Sequential Updating:
A Behavioral Model of Belief Change

Korhan Koçak*

First draft: March 14, 2016
Current draft: April 7, 2018

*Princeton University. Email: kkocak@princeton.edu. The author is grateful to Roland Bénabou, Romain Ferrali, Germán Gieczewski, Matias Iaryczower, Amanda Kennard, Asya Magazinnik, Sepehr Shahshahani, Sondre Solstad, Dustin Tingley, and conference and workshop participants at Princeton University, Midwest Political Science Association Conference, European Political Science Association Conference, and Empirical Implications of Theoretical Models Summer Institute.

Abstract

The way information is processed by decision makers is frequently at odds with the standard theoretical framework. On many contentious topics, decision makers' beliefs often fail to converge despite exposure to a high number of signals, or still continue to polarize. Individuals have been shown to exhibit the backfire effect, holding onto their initial beliefs even more strongly when presented with evidence to the contrary. The first piece of information in a series often acts as an anchor, and thus has an excessively high weight in posteriors. Decision makers' beliefs continue to shift after repeated signals, although no learning occurs. Based on recent findings on working memory, this paper proposes a cognitively less demanding yet ultimately incorrect way to update beliefs that accounts for these cognitive biases. If a signal is informative on multiple dimensions, and if individuals can hold only a limited number of items in working memory, then the order with which beliefs are updated across different dimensions leads to different posteriors. The unifying model advanced by this paper shows that polarization is possible even under common priors and common signals, and suggests new avenues of research for a better understanding of how information is processed by individuals, as well as how persuasion occurs.

1 INTRODUCTION

Picture an audience watching a presidential debate before an election. Audience members are uncertain about the merits of the candidates, as well as that of their positions. The candidates wish to reduce these uncertainties in their favor: to show that they are the more qualified candidate because they have superior policies planned, and that their policies are superior because they are the more qualified candidate. If the audience members had common priors on both dimensions and were perfectly Bayesian, their beliefs would move together. The order of speakers would not matter, as information that arrived earlier would have the same weight on posteriors as if it arrived later instead. Candidates would never repeat their messages; given the limited time they have they would like to cover as many topics as possible.

Yet casual observation shows us time and again that standard theory fails us in these presumptions, especially in the realm of politics, where feedback often do not arrive fast enough for agents to correct their cognitive biases.[1] Indeed, few would be surprised to see the audience leave the room more polarized after having watched the debate; that the

[1]Consider for example beliefs regarding climate change: a Pew study found that only 50% of Americans believed in human induced climate change in contrast to 87% of American scientists. This, in part, arises from the fact that the more obvious adverse effects of climate change will only materialize far into the future.

first answer to each question swayed some in the audience more than could be accounted for; that candidates did spend precious air time repeating the same messages over and over again and that this had an impact on the audience's beliefs. The evidence that decision makers frequently exhibit behavior of this sort is well established. The question remains as to what are the mechanisms behind them.

The present paper proposes that this stems from complexity of the problem of belief change under multiple dimensions of uncertainty, and the limits of working memory. In particular, I propose that some agents are prone to making an error in the order with which items are discarded from the working memory and placed in long-term memory. Recent research on memory provides evidence that working memory updating involves an item-specific active removal process of outdated information (Ecker et al., 2014). Because of the capacity limits of working memory, discarding irrelevant information is crucial for cognition to function properly. However, little is known in political science and economics about how the mind determines when an item is no longer relevant and how this discarding of information affects decision making. This paper explores a problem of this class. Specifically, I investigate what happens when there are multiple items to be updated, and an item is discarded before its full information value is extracted.

Consider the following problem. There is a decision maker who has to take an action after observing a signal from an information source whose type is unknown. There is a payoff relevant state of the world which is not observed by the decision maker. When the source sends a signal about the state, this signal affects the decision maker's beliefs about the state of the world and the type of the signal source. For the decision maker to update her beliefs, three items need to be in her working memory: two prior beliefs and the signal.[2] If she had no cognitive limits, she could update both parameters simultaneously, and then discard the priors after forming the posteriors (Figure 1(a)). However, given limits on working memory, the decision maker needs to update the parameters sequentially instead, updating one parameter before the other.[3] Whichever order she does her updating, after she updates the first belief she now needs to hold four items in her working memory: two priors, one posterior and the signal. If the limit of her working memory only allows for three items to be held simultaneously, then the correct way to update would be to store the updated belief in long-term memory, and then use the two priors and the signal to update the other belief. Only after updating the other item can

[2]Roughly put, one can think of long-term memory as an infinitely big table, and working memory as the agent's hands. Theoretically, the table can hold an infinite number of items, although some items are easier to access than others. In contrast, the hands can hold only a limited number of items. If the agent has her hands full (i.e. her working memory is at its capacity) and she needs to access another item, she must put down an item in her hands on the table (i.e. store in long-term memory to access later).

[3]Experimental evidence from the psychology literature suggests that people do cognitive tasks serially, and not in parallel. See for example (Oberauer, 2002).

Figure 1: Description of the Updating Process

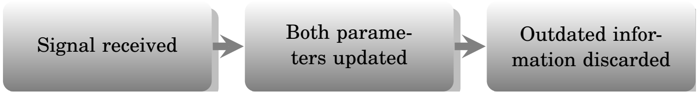

(a) Rational agent without cognitive limits

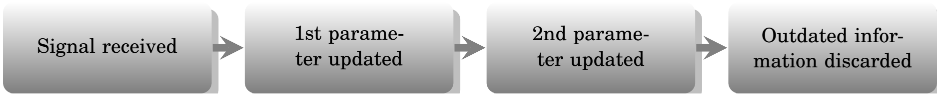

(b) Rational agent with cognitive limits

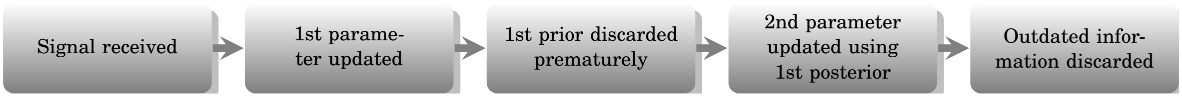

(c) Behavioral agent

she finally discard the priors that are now outdated (Figure 1(b)).

Now consider an error in this process. Suppose that after updating the first item, the decision maker discards the prior of the now updated parameter, presuming that it is no longer useful since it is now outdated. When she then wants to update the other parameter, she no longer has the prior, but only the posterior. Hence she must use the only tool at her disposal, the posterior of the first item, to form her posterior on the other item. Depending on the order with which parameters are updated, this error leads to a distinct set of posteriors (Figure 1(c)).

This proposed bias has a number of potential implications that are corroborated in the experimental literature. The model shows how agents with the same priors but different orders of updating end up with different posteriors after observing the same signal, a phenomenon known as *attitude polarization*. For some subset of parameters, this leads to the *backfire effect* where some agents update in the 'wrong' direction. When there are competing information sources, this model shows *anchoring* is possible, where earlier signals have more weight in the posteriors than later signals. Finally, I show that agents who update sequentially are susceptible to persuasion by repetition or to *truth effect*, due to the inconsistency in their beliefs.

To get an intuition on how sequential updating leads to different posteriors with common priors and signals, consider a brief example borrowed from Morris (2001). Suppose there is uncertainty about whether affirmative action would be a beneficial policy. There is a social scientist who has information about the welfare implications of the policy, independent of that of the policy maker. Specifically, the social scientist may be the good type who has the correct information, or the bad type who has wrong information. The policy maker believes that affirmative action would be beneficial with probability 2/3. She also believes that the social scientist is good with probability 3/4. Suppose the social scientist says that affirmative action would not be beneficial. If the policy maker makes no error in updating, her posteriors about the two parameters would be as follows:

$$\Pr(\text{scientist is good} \mid \text{scientist against policy}) = \frac{1/3(3/4)}{1/3(3/4)+2/3(1/4)} = \frac{3}{5} \tag{1}$$

$$\Pr(\text{policy beneficial} \mid \text{scientist against policy}) = \frac{2/3(1/4)}{2/3(1/4)+1/3(3/4)} = \frac{2}{5} \tag{2}$$

Consider now two policy makers who update sequentially as described above. First, consider a 'naive' agent who updates her belief on the probability of policy being beneficial first, then discards the prior before forming a posterior on the probability of the social scientist being the good kind. Since her belief on the policy is formed using both priors, by Equation (2) we know that her posterior on the probability of this being a beneficial policy is 2/5, same as the policy maker that makes no error in updating. But because she erroneously thinks that her prior on the probability of policy being good no longer matters, she discards it, and is only left with her posterior when she goes on to update her belief on the type of the social scientist. Hence, her posterior on the type of the social scientist is:

$$\overset{n}{\Pr}(\text{scientist is good} \mid \text{scientist against policy}) = \frac{3/5(3/4)}{3/5(3/4)+2/5(1/4)} = \frac{9}{11} > \frac{3}{5}$$

Note that the posterior of the naive policy maker on the type of the social scientist is greater than that of a rational policy maker. It is in fact greater than the prior, meaning that the naive policy maker and the rational policy maker update in opposite directions upon observing the same signal, though having started with identical priors.

Finally consider a 'dogmatic' policy maker who updates the belief about the type of the social scientist first, as in Equation (1), and then discards the prior. She uses her posterior on the probability of the social scientist being the good type when evaluating the probability of policy being good:

$$\overset{s}{\Pr}(\text{policy beneficial} \mid \text{scientist against policy}) = \frac{2/3(2/5)}{2/3(2/5)+1/3(3/5)} = \frac{4}{7} > \frac{2}{5}$$

The posterior of the dogmatic policy maker on the state of the world is slanted towards her prior relative to the Bayesian benchmark. Thus, by one simple yet crucial tweak we have three policy makers who all start with the same priors, and observe the very same information, and reach different conclusions.

In independent and parallel work, Cheng and Hsiaw (2016) examine a bias very similar to the one presented in this paper. In their model, agents "pre-screen" any signal they receive by updating their beliefs on the type of the expert (signal source) first, and then use that posterior to update their belief on the state.[4] The two papers, however, take different routes in dealing with the question of sequences of signals. In Cheng and Hsiaw (2016), decision makers reevaluate all past signals under the light of the most current posterior belief about the quality of the source, which allows the authors to extend their results to arbitrary sequences of signals. In contrast, in the present paper I do not require decision makers to revisit past observations, but I restrict focus to certain classes of sequences. The two papers have comparable results on divergence of beliefs and anchoring; I also examine how sequential updating may lead to persuasion by repetition and backfire effect, as well as examining the implications of this bias in a political agency model.

In terms of mechanics, this paper can be considered a part of the literature on two-sided updating, as there is uncertainty regarding both the state of the world, and quality or the identity of the information source. In a 1985 paper, Sobel examines a repeated cheap talk game where the sender's preferences are either congruent with or exactly opposite of the receiver's preferences (Sobel, 1985). The sender observes the state of the world, and sends a signal to the receiver, who takes an action relevant to both players' payoffs. A "friendly" sender always sends the correct signal, so upon observing an inaccurate signal the receiver becomes certain the sender is an enemy, and completely ignores any future signals henceforth. This gives the sender only one chance to fool the receiver. Benabou and Laroque (1992) extend this to a setting where the sender's information is imperfect, therefore a wrong signal could be attributed to an honest mistake from a good sender as well. In their paper, the reputation of the sender fluctuates, improving with each correct message, and decreasing with each wrong message. Morris (2001) examines a political agency model in which an adviser with private information about the state of the world, whose preferences are either congruent with a decision maker or biased towards a known policy, gives policy recommendations to the decision maker. If the good adviser's private information is the state the biased adviser doesn't like, she reveals her information truthfully. If, however, the state is the one the biased adviser does like, the good adviser faces a trade-off between telling the truth and losing some reputation versus lying and improving her reputation at the cost of leading the decision maker to-

[4]Pre-screeners in their terminology corresponds to the dogmatic agents in mine. They do not study a counterpart of the 'naive' agents in the present paper, as they consider only one sequence of updating: source first and state second.

wards the wrong policy. Morris shows that when the reputation concerns are sufficiently salient, there is no informative equilibrium. In all these above mentioned papers the focus is the optimal strategy of the sender, who cares about the action taken by a Bayesian receiver as well as her own reputation. The present paper, on the other hand, studies the reputation of a signal source who is not strategic, and where signals are produced mechanically. My focus is the exposition of the implications of sequential updating, and therefore I abstract away from strategic considerations of the sender.

There is a growing literature on why people's beliefs polarize despite the abundance of information. Ortoleva and Snowberg (2015) attribute this discrepancy between standard economic theory and actual empirical observations to what they call "correlational neglect"; the tendency of people to underestimate the degree of correlation in the information they receive. Similarly, Eyster and Rabin (2010) write about herding in which people naively believe that others' actions are based solely on their own private information, therefore effectively over-weighing the first movers' beliefs. DeMarzo et al. (2003) call this "persuasion bias," and examine how boundedly rational agents form opinions in a network. They show that when agents fail to adjust for repetitions of information, an individual's influence on the consensus beliefs depends not only on accuracy, but also on her place in the network. In all these papers the bias stems from decision makers underestimating the correlation between the signals they receive from different sources. In contrast, in the present model polarization is possible even when there is a single signal source. In that sense, this paper extends to broader settings where the bulk of information comes from only one source, such as advertising or state propaganda.

Baliga et al. (2013) show that polarization is possible when agents are ambiguity averse and dynamically consistent. This happens when agents stick with their prior plans that hedge against ambiguity, even when that doesn't maximize their expected utility conditional on the signals they have received. Therefore, decisions of agents with different priors may diverge after observing the same signal. Acemoglu et al. (2016) instead focus on a setting where different individuals who start with different priors are uncertain about how to interpret the signals they receive. Andreoni and Mylovanov (2012) show polarization is possible by assuming that decision makers receive differential private signals along with common signals. Another relevant paper by Benoit and Dubra (2014) shows that polarization with rational agents with different priors is possible. In their model there is a proposition about an issue that might be true or false, and an ancillary matter regarding which of the attributes of the issue are more important for the proposition. Agents' beliefs about the likelihood of the proposition being true depend on which ancillary matter they believe to be more relevant with respect to the proposition. In this setting, a signal that is good news about one matter and bad news about another may polarize agents' posteriors, in the sense that both agents who believe the proposition to be true and those who do not become more confident of their beliefs after observing such a signal. The present model diverges from these papers as polarization is possible with common priors and common signals, if the agents differ in their order of updating across different parameters.

Rabin and Schrag (1999) explore a setting where agents with some positive probability misread signals that go against their priors, and show that anchoring bias (or confirmation bias) is possible. Similarly, Fryer et al. (2016) study a setting where signals may be ambiguous, but cannot be recorded as such. The agents who observe ambiguous signals interpret them in accordance with their priors, thus "double updating" and exhibiting confirmation bias.

These models provide various explanations for the truth effect, attitude polarization or anchoring, whereas the updating rule employed in the current paper leads to all three cognitive biases.

2 Model

2.1 Benchmark

There is a decision maker, two states $S = \{A, B\}$, and two actions $s = \{a, b\}$. The utility of the decision maker is given by

$$u(s,S) = \begin{cases} 1 \text{ if } s = a \text{ and } S = A \text{ or } s = b \text{ and } S = B \\ 0 \text{ if } s = b \text{ and } S = A \text{ or } s = a \text{ and } S = B \end{cases} \tag{3}$$

The decision maker does not observe the actual state of the world, but knows that the prior probability that the state is A is $x_0 \in (0, 1/2]$, so that a priori state B is more likely. The decision maker also receives a signal that comes from a signal source. The type of the signal source γ can be either high, H in which case it publishes the correct signal $\sigma = S$ with probability h, or it can be low L, so that it publishes $\sigma = S$ with probability $l < h$. The prior probability that the source is the high type is given by $y_0 \in (0, 1)$ and type of the source is drawn independently from the state. Without loss of generality, I assume that $y_0 h + (1 - y_0)l \geq 1/2$, so that signals are more likely to be correct than not - otherwise the decision maker could revert the signals. I denote the signal published by the source with the corresponding Greek letter, i.e. α for A, and β for B. When there are multiple periods, I follow the convention of denoting by σ_t the message in period t, and I use σ^t to refer to the history of all messages up to and including period t.[5]

There are three types of agents. First is the "rational" type; the type that updates both parameters correctly. The second type, which I call "naive", updates the posterior about the state of the world first, and uses that posterior to update her belief about the type of the signal source. This is the type that primarily pays attention to the state of the world and secondarily to sources of signals. Finally, I call the type that updates the posterior about the type of the source first, and uses that posterior to update beliefs about the state of the world the "dogmatic" type. This is the type that is primarily concerned about what a signal says about its source.

For ease of exposure going forward, let us define the following. Let $d_t(\sigma^t) = y_t(\sigma^t)h + (1-y_t(\sigma^t))l$ denote the probability of observing the correct signal, $r_t(\sigma^t) = x_t(\sigma^t)(2h-1) + 1-h$ the probability of observing the signal α when the sender is type H, and $q_t(\sigma^t) = x_t(\sigma^t)(2l-1)+1-l$ the probability of observing the signal α when the sender is type L. Furthermore, let λ denote the likelihood ratio of state being A versus B, and similarly let ω denote the likelihood ratio of type of the source being high versus low. Finally, let δ denote the likelihood ratio of observing the correct signal versus not. :

$$\lambda_t(\sigma^t) = \frac{x_t(\sigma^t)}{1-x_t(\sigma^t)} \quad \omega_t(\sigma^t) = \frac{y_t(\sigma^t)}{1-y_t(\sigma^t)} \quad \delta_t(\sigma^t) = \frac{d_t(\sigma^t)}{1-d_t(\sigma^t)} \tag{4}$$

Suppose the first signal is α, then λ_1 and ω_1 become:

$$\lambda_1(\alpha) = \lambda_0\delta_0 = \lambda_0\left(\frac{y_0(h-l)+l}{1-y_0(h-l)-l}\right) = \frac{x_0(y_0(h-l)+l)}{(1-x_0)(1-y_0(h-l)-l)}$$

$$\omega_1(\alpha) = \omega_0\frac{r_0}{q_0} = \omega_0\left(\frac{x_0(2h-1)+1-h}{x_0(2l-1)+1-l}\right) = \frac{y_0(x_0(2h-1)+1-h)}{(1-y_0)(x_0(2l-1)+1-l)}$$

Similarly, when the signal observed is β, we have that $\lambda_1(\beta) = \frac{\lambda_0}{\delta_0}$, and $\omega_1(\beta) = \omega_0\frac{1-r_0}{1-q_0}$.

When a biased decision maker is faced with the same problem, however, his posteriors evolve differently. If the parameter first updated was x, y is updated to y_1^n instead of y_1, and I write ω_1^n instead of ω_1 where the superscript n refers to the type of the naive agent:

$$\omega_1^n(\alpha) = \omega_0\frac{r_1(\alpha)}{q_1(\alpha)} = \frac{y_0(x_1(\alpha)(2h-1)+1-h)}{(1-y_0)(x_1(\alpha)(2l-1)+1-l)}$$

[5] Note that I do not restrict l to be greater than half, that is, the low quality signal source may send the inaccurate signal more often than it does the correct signal. Substantively, this would be the case where there is a source of information that benefits from increased polarization, This is to preserve symmetry in the model for sake of exposition, and all results follow through in a slightly modified model where the low quality source is biased towards one signal.

Conversely, if the parameter first updated was y, x is updated to x_1^s instead of x_1, and so:

$$\lambda_1^s(\alpha) = \lambda_0 \delta_1(\alpha) = \frac{x_0(y_1(\alpha)(h-l)+l)}{(1-x_0)(1-y_1(\alpha)(h-l)-l)}$$

where the superscript s refers to the type of the dogmatic agent.[6]

With this setup, we are now ready to state the first result of the paper.

Proposition 1. *When a priori signals are more (less) likely to be correct than not, naive agents have higher (lower) posteriors about the type of the signal source than rational agents with the same information. Dogmatic agents' posteriors slant towards the state they initially believe is more likely, relative to rational agents with the same information.*

Proof. *All proofs are in the Appendix.* □

To see get an intuition to why these differences in posteriors matter, suppose for ease of exposition that $h = 1$ and $l = 1/2$, that is, the high type of source is always accurate, whereas the low type produces pure gibberish. Since the decision makers want to match their action to the state, they take action a if and only if their posterior on x, denoted by x_1, is greater than a half (or if $\lambda_1 > 1$). Suppose that the source publishes the signal $\sigma = A$. Then, the values of x_0 and y_0 that leave the decision makers of different types indifferent are stated in Corollary 1 and presented in Figure 2.

Corollary 1. *Suppose $h = 1$, $l = 1/2$, and $\sigma = \alpha$. Then, naive and rational agents take action a whenever $2x_0 + y_0 > 1$, and action b whenever $2x_0 + y_0 < 1$. dogmatic agents take action a if $y_0 > \frac{1-2x_0}{4x_0^2-2x_0+1}$, and action b if $y_0 < \frac{1-2x_0}{4x_0^2-2x_0+1}$. Thus, agents whose priors lie in the interval $y_0 \in \left(1-2x_0, \frac{1-2x_0}{4x_0^2-2x_0+1}\right)$ disagree on the optimal action after observing the same signal α, with dogmatic agents less likely to follow the advice of the source.*

As can be seen, there is a strictly larger set of parameters in which the dogmatic type agents pick action b, even when it is not, strictly speaking, rational to do so. Put another way, when all agents have common priors and those priors are between the blue line and the red line on Figure 2, all rational and naive agents take the action a, and all dogmatic agents take the action b after the source publishes the signal $\sigma = A$.

Divergence in beliefs may happen in two different ways: either the biased agents update in the same direction as the rational agents but by different magnitudes, or they

[6]Implicit in this setup is the assumption that biased decision makers may hold conflicting beliefs as long as they are unaware of the conflict. Henceforth I maintain the assumption that decision makers' beliefs remain inconsistent until this is revealed via external stimuli, e.g. reminders about past signals.

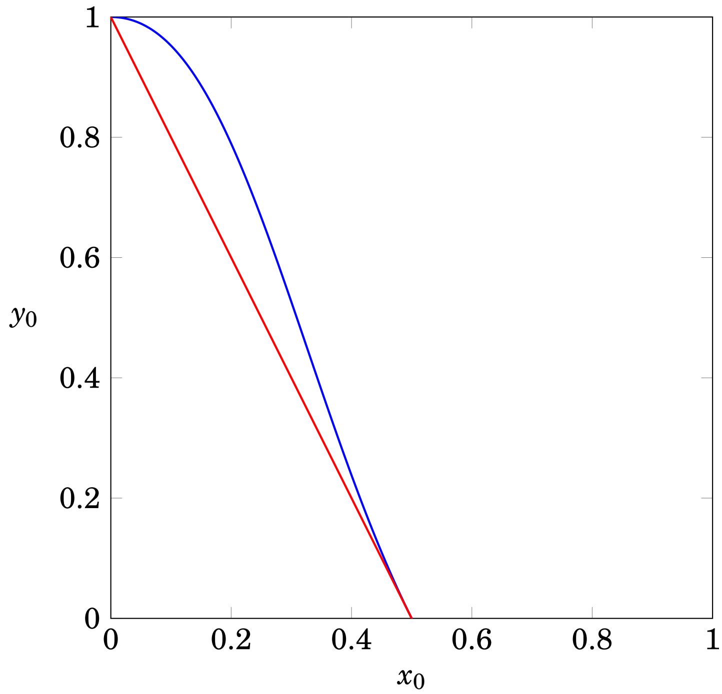

Figure 2: Set of priors that leave agents indifferent between the two actions after one signal of A. Red line refers to rational and naive agents who are indifferent between two actions, and blue line to dogmatic agents.

update in opposite directions. The latter pertains to a stronger form of polarization, where biased agents update in the 'wrong' direction, doubling down on their priors in the face of feasible information that contradicts it.

Lemma 1. *Suppose $h > 1/2 > l$ and the source sends the signal α. Then, whether the naive agent updates in favor or against the source depends on how her prior belief on the state compares to the prior probability that the message is correct. In particular, she trusts the source more after the signal if and only if her prior belief that the state is A and the prior probability that the message is correct are sufficiently high:*

$$\omega_1^n(\alpha) > \omega_0 \text{ if and only if } \lambda_0\delta_0 > 1$$

Conversely, whether the dogmatic agent updates towards or away from A depends on how her prior belief on the quality of the signal compares to the prior likelihood that a signal α comes from the high type. In particular, she updates her belief towards A after the signal if and only if her prior belief that the source is high quality and the prior likelihood that the signal α comes from the high type are sufficiently high:

$$\lambda_1^s(\alpha) > \lambda_0 \text{ if and only if } \omega_0\frac{r_0}{q_0} > \frac{1/2 - l}{h - 1/2}$$

The next result pertains to a stronger form of polarization, where beliefs of agents with common priors move in opposite directions upon observing the same signal.

Proposition 2. *Suppose that $l < 1/2 < h$. That is, suppose there exists a sender that's more often accurate than half the time, and one that is less. Then, for the naive type:*

$$\frac{1-y_0(h-l)-l}{y_0(h-l)+l} < \lambda_0 < 1 \Longleftrightarrow \omega_1(\alpha) < \omega_0 < \omega_1^n(\alpha)$$

and:

$$\frac{1-y_0(h-l)-l}{y_0(h-l)+l} > \lambda_0 > 1 \Longleftrightarrow \omega_1(\alpha) > \omega_0 > \omega_1^n(\alpha)$$

And for the dogmatic type:

$$\frac{1/2-l}{h-1/2} < \omega_0 < \frac{q_0(1/2-l)}{r_0(h-1/2)} \Longleftrightarrow \lambda_1^s(\alpha) < \lambda_0 < \lambda_1(\alpha)$$

Conversely:

$$\frac{1/2-l}{h-1/2} > \omega_0 > \frac{q_0(1/2-l)}{r_0(h-1/2)} \Longleftrightarrow \lambda_1^s(\alpha) > \lambda_0 > \lambda_1(\alpha)$$

This result states the necessary and sufficient conditions for the parameters of the model that lead to posteriors of biased agents diverging from those of rational agents with same priors after one signal, in the sense that while rational agents update their beliefs towards one direction, biased agents update in the opposite direction. naive agents update their beliefs about the information source in the opposite direction of rational and dogmatic agents when (a) a signal that is sufficiently likely to be correct sends a signal against the prior on the state, or (b) a signal that is sufficently likely to be wrong sends a signal in line with the prior on the state. In the former case, naive agents update up on the type of the source whereas the others update down, and vice versa in the latter case.

As for the divergence of the posterior of the dogmatic type, this would for example happen when agents initially believe that B is more likely than A ($x_0 < 1/2$), and the signal is more likely to be correct than not (i.e. $h(y_0) + l(1-y_0) > 1/2$). Rational and naive agents would believe that A is more likely than what their priors suggested after observing the signal α, because a signal that's likely to be correct said so. Dogmatic agents would, on the other hand, first update on the type of the source, and because they believe B is more likely than A, their posterior on the quality of the signal source would fall when they observe it transmitted α. In the range described in the proposition, this leads to their posterior belief about the signal being that it is more likely to be incorrect than not (i.e. $h(y_1(\alpha)) + l(1-y_1(\alpha)) < 1/2$). When they update their belief on the state of the world, they use this posterior rather than their prior, which leads them to conclude that state is even less likely to be A than they initially believed. This is known as the backfire effect: when decision makers react to evidence against their priors by holding onto their initial beliefs even more strongly.

The opposite case where the dogmatic agents update towards a signal that is likely false but is in line with priors also holds. That is, suppose the prior is in favor of A, but the signal is more likely to be wrong than right. Then, when the source transmits signal α, rational and naive types have posteriors on A lower than their priors. The dogmatic agents, however, update on the type of the source first. And because the signal agrees with their prior, they believe that the signal is more likely to be correct, and therefore update towards A their belief on the state of the world.

This provides a new intuition as to why some voters in various contexts seem very eager to give credence to news sources they had likely never heard of before, and at the same time scorn reputable media outlets and journalists. If these voters exhibit the dogmatic type of bias described in this paper, they likely update down on outlets which transmitted signals that go against their priors, and update up on those that agree; and use these posterior beliefs to update on the state of the world, both of which serve to strengthen their priors.

This divergence may also occur when two dogmatic agents agree on the type of the source but have different priors on the state of the world. Denoting the priors of one agent with hats, and the other with tildes, if their priors on the state of the world satisfy:

$$\frac{\hat{q}_0}{\hat{r}_0} < \left(\frac{h - 1/2}{1/2 - l}\right)\omega_0 < \frac{\tilde{q}_0}{\tilde{r}_0}$$

then these agents update in opposite directions, although they have identical priors about how likely the message they observe is to be correct.

2.2 Order Effects

Suppose now there are two sources of information and two periods. In the first period sender 1 sends a signal, and in the second period sender 2 sends a signal. For ease of exposition, I assume that the sources are either perfectly accurate ($h = 1$), or perfectly inaccurate ($l = 0$). Thus, the sources are either the honest type who publishes the true signal, or the sinister type that publishes the wrong signal. The types of the sources are drawn independently. Let x_0 as before denote the prior on the state, and let y_0 denote the probability that the first signal sender is the honest type, and let z_0 denote the probability that the second sender is the honest type. Let $\omega_t = \frac{y_t}{1-y_t}$, and $\mu_t = \frac{z_t}{1-z_t}$.

The cases in which two senders send the same signal are not of particular interest; if they are jointly precise with probability greater than a half the posterior on the state will be closer to the state they published regardless of the type of the decision maker, and otherwise the posterior will move away from the signal. I will therefore restrict attention to the case when the two sources publish different signals. Clearly, in this case exactly one of the signal senders must be the honest type. Suppose without loss of generality the first source sends $\sigma_y = \alpha$, and the second source sends $\sigma_z = \beta$. Note that there are two potential scenarios, either the state is A, the first source is honest and second sinister; or the state is B, in which case the first source is sinister and second honest.

Suppose further that these signals come in sequence, with first source sending their message first followed by the second: $\sigma^t = (\sigma_{y,1} = \alpha, \sigma_{z,2} = \beta)$, where the first item in the subscript refers to the identity of the source and the second item refers to the period in which the signal is observed. After both signals, the decision maker makes the decision. With a rational decision maker the order in which the signals arrive won't matter. The posterior LR in either case will be:

$$\lambda_2(\sigma^t) = \omega_2(\sigma^t) = \frac{1}{\mu_2(\sigma^t)} = \frac{\lambda_0\omega_0}{\mu_0} \tag{5}$$

Consider now the naive type, who updates his belief about the state of the world first:

$$\lambda_1(\sigma_{y,1} = \alpha) = \lambda_0\omega_0$$

and then updates his belief about the first source:

$$\omega_1^n(\sigma_{y,1} = \alpha) = \lambda_1(\sigma_{y,1} = \alpha)\omega_0 = \lambda_0\omega_0^2 \tag{6}$$

When the second signal is received, again it is the state of the world that is first updated:

$$\lambda_2^n(\sigma^t) = \frac{\lambda_1(\sigma_{y,1} = \alpha)}{\mu_0} = \frac{\lambda_0\omega_0}{\mu_0} \tag{7}$$

and finally,

$$\frac{1}{\mu_2^n(\sigma^t)} = \frac{\lambda_2^n(\sigma^t)}{\mu_0} = \frac{\lambda_0\omega_0}{\mu_0^2} \tag{8}$$

The assumption that the biased decision makers may hold inconsistent beliefs implies that until the inconsistency is revealed, the naive agent holds the beliefs regarding the two sources and the state described in Equations (6), (7) and (8), although consistency would require these three beliefs be equal. According to these equations, the naive agent's posteriors on the sources are different from those of a Bayesian agent, although the direction and the magnitude of the bias depends on the priors on the sources. The naive agent however holds correct posteriors regarding the state of the world. This is in line with the findings of Section 2.1; unless the naive agent observes multiple signals from one source, his inaccurate beliefs about the sources do not influence his beliefs regarding the state of the world, and he has the same beliefs regarding the state of the world as a Bayesian decision maker would.

This however is not true for the dogmatic type, who updates his belief about the type of the signal source first.

$$\omega_1(\sigma_{y,1}=\alpha)=\lambda_0\omega_0 \tag{9}$$

Given those posteriors about the first source, the posterior belief about state of the world prior to observing the latter signal is:

$$\lambda_1^s(\sigma_{y,1}=\alpha)=\lambda_0\omega_1(\sigma_{y,1}=\alpha)=\lambda_0^2\omega_0$$

The second source's signal will then be evaluated under the light of this new posterior:

$$\frac{1}{\mu_2^s(\sigma^t)}=\frac{\lambda_1^s(\sigma_{y,1}=\alpha)}{\mu_0}=\frac{\lambda_0^2\omega_0}{\mu_0} \tag{10}$$

And finally, the posterior about the state of the world:

$$\lambda_2^s(\sigma^t)=\frac{\lambda_1^s(\sigma_{y,1}=\alpha)}{\mu_2^s(\sigma^t)}=\left(\lambda_0^2\omega_0\right)\frac{\lambda_0^2\omega_0}{\mu_0}=\frac{\lambda_0^4\omega_0^2}{\mu_0} \tag{11}$$

From Equations (9), (10) and (11) one can see that the dogmatic decision maker has correct beliefs about the first source, has slanted up (down) posteriors about second source if second source agreed (disagreed) with his prior, and overweights both his prior and the first source's message relative to a Bayesian decision maker in his posterior belief about the state of the world.

Alternatively, if the second source sent a signal first, the posterior on the state of the world would be:

$$\lambda_2^s(\sigma_{z,1}=\beta,\sigma_{y,2}=\alpha)=\frac{\lambda_0^4\omega_0}{\mu_0^2}$$

Therefore, the source that sends the first signal has a greater influence on the dogmatic type's posterior on the state of the world.

Therefore, the order in which the signals are observed matters in determining the final posteriors. There is an abundance of empirical evidence showing that people exhibit a primacy effect, over-weighing earlier information relative to latter information. This model provides a novel explanation as to why people may exhibit this bias.

2.3 Repetition and Persuasion

One empirical observation that is quite a curiosity for economic theory is the apparent utility of repetition of information. According to theory, learning of a piece of information must be sufficient for an agent to update his beliefs, and repeated exposure to same information should have no influence on beliefs, as all that can be incorporated is already incorporated to posteriors.[7] Yet various strategists whose job description is persuading others, from politicians to marketers, prescribe repetition for successful conversion.[8]

While there are models that explain the impact of repeated signals from different sources when the decision maker fails to adjust for their correlation, scant attention has been paid to the case when repeated signals come from the same source.[9] Unless the agent's memory is so limited that they forget the sources from which they received past signals, these models fail to account for the large body of evidence of persuasion by repetition from a single source. This is the case for advertisement, where consumers know it's the company that make the advertisement; presidential debates where the audience sees the same candidate repeat the same arguments; and in any authoritarian regime with state control over media sources, which repeat the same information over and over again.[10]

Consider the following: there is again a state of the world which is unknown, and a single signal source whose type is also unknown, with $h = 1$ and $l = 0$ as before. Suppose without loss of generality the source says that the state is A in each period: $\sigma^t = (\alpha, \alpha, \ldots, \alpha)$. The posterior likelihood ratio of a rational agent will then be:

$$\lambda_1(\sigma_1 = \alpha) = \lambda_0 \omega_0$$

where subscripts denote the number of times the signal is received. In fact, since the rational agent incorporates all the information in the signal after the first iteration, we have for all t:

$$\lambda_t(\sigma^t) = \lambda_0 \omega_0$$

Suppose now instead this agent is naive. Upon receiving the signal from the source, she first updates her belief about the state of the world:

$$\lambda_1^n(\sigma^1) = \lambda_0 \omega_0$$

[7] See DeMarzo et al. (2003) for an account.

[8] See for example Lewandowsky et al. (2012) and Fernandes (2013).

[9] See for example (DeMarzo et al., 2003; Ortoleva and Snowberg, 2015).

[10] George W. Bush infamously said "See in my line of work you got to keep repeating things over and over and over again for the truth to sink in, to kind of catapult the propaganda."

And then her beliefs about the type of the signal source:

$$\omega_1^n(\sigma^1) = \lambda_1^n(\sigma^1)\omega_0 = \lambda_0\omega_0^2$$

where superscripts denote the type of the agent.

Note that the beliefs of the agent on the two parameters are inconsistent: given the structure of the game he cannot rationally believe that the probability of the sender being honest and the probability that the state is what the sender signaled could be different. However, our agent is biased in such a way that he takes the signal to be relevant for one parameter at a time instead of being jointly relevant for both parameters. If he were to note the inconsistency at some point, he could correct it by fixing either parameter and achieve consistency. If the inconsistency in his beliefs eludes the agent, however, and he fails to 'fix' one of his beliefs to conform with the other by himself, upon being reminded of the same signal at a later date he will fix it then. That is, once he observes the signal again, being a naive agent he will first update his belief on the state of the world by adjusting his belief on that parameter to conform with his posterior on the type of the signal source. In other words, when the naive agent notes the inconsistency between two beliefs, one correct and the other not, because of the way his bias works he "fixes" this discrepancy by changing his correct belief to be in line with his wrong belief. Thus:

$$\lambda_2^n(\sigma^2) = \omega_1^n(\sigma^1) = \lambda_0\omega_0^2$$

Since according to the bias in this model a biased agent takes parameters into consideration sequentially, he will, after updating one parameter upon receiving a signal, update the other parameter immediately after, basing beliefs on the already updated parameters. This is the same error the naive agent committed after the first iteration of the signal, namely, the posterior on the source is formed based on the posterior on the state of the world and the prior on the source, only this time the posterior on the state of the world is "wrong". Thus, the naive agent then reevaluates his beliefs about the type of the signal source as follows:

$$\omega_2^n(\sigma^2) = \lambda_2^n(\sigma^2)\omega_0 = \lambda_0\omega_0^3$$

Note that $\lambda_2^n(\sigma^2)$ and $\omega_2^n(\sigma^2)$ are also inconsistent. This inconsistency, if revealed to the biased agent by another iteration of the same signal, is resolved in the same way as before. Going this way, after t iterations of the same signal from the same signal source, the naive agent has the following beliefs:

$$\lambda_t^n(\sigma^t) = \lambda_0\omega_0^t \qquad \text{and} \qquad \omega_t^n(\sigma^t) = \lambda_0\omega_0^{t+1}$$

As the number of signals go to infinity, the posterior of the naive agent about both parameters converge to 0 for all $y_0 < 1/2$ and converge to 1 for all $y_0 > 1/2$. Thus, as the number of signals increase the weight of the prior on the state of world diminish, and the prior on the source is sufficient to determine asymptotic beliefs regarding both parameters. Note that this happens despite lack of learning. The bias stems from the inability of the naive agent to resolve inconsistencies in his belief without committing the same errors that lead to the inconsistencies in the first place.

Consider now the dogmatic agent. He first updates his belief about the type of the signal sender:

$$\omega_1(\sigma^1) = \lambda_0\omega_0$$

So his posterior about the state of the world are:

$$\lambda_1^s(\sigma^1) = \lambda_0^2\omega_0$$

By a symmetric argument with the naive agent, we have:

$$\lambda_t^n(\sigma^t) = \lambda_0\omega_0^{t+1} \qquad \text{and} \qquad \omega_t^n(\sigma^t) = \lambda_0\omega_0^t$$

Thus, as signals are repeated the weight of the initial trust in the signal source goes to zero for the dogmatic agent, and her beliefs about both parameters go to 1 if $x_0 > 1/2$, and go to 0 if $x_0 < 1/2$.

To fix ideas with one last example, suppose there are two agents, one naive and one dogmatic with identical priors such that: $x_0 \in (1/2, 1)$, $y_0 \in (0, 1/2)$. After receiving sufficiently many signals, the naive agent's posteriors can be arbitrarily close to $x_\infty^n = y_\infty^n = 0$, whereas the dogmatic agent's posteriors would be $x_\infty^s = y_\infty^s = 1$. Thus, although agents start with the same priors and see the same signal repeated infinitely many times, they end up with diametrically opposite beliefs about both parameters of the world; one with arbitrarily high trust in the signal source and what it says to be the state of the world, other with no trust whatsoever, and believing the state of the world to be the exact opposite of what the source says it is. This is in line with empirical findings that show that repeated communications about a subject can lead to greater attitude polarization on that subject. (Judd and Brauer, 1995)

It is worth noting here that the results in this section are driven by the assumption that the source is either the good type which reports truthfully with probability one or the bad type which reports truthfully with probability zero. If this wasn't the case, having different beliefs on the state of the world and the quality of the sender would not be inconsistent, and the agent would update as before. This is because in such a case the decision maker is uncertain whether an unlikely signal is due to the type of the sender or an unlikely draw.[11]

2.4 LONG-RUN

Now consider the case when the game is played infinitely many periods. In this section I assume that after receiving signals and updating their beliefs, agents forget the signals themselves. Therefore, it is as if each period they start the game from scratch, only with their posteriors having replaced their priors.[12] In other words, decision makers exhibit *source amnesia*, the tendency for individuals to recall the information itself, but not its source. As Schacter et al. (1984) put it: "One of the central characteristics of the phenomenon is that in a situation in which memory for two attributes of an event is probed, subjects demonstrate knowledge of one attribute but not the other. The remembered attribute corresponds to a fact or item that has been presented to the subject; the forgotten one is the source of the recalled information." (p.594)

Consider the case when the state is A and the source is either perfectly right ($h = 1$) or perfectly wrong ($l = 0$), that is, either it always tells the truth, or always lies. Therefore, we have $\lambda_t(\sigma^t) = \rho_t(\sigma^t)$ and $\omega_t(\sigma^t) = \delta_t(\sigma^t)$ for all t.

Proposition 3. *There exists a set of priors for which naive agents and dogmatic agents starting with the same priors have opposite beliefs in the limit. Further, this set can be partitioned into two subsets where Bayesian agents side with either group.*

The intuition for Proposition 3 is best understood via Figure 3. In that figure, agents whose priors are to the northeast of the lines are those whose beliefs converge to state being A and source being trustworthy, that is $\lim_{t\to\infty} x_t(\sigma^t) = \lim_{t\to\infty} y_t(\sigma^t) = 1$ and those on the southwest have the opposite asymptotic beliefs of $\lim_{t\to\infty} x_t(\sigma^t) = \lim_{t\to\infty} y_t(\sigma^t) = 0$, where $\sigma^t = (\alpha, \alpha, \ldots, \alpha)$. Red line refers to rational agents, green to naive agents, and blue to dogmatic agents. It is clear from the figure that the asymptotic beliefs of naive agents are swayed more by their priors about the signal source. Conversely, the asymptotic beliefs of the dogmatic agents are influenced more by their priors on the state of the

[11]This distinction parallels the one between Sobel (1985) and Benabou and Laroque (1992).

[12]Without this assumption, the results of Blackwell and Dubins (1962) apply, and beliefs converge on the truth as the number of signals goes to infinity. This assumption may be especially relevant in the age of information. With the advent of the Internet, people receive countless signals every day, and it might be reasonable to assume that while these signals influence their beliefs, people might not be able to remember each and every individual signal they have received. An equivalent interpretation is through the lens of an intergenerational transmission mechanism of beliefs, where parents forward their posterior beliefs without passing on the history of signals.

world.

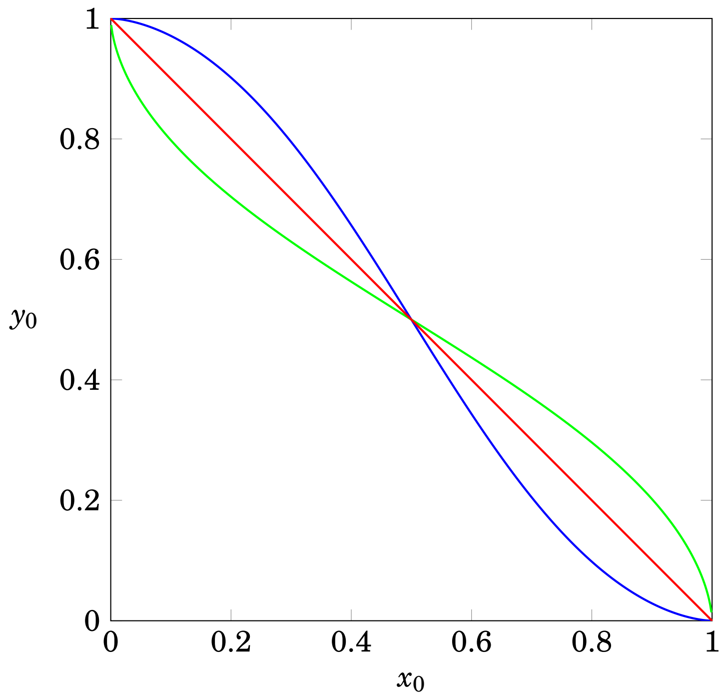

Figure 3: Set of priors that leave agents indifferent between the two actions after infinite signals of A. Red line refers to rational agents who are indifferent between two actions after an infinite sequence of signals, green line to naive agents, and blue line to dogmatic agents.

One implication of this result is that agents with different biases but identical priors on both parameters may have diametrically opposite beliefs after observing the same arbitrarily long sequence of signals. This results provides a novel intuition as to why different individuals observing the same sequence of signals may polarize, relating this phenomenon to the question of whether decision makers primarily pay attention to the state of the world or to signal source.

3 Agency

In this section, I study an agency model in the spirit of Maskin and Tirole (2004). In each of two periods there is an unknown state of the world denoted by $S_t = \{A, B\}$, $t \in \{1, 2\}$ drawn independently, where the prior probability that the state is A is $x_0 \in (1/2, 1)$. There is a voter, whose utility is the same as that of the decision maker given in Equation 3. There is also an incumbent politician who knows the state and is one of two types: with probability $y_0 > 1/2$ she is the congruent type who has the same utility regarding policy as the voter (i.e. gets utility one if action matches the state, zero otherwise), and with the complementary probability she has the reverse preferences as the voter (i.e. gets utility zero if action matches the state, one otherwise). I assume that the politician derives an ego rent of $R > 0$ for every period she is in office.

In the first period the politician chooses an action s from the set $\{a, b\}$. At the end of period one, after the voter observes the politician's policy choice but before the state is revealed, the voter decides whether to reelect the incumbent or elect a challenger of unknown type drawn from the same pool as the incumbent. In the second period the winner of the election picks an action again, after which the game ends and payoffs are realized.

The timings of the game is as follows:

1. Nature chooses the state and incumbent's type;
2. Incumbent observes the state and own type, chooses policy;
3. Voter observes policy choice, decides to vote for the incumbent or a challenger of unknown type;
4. Incumbent chooses the policy for the second period;
5. Voter observes state, the payoffs are realized.

Again, I use the notation $\lambda_t(\sigma^t) = \frac{x_t(\sigma^t)}{1-x_t(\sigma^t)}$ and $\omega_t(\sigma^t) = \frac{y_t(\sigma^t)}{1-y_t(\sigma^t)}$ for the likelihood ratios.

For $R < 1 - y_0$, note that it is a dominant strategy for either type of incumbent to act on her policy preference in the first period, even if this means that she won't be reelected. This is because choosing the policy that yields higher utility yields a minimum of $1 + R + (1 - y_0)$ for the incongruent type, which is higher than the best outcome she can get by pandering, $1 + 2R$. So for these parameter values, the politician acts on her signal in the first period, matching the action to state if she is congruent and mismatching otherwise. I will henceforth focus on the case when $R > 1 - y_0$.

To derive the best response of the voter, I will look at the two types separately. For a Bayesian voter, the problem is straightforward. If he observes that the politician picked s, his posteriors are:

$$\lambda_1(a) = \lambda_0\omega_0 \qquad \omega_1(a) = \lambda_0\omega_0 > \omega_0 \qquad \lambda_1(b) = \frac{\lambda_0}{\omega_0} \qquad \omega_1(b) = \frac{\omega_0}{\lambda_0} < \omega_0$$

However, with a naive voter, the posteriors are:

$$\lambda_1^n(a) = \lambda_0\omega_0 \qquad \omega_1^n(a) = \lambda_0\omega_0^2 > \omega_0 \qquad \lambda_1^n(b) = \frac{\lambda_0}{\omega_0} \qquad \omega_1^n(b) = \frac{\omega_0^2}{\lambda_0}$$

Note that how ω_0 compares with $\omega_1^n(b) = \frac{\omega_0^2}{\lambda_0}$ depends on the relationship between λ_0 and ω_0, or equivalently, between x_0 and y_0. Specifically, if $y_0 > x_0$, the naive voter thinks that the politician who took the unpopular action is still more likely to be congruent than a challenger of unknown type. This is because he updates on the state of the world first, which is closer to the state that matches the politician's policy; and uses that posterior to evaluate the probability that incumbent is the congruent type.

There is no turnover of politicians when the voter is naive; the incumbent is always reelected, even after she takes the unpopular action. While this helps the congruent type of politician keep her job when the state is B, so does the incongruent politician when the state is A. Since the state is more likely to be A than not, the latter effect dominates the former in terms of welfare. The utility of the Bayesian voter is

$$W^b = y_0(2x_0(1 - y_0) + 1 + y_0) \tag{12}$$

whereas the naive voter has utility $W^n = 2y_0$, less than the utility of the Bayesian voter in Equation 12.

On the other hand, when $x_0 > y_0$, and therefore the voters are more confident about the state than they are of politicians' types, the bias of the naive type has no bite, and both types of voters act the same way.

4 Discussion

So far, I ignored the question of why some agents are naive and some are dogmatic. While the model is agnostic about this, a brief discussion here is in order. One answer has to do with the framing of signals. For example, a presidential debate is primarily about forming beliefs about candidates, and policy discussions in a debate are just ways for voters to get to know the candidates better. This gives, therefore, a natural heuristic to observers regarding what they should be focusing on during the debate: candidates. This does not imply beliefs about substantive issues remain unchanged during the course of the debate; beliefs, especially those that are not very strong, are likely to shift, partly based on the positions of candidates that are deemed favorable or unfavorable by the agent. Thus, an agent who updates sequentially would be more likely to be the dogmatic type when watching a presidential debate, the type who primarily cares about getting the beliefs about the type of the source right, and pays relatively less attention to forming correct posteriors regarding substantive issues.

Conversely, if the signal is framed as a policy debate, then we would expect the audience to be primarily focused on the substantive issue at hand, and less so on the speakers. This could result in a greater propensity to exhibit the naive kind of behavior, where agents primarily want to get their posteriors on the state of the world correct, potentially at the expense of biased posteriors on the speakers.

The idea that frames might make certain kinds of bias more prevalent has a number of substantive implications. For instance, if election campaigns indeed lead some voters to be more dogmatic -in the sense of this paper- this would lead to biased beliefs regarding policies. Strong priors about candidates rooted in strong partisanship would exacerbate polarization on substantive matters. Similarly, a politician who wants to improve her approval ratings might find it optimal to frame populist discourses as policy discussions with the aim that voters would associate her with positions they already deem plausible.

Whether agents exhibit this bias or not may also relate to cognitive load. The consensus in the psychology literature is that the working memory can hold the magical number ± 7 items, and some of these slots can be occupied by salient negative emotions like anxiety (Darke, 1988). For example, it was shown that poorer subjects were able to hold fewer objects in their working memory relative to subjects who had fewer financial concerns (Mani et al., 2013). Another set of experiments show that subjects who were more anxious and felt threatened were more likely to exhibit the backfire effect (Nyhan et al., 2014). The model in the present paper suggests that these two findings may be related. If indeed sequential updating stems from an error due to limits of working memory, then changes in emotions invoked in the way signals are presented may influence the number of empty slots on the working memory. This could, in turn, influence whether the agents commit the error proposed here.

This provides a novel, albeit cynical answer as to why senders may want to load their messages with salient emotions. Information that is loaded with negative imagery meant to invoke fear or anger may lead receivers to update their beliefs in a biased way by impairing working memory. This is in line with the observation that politicians and news outlets are often more successful at persuasion when they rouse negative emotions in their audiences.

5 Conclusion

When there are more degrees of uncertainty than there is information, the cognitive task of a decision maker who wishes to update his beliefs is nontrivial. Drawing on evidence from the psychology literature on working memory, I proposed an intuitive and tractable way to update beliefs across multiple dimensions that lead to several salient cognitive biases in the realm of politics. By assuming that some agents update their beliefs sequentially across multiple parameters instead of simultaneously, I have shown that attitude polarization is possible, particularly with common priors and information. This is true in two senses of polarization, not only can posteriors be more distant from each other than priors—which can happen when agents update in the same direction but by different magnitudes—beliefs may in fact go in opposite directions altogether. This suggests a novel interpretation to the observation that voters seem willing to give too much credence to—and allow themselves to be swayed by—dubious sources that agree with their world views, while rejecting credible news sources that challenge them. Relatedly, the backfire effect, where decision makers react to evidence that goes against their priors by strengthening their initial beliefs is explained via the bias proposed in this paper.

In extensions to the benchmark model I show that sequential updating leads to other widely studied cognitive biases such that anchoring and persuasion by repetition from a single source, known in the psychology literature as the truth effect. For the former I showed that decision makers may put more weight on the earlier messages they receive in a sequence of messages. The truth effect, on the other hand, stems from the inconsistency in the beliefs of the decision maker across dimensions when he updates sequentially. I then presented a canonical agency model and showed that sequential updating leads to less turnover of bad politicians, and lower welfare to underline the implications of such a bias in an election setting. Finally, I elaborated on some potential mechanisms that may drive this bias, and whether the framing of the signal may be used to manipulate the sequence with which decision makers update their beliefs. Future research should study whether different orders of updating can be induced by senders, as well as the implications of sequential updating under various settings.

Appendix

Proof of Proposition 1

Let's start with the first part and focus on the case when $y_0 h + (1-y_0)l > 1/2$, the other case is symmetric. Recall that $\omega_1^n(\alpha) = \omega_0\rho_1(\alpha)$ and $\omega_1^s(\alpha) = \omega_1(\alpha) = \omega_0\rho_0$. Therefore, we have:

$$\begin{aligned}\omega_1^n(\alpha) > \omega_1(\alpha) \Leftrightarrow \rho_1(\alpha) > \rho_0 \Leftrightarrow (x_1(\alpha) - x_0)(h-l) > 0\\ \Leftrightarrow x_1(\alpha) > x_0 \Leftrightarrow \lambda_1(\alpha) > \lambda_0 \Leftrightarrow \delta_0 > 1 \Leftrightarrow y_0(h) + (1-y_0)l > 1/2\end{aligned}$$

Similarly:

$$\begin{aligned}\omega_1^n(\beta) > \omega_1(\beta) \Leftrightarrow \rho_1(\beta) < \rho_0 \Leftrightarrow (x_1(\beta) - x_0)(h-l) < 0\\ \Leftrightarrow x_1(\beta) < x_0 \Leftrightarrow \lambda_1(\beta) < \lambda_0 \Leftrightarrow \delta_0 > 1 \Leftrightarrow y_0(h) + (1-y_0)l > 1/2\end{aligned}$$

A symmetric argument establishes the opposite case.

For the second part, $\lambda_1^s(\sigma) < \lambda_1(\sigma) = \lambda_1^n(\sigma)$ if $x_0 < 1/2$:

$$\lambda_1^s(\alpha) < \lambda_1(\alpha) \Leftrightarrow \delta_1(\alpha) < \delta_0 \Leftrightarrow y_1(\alpha) < y_0 \Leftrightarrow x_0 < 1/2$$

Similarly,

$$\lambda_1^s(\beta) < \lambda_1(\beta) \Leftrightarrow \delta_1(\beta) > \delta_0 \Leftrightarrow y_1(\beta) > y_0 \Leftrightarrow x_0 < 1/2$$

By a symmetric argument, $\lambda_1^s(\sigma) > \lambda_1(\sigma) = \lambda_1^n(\sigma) \Leftrightarrow x_0 > 1/2$. □

Proof of Corollary 1

First, note that $E[u(a,S)] = x_1(\sigma) > 1 - x_1(\sigma) = E[u(b,S)] \Leftrightarrow x_1(\sigma) > 1/2$, and so decision makers choose a whenever they believe state A to be more likely than state B. For rational and naive agents, we have: $x_1(\alpha) = \frac{x_0(1+y_0)}{2x_0y_0+1-y_0} > 1/2 \Leftrightarrow 2x_0 + y_0 > 1$. For dogmatic agents,

$$\begin{aligned}x_1^s(\alpha) = \frac{x_0(1+y_1)}{2x_0y_1+1-y_1} > 1/2\\ \Leftrightarrow y_0(4x_0^2 - 2x_0 + 1) > 1 - 2x_0\\ \Leftrightarrow y_0 > \frac{1-2x_0}{4x_0^2 - 2x_0 + 1}\end{aligned}$$

□

Proof of Lemma 1

When the signal is α, the naive type's posterior on the sender's type is greater than his prior when:

$$\begin{aligned}\omega_1^n(\alpha) > \omega_0 &\Leftrightarrow x_1(\alpha)(2h-1) + 1 - h > x_1(\alpha)(2l-1) + 1 - l\\ &\Leftrightarrow 2x_1(\alpha)(h-l) > h - l\\ &\Leftrightarrow x_1(\alpha) > 1/2\\ &\Leftrightarrow \lambda_0\delta_0 > 1\end{aligned}$$

where δ_0 is as defined as in Equation (4). Similarly, the dogmatic type's posterior on the state is greater than his prior when:

$$\begin{aligned}
\lambda_1^s(\alpha) > \lambda_0 &\Leftrightarrow y_1(\alpha)(h-l)+l?1-y_1(\alpha)(h-l)-l \\
&\Leftrightarrow y_1(\alpha) > \frac{1/2-l}{h-l} \\
&\Leftrightarrow \omega_1(\alpha) > \frac{1/2-l}{h-1/2} \\
&\Leftrightarrow \omega_0 \frac{r_0}{q_0} > \frac{1/2-l}{h-1/2}
\end{aligned}$$

where $r_0 = x_0(2h-1)+1-h$, and $q_0 = x_0(2l-1)+1-l$. Note that in the latter case, the converse is only possible when $h > 1/2 > l$. That is, the dogmatic type can only update away from the signal when it is possible that the signal is likely wrong.

Proof of Proposition 2

Suppose $l < 1/2 < h$. Then, note that:

$$\omega_0 < \omega_1(\alpha) \Leftrightarrow \lambda_0 > 1$$

and

$$\frac{1/2-l}{h-1/2} < \omega_0 \Rightarrow \delta_0 > 1 \Rightarrow \lambda_1(\alpha) > \lambda_0$$

With Lemma 1, these give us the desired result.

□

Proof of Proposition 3

Suppose without loss of generality the signal source transmits A every period. Then, the Bayesian agent's posterior likelihood ratios are characterized as follows:

$$\begin{aligned}
&\lambda_1 = \omega_1 = \lambda_0\omega_0 > 1 \\
&\Leftrightarrow \lambda_2 = \omega_2 = \lambda_1\omega_1 = (\lambda_0\omega_0)^2 > \lambda_0\omega_0 = \lambda_1 = \omega_1 \\
&\Leftrightarrow \lambda_t = \omega_t = (\lambda_0\omega_0)^t > (\lambda_0\omega_0)^{t-1} = \lambda_{t-1}\omega_{t-1} \\
&\Longrightarrow \lim_{t\to\infty} \lambda_t = \lim_{t\to\infty} \omega_t = \infty \Longleftrightarrow \lambda_0\omega_0 > 1 \\
\text{and symmetrically} \quad &\lim_{t\to\infty} \lambda_t = \lim_{t\to\infty} \omega_t = 0 \Longleftrightarrow \lambda_0\omega_0 < 1
\end{aligned}$$

Since the naive agent first updates the belief on the state of the world in each period, and then uses that posterior to form beliefs about the type, his beliefs are characterized as follows:

$$\begin{aligned}
\lambda_t^n = \lambda_{t-1}^n \omega_{t-1}^n \quad &\text{and} \quad \omega_t^n = \omega_{t-1}^n \lambda_t^n \quad \text{or} \\
\lambda_t^n = \lambda_0^{F_{2t-2}} \omega_0^{F_{2t-1}} \quad &\text{and} \quad \omega_t^n = \lambda_0^{F_{2t-1}} \omega_0^{F_{2t}}
\end{aligned}$$

where F_x refers to the xth number on the Fibonacci sequence. First of all, note that if $\lambda_0 > 1$ and $\omega_0 > 1$, then $\lim_{t\to\infty} \lambda_t^n = \lim_{t\to\infty} \omega_t^n = \infty$. Similarly, if $\lambda_0 < 1$ and $\omega_0 < 1$, then $\lim_{t\to\infty} \lambda_t^n = \lim_{t\to\infty} \omega_t^n = 0$. When one of the ratios is greater and the other less than one, the asymptotic beliefs are slightly more involved to calculate. Suppose $\log(\lambda_0)\log(\omega_0) > 0$ and take some $z \in \mathbb{R}$:

$$\begin{aligned} \log(\lambda_t^n) = F_{2t-2}\log(\lambda_0) + F_{2t-1}\log(\omega_0) > z \\ \Longleftrightarrow \log(\lambda_0) + \frac{F_{2t-1}}{F_{2t-2}}\log(\omega_0) > \frac{z}{F_{2t-2}} \end{aligned}$$

Let F be the limit of the Fibonacci sequence, then, as $t \to \infty$:

$$\begin{aligned} &\log(\lambda_0) + F\log(\omega_0) > 0 \Longleftrightarrow \lambda_0\omega_0^F > 1 \\ &\Longrightarrow \lim_{t\to\infty} \lambda_t^n = \lim_{t\to\infty} \omega_t^n = \infty \\ &\log(\lambda_0) + F\log(\omega_0) < 0 \Longleftrightarrow \lambda_0\omega_0^F < 1 \\ &\Longrightarrow \lim_{t\to\infty} \lambda_t^n = \lim_{t\to\infty} \omega_t^n = 0 \end{aligned}$$

By symmetric arguments for the dogmatic agent

$$\begin{aligned} \lambda_0^F\omega_0 > 1 \Longrightarrow \lim_{t\to\infty} \lambda_t^s = \lim_{t\to\infty} \omega_t^s = \infty \\ \lambda_0^F\omega_0 < 1 \Longrightarrow \lim_{t\to\infty} \lambda_t^s = \lim_{t\to\infty} \omega_t^s = 0 \end{aligned}$$

□

REFERENCES

Acemoglu, Daron, Victor Chernozhukov, and Muhamet Yildiz, "Fragility of asymptotic agreement under Bayesian learning," *Theoretical Economics*, 2016, *11* (1), 187–225.

Andreoni, James and Tymofiy Mylovanov, "Diverging Opinions," *American Economic Journal: Microeconomics*, 2012, *4* (1), 209–232.

Baliga, Sandeep, Eran Hanany, and Peter Klibanoff, "Polarization and ambiguity," *The American Economic Review*, 2013, *103* (7), 3071–3083.

Benabou, Roland and Guy Laroque, "Using privileged information to manipulate markets: Insiders, gurus, and credibility," *The Quarterly Journal of Economics*, 1992, pp. 921–958.

Benoit, Jean-Pierre and Juan Dubra, "A Theory of Rational Attitude Polarization," *Available at SSRN 2529494*, 2014.

Blackwell, David and Lester Dubins, "Merging of Opinions with Increasing Information," *The Annals of Mathematical Statistics*, September 1962, *33* (3), 882–886.

Cheng, Ing-Haw and Alice Hsiaw, "Distrust in Experts and the Origins of Disagreement," SSRN Scholarly Paper ID 2864563, Social Science Research Network, Rochester, NY December 2016.

Darke, Shane, "Anxiety and working memory capacity," *Cognition and Emotion*, April 1988, *2* (2), 145–154.

DeMarzo, Peter M., Dimitri Vayanos, and Jeffrey Zwiebel, "Persuasion bias, social influence, and unidimensional opinions," *The Quarterly Journal of Economics*, 2003, *118* (3), 909–968.

Ecker, Ullrich KH, Klaus Oberauer, and Stephan Lewandowsky, "Working memory updating involves item-specific removal," *Journal of Memory and Language*, 2014, *74*, 1–15.

Eyster, Erik and Matthew Rabin, "NaÃŕve Herding in Rich-Information Settings," *American Economic Journal: Microeconomics*, November 2010, *2* (4), 221–243.

Fernandes, Juliana, "Effects of Negative Political Advertising and Message Repetition on Candidate Evaluation," *Mass Communication and Society*, March 2013, *16* (2), 268–291.

Fryer, Roland G., Philipp Harms, and Matthew O. Jackson, "Updating Beliefs when Evidence is Open to Interpretation: Implications for Bias and Polarization," 2016.

Judd, Charles M. and Markus Brauer, "Repetition and evaluative extremity," in "Attitude strength: Antecedents and consequences," Vol. 4 1995, pp. 43–72.

Lewandowsky, Stephan, Ullrich KH Ecker, Colleen M. Seifert, Norbert Schwarz, and John Cook, "Misinformation and its correction: Continued influence and successful debiasing," *Psychological Science in the Public Interest*, 2012, *13* (3), 106–131.

Mani, Anandi, Sendhil Mullainathan, Eldar Shafir, and Jiaying Zhao, "Poverty Impedes Cognitive Function," *Science*, August 2013, *341* (6149), 976–980.

Maskin, Eric and Jean Tirole, "The politician and the judge: Accountability in government," *The American Economic Review*, 2004, *94* (4), 1034–1054.

Morris, Stephen, “Political Correctness,” *Journal of Political Economy*, 2001, *109* (2), 231–265.

Nyhan, Brendan, Jason Reifler, Sean Richey, and Gary L. Freed, “Effective Messages in Vaccine Promotion: A Randomized Trial,” *Pediatrics*, April 2014, *133* (4), e835–e842.

Oberauer, Klaus, “Access to information in working memory: exploring the focus of attention.,” *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 2002, *28* (3), 411.

Ortoleva, Pietro and Erik Snowberg, “Overconfidence in political behavior,” *The American Economic Review*, 2015, *105* (2), 504–535.

Rabin, Matthew and Joel L. Schrag, “First Impressions Matter: A Model of Confirmatory Bias,” *The Quarterly Journal of Economics*, February 1999, *114* (1), 37–82.

Schacter, Daniel L., Joanne L. Harbluk, and Donald R. McLachlan, “Retrieval without recollection: An experimental analysis of source amnesia,” *Journal of Verbal Learning and Verbal Behavior*, 1984, *23* (5), 593–611.

Sobel, Joel, “A theory of credibility,” *The Review of Economic Studies*, 1985, *52* (4), 557–573.