

William J. Berger, Daniel J. Singer, Aaron Bramson, Patrick Grim, Jiin Jung, and Bennett Holman

Philosophical Considerations of Political Polarization

Abstract: This chapter illustrates how philosophy and political science can inform one another by providing an overview of philosophical contributions the authors have made on the topic of political polarization. The authors outline three contributions they have made to understanding political polarization, particularly of the epistemic kind, discussing work that gives clearer terminology for and ways of measuring polarization, precise mechanistic accounts of polarization, and a novel normative view about a possible source of polarization that casts polarization as a possible outcome of rational, but limited, agents interacting. This last contribution illustrates, contra recent work, how dynamics akin to epistemic bubbles and echo chambers can develop without associated epistemic vices. Taken together, these projects can serve as a guide for producing philosophical work which both contributes to a mainstream disciplinary literature as well as informs cross-disciplinary, empirical literatures in the social sciences.

1 Introduction

The study of community fragmentation has a long history in political science (see Duverger 1954; Lipset 1960; and Fearon 2003), though the topic has become more pressing for the study of American politics in recent decades, as the country has experienced growing polarization, and at the same time, rising discord and acrimony (Iyengar 2019). The attention political scientists have given it at both elite and mass levels has led to broad and sustained disciplinary scrutiny, indicated at least in some part by five major review articles on the topic (Hetherington 2009; Layman et. al. 2006; Fiorina and Abrams 2008; Prior 2013; and Iyengar 2019). As this literature has grown, however, so has terminological confusion, as the expanse of papers has come to discuss polarization in different terms and offer different dynamical processes to explain the trends. This chapter illustrates how philosophy and political science can inform one another. We do this by providing an overview of philosophical contributions we have made on the topic of political polarization. We will show how our contributions to social epistemology can help fill gaps and correct inconsistencies in the empirical political science literature. The body of work presented here provides clear exam-

ples of how analytic philosophy can borrow from and inform empirical social science, and political science more specifically. We outline three contributions to understanding political polarization, particularly of the epistemic kind. We will discuss work that gives clearer terminology for and ways of measuring polarization, precise mechanistic accounts of polarization, and a novel normative view about a possible source of polarization that casts polarization as a possible outcome of rational, but limited, agents interacting. This last contribution is particularly interesting given recent developments in a literature around epistemic bubbles and echo chambers (Flaxman et. al. 2016; Jasny et. al. 2015; Levy and Razin 2019; Nguyen 2020; and Santos 2020), which understand echo chambers as due to epistemic vice. As we will discuss further, the work of Daniel J. Singer et. al. (2018) shows polarization can arise without such epistemic vices. By gaining a better understanding of how to discuss and measure polarization, what dynamics might give rise to it, and the normative status of those mechanisms, this piece draws connections between the philosophical and social scientific literatures on social justice and oppression.

The next section illustrates how analytic philosophy can contribute more precisely defined concepts and measures to political science. We offer added resolution to research on polarization by presenting our work that provides new senses and measures of polarization. Section 3 uses the conceptual apparatus from section 2 and looks at how approaches in computational modeling shed light on process-based accounts of polarization. Here, we survey plausible mechanisms that give rise to the kinds of polarization we observe in society. Section 4 then offers an additional model with distinctive normative upshots, arguing that under certain conditions polarization can occur even when everyone is listening to one another in good faith. Section 5 concludes by reflecting on the ways that the philosophical work here can contribute to empirical political science. While by no means exhaustive, these accounts demonstrate how philosophy might be informed by political science and contribute to its empirical process in turn. And since our research primarily contributes to social epistemology, we focus here on polarization that looks at political beliefs rather than the fascinating and influential literature on affective polarization, though the latter is by no means less important.

2 Senses of Polarization

Though Maurice Duverger (1954) argues that American politics cannot not support a multiplicity of parties, the two dominant parties that it did produce were famously thought not to have “a dime’s worth of difference” between

them (the phrase coming from 1968 third-party presidential hopeful George Wallace). That began to change by the 1980s, with Keith Poole and Howard Rosenthal showing diminishing ideological overlap of Members of Congress from Democratic and Republican states (Poole and Rosenthal 1984). While the 1950s and 1960s were marked by considerable legislative common ground, that overlap was slipping away by the late 1970s. Poole and Rosenthal's NOMINATE scores (Figure 1) are one way of measuring partisan drift, providing a standardized location for each congressional member across time, and allowing for a clear articulation of the increasing partisan divide (McCarty et. al. 2006).

Poole and Rosenthal (1984) uses the one-dimensional alignment of congressional votes to estimate elite political drift, but polarization might also be understood differently, say, by looking at the consistency with which disparate political issues get bundled together. This is the approach of Delia Baldassarri and Andrew Gelman (2008), which measures the degree of mass political polarization, operationalizing the concept by assessing the level of correlation between issue pairs. It argues that, in a completely unpolarized population, one's position on one policy issue would not be indicative of one's position on any other—voters' political attitudes would be totally uncorrelated with one another. In a polarized society, however, taking a position on any issue—for or against—would perfectly predict every other political position one holds. Baldassarri and Gelman find that over time there is an increasing correlation between issue pairs in the U.S., which is itself strongly associated with partisanship and ideology. The diminishing heterogeneity of political attitudes among Democrats and Republicans further demonstrates the exacerbating political rift in the U.S.

The work of Poole and Rosenthal along with that of Baldassarri and Gelman are exemplary treatments of an epistemic kind of polarization in political science, among notable others (see Hetherington 2009; Layman et. al. 2006; Fiorina and Abrams 2008; and Prior 2013 for more expansive literature reviews). From even this brief discussion, however, it is evident that identification of polarization can be operationalized in different ways using different measures and conceptual apparatuses. In the absence of careful taxonomy, it is difficult to assess the extent of polarization or establish that accounts are not at odds with one another. The distance between median members of groups is one way of measuring polarization, but intra-group belief uniformity is another. Importantly, these ways of conceptualizing polarization are independent of one another, making it possible for polarization to increase in one sense and decrease in another. The signature elite-level polarization illustrated by NOMINATE scores is itself the confluence of different trends. The ideological middle-ground between the parties evaporates, while partisan extremes increase as well. So, while the fact of polarization is not in dispute, we draw attention to the confusion that can re-




Figure 1: Congressional NOMINATE scores, 1963–2013.

sult from failing to distinguish among a variety of conceptual modes and measures.

Jamie Luguri and Jamie Napier (2013) serves as an example of how muddled discussions of polarization impede research on the topic. In addressing whether abstract or concrete ideation generates polarization, they point to competing evidence on the subject, arguing that though Daniel Yang et. al. (2013) finds abstract thinking decreased polarization, Alison Ledgerwood et. al. (2010) demonstrates that an abstract framing of an issue increases polarization. However, Ledgerwood et. al. (2010) shows that abstract thinking creates greater polarization in one sense (group size parity), while Yang et. al. (2013) shows greater polarization in a completely different sense (distinctness). Such confusion confounds the study of polarization, inhibiting the production of coherent research and illustrating the importance of terminological precision even when just assessing whether or not a given system elicits polarization.

In Aaron Bramson et. al. (2016), we offer a number of distinct senses of polarization in order to enhance conceptual clarity and keep our accounting straight regarding whether polarization, in its various senses, is growing or shrinking. ‘Polarization,’ we argue, acts like an umbrella term, containing differ-

ent elements or kinds, rather than as one monolithic concept. In that paper, we offer a taxonomy of and measurement tools for distinct senses of polarization, tracking each type through empirical cases. The thinking here is that in order to understand polarization we must have conceptual clarity fostered by formal measures of the properties we hope to track. Only then can polarization (of this or that type) be evaluated as either increasing or decreasing.

The first and simplest sense of polarization offered in Bramson et. al. (2016) is ‘spread,’ how far apart the extremes of the distribution are from one another. In the one-dimensional case, where the relevant sense of polarization can be understood by placing people on a line, spread measures how broadly elements (e.g. votes or beliefs) range. As Figure 1 shows, one of the distinctive ways Congress has polarized is through the increased spread of political elites. While Democrats have maintained a left flank around -0.75 for most of the timeseries, the right-most Republicans have moved from a value of 0.75 to greater than 1 over the same period.




Figure 2: Group Divergence measure, with (b) showing more consensus than (a).

‘Group divergence’ and ‘group consensus’ are two further senses of polarization that cleanly pertain to polarization seen in the NOMINATE project. Whereas group divergence measures how far the median members of each group are from each other (Figure 2), group consensus identifies the within-group homogeneity (Figure 3). While spread measures the extremeness of group members, divergence measures the distance of median members of groups, increasing as median members separate, and consensus measures the internal group coherence. An example of consensus is the way party caucuses can deviate from the interests of party leadership. As the ideal points of legislators stray from the parties’ centers (and frequently from party leadership), like members of the Freedom Caucus on the right and Progressive Caucus on the left, the consensus of the groups decreases, thereby decreasing the level of polarization.

‘Distinctness’ is another important conceptualization of polarization (Figure 4). As the overlap between groups diminishes, the level of polarization rises. As indicated above, because in the U.S. Members of Congress have the ability to




Figure 3: Group Consensus measure, with (b) showing more polarization than (a).

vote against party leadership, it was common for some Democrats to vote with Republican leadership and vice versa. As this tendency diminishes, distinctness increases, leaving less overlap between the parties and thus the prospect for compromise. Political parties, particularly in a two-party system like the U.S., are a coalition of divergent interests. In U.S. representative bodies (unlike in parliamentary systems), members are allowed to vote against their party leadership and often do. Party discipline evaluates how regularly party members vote with leadership (Krehbiel 2000). Increasing party discipline is thereby one realization of increased distinctness polarization. As can be seen in Figure 1, the considerable partisan overlap that existed in the 1960s and 1970s has all but evaporated, leaving two distinct and differentiated parties.




Figure 4: Distinctness measure, with (b) showing more polarization than (a).

A further marker of polarization we discuss in Bramson et. al. (2016) is ‘size parity,’ or the membership ratio between two groups (Figure 5). Intuitively, lower parity indicates less polarization. Were all but one individual to be a member of the dominant group, we would take that community to be nearly homogenous. Interestingly, Frances Lee (2016) argues, somewhat counterintuitively, that partisan compromise was easier during the middle of the twentieth century when Republicans tended to be a persistent minority in Congress, since vote trading was needed for Republican members to advance their agenda. Her argument gives

reason to think that diminished size parity has fundamentally affected political gridlock and polarization.

Not all of the measures we outline in Bramson et. al. (2016) require exogenous identification of groups with predetermined constituencies. ‘Community fragmentation,’ for instance, looks to endogenously identify the groups by looking at natural splits in the distribution of the members of a population (Figure 6). Often it is perfectly straightforward to mark groups from the start, like Democrats and Republicans in Congress. However, where one seeks to identify congressional factions, like the Tea Party or the Freedom Caucus, community fragmentation is a useful tool to evaluate the number and boundaries of groups in a population, rather than assume the result.




Figure 5: Size parity measure, with (a) showing more polarization than (b).

We see polarization as an umbrella term, containing distinct and independent senses, the measures of which do not always move in the same direction. Increased coherence can often entail members moving inward, which then diminishes spread (as people pull back from extremes) and induces more overlap. Using respondents’ beliefs from the General Social Survey (GSS), Bramson et. al. (2016) proceeds to show how the distinct measures of polarization vary over time and relative to one another. Religion turns out to be a good domain to illustrate how these concepts come apart. The paper specifically looks at respondents’ belief in the divinity of the Bible as a vivid illustration of this. While there is some correlation between the measures, there also exist clear instances when the measures are anti-correlated.

Analytic philosophy is well suited for this sort of conceptual work. As social scientific literature becomes expansive, it can lose the consistency that initially prevailed. These sorts of definitional inquiries and operationalizations are conducive for philosophers to take on, filling a crucial and treacherous gap in an empirical literature. Though the work in Bramson et. al. (2016) omits certain conceptualizations, such as multi-dimensional approaches like those of Baldassarri and Gelman (2008), it serves as a guide for further clarificatory work, not only on polarization, but in other domains within political science as well.




Figure 6: Community fragmentation measure, with (b) illustrating the latent constituencies of (a).

3 Polarization Dynamics

While Poole and Rosenthal (1984) and Baldassarri and Gelman (2008) are primarily concerned with measuring polarization, political science has also offered causal accounts to explain why it has happened. Larry Bartels (2000) demonstrates how partisanship has played an increasing role in U.S. presidential voting since the 1950s, becoming particularly contentious of late. The middle ground among partisans has slowly slipped away, leading to the disappearance of moderate voters and the rise of party extremism. Partisanship spiked similarly around moments like the passage of the Voter Rights Act and the Vietnam War, but the steady trend observed since the 1980s is distinctive. Motivated partisanship is no longer an exception, but a rule. Geoffrey Layman and Thomas Carsey (2002) find similar trends, showing the widening of American political attitudes over the course of three decades, as well as demonstrating that mass partisan polarization is primarily due to increasing issue salience rather than a party realignment. In support of their hypothesis, they find that on issues of race, social welfare, and morality the divide between partisans has intensified over the last half of the twentieth century.

This is not the only causal story on offer for the rise of polarization, though. McCarty, Poole and Rosenthal (2006) argue that the trend is associated with increased congressional turnover, where younger members take the seats of older members. Similar to Geoffrey Layman and Thomas Carsey (2002), Richard Fleisher and John Bond (2004) and Christian Grose and Antoine Yoshinaka (2003) argue that polarization is linked to the increasing ideological purity of Members of Congress. Others point to historical shocks stemming from the Civil Rights Act and the partisan realignment of Southern Democrats to explain the increased homogeneity within the parties (Hetherington 2009; Rae 1994; and Han and Brady 2007).

Some discussions of political polarization build on social psychological research that offers finer mechanistic accounts of polarization that might well motivate the political dynamics we observe. Notably, Lord et al (1979) shows how biased assimilation of evidence can strengthen belief polarization. Sustained interest in the processes of group and political polarization, however, is more recent. Hetherington and Weiler (2018) have newly made the case that cultural polarization has cleaved Americans right down to the cars they drive and the beers they drink. Relatedly, Hopkins (2018) has shown that polarization is a national phenomenon—Democrats in Georgia and Oregon are more similar than cross-partisans in the same state. Studies such as these provide evidence of mass polarization, indicating that political polarization is occurring through generic pathways that do not directly relate to the structure or incentives of formal political institutions.

Beyond the historical contingencies that induced specific instances of polarization, we might also want to know the dynamical processes at work “under the hood” that brings about these cleavages. Polarization might, for instance, be brought about by “stretching” existing groups or by “pulling” coalitions apart. Philosophy, and computational philosophy in particular, is well suited to explore these mechanisms. What dynamical processes give rise to the kinds of political polarization above? In Patrick Grim et. al. (2012), we provide an illustration of how such dynamical processes might be designed, while, in Bramson et. al. (2017), we offer a more comprehensive look at computational models of polarization, including Robert Axelrod’s dissemination of culture model (1997), Rainer Hegselmann and Ulrich Krause’s bounded confidence model (2002), and Frank Harary’s structural balance project (1959). Bramson et. al. (2017) looks at these three broad categories of models to interrogate distinctive approaches to polarization that have been taken, as well as offering corresponding empirical examples.

Axelrod (1997) offers a multidimensional approach to the dynamics of polarization. In the model, agents are arranged on a grid, each with an identity vector consisting of elements which take values from 0–9. Agents are thought to be close to one another insofar as they share element values in the same position of their vectors. If agent 1 has vector {1, 2, 5} and agent 2 {1, 1, 5}, they share two points of similarity. The model operates by randomly pairing neighbors and allowing for harmonization of agents with disparate features. When similar agents are paired, one alters their vector in order to become more similar to the other, thereby moving towards cultural consensus. So, in this example, agent 1 would change its second element from ‘2’ to ‘1’ to be more similar to agent 2. What generally results, however, is a kind of polarization where distinct and stable communities emerge that have no overlapping traits. Those communities,

however, are not well-balanced, which is to say that they do not have size parity. There is only a narrow range of initial parameter settings which produce moderately sized groups, and even then, the results are not robust.

Because the project is multidimensional, its outcomes do not neatly conform to our measures of polarization. We can clearly see where the senses of group size, distinctness, and coherence are all relevant for interpreting the model's result with the same measures. For other senses, the features of the model may require a new measure in order to capture these concepts of polarization aside from those described here. The Axelrod model did not set out to study polarization specifically, but when we see polarization emerge as it does here, we are led to ask how to describe it. To that end, the Axelrod model is an illustrative example of how the kind of polarization that social scientists observe might emerge.

It is important that a plausible model of polarization not result in an unrealistic picture in which groups move to the absolute extremes: converging to 0 and 1, or the highest and lowest values that the model offers. The reason for this is primarily empirical—the kind of social and political polarization we see in the world does not paint society into opposite corners, but depicts constituencies pulling away from one another. Figure 1's NOMINATE scores illustrate how U.S. political parties have polarized in a number of senses (spread, distinctness, and group convergences, among others), though groups never pool at the poles, but somewhere in between. Such realistic desiderata, which do not force equilibria to extremes, are particularly relevant for political contexts in which the metrics we have—Likert scales of political ideology, NOMINATE scores, or indices like those in Baldassarri and Gelman (2008)—do not document such stark outcomes.

Bounded confidence models aim to produce polarization with this property. Hegselmann and Krause's model (2002, 2005, and 2006) is an example of exactly this sort of effort, observing how polarization might arise without agents eventually coalescing at the poles (Figure 7). The Hegselmann and Krause model produces polarization through local updating, where agents update their beliefs according to those in their peer neighborhood: those who have beliefs that are sufficiently close. Belief updating occurs here by taking a weighted average of the opinions of those that are "close enough" to one's own—where close enough is determined by a social distance threshold ϵ .

Hegselmann and Krause position beliefs on a one-dimensional scale from 0 to 1 and vary the social distance threshold to determine how polarization might take hold. As the model runs, agents weigh their beliefs against those of others in their neighborhood, softening their views to take one more in line with those around them. This produces the kind of pattern seen in Figure 7 where low-level clustering occurs at first, followed by higher level clustering. At threshold values




Figure 7: Example changes in opinion across time from single runs with different threshold values $\epsilon \in \{0.01, 0.15, 0.25\}$. Hegselmann and Krause (2002).

of 0.25 a single consensus opinion emerges, while at a value of 0.01 individuals cluster into a small number of local groups, exhibiting the kind of community fragmentation seen in Figure 6. At intermediate values, however, of say 0.15, two distinct groups emerge from the whole, clustering at beliefs around 0.25 and 0.75. This result is contingent on the number of agents present and their initial position along the 0–1 measure. If there were too few agents or they were non-uniformly distributed along the range, this would alter the likelihood of polarization as well as the final values at which the system converges.

The Hegselmann-Krause model is distinctive in that polarization *increases* over the runs only in the senses of group distinctness and group consensus, whereas the population-based measures (such as spread) as well as community fragmentation and group divergence measures *decrease*. This contrasts with the Axelrod model in which polarization increases across most of the senses. Though the Hegselmann-Krause model produces outcomes that intuitively strike one as highly polarized, they are polarized in only a few of senses. In some ways, it is right to think of the Hegselmann-Krause model as a model of polarization. For

example, the model's distance parameter ϵ might help us consider how narrow political information bubbles produce greater degrees of political polarization. However, in other senses the model creates greater consensus by decreasing spread, for instance, while leaving distinctness unchanged. The model feature that leads to this outcome—"peeling back from the edges," where groups are formed by the most extreme members of the society being pulled inward—does not conform well to observation where we can see that groups often form and repel one another (e.g. Sunstein 2002) or congressional polarization as observed through NOMINATE scores. So though bounded confidence models offer insight into potential dynamics of polarization, they do not robustly correspond to observed instances of political polarization.

The last category of models addressed by Bramson et. al. (2017) is structural balance models. These systems capture a common sensibility that "the enemy of my enemy is my friend" and have recently been employed to assess the dynamics of congressional polarization (Neal 2020). Structural balance models consist of agents on networks. The links between the agents are either positive (if they are friends) or negative (if they are enemies). A network structure is considered balanced whenever all paths—all unique sequences of links connecting each pair of nodes—have an odd number of negative links connecting enemies. In dynamical versions of these models, links in connected groups will rearrange themselves when the structure is unbalanced, and they continue to do so until the structure becomes balanced.

A clear way of assessing polarization in structural balance is by using a measure of community fragmentation, since we can count up the number of stable clusters in the network after the system has become balanced. As we discuss in Bramson et al (2017), other measures can also be employed to assess the system's polarization prior to it becoming balanced. While structural balance models give us an insightful look at the possible dynamics of polarization, they do not tend to produce the kind of polarization that social systems commonly elicit, where groups are divided, but not maximally so.

Taken together the cultural consensus, bounded confidence, and structural balance models offer promising insights into plausible dynamical processes that bring about polarization, potentially shedding light on empirical accounts like those of Layman and Carsey's findings regarding conflict extension and conflict displacement (2002). It is, of course, possible to achieve polarization simply by pulling groups apart from one another, but the models above each provide more sophisticated dynamics for the process. Simpler models would merely force an outcome of two groups in complete opposition, but such a configuration does not conform to our observational evidence of politics.

4 Rational Polarization

The approaches considered so far are predominantly positive or descriptive, looking at how polarization can be understood, measured, and operationalized. A further consideration, which has received less attention in political science, is whether polarization can be the outcome of a rational process. Though political science is primarily empirical, normative work also plays a role in the discipline (e.g. Zaller 1992; Carsey and Layman 2006; and Druckman et. al. 2012). In Singer et. al. (2018), we provide a new model of a potential mechanism of polarization as well as a reason to think that the polarization in this model is produced by fully epistemically rational agents trying to share their information.

Singer et. al. (2018) offers a minimal agent-based model (ABM) of group deliberation to track key elements of groups sharing information. Like Hegselmann and Krause (2002) and Axelrod (1997), Singer et. al. (2018) provides an agent-based computer simulation that considers how simple, lower-order dynamics interact to produce higher-order phenomena. In Singer et. al.'s model, agents are seeded with a set of reasons for or against some proposition. Reasons are modelled as pairs of numerical weights and valences—positive or negative—for or against some matter p . Each time the model is run, a different set of reasons are randomly given to agents. An agent's belief (at any given moment of the model) is whatever their evidence for and against p supports.

The agents deliberate by randomly sharing the reasons they have. When agents have unlimited memory, eventually all the agents converge on the same belief since they all end up having the same reasons. The paper also considers more realistic cases in which agents' memories are constrained so that they can only hold a limited number of reasons. For agents with limited memory, three different memory management strategies are considered where evidence is dropped when agents exceed their memory capacity: one where they randomly forget a reason ('simple-minded'), one where they forget the reason with the lowest weight ('weight-minded'), and one where they forget the lowest weighted reason contrary to the view they support at the time 'coherence-minded'.

The paper measures four senses of polarization in the model output: the time it takes for the group to uniformly agree on a position (for or against p), the divergence of the groups for and against p measured by the distance between group means, the internal subgroup consensuses, and the size parity of the groups. Whereas simple-minded agents, weight-minded agents, and agents with unlimited memory all eventually converge to a single view, coherence-minded agents almost never converge. Moreover, groups of coherence-minded agents typically fracture into evenly sized groups for and against p , and group diver-

gence and subgroup consensus both increase through time (Figure 8). So, groups of coherence-minded agents polarize when they are sharing their reasons for their beliefs.

One might think that being coherence-minded is biased in the same way as Lord et. al.'s (1979) agents who exhibit biased assimilation of new evidence. That is not how Singer et. al.'s coherence-minded agents work, though. Lord et. al.'s subjects are biased because they pre-judge incoming evidence as a function of their previously held beliefs. But coherence-minded agents objectively consider all of their evidence, including the new incoming evidence, before discarding any of it. So, it is not that the new evidence is treated differently or less weighty than it ought to be treated. Coherence-minded agents are doing their best to manage their memory in light of what all of the evidence supports.

Another impetus for thinking that coherence-minded memory management is rational comes from its balancing the epistemic value of an agent having coherent doxastic states against a respect for the weight of their evidence. For agents with limited memories, coherence considerations support favoring evidence that supports the apparently true view and respecting the weight of one's evidence supports keeping the strongest evidence. This leads agents to first forget the weakest reasons in support of apparently (from where they stand) mistaken views.




Figure 8: Histogram of beliefs and strengths for a typical run with coherence-minded agents. Singer et. al. (2018).

Treating evidence in this way leads groups of coherence-minded agents polarize in ways that look more like the kinds of polarization we typically see. Groups form endogenously, without dependence on a tuning parameter (like the value

ε in the Hegselmann and Krause model) and without polarization that crashes into the extreme values of the model.

One upshot of thinking about polarization in the way that Singer et. al. (2018) does is that it focuses us on the distinctive dynamical effects of echo chambers to polarization, similar to extant work in political science (Prior 2007; Stroud 2010; and Tucker et. al. 2018). If Singer et. al.'s mechanism of polarization is real, then we should expect epistemic bubbles to be particularly debilitating to society, since agents will have their memory filled by reasons that support their view, thereby making it very difficult for them to be swayed by contrary evidence. This means that, insofar as we think that polarization harms political discourse, we would have reason to be particularly wary of echo chambers. Moreover, our work highlights the fact that polarization can happen even when rational agents are sharing their information and listening to others, adding to a growing literature on the rationality of polarization (Lodge and Taber 2000; Redlawsk 2002; Sunstein 2002; Taber and Lodge 2006; and Druckman et. al. 2012). Recent work characterizes echo chambers as groups that selectively exclude people and evidence from epistemic communities (Nguyen 2020 and Santos 2020). There are obvious epistemic vices associated with this, and some cases of polarized groups might include echo chambers. As we demonstrate in Singer et. al. (2018) however, polarized groups can also emerge without those epistemic vices. This shows how polarization and echo chambers can come apart and also illustrates how computational models can be used to generate "how possible" explanations with distinctive normative upshots. This finding gives good reason for political science and public policy to reckon more fully with polarization which is not the product of any malfeasance. Rather than take aim only at irrationality or malice in public discourse, we ought also to address how to avoid polarization caused by how we listen to (and remember) what others say.

5 What Philosophy and Political Science and Learn from Each Other

While political science is increasingly an empirical field, philosophy maintains a commitment to theoretical inquiry, suggesting a basic obstacle for interdisciplinary research. Our work presented here, however, provides a sample of philosophical projects that turn to political science for inspiration and motivation, while producing results that are equally of interest for empirical social science as they are for social epistemology and computational philosophy.

The projects presented here offer conceptual, dynamical, and normative traction on the problem of political polarization, specifically germane to instances relating to people's beliefs. Bramson et. al. (2016) provides a series of definitions and measures of different kinds of polarization. Rather than working with a single concept, we argue for conceptualizing polarization as an umbrella term containing distinct kinds, measured differently, and importantly, which can vary independently of one another. As that paper demonstrates, it is both conceptually possible and empirically commonplace for some measures to increase while others decrease. This would be impossible were we to treat polarization as a singular concept.

Using those distinct senses, Bramson et. al. (2017) pursues a more comprehensive account of the dynamical models that have been offered for polarization. The paper considers three broad classes of models—cultural dissemination, bounded confidence, and structural balance—to probe possible underlying dynamics of polarization and the wider applicability of the polarization measures. These classes of models are not meant to be exhaustive, but broadly indicative of the mechanistic approaches to polarization in which the outcome does not merely result from two groups being pulled away from one another. Like empirical instances of polarization, the measurements of polarization in these computational models do not move in lock step, both indicating that no one model can be considered comprehensive and illustrating a more general lacuna in the literature owing to lack of strong correspondence between any of the three models and observed instances of polarization in politics.

Singer et. al. (2018) goes some distance towards offering a better model of polarization, while also interrogating its descriptive and normative rational features. When rational, forgetful agents in this model exchange information they polarize in ways analyzed in Bramson et. al. (2016). Coherence-minded agents—those that manage their memory by forgetting the least consequential reason that does not conform with their prevailing belief—consistently polarize in natural ways and do so for rational reasons. This sort of modeling exercise buttresses the work of Bramson et. al. (2017) to show further cases of model classes that plausibly produce the kinds of polarization observed by political science. Though the potentially rational aspects of polarization have been largely overlooked in the empirical literature, those aspects have clear and demonstrable implications for both how we understand the phenomenon and the policies offered to countermand the trends. Idealized computational models offer a means of providing explanatory accounts of empirical phenomena which would otherwise be difficult to obtain. Singer et. al. (2018) illustrates how polarization can still occur when everyone is behaving rationally and draws attention to the shortcomings of accounts which rule out rational processes.

The work we have done exemplifies a trend in philosophy to look to nearby fields, in this case the social sciences, to motivate philosophical questions that simultaneously offer clarity on how other fields might theorize and operationalize concepts. Analytic philosophy is well suited for taking on definitional, process-oriented, and normative puzzles. A lack of coherent concepts and measurement undermines good science by muddying conclusions and propagating incompatible results. Bramson et. al. (2016) shows not only that there are numerous kinds of polarization, but that a particular case can yield measurements that move in opposite directions. Incorporating political science in the process of developing philosophical research can make philosophy richer and more naturalistic, while offering stronger interdisciplinary pull on empirically oriented disciplines.

Bibliography

- Axelrod, Robert (1997): "The Dissemination of Culture: A MODEL WITH LOCAL CONVERGENCE AND GLOBAL POLARIZATION." In: *JOURNAL OF CONFLICT RESOLUTION* 41. No. 2, pp. 203–226.
- Baldassarri, Delia, and Andrew Gelman (2008): "Partisans without Constraint: Political Polarization and Trends in American Public Opinion." In: *American Journal of Sociology* 114. No. 2, pp. 408–446.
- Bartels, Larry M. (2000): "Partisanship And Voting Behavior, 1952–1996." In: *American Journal of Political Science* 44. No. 1, pp. 35–50.
- Bramson, Aaron, Patrick Grim, Daniel J. Singer, Steven Fisher, William Berger, Graham Sack, and Carissa Flocken (2016): "Disambiguation of social polarization concepts and measures." In: *The Journal of Mathematical Sociology* 40. No. 2, pp. 80–111.
- Bramson, Aaron, Patrick Grim, Daniel J. Singer, William J. Berger, Graham Sack, Steven Fisher, Carissa Flocken, and Bennett Holman (2017): "Understanding Polarization: Meanings, Measures, and Model Evaluation." In: *Philosophy of Science* 84. No. 1, pp. 115–159.
- Druckman, James N., Jordan Fein, and Thomas J. Leeper (2012): "A Source of Bias in Public Opinion Stability." In: *American Political Science Review* 106. No. 2, pp. 430–454.
- Duverger, Maurice (1959): *Political Parties: Their Organization and Activity in the Modern State*. New Jersey: John Wiley & Sons.
- Fiorina, Morris P., and Samuel J. Abrams (2008): "Political Polarization in the American Public." In: *Annual Review of Political Science*. 11, pp. 563–588.
- Flaxman, Seth, Sharad Goel, and Justin M. Rao, Justin (2016): "Filter Bubbles, Echo Chambers, and Online News Consumption." In: *Public Opinion Quarterly* 80. No. S1, pp. 298–320.
- Fleisher, Richard, and John R. Bond, (2004): "The Shrinking Middle in the US Congress." In: *British Journal of Political Science* 34. No. 3, pp. 429–451.
- Grim et. al. (2012): "Philosophical Analysis in Modeling Polarization: Notes from a Work in Progress." In: *APA Newsletters* 12. No. 1, pp. 7–15.

- Grose, Christian R., and Antoine Yoshinaka (2003): "The Electoral Consequences of Party Switching by Incumbent Members of Congress, 1947–2000." In: *Legislative Studies Quarterly* 28. No. 1, pp. 55–75.
- Han, Hahrie, and David W. Brady (2007): "A Delayed Return to Historical Norms: Congressional Party Polarization after the Second World War." In: *British Journal of Political Science* 37. No. 3, pp. 505–531.
- Harary, Frank (1959): "On the Measurement of Structural Balance." In: *Behavioral Science* 4. No. 4, pp. 316–23.
- Hegselmann, Rainer, and Ulrich Krause (2002): "OPINION DYNAMICS AND BOUNDED CONFIDENCE MODELS, ANALYSIS, AND SIMULATION." In: *Journal of Artificial Societies and Social Simulation* 5. No. 3.
- Hegselmann, Rainer, and Ulrich Krause (2005): "Opinion dynamics driven by various ways of averaging." In: *Computational Economics* 25. No. 4, pp. 381–405.
- Hegselmann, Rainer, and Ulrich Krause (2006): "Truth and Cognitive Division of Labor: First Steps towards a Computer Aided Social Epistemology." In: *Journal of Artificial Societies and Social Simulation* 9. No. 3, p. 10.
- Hetherington, Marc J. (2009): "Putting Polarization in Perspective." In: *British Journal of Political Science* 39. No. 2, pp. 413–448.
- Hetherington, Marc J., and Jonathan Weiler (2018): *Prius Or Pickup?: How the Answers to Four Simple Questions Explain America's Great Divide*. New York: Houghton Mifflin.
- Hopkins, Daniel J. (2018): *The increasingly United States: How and Why American Political Behavior Nationalized*. Chicago, London: University of Chicago Press.
- Iyengar, Shanto, Yphtach Lelkes, Matthew Levendusky, Neil Malhotra, and Sean J. Westwood (2019): "The Origins and Consequences of Affective Polarization in the United States." In: *Annual Review of Political Science* 22, pp. 129–146.
- Jasny, Lorien, Joseph Waggle, and Dana R. Fisher (2015): "An empirical examination of echo chambers in US climate policy networks." In: *Nature Climate Change* 5. No. 8, pp. 782–786.
- Krehbiel, Keith (2000): "Party Discipline and Measures of Partisanship." In: *American Journal of Political Science* 44. No. 2, pp. 212–227.
- Layman, Geoffrey C., and Thomas M. Carsey (2002): "Party Polarization and 'Conflict Extension' in the American Electorate." In: *American Journal of Political Science* 46. No. 4, pp. 786–802.
- Layman, Geoffrey C., Thomas M. Carsey, and Juliana M. Horowitz (2006): "Party Polarization in American Politics: Characteristics, Causes, and Consequences." In: *Annual Review of Political Science* 9, pp. 83–110.
- Ledgerwood, Alison, Yaacov Trope, and Shelly Chaiken (2010): "Flexibility Now, Consistency Later: Psychological Distance and Construal Shape Evaluative Responding." In: *Journal of Personality and Social Psychology* 99. No. 1, pp. 32–51.
- Lee, Frances E. (2016): *Insecure Majorities: Congress and the Perpetual Campaign*. London: University of Chicago Press.
- Levy, Gilat, and Ronny Razin (2019): "Echo Chambers and their Effects on Economic and Political Outcomes." In: *Annual Review of Economics* 11, pp. 303–328.
- Liberman, Akiva, and Shelly Chaiken (1992): "Defensive Processing of Personally Relevant Health Messages." In: *Personality and Social Psychology Bulletin* 18. No. 6, pp. 669–679.

- Lipset, Seymour M. (1960): *Political Man: The Social Bases of Politics*. New York: Garden City.
- Lodge, Milton, and Charles S. Taber (2005): "The Automaticity of Affect for Political Leaders, Groups, and Issues: An Experimental Test of the Hot Cognition Hypothesis." In: *Political Psychology* 26. No. 3, pp. 455–482.
- Lord, Charles G., Lee Ross, and Mark R. Lepper (1979): "Biased Assimilation and Attitude Polarization: The Effects of Prior Theories on Subsequently Considered Evidence." In: *Journal of Personality and Social Psychology* 37. No. 11, pp. 2098–2109.
- Luguri, Jamie B., and Jamie L. Napier (2013): "Of two minds: The interactive effect of construal level and identity on political polarization." In: *Journal of Experimental Social Psychology* 49. No. 6, pp. 972–977.
- McCarty, Nolan, Keith T. Poole, and Howard Rosenthal (2006): *Polarized America: The Dance of Ideology and Unequal Riches*. Cambridge: MIT Press.
- McHoskey, John W. (1995): "Case Closed? On the John F. Kennedy Assassination: Biased Assimilation of Evidence and Attitude Polarization." In: *Basic and Applied Social Psychology* 17. No. 3, pp. 395–409.
- Munro, Geoffrey D., Peter H. and Ditto (1997): "Biased Assimilation, Attitude Polarization, and Affect in Reactions to Stereotype-Relevant Scientific Information." In: *Personality and Social Psychology Bulletin* 23. No. 6, pp. 636–653.
- Neal, Zachary P. (2020): "A sign of the times? Weak and strong polarization in the US Congress, 1973–2016." In: *Social Networks* 60, pp. 103–112.
- Nguyen, C. Thi (2020): "ECHO CHAMBERS AND EPISTEMIC BUBBLES." In: *Episteme* 17. No. 2, pp. 141–161.
- Rae, Nicol C. (1994): *Southern Democrats*. New York: Oxford University Press on Demand.
- Redlawsk, David P. (2002): "Hot Cognition or Cool Consideration? Testing the Effects of Motivated Reasoning on Political Decision Making." In: *The Journal of Politics* 64. No. 4, pp. 1021–1044.
- Plous, Scott (1991): "Biases in the Assimilation of Technological Breakdowns: Do Accidents Make Us Safer?" *Journal of Applied Social Psychology* 21. No. 13, pp. 1058–1082.
- Poole, K. T., and H. Rosenthal (1984): "The Polarization of American Politics." *The Journal of Politics* 46. No. 4, pp. 1061–1079.
- Prior, Markus (2007): *Post-Broadcast Democracy: How Media Choice Increases Inequality in Political Involvement and Polarizes Elections*. New York: Cambridge University Press.
- Prior, Markus (2013): "Media and Political Polarization." In: *Annual Review of Political Science* 16, pp. 101–127.
- Santos, Breno R. (2021): "Echo Chambers, Ignorance and Domination." In: *Social Epistemology*, 35, No. 2, pp. 109–119.
- Singer, Daniel J., Aaron Bramson, Patrick Grim, Bennett Holman, Jiin Jung, Karen Kovaka, Anika Ranginani, and William J. Berger. (2019): "Rational social and political polarization." In: *Philosophical Studies* 176. No. 9, pp. 2243–2267.
- Stroud, Natalie J. (2010): "Polarization and Partisan Selective Exposure." *Journal of Communication* 60. No. 3, pp. 556–576.
- Sunstein, Cass R. (2002): "The Law of group polarization." In: *Journal of Political Philosophy* 10. No. 2, pp. 175–195.
- Taber, Charles S., and Milton Lodge (2006): "Motivated skepticism in the evaluation of political beliefs." In: *American Journal of Political Science* 50. No. 3, pp. 755–769.

- Tucker, Joshua A., Andrew Guess, Pablo Barberá, Cristian Vaccari, Alexandra Siegel, Sergey Sanovich, Denis Stukal, and Brendan Nyhan (2018): "Social media, political polarization, and political disinformation: A review of the scientific literature." In: *SSRN Electronic Journal*. <https://dx.doi.org/10.2139/ssrn.3144139>, accessed 28 October 2021.
- Yang, Daniel Y. J., Jesse L. Preston, and Ivan Hernandez (2013): "Polarized Attitudes Toward the Ground Zero Mosque Are Reduced by High-Level Construal." In: *Social Psychological and Personality Science* 4. No. 2, pp. 244 – 250.
- Zaller, John R. (1992). *The Nature and Origins of Mass Opinion*. New York: Cambridge University Press.