Foundations for an Economic Model of Echo Chambers

by

Ryan Scott

A Thesis Submitted in Partial Fulfillment of the Requirements for the
Degree of Bachelor of Science, Honours
in the Department of Economics
University of Victoria
2019

Supervised by Dr. Chris Willmore

For

Dr. Elisabeth Gugl, Honours Co-Advisor

Dr. Martin Farnham, Honours Co-Advisor

Abstract

Many social scientists and news articles warn of the negative effects of echo chambers. Economists have built upon the psychological concepts of confirmation bias and cognitive dissonance to construct models that describe “echo-chamber “-like behavior. Economists find that the behaviors observed in echo chambers are perfectly rational and utility maximizing. Further work on modelling echo chambers should focus on negative externalities arising from actions undertaken by agents in echo chambers and by extending network based Bayesian models to include agents with utility functions.

I – Introduction

In the wake of the 2016 American Presidential election, Great Britain's Brexit vote, and the recent outbreak of measles in Oregon and Vancouver there has been increasing discourse about echo chambers and the potential harm they may cause. Increasingly, polarization and segregation are being viewed as a serious impediment to public policy decision-making and therefore may have adverse effects on societal welfare. The Brexit vote and subsequent uncertainty about exit terms and agreements have already caused significant contractions in the British economy with both losses to the value of the pound and major firms relocating outside the country. The recent spread of measles in Vancouver has re-sparked a public debate about mandatory vaccinations for school-aged children in the face of a public health crisis.

"An *echo chamber* is a sociological setting where peoples' prior beliefs are 'echoed back' giving the impression that their beliefs are correct" (Baumgaertner, 2014). Under this definition an echo chamber encompasses both online forums/communities, but also includes the "conventional wisdom" of Galbraith (1958) that even academic fields occasionally struggle with. However, echo chambers are particularly prevalent online, on websites such as YouTube, on Reddit (a news/information forum in which content is driven upwards by a community voting mechanism), and on Facebook (Bakshy, 2015). Echo chambers are prime grounds for the propagation and spread of inaccurate, misleading, or false information, because they tend to discourage or silence dissenting viewpoints and evidence. Quattrociocchi *et al*. (2016) & Del Vicaro *et al*. (2017) investigate the spread of misinformation in the context of the Brexit vote. Their research finds strong

evidence that users choose to selectively consume information and tend to join communities of like-minded individuals.

Economists tend to view more information as good, and assume that inaccurate information will simply be discarded as agents make decisions. This is not to say that economists argue that agents are always better off revealing the true state of the world. As I will discuss below agents may be better off not seeking the true state of the world. In light of empirical evidence, economists might be able to contribute by generating a theoretical model that addresses why individuals choose to remain in echo chambers and allow for analysis of societal welfare under echo chambers and without.

In this paper I outline some different approaches economists take to model seemingly irrational behaviors in the areas of information consumption and the formation of beliefs. Research in this area builds heavily on insights from psychology. In Section III, I briefly review confirmation bias, cognitive dissonance, and the sunk-cost fallacy, three major insights in the psychological literature that economists have built upon, each capturing a different reason why agents modify their information consumption and belief formation behavior. In Section IV, I briefly review the basic principles of Bayesian updating and signalling. In Section V, I detail the economic approaches to models of information consumption and belief formation. In Section VI, I discuss my findings, and conclude with some thoughts about future work in this area.

II – The Market for Ideas

Economists have long investigated markets. Therefore, it would seem natural to begin a discussion of an economic model of echo chambers with a market metaphor. As I asserted above, a hallmark of echo chambers is the repetition of inaccurate or false information within a chamber and the silencing of dissenting opinions. One classic conception of the spread and share of information is the marketplace for ideas. I argue that market failure in the marketplace for ideas provides an intuitive conceptual starting point for thinking about echo chambers and why, from a structural standpoint, they may exist.

The market for ideas intuition is very simple: through public debate correct or good ideas will "outcompete" bad ideas and only good information will remain in the public discourse. In liberal democracies public decisions are made on the basis of this debate. To this end, a properly informed public that makes well-informed decisions can be thought of as a positive externality.

However, there may be market imperfections in the marketplace for ideas that lead to market failure. Imperfect competition reigns in the market for ideas. One facet of this is the decreasing number of media firms each holding an increasing number of subsidiaries. If consumers are getting their information from an ever-smaller number of sources, those firms have a greater ability to drive and frame public debate (Herman & Chomsky, 1988). Market failure will result in this case when media firms work to ensure that incorrect or bad viewpoints are not driven from the public discourse. Napoli (2018) makes the case that journalism is a public good in a liberal democracy and as such the under-provision of journalism represents a market failure. Moore (2015) notes that in the absence of a

credible mechanism to discipline those who propagate falsehoods the market's truth seeking mechanism can break down.

On the other hand, the Internet has greatly decreased the costs of distributing information for individuals and smaller firms. The ways in which consumers come across this new content then becomes of great importance. Google and Facebook have both come under increased scrutiny for how their services organize search results, facilitate the sharing of content, and suggest additional content for consumption (Napoli, 2018). The algorithms by which companies and consumers can limit their access to information are one type of a filter bubble. Similarly to imperfect competition, filter bubbles allow incorrect information to survive the public discourse by allowing some of the public to segregate themselves from the debate.

Finally, if consumers are unable to differentiate between information of high and low quality, market failure may occur as a result of asymmetric information. This is because it is presumably more expensive to produce information of high quality, but if consumers are not able to distinguish high and low quality information they may be unwilling to pay a price high enough to cover the cost for high quality firms. If this is the case the high quality firms will exit the market and only low quality firms will remain. This completely undermines the market for information as it rests on the assumption that consumers can perfectly observe the quality of information being supplied in the market and would presumably be willing to pay a higher price for higher quality information. Along these lines Napoli (2018) argues that asymmetric information results from the inability of consumers to distinguish information of differing qualities, which he describes as boundedly rational.

I argue that each of these market imperfections could cause an echo chamber to form, or compound the effect of an existing one, if the market for ideas analogy is correct. But these market imperfections may not be sufficient in terms of working towards an economic model of echo chambers. What certainly is necessary is to take into account the psychologies of the consumers in the market for ideas themselves.

III – Contributions from Psychologists

Psychologists explain behaviors such as those observed in echo chambers primarily as cognitive dissonance, confirmation bias, and the sunk cost fallacy. Leon Festinger first proposed a theory of cognitive dissonance in 1957. He proposed that a "pair of cognitions is inconsistent if one cognition follows from the [opposite] of the other" (Cooper, 2007) and that this caused people psychological distress. Festinger first gave the example of doomsday prophets whose predicted apocalypse did not come to pass. Very quickly these people set new conditions that reconciled their previous predictions with their new (continued) existence. Festinger concluded that people act to minimize or resolve dissonance in themselves. In the case of echo chambers, I consider the suppression of dissenting views to be an example of the community working to ensure that cognitive dissonance is minimized. Alternatively, one example of the minimization of cognitive dissonance is the recent comparisons by American Evangelical Christians of President Trump to King Cyrus to reconcile his personal life with the doctrine of Christ. Cooper (2007) provides a useful overview on this subject.

Confirmation bias refers to one's tendency to rely more heavily on information that supports one's prior beliefs than on information that contradicts one's prior beliefs. If one believes (incorrectly for the record) that vaccines cause autism one is far more likely to accept a study that finds this to be the case than a study identical in all other aspects that finds no relation. Nickerson (1998) provides an extensive literature review and discusses evidence in a variety of applied contexts.

The sunk cost fallacy is the tendency of agents to take into account costs that are non-recoverable, and therefore irrelevant to optimal decision making, when deciding to continue to invest in a project. In the context of echo chambers I refer to this as "post-rationalization". Consider someone who has made a decision that comes to pass as sub-optimal. A person who seeks to post-rationalize will make a further decision that makes the previous decision utility maximizing in light of the sunk cost of the original decision. For example, an undergraduate student who regrets undertaking the honours program may continue on to graduate school, so as to make undertaking the honours program optimal. Arkes and Blumer (1985) review the subject and provide experimental evidence suggesting that having undertaken prior courses in economics was not enough to overcome the sunk cost fallacy. Even economists (albeit undergraduate ones) are not immune from cognitive biases.

All of this implies that to some degree agents may desire to limit or control the amount of or type of information they receive. Economists such as Golman *et al.* (2017) identify physical avoidance, inattention, biased interpretation of information, forgetfulness, and what the authors term "self-handicapping", which is described as the tendency for agents to pick tasks poorly matched with their abilities in order to prevent themselves

from learning about their own limits, as the major methods by which agents avoid information acquisition. I focus on biased interpretation in this paper, and I believe that "self-handicapping" could be considered an example of agents minimizing cognitive dissonance. Agents would suffer disutility if the limits on their abilities were revealed to them and thus seek to remain ignorant of those limits. The authors give many examples of why they believe agents choose to ignore or avoid information, ranging from ignoring the results of a medical test to an inventor who ignores his financial portfolio when the stock market is down.

IV – Bayesian Updating and Signalling

Bayesian updating is the major paradigm in information acquisition. In a typical model of Bayesian updating, a rational agent has a given prior belief about the likelihood of an event and receives information as signals, often of varying degrees of quality. The agent updates their prior beliefs using Bayes' Rule.

$$P(\Omega = A|\sigma = a) = \frac{P(\Omega = A)P(\sigma = a|\Omega = A)}{P(\sigma = a)}$$

The basic implementation of Bayes' Rule assumes that the agent can perfectly keep track of and weigh all the information they receive and that updating is instantaneous.

Economists have made great use of signalling in this field. In most models the true state of the world is an element in a set containing two elements.

$$\Omega \epsilon \{A, B\}$$

Each period, an agent receives a signal correlated with the true state of the world with some degree of accuracy.

$$\sigma \in \{a, b\} \; with \; accuracy \; p$$

The agent updates their beliefs about the true state of the world (say in this case the probability that the true state of the world is A) using Bayes' Rule. Signals received with high accuracy change beliefs about the state of the world more than signals with low accuracy. Consider the following example: Let there be two states of the world, A and B, with probabilities .5 and .5 respectively. Let the probability that a signal, σ, be equal to *a* given the state of the world is A be .7 across all periods. Let the agent receive three signals all of which are equal to *a*. I show in the following table how the belief of the agent that the true state of the world is A can be calculated:

$P(\Omega=A, B)$	$P(\sigma_t=a\|\Omega)$	$P(\Omega)P(\sigma_1=a\|\Omega)$	$P(\Omega)P(\sigma_1=a\|\Omega)P(\sigma_2=a\|\Omega)$	$P(\Omega)P(\sigma_1=a\|\Omega)P(\sigma_2=a\|\Omega)P(\sigma_3=a\|\Omega)$	$P(\Omega=A\|\sigma_1=a, \sigma_2=a, \sigma_3=\alpha)$
.5	.7	.35	.245	.171	.924
.5	.3	.15	0.045	0.014	.076

After receiving 3 signals the agent believes the true state of the world to be A with 92.4% probability.

V – Economic Models

I will focus on 6 models that capture different aspects of echo chamber-like behavior.

Table 1 - Model Summary

Author	Year	Model Type	Findings
Akerlof & Dickens	1982	Utility derived from beliefs, Agent suffers from cognitive dissonance	Agent chooses to forgo safety equipment to avoid recognizing a dangerous work environment
Yariv	2002	Utility derived from beliefs, agent has preference for consistency, updates belief in a Bayesian fashion	Agents may wish to ignore information to maximize utility and beliefs formed early on may be impossible to shake
Eyster	2001	Utility derived from beliefs, agent works to justify past decisions	Agents minimize cognitive dissonance and act to post-rationalize
Kovach	2015	Bayesian consumer of information with utility dependent on beliefs	Agents over-weight acceptable information and under-weight the bad. Agents with same initial information may come to believe in opposite hypotheses.
Eliaz & Spiegler	2006	Utility derived from beliefs, agent is Bayesian information consumer	Model fails to account for intuitive examples given by authors
Rabin and Schrag	1999	Bayesian agent who misreads signals due to confirmation bias	Agent may come to believe in a false hypothesis with near certainty despite an infinite amount of information
Bowden	2014	Bayesian agent who misreads signals due to confirmation bias, in network setting	Confirmation bias slows the spread of new information in a network
Jann & Schottmuller	2018	Agents derive utility from information, actions, and belief	Agents choose to segregate themselves into communities with similar beliefs.

The basic Bayesian model has been extended to account for human failings. Rabin and Schrag (1999) model a Bayesian updater who suffers from confirmation bias who in each period receives a signal that is correlated with the true state of the world. It is initially assumed that these signals are of equal accuracy:

$$P(\sigma_t = a|A) = P(\sigma_t = b|B) = \theta$$

$$for\ \theta\ \epsilon\ (.5, 1)$$

With positive probability q, the agent falsely interprets signals in support of one state of the world as actually supporting a different state of the world. This leads to an updated belief that is incorrect. Consider the recent case of news that the Special Counsel's investigation into Russian election meddling was coming to a close. Many on the political left took this as a signal of immanent impeachment. This was not the case and shows that a mistaken interpretation captures the confirmation bias effect. The probability of mistaken interpretation occurring (q) is not dependent on how strongly the agent believes in one state of the world. In this case the number of signals perceived to be correlated with A and B are given by n_a and n_b respectively. Rabin and Schrag (1999) define the probability that the agent misinterprets the signal as θ^* and the probability that the agent correctly perceives the signal as θ^{**}. These are given by:

$$\theta^* = (1 - \theta) + q\theta$$

$$\theta^{**} = \theta + q(1 - \theta)$$

Thus the probability that the true state of the world is A is given by the following equation:

$$P(x = A|n_a, n_b) = \frac{\theta^{\lambda}}{\theta^{\lambda} - (1-\theta)^{\lambda}} \; where \; \lambda = n_a(\theta^*, \theta^{**}) - n_b(\theta^*, \theta^{**})$$

This probability is conditional on the number of signals perceived to be correlated with each state of the world. As the probability of a mistaken interpretation increases, n_a and n_b change respectively. This modeling choice is slightly different from other models I present here, in which confirmation bias entails the overlooking of countering evidence as opposed to a simple misunderstanding about the state of the world the signal supports. The authors demonstrate that this mistaken interpretation is enough to bias the beliefs of the agent such that the agent may believe very strongly in a false state of the world even if the agent receives an infinite number of additional signals. This directly contradicts the conventional wisdom that enough correct information will correct for incorrect information, allowing the correct belief to be formed.

Bowden (2014) examines how confirmation bias affects the spread of information in a network, building on Rabin and Schrag's work. Bowden assumes a network of Bayesian agents in which information flows from one agent to the next and each agent is indirectly connected to all other agents via third parties. In his simplest specification, this takes the form of a ring of agents.

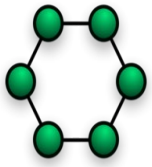

Fig 1. Ring Network. From Wikipedia Commons, 2006,

https://en.wikipedia.org/wiki/Ring_network#/media/File:NetworkTopology-Ring.png

Each period, individuals receive a signal and compare it to the last period belief of their neighbours, updating their belief with a modified Bayes' Rule. The higher the bias of agents, the more likely society is to retain the current outlook as agents rely on the beliefs of their neighbours and may misinterpret their own private signal as supporting the public view when it actually does not. Bowden concludes that confirmation bias slows the rate of the spread of information in the network of agents.

Jann and Schottmuller (2018) propose a model in which agents play a cheap-talk game. Jann and Schottmuller build on the simple intuition that it is easier to share your beliefs with those who agree with you. Agents choose to enter one of a range of available rooms, observe a private signal, and send a message to only the other occupants of the room. Agents then undertake an action realizing a payoff. Payoffs are a function of actions, preferences, and information and explicitly depend on the weighted payoffs of other agents.

$$u_i = -(a_i - b_i - \sum_{k=1}^{n} \theta_k)^2 - \alpha \sum_{j \neq 1} (a_j - b_j - \sum_{j=1}^{n} \theta_j)^2$$

Utility increases in information, θ, as it is either equal to 0 or 1, and the action they undertake. Beliefs are simply a real number that will be discussed below. Actions are not

strictly defined by the authors, other than through a first order condition that shows the optimal action for agent I is agent I's belief plus a weighted sum of all other agent's beliefs dependent on the value of α. I argue that the most important extension to this model would be a better definition of what it means to undertake an action. However, the authors focus on the messaging and room choice aspects of their model. In this model agents can lie by sending messages that are the opposite of their private signal. The authors focus on the most informative equilibrium of this game in which agents do not lie. Agents truthfully reveal their information (send a message) if the difference between their belief b_i and the mean belief in the room they are located in is sufficiently small, otherwise they "babble" in which they send messages that are independent of their private signal.

$$b_i \in [\bar{b} - \frac{n-1}{n}(p - \frac{1}{2}), \bar{b} + \frac{n-1}{n}(p - \frac{1}{2})]$$

The above truth-telling interval increases in size as the accuracy of private signals increase and the number of people in the room. As each agent has unique information about the true state of the world and the payoffs are linear, Jann and Schottmuller find that segregation is welfare increasing, because the total amount of information available to each agent is higher under segregation than without. The ability of agents to segregate into communities with similar beliefs allows agents to tell the truth about their signal. This encourages the exchange of information and reduces babbling. Most interestingly, as the authors show agents send true messages when close to a mean belief in a room, they argue that segregation arises from polarization and not the other way around. Communities with high levels of polarization opt to segregate and low polarization communities opt to remain

integrated. However, the authors are clear that polarization itself is not welfare maximizing, only segregation in response to it is.

While some have attempted to model information as a consumption good to take advantage of basic economic models, Dubra and Echenique (2001) prove that if the state space is uncountable, monotonic preferences over information cannot be represented by a utility function. As such, I do not review this strand of the literature at length.

Beliefs have been somewhat mistreated in economics. Akerlof and Dickens (1982) propose a model in which beliefs directly enter the utility function. The authors show that fully informed workers in dangerous occupations who suffer disutility from cognitive dissonance may choose to forgo safety equipment. Yariv (2001) contributes a theoretical underpinning for direct inclusion of beliefs in the utility function.

Yariv (2002) presents a model in which agents have a preference for and directly derive utility from the consistency of beliefs over time. Her model attempts to capture the tradeoff between instantaneous utility and consistency over time. As in Yariv (2001), utility is given by the sum of discounted utilities for all periods with utility in each period being derived from beliefs and actions in that period only:

$$U_\tau = \sum_{\tau=t}^{T} \delta^{\tau-t}(u(a_\tau) + \gamma v(\mu_{\tau-1}, \mu_t))$$

Actions are assumed to be fully consistent with beliefs. If the action undertaken is consistent with the true state of the world agent receives $u(a_\tau) = 1$ otherwise the agent receives zero utility. Observing a signal each period, the agent updates their beliefs using Bayes' rule.

$$\mu_{\tau+1}^{B} = \begin{cases} \dfrac{\mu_\tau q_\tau}{\mu_\tau q_\tau + (1-\mu_\tau)(1-q_\tau)} \\ \dfrac{\mu_\tau (1-q_\tau)}{\mu_\tau (1-q_\tau) + (1-\mu_\tau) q_\tau} \end{cases}$$

She shows that under these preferences agents can become over- and under-confident when compared to a Bayesian observer, that first-period beliefs can self-reinforce if extreme, and that agents may wish to selectively consume or ignore information to maximize utility. The degree to which the agent is accurately able to predict their future preference for consistency determines the degree to which the agent displays the above results.

Eyster (2001) develops a model in which agents have a taste for consistency; however, agents do not seek to reduce disutility arising from dissonance rather they wish to avoid recognizing they have made a suboptimal decision. To this end, they undertake a further action to rationalize the former. In the simplest case, the decision maker chooses an action from a set of actions in period one. In period two, the true state of the world is revealed to them, and they choose an action from a set dependent on the action they undertook in period one and the true state of the world. Eyster defines the regret that the agent incurs as the difference between the maximum payoff she could receive if she could choose action 1 again, holding action 2 and state of the world fixed, and her actual payoff. This is defined by the function:

$$regret = -\big(max_{a_2}\{g(a_1, a_2, \omega)\} - g(a_1, a_2, \omega)\big) - \rho\big(max_{a_1}\{g(a_1, a_2, \omega)\} - g(a_1, a_2, \omega)\big)$$

Here, ρ represents the degree to which the agent displays the need to post-rationalize. Eyster shows that agents who have stronger tastes for consistency work to minimize their regret, choosing actions in period 2 that justify actions taken in period 1. Eyster generalizes his model to a many-period analysis. Kovach (2015) presents a similarly motivated Bayesian model where the agent's response to information is a function of past actions and the Decision Maker displays from cognitive dissonance. He finds agents will over-update when presented with good news and under-update when presented with bad news. Kovach considers polarization of belief in the context of his model. He shows that two agents with identical prior beliefs and who receive the same information can have different posterior beliefs.

Direct inclusion of beliefs in utility functions has been met with some criticism. Eliaz and Spiegler (2006) propose their own Bayesian model in which beliefs directly enter the utility function. The authors seek to capture the effect of an anticipatory attitude to information. They give the example of a manger that seeks counsel when he expects his advisors will agree with him, and avoids them when he is uncertain as an example of such an anticipatory attitude. They struggle with ordinal utility representation, finding that agents with vastly different beliefs may choose the same information source and that non-monotonic transformations of preferences exist such that those who disagree before the transformation may end up with identical preferences afterwards. The authors believe this to be a serious shortcoming as they argue that a belief that the state of the world is "good" is preferable to a belief that the state of the world is "bad". Agents who hold differing beliefs about what is "good" should have different preferences. Transformations of these preferences show that these differing agents may be

indistinguishable from one another in that they choose the same signals. The authors give an example in which two sets of preferences rationalize the same attitude towards information, but the preferences themselves generate utility rankings that are the opposite of each other. Ordinal utility may not be enough when building an economic model of echo chambers, at least under Eliaz and Spiegler's setup. The authors also fail to capture the effect anticipatory attitudes to information that they set out to. However, the authors themselves defend the approach, disagreeing that their failures suggest economists abandon the effort to integrate anticipatory attitudes and beliefs into agent's utilities. The authors suggest that abandoning the Bayesian component of their model may be the best way to proceed. This will also allow for preferences to be robust to affine transformations and preferences become cardinal.

VI – Conclusions and Directions for Future Modelling

Economists have built upon the psychological insights of confirmation bias, cognitive dissonance, and post-rationalization to model seemingly irrational behavior. Bayesian updating and signaling feature heavily in this literature and seem to be useful features to retain in future models as they are prevalent in economics at large and capture belief formation and information consumption adequately. Assuming agents derive utility from actions, beliefs, information, or some combination of these seem to be reasonable if economists wish to model echo chambers. It indeed seems likely that agents have preferences over their beliefs. Very few people wish to be wholly honest with themselves about their flaws. It may be difficult to admit that the world in which we live does not

match our previous perception of it. In order to maintain the illusion that the world is the way we believe it to be, it would seem perfectly reasonable to behave somewhat irrationally. Economists, however, argue that the agents who "suffer" from confirmation bias, cognitive dissonance, or post-rationalization are far from irrational. All of the behaviors that are observed in echo chambers an economist would call perfectly rational, given that the agent react as the models here suggest. If beliefs are included in the utility function, an agent rationally maximizes their utility. If the agent consumes information and wishes only to know what they already believe they rationally ignore it. From a purely utilitarian perspective an outside attempt to force the agent to change their behavior would be welfare reducing in the absence of compensation.

I consider this somewhat distressing, in light of empirical evidence that suggests echo chambers increase polarization and are difficult to dismantle. I suggest that future work on multi-agent models incorporating externalities may reconcile the obvious disutility of individual agents being forced out of their echo chambers with the societal disutility that arises with echo chambers. Two models stand out as candidates for foundations of future models Bowden's (2014) network based model and Jann and Schottmuller's (2018) room based model.

The network setting of Bowden's model is attractive as it captures the persistence of societal beliefs over time when agents take into account the beliefs of agents around them. This definitely seems to model the "echo" of the echo chamber. However, agents in his model are simply Bayesian updaters which seems a little mechanical, even for economists. Integrating a utility function (if one can be generated from underlying preferences) into the model seems to be a natural direction to go. Yariv's (2002) model seems to be very flexible,

and is at least supported by axiomatic foundations laid out in Yariv (2001). However, Eyster's (2001) model of minimized regret may be more tractable and still captures a desire for consistency. Both Bowden and Eyster's models are many-period and it seems reasonable to assume that an agent who seeks to minimize regret may take into account the perceptions of agents around them. In this way, it might be optimal for communities to post-rationalize actions they undertook in an earlier period.

Jann and Schottmuller's model captures the important segregation aspect of echo chambers. However, the authors define the actions that agents undertake in the model extremely loosely. Actions have consequences and future models should address this. To this end, defining actions in such a way as to allow for externalities to arise in the model would be a great improvement. Just because agents choose to segregate into a room does not mean that the ramification of actions they undertake stay inside with them. Furthermore, beliefs are assumed to be static in this model and known to other agents. Incorporating Bayesian updating of beliefs in a multi-period setting also would be a good extension of the model. Finally, Jann and Schottmuller focus on the most informative equilibrium in their model, eliminating equilibriums in which agents spread false information. This is an observed trend in echo chambers and as such should be considered in future work.

Economists have much to contribute in this field. I hope that this summary of existing models and the suggestions I have made for future work here are a part (however small) of that contribution.

References

Akerlof, G. & Dickens, WT. (1982) The economics consequences of cognitive dissonance. The *American Economics Review. 73(3)*, 307-319

Allahverdyan AE, Galstyan A (2014) Opinion Dynamics with Confirmation Bias. PLoS ONE 9(7): e99557. Doi:10.1371/journal.pone.0099557

Arkes, HR. & Blumer, C. (1985) The psychology of sunk cost. *Organizational Behavior and Human Decision Processes, 35.* 124-140
Arrow, KJ. (1996) The economics of information: an exposition. *Empirica, 23* 199-128

Bakshy, E. Messing, S & Adamic, L. A. (2015). Exposure to ideologically diverse news and opinion on Facebook. *Science, 348*(6239), 1130-1132. Retrieved from http://science.sciencemag.org.ezproxy.library.uvic.ca/content/348/6239/1130

Baumgaertner, B. (2014). Yes, no, maybe so: A veritistic approach to echo chambers using a trichotomous belief model. *Synthese,191*(11), 2549-2569. Retrieved from http://www.jstor.org/stable/24019977

Bowden, M. (2014) Information contagion within social networks in the presence of confirmatory bias. *Malaysian Journal of Economic Studies. 51(2).* 151-166

Braman, S. (2005) The micro- and macroeconomics of information. *Annual Review of Information Science and Technology.* 3-52
Cooper, J. (2007) Cognitive dissonance: fifty years of a classic theory. *SAGE Publications LTD.* Retrieved from: http://dx.doi.org.ezproxy.library.uvic.ca/10.4135/9781446214282

Del Vicario et. al (2017) Mapping social dynamics on facebook: the brexit debate. *Social Networks, 50.* Retrieved from: https://reader.elsevier.com/reader/sd/pii/S0378873316304166?token=69DF86991FF1B870B87D6BD07858A348C4E5FC964C281044634A6FF287B008F8469221B2B9986290E7DCA4069DB87102

Dubras, J. & Echenique, F. (2001) Monotone preferences over information. *Cowles Foundation Discussion Paper.* Retrieved from http://www.bvrie.gub.uy/local/File/JAE/2001/iees03j3210701.pdf

Eliaz, K. & Spiegler, R. (2006) Can anticipatory feelings explain anomalous choices of information sources? *Games and Economic Behavior, 56.* 87-104.

Eyster, E. (2002). Rationalizing the past: a taste for consistency. Retrieved from http://www.lse.ac.uk/economics/Assets/Documents/personal-pages/erik-eyster/rationalising-the-past-a-taste-for-consistency.pdf

Gailbraith, JK. (1958) *The affluent society.* New York: Penguin Books

Golman, R. *et al.* (2017) Information Avoidance. *Journal of Economic Literature. 55(1),* 96-135

Herman, E. S., & Chomsky, N. (1988). *Manufacturing consent: The political economy of the mass media*. New York: Pantheon Books

Jann, O & Schottmuller, C. (2018) Why echo chambers are useful. Retrieved from: https://olejann.net/wp-content/uploads/echo_chambers.pdf

Kovach, M. (2015) Twisting the truth: a theory of cognitive dissonance. Available at: https://ssrn.com/abstract=2567080

Moore, C. (2015) Failures in the marketplace of ideas: misinformation, disinformation, and the affordable care act. Retrieved from: https://jscholarship.library.jhu.edu/bitstream/handle/1774.2/38082/MOORE-THESIS-2015.pdf

Napoli, PM. (2018) What if more speech is no longer the solution? First amendment theory meets fake news and filter bubbles. *Federal Communications Law Journal 70(1).* 55-104

Nickerson, R. (1998) Confirmation bias: a ubiquitous phenomenon in many guises. *Review of General Psychology, 2(2).* 175-220

Rabin, M. & Schrag, JL. (1999) First impressions matter: a model of confirmatory bias. *The Quarterly Journals of Economics, 114(1).* 37-82. Retrieved from: https://www.jstor.org/stable/2586947

Quattrociocchi et. al (2016) Echo chambers on facebook. Working paper version. Available at: http://www.law.harvard.edu/programs/olin_center/papers/pdf/Sunstein_877.pdf

Yariv, L. (2001) Believe and let believe: axiomatic foundations for belief dependent utility functionals. *Cowles Foundation Discussion Paper No. 1344.* Retrieved from: http://papers.ssrn.com/abstract=295729

Yariv, L. (2002) I'll see it when I believe it – A simple model of cognitive consistency. *Cowles Foundation Discussion Paper No. 1352.* Retrieved from: http://papers.ssrn.com/abstract=300696