# AGREEING TO DISAGREE POLITICALLY

Brad R. Taylor
University of Southern Queensland
Brad.Taylor@usq.edu.au

William Bosworth
London School of Economics

Forthcoming (2020) in *New Political Economy*

*Abstract*

Robert Aumann's agreement theorem and subsequent work shows that people who are rational in a certain Bayesian sense cannot agree to disagree on matters of fact, as long as there is common knowledge of this common rationality. This result hinges on a type of epistemic impartiality: a rational person will not give extra weight to a piece of evidence simply because they themselves discovered it rather than someone else. Of course, "Good Bayesians" who agree on matters of fact can nevertheless disagree politically due to differences of preference and value. We question the possibility of reasonable political disagreement for liberals. On our reading, a "Good Liberal" must not give extra weight in public deliberations to their own preferences or values simply because they are their own. This political impartiality mirrors the epistemic impartiality of Aumann's theorem and we argue that disagreement on policy is impossible in a world of "Good Liberal Bayesians," assuming common knowledge of both Bayesian rationality and Liberal reasonableness. The persistence and predictability of disagreement in the real world provides support for expressive accounts of political behaviour and points to the important role of epistemic trust in politics. This issue of epistemic trust provides insight into recent trends in political polarization in the United States and elsewhere.

## I. INTRODUCTION

Robert Aumann's (1976) agreement theorem and its extensions have shown that two rational actors with common priors and common knowledge of posteriors cannot persistently agree to disagree over factual judgements. This result hinges on a type of epistemic impartiality: a rational person will not give extra weight to a piece of evidence simply because they themselves discovered it rather than someone else. Aumann's work and various extensions of it, we argue, provide a compelling case that in a world of "Good Bayesians" with common knowledge that others are also Good Bayesians, individuals will tend towards agreement in public deliberation over matters of fact. Since this is not always true in the real world, Aumann's finding gives us a place to look for deliberative failings.

Perfectly rational people can differ on values and preferences. But just as Bayesian rationality requires epistemic impartiality, the idea of "political liberalism" championed by Rawls (2005) requires impartiality in political argument. Although people can reasonably differ in their moral beliefs, in order to form a legitimate and stable political association "Good Liberals" must set these differences aside to reach an 'overlapping consensus' on the basic set of rules which are to govern society. Rawls's argument and much of the work building on it has been rather vague on just how impartial a political liberal needs to be. In this paper we argue that the most

plausible version of political liberalism commits us to an extension of Aumann's theorem under which all agreements to disagree over policy is deemed unreasonable by violating the principles of either Bayesian rationality or Rawlsian reasonableness. "Good Bayesian Liberals" thus cannot persistently and predictably disagree over policy in a substantial way.

## II. BAYESIANISM AND EPISTEMIC IMPARTIALITY

Our intention in this section is not to provide a comprehensive review of the literature or a technical presentation of Aumann's result. Rather, we wish to establish the proposition that substantial, persistent, and predictable disagreement on questions of fact among Bayesian rationalists is only possible when at least one party puts undue weight on their own judgments, experiences, or mental ability. As Aumann (1976, p. 1236) admits, the agreement theorem is mathematically trivial once the Bayesian framework is adopted, but the framework and the assumptions upon which the conclusion rests require some explanation.

Aumann (1976) shows that two rational individuals with identical priors, common knowledge of rationality, and common knowledge of posterior probability judgements must have identical posterior judgements. The intuition behind this result is that somebody else's posterior belief provides a summary of all the information they have been exposed to. Suppose that two players are asked to take turns privately flipping the same coin ten times and then publicly estimating the probability that their next toss will land heads. A long series of tosses which landed predominantly on heads would provide some evidence that the coin is biased, and so a player seeing such a series would estimate the chance of the next toss landing heads higher than someone not exposed to that information. However, if the other player by chance landed mostly tails, they would similarly have reason to believe that the coin is biased in the opposite direction. The players would disagree because they have been exposed to different information. However, if public probability estimates are made (and both players believe the other to be honest and rational, as well as being honest and rational themselves) each player can conclude that the other player saw evidence quite different from their own. Since neither player has any reason to favour the evidence of their own eyes over that of that of the other player's, rationality requires that each player take into account the evidence provided by the other's probability judgment and change their own judgement accordingly. Only once all disagreement has been resolved will equilibrium be reached.[1]

Subsequent work has clarified and extended this result.[2] The original theorem was formulated in terms of probability judgments, but this result has been generalised to non-probabilistic beliefs (Bacharach, 1985; Samet, 2010). Aumann's logic applies to any unambiguous factual claim which can be represented in terms of possible worlds.

Agreement will not be reached instantly, however, since once two Bayesians become aware of one another's beliefs they both need to update to account for the disagreement, and then update based on one another's updated belief, and so on. This process of reaching complete agreement might take a very long time indeed, but as long as information is being transmitted the time required will be finite (Geanakoplos & Polemarchakis, 1982). Aaronson (2005) builds on this insight to show when disagreement follows a random walk pattern, approximate

[1] This assumes that agents are Bayesian truth-seekers. In reality, some questions may be deemed so unimportant that any costly effort to answer them correctly is irrational. That is, instrumentally rationally agents can be ignorant or even epistemically irrational (Downs, 1957; Caplan, 2007; Taylor 2020).

[2] See Bonanno and Nehring (1997) for a more detailed, albeit somewhat dated, review of this literature.

convergence can be reached in a reasonable time frame. Specifically, Bayesian agents will agree within $\varepsilon$ with high probability after exchanging around $1/\varepsilon^2$ messages. Crucially, the number of steps required does not depend on the amount of relevant knowledge possessed by each agent, but only the degree of disagreement we are willing to abide. These results, however, assume that such messages transmitted between the agents are the only source of new information; with other new evidence arriving the results no longer hold. A more general result which does allow for new outside information is that of Hanson (2002), who shows that at any point during this process of communication the direction of disagreement will be unpredictable to each participant. That is, an agent's best estimate of the truth and their best estimate of the other agent's future belief must be equal. These can be different *ex post*, but it must be impossible *ex ante* for the agent to predict whether their counterpart will give a higher or a lower estimate.

These results suggest that Bayesians will often move towards agreement over time, and if no new outside information is received they should be able to reach approximate agreement within a reasonable timeframe regardless of how complicated the question is or the extent to which their beliefs initially differ. Moreover, even if new information does arrive, the direction of disagreement should be unpredictable. The existence of some disagreement at any point in time is not troubling from a Bayesian perspective, but *persistent*, *substantial,* and *predictable* disagreement is.

The assumption of common knowledge does a lot of work in Aumann's original formulation. It is not enough for the two parties to know each other's posterior beliefs; they must also know that the other knows their posterior beliefs, know that they know this, and so on. If Alice and Bob know each other's posterior beliefs but Alice doesn't know that Bob knows Alice's posterior belief, Alice might think Bob's judgement is made without all the relevant evidence (i.e. Alice's posterior belief). Subsequent work, however, has shown that common knowledge is not a necessary condition for the result to hold. Common knowledge requires certainty that other parties are rational and honest, that they know our beliefs, that they know we know their beliefs, and so on. Common belief instead requires only that the relevant parties believe these things, even if there's some chance they are mistaken. This opens the door for some disagreement, but the extent is restricted by the level of credence and approximations of common knowledge will approximate the agreement result (Monderer & Samet, 1989).

For the purposes of our argument, the key point to emerge from the Bayesian literature on agreeing to disagree is that disagreement can only persist if there is epistemic partiality – the tendency of people to favour their own beliefs, experiences, or reasoning on the basis that these are their own. There is in general no reason to think that we have a firmer grasp on the truth than any other person. Some people are more informed or better at reasoning than others, but in a world of Bayesian rationalists this would lead to just as many people putting greater weight on others' judgements than their own and the convergence result would persist.

The most contested assumption of Aumann's result is that of uncommon priors. Aumann (1998) sees this as formalising the obvious constraint on rationality that difference in belief must ultimately derive from differences in information. Two people exposed to exactly the same information should have the same beliefs. In reality differences in genes and upbringing will mean this is not the case, but such differences cannot be justified rationally. If nature or society randomly allocates personality traits which alter the way we process information, there is no reason to systematically favour our own priors over others. As Hanson (2006) says, "uncommon priors require origin disputes."

Rational people can come to different factual conclusions because they have different experiences and are thus exposed to different evidence. However, the factual judgement of another person you know to be rational provides a powerful piece of evidence in itself, since it is based on their rational judgement of the evidence they have collected. Unless you have reason to believe that your evidence is better than theirs, you have no reason to privilege your own judgement over theirs. The point relevant to our argument is that disagreement is only possible when Bayesian agents give their own beliefs – whether prior or posterior – priority over the beliefs of others.

## III. POLITICAL LIBERALISM AND AXIOLOGICAL IMPARTIALITY

Political theorists have also struggled with the phenomenon of disagreement. One influential response has been the "political liberalism" literature stemming from the work of Rawls (2005). Rawls seeks to find rules which can be generally accepted by all reasonable people, even when those people have differing religious and ethical beliefs, which Rawls calls "comprehensive doctrines." Although individuals differ in their comprehensive doctrines, a reasonable citizen will not seek to impose their own views on others without justification, and such justification must proceed on the basis of reasons which are acceptable to all other reasonable people, regardless of their own comprehensive ethical and religious views. Thus, liberalism is a non-comprehensive political doctrine which requires an outward-looking form of justification. In order to create a legitimate and stable set of institutions, Rawls insists that reasonable public justification needs to hit on an "overlapping consensus" of basic principles of justice which all reasonable people can be expected to accept.

Political liberals assume a kind of epistemological impartiality similar to Aumann's theorem. Rawls's argument here has some similarity to Aumann's agreement theorem. In both cases a certain type of impartiality leads to a certain agreement among agents who would otherwise disagree. In early versions of this argument, Rawls (1971) uses "the veil of ignorance" as a hypothetical way of inducing impartiality between different conceptions of the good. Individuals are asked to choose the basic principles to govern society while being denied knowledge of their own place in society. The poor, for example, cannot favour a generous social welfare system for selfish reasons since they don't know whether they are rich or poor, talented or untalented.

It is not coincidental that the veil of ignorance has also been used to justify the common prior assumption. Binmore (2007: 395) describes the so-called Harsanyi doctrine,

> "Imagine that a veil of ignorance conceals all the information you have ever received. Harsanyi thinks that ideally rational folk in this state of sublime ignorance would all select the same prior."

Put another way, the posteriors from a prior position of sublime ignorance should be equivalent for rational individuals. The devices used to simulate impartiality for choice conditions also appear to capture the conditions assumed by Aumann's theorem.

While the veil is still important for simulating impartial choice conditions, political liberalism has now turned to the grounds for justifying the veil in the first place with the development of the idea of 'public reason'. Again, this justification is supposed to somehow approximate an impartial form of reasoning. Thomas Nagel has described this impartiality as a kind of epistemological restraint. He says,

> "The idea is that when we look at certain of our convictions from outside, however justified they may be from within, the appeal to their truth must be seen merely as an appeal to our beliefs, and should be treated as such unless those beliefs can be shown to be justifiable from a more impersonal standpoint." (Nagel, 1987: 230)

John Rawls also endorses a similar kind of restraint (Nagel, 1987: 217-8, although see Rawls, 1993: 116). This is precisely the same sort of justification Hanson (2006: 326) gives for the irrationality of uncommon priors. He suggests that

> "If priors are transmitted culturally via children copying visible adults, standard theories about individual variations in such culturally transmitted belief tendencies offer little support for the idea that some children are better able to select the most truth-tracking cultural elements from among the available cultural transmissions."

The likelihood that our priors track the truth better than others likewise requires evidence from, in Nagel's terms, an "impersonal standpoint". In fact, Nagel follows this by claiming

> "it must be possible to present to others the basis of your own beliefs, so that once you have done so, *they have what you have*, and can arrive at a judgment on the same basis."

So the Good Liberal cannot give extra weight to their own values simply because they are their own. This is not to say that Good Liberals will have identical preferences or identical values. The manufacturer and the consumer can reasonably continue to prefer different policies because of their divergent interests; the libertarian and the socialist can reasonably continue to disagree on which policy is morally superior because of their different value judgements. However, Good Liberals will tend towards agreement on the degree to which this consideration is publicly justified as relevant to policy. We can thus distinguish between publicly justifiable political positions and privately held political preferences and values.[3]

Just as Aumann's theorem requires common knowledge of rationality and of posterior beliefs, we can only expect agreement if there is common knowledge of liberalism and common knowledge of honest policy judgements. That is, all of those in the discursive group must be Good Liberals and this must be common knowledge. We can no longer be sure of agreement if some fanatically put extra weight on their own values, if people believe others might be fanatics, or if people think others might think others might be fanatics, etc. Similarly, divergence will be possible if political judgements are not clearly presented to others, or if there is doubt about the clarity of this presentation.

## IV. GOOD BAYESIAN LIBERALS WILL AGREE

People prefer different policies due to differences in factual judgements, preferences, and values. All three sources of disagreement, if persistent, substantial, and predictable, rest on some sort of partiality which can be deemed illegitimate if we accept the Bayesian position of

[3] We are making a distinction between what Harsanyi (1955, p. 315) calls "subjective preferences" (what individuals actually prefer based on their self-interest and any other factors taken into consideration) and "ethical preferences" (what individuals think they ought to prefer if they considered the interests of all affected parties equally). This captures the intuition that we may disapprove of certain states of affairs, while still preferring them for selfish reasons. Ethical preference are less partial than subjective preferences, since they remove one source of bias. However, from the perspective of political liberalism the particular moral assumptions built into ethical preferences remain as a source of partiality.

section II and the liberal position of section III. A Good Bayesian Liberal cannot abide disagreement over which policies are most justified given the preferences and values of those in society. A policy disagreement arising from differing interests is partial in an obvious and familiar sense. A policy disagreement arising from differences in factual judgement implies that some are favouring their own beliefs or judgments over others, or suspect others of doing so. A policy disagreement arising from differences in values implies that some are favouring their own value judgements over others, or suspect others of doing so.

In our specification of impartiality, the reasoning behind a public policy position (as opposed to a private policy preference or value judgement) must be non-indexical in the sense that it does not matter who is doing the reasoning. A manufacturer who benefits from a protectionist policy must agree with a consumer harmed by it if both are being impartial in the sense of appealing only to non-indexical reasons. Similarly, a libertarian and a socialist must agree on the appropriate level of redistribution if both are committed to a fair compromise between different value judgements when it comes to making policy. If not, they must differ on the extent to which they think their own value judgements should be the basis for policy. In other words, at least one must be reasoning in an indexical way and so not relying on reasons which are not acceptable to all Good Bayesian Liberals.

Of course, the real world is replete with persistent policy disagreements. The Good Bayesian Liberal must be in some sense troubled by this, but there is no obvious practical path forward since we know that others are not playing by the rules of reasonable and rational discourse. We cannot move towards the consensus view because there is no consensus. Our argument then, does not provide practical advice on how to behave rationally and impartially in our non-ideal world, but we do think it has important individual and institutional implications for the idea of public reason. On one hand, it reinforces the notion that public judgements must be made impartially on the basis of reasons acceptable to all. On the other hand, the Aumannian way of thinking, and by extension the liberal way of thinking, does not require any actual public justification of beliefs or policy preferences if there is common knowledge that everyone is being a Good Bayesian Liberal. Since the factual or policy judgement represents a summary of all the relevant considerations available to the individual, there is no need for the details to be communicated. Given that we do not live in an ideal world, public justification based on reasons acceptable to all acts as a substitute for common knowledge and has an important role in fostering impartiality.

As we argue in the following section, we see the above analysis as providing evidence that political argument in the real world does involve a regrettable degree of epistemic and axiological partiality. However, it is also important to note that this is not the only possible interpretation. One man's *modus ponens* is another man's *modus tollens,* and our argument may be taken as a *reductio ad absurdum* of the notion of impartiality as we specify it here. Bosworth and Taylor (2020) argue we can make sense of political disagreement among political liberals by treating moral values as indexical statements. Taking this position would help clarify the grounds for reasonable disagreement for political liberals. If it is reasonable to agree to disagree on policy, reasoning must in some sense be indexical – it must be acceptable to favour our own beliefs, interests, or values because they are our own. If this is so, we should ultimately be able to identify the indexical reasons on which a given political disagreement rests by making political statements precise using, for example, the method of elimination (Bosworth, 2016; Chalmers, 2011). If we are successful in removing all vague language, the differences we end up

with must be indexical, and this will reflect either a source of bias which must be removed or the true grounds of reasonable political disagreement.

## V. CONCLUSION: WHY DOES POLITICAL DISAGREEMENT PERSIST?

Before thinking about the implications of the above analysis to political discourse, it is worth clarifying the scope of our argument. Our finding that Good Liberal Bayesians will not agree to disagree when seeking the truth and a fair compromise on policy differences does not preclude agreement to disagree motivated by other concerns such as the avoidance of interpersonal or civil strife. Although our argument suggests that a young progressive activist and her conservative uncle should come to agreement on immigration policy over Christmas dinner, we fully accept that in reality the benefits of such belief convergence for the parties involved are outweighed by the costs of such an exchange in terms of familial harmony. Likewise, an agreement to disagree on basic values among members of a multicultural society may be a prerequisite for peace. It may be instrumentally rational but epistemically irrational to agree to disagree on facts, values, and political positions. Our analysis applies only when the parties to political argument are attempting to be Good Bayesian Liberals.

Although imperfections in information and communication may in the real world allow some disagreement to persist among Good Bayesian Liberals, the nature of disagreement we see in the real world are sharply at odds with the requirements of Bayesian rationality and liberal impartiality (Cowen & Hanson, 2004). Any two people, no matter how apparently committed to truth-seeking and liberal impartiality, will easily be able to find political and factual questions on which they significantly, persistently, and predictably disagree. Such disagreement is not normally seen as cause for concern or embarrassment, and mutual awareness of disagreement very often fails to shift beliefs at all or may even reinforce disagreement. Aumann's restrictive assumptions guarantee absolute agreement among Bayesian agents with common priors, but as we showed in section II more realistic assumptions suggest that disagreement will be modest and should tend to reduce over time (Aaronson, 2005; Cowen & Hanson, 2004; Geanakoplos & Polemarchakis, 1982; Hanson, 2002, 2003; Monderer & Samet, 1989).

The phenomenon of political disagreement has become increasingly important in recent years with the rise of right-wing populism and the resurgence of the democratic socialist left in the United States, not to mention the more extreme fringe movements on both the left and right. When we look at public opinion on particular issues there is no great polarization to be seen. Instead, we've seen a dramatic rise in partyism and affective polarization (Iyengar, Lelkes, Levendusky, Malhotra, & Westwood, 2019; Sunstein, 2015). Those identifying as a Republican or Democrat have come to view their opponents in increasingly negative terms. The 2016 US presidential election saw a dramatic increase in extremely negative ratings of both Trump and Clinton (Christenson & Weisberg, 2019), and tests of implicit association have suggest that partisan bias is now stronger than racial bias in America (Iyengar & Westwood, 2015). In 1960 four percent of Democrats and five percent of Republicans responded that they would be displeased if their son or daughter married someone from the opposing political party; in 2010 this had increased to 33 percent and 49 percent (Iyengar, Sood, & Lelkes, 2012). Again, it appears that this partisan bias is more widespread than racial bias (Sunstein, 2015, p. 5). Such vicious and fundamental disagreement does not simply increase incivility and prevent political compromise; it threatens to undermine the liberal democratic social contract itself. When opposing views are seen not simply as wrong but as corrupt, dangerous, and illegitimate the imperative to accept electoral defeat and peacefully transfer power is called into question.

How are we to explain such disagreement in light of the argument above that Good Liberal Bayesians cannot agree to disagree politically? Obviously, at least one of the assumptions underlying the argument must be seriously incorrect (i.e. not even approximately true). One possibility is that the role of social identity in partisanship is to blame here. Humans are a naturally group-forming animal, and social identity is a strong force (Tajfel, 2010). Even arbitrary group affiliations can have profound effects (Billig & Tajfel, 1973). Party affiliation, particularly in the United States, is an important form of social identity reflecting underlying values and beliefs reproduced regularly through election campaigns (Bartle & Bellucci, 2014; Green, Palmquist, & Schickler, 2002; Greene, 2004).

Social identity can explain political disagreement in two ways. First, individuals may form and express political views not as a way of influencing policy outcomes but rather as a way of signalling group loyalty (Brennan & Hamlin, 1998; Brennan & Lomasky, 1993; Hamlin & Jennings, 2004, 2011; Hillman, 2010; Taylor, 2015) or even to anger and please others (Glazer, 2008). If we accept this hypothesis, political disagreement is easily explained by the low-stakes nature of political choice: it may simply be too costly in terms of attention and cognitive effort to update beliefs which are instrumentally unimportant. The logic of rational ignorance (Downs, 1957; Hindmoor & Taylor, 2015, Chapter 8; Taylor, 2020) could thus be extended to disagreement, and we would have no reason to think voters will be either Liberal or Bayesian in the relevant sense. A second possibility is that partisans are for the most part Good Liberal Bayesians (i.e. they seek the truth and are willing to make compromises on liberal grounds), but they do not trust their ideological opponents to follow the same rules. In Aumann's terms, this would be a violation of the common knowledge assumption. Views will not converge without a general belief that others are politically liberal truth-seekers. If this is the case, the most pressing issue is trust rather than partiality, although of course the fact that people view their co-partisans as more trustworthy than those across the ideological aisle is itself an expression of epistemic partiality. This would lead us to expect convergence within parties or social identities and divergence between them. If each party becomes its own epistemic community within which beliefs are rationally updated, we would see intracommunity but not intercommunity convergence. This is consistent with the recent pattern of party polarization in the United States, in which the major parties have each become more ideologically homogeneous and distant from one another (Layman et al, 2006; Fiorina 2017).[4]

The problem of epistemic trust may have been worsened by recent changes to the media and communication landscape. As Suntein (2018) argues, the internet provides consumers of media with content which more closely matches their personal interests, including their ideological positions. Social media further exacerbates this problem as individuals weave themselves 'information cocoons' which shut out opposing views and amplify their pre-existing biases. They are helped in this by algorithms such as Facebook's news feed and YouTube's system of recommending videos based on past viewing habits (O'Callaghan, Greene, Conway, Carthy, & Cunningham, 2015). As users deny themselves access to opposing voices they lose the information required to update their beliefs rationally. Again, this requires a certain degree of epistemic partiality insofar as Good Liberal Bayesians have no reason to think their side of politics has more direct access to the truth than the other side.

[4] Although the Bayesian approach can make sense of convergence within epistemic communities, it would not predict that the assertion of a political position or belief could lead a Bayesian agent to move in the opposite direction. Since this appears to happen to some extent (see e.g. Samuels & Zucco, 2018) more investigation from a Bayesian approach is required.

Polarization also has the potential to undermine political bargaining. Even without a general commitment to political liberalism, self-interest can produce compromise through bargaining. Here it is useful to consider the work of Buchanan (1975; Buchanan & Tullock, 1962; Brennan & Buchanan, 1985), who sees politics as a complex form of exchange. Just as well-functioning markets with low transaction costs allow buyers and sellers to make mutually-beneficial trades, a representative democracy which allows political bargaining and exchange through log-rolling allows for mutually-beneficial political exchanges. Those with weak preferences could 'sell' their votes to those with strong preferences, and in the absence of transaction costs the equilibrium in the political market would be welfare-maximising. Thus we would effectively see a political analogue of the Coase theorem (Coase, 1960). Applying the same assumptions of rationality, competition, and low transaction costs to government that we typically apply to markets, democracy can be shown to produce efficient results (Wittman, 1989, 1995). Mistrust is an important source of transaction costs in general (Bohnet, Frey, & Huck, 2001; Gambetta, 1990; Lindenberg, 2000). A mistrust of partisan opponents, and particularly opposing political elites, due to dishonesty or irrationality creates barriers to political exchange.

One outstanding question here is how someone who wishes to be epistemically and politically impartial – a Liberal Bayesian wannabe, to adapt Hanson's (2003) phrase – should respond to our argument in practical terms. The preponderance of persistent disagreement in the real world prevents any attempt to move towards consensus for any salient policy question. A naïve interpretation of our argument would result in a Liberal Bayesian wannabe agreeing with whomever they happened to talk to last. This is clearly just as irrational as sticking with a fixed set of beliefs no matter what. Another option is to seek out smaller epistemic communities able to foster greater epistemic trust. The success of this approach, of course, is contingent on the epistemic standards of this community tracking truth and value acceptably. Moreover, if things go wrong such intentional epistemic communities can produce ideological amplification and radicalisation. Indeed, the contemporary alt-right – the customary example of dangerous online radicalisation – has roots in the rationalist and sceptic communities which sought truth by casting off the constraints of political correctness and religious orthodoxy (Rozner, 2018; Torres, 2017).

There is no simple solution to be problem of political disagreement. People are prone to form group identities and to treat outsiders with suspicion. As a normative standard, Bayesian rationality is compelling but as a description of human behaviour, it falls short. The persistence of disagreement, however, makes political liberalism all the more important. If human nature precludes the possibility of political consensus, peace requires compromise and tolerance. Since no epistemic or moral community has a monopoly on truth or justice, the liberal society must respect freedom of association (Kukathas, 2003). However, political liberalism also requires an overlapping consensus on these basic liberal principles. There is a tension here: greater freedom of association allows illiberal groups to isolate themselves from the wider community and reinforces illiberalism (Taylor & Crampton, 2010). In attempting to deal with the consequences of disagreement through decentralization and freedom of association, we may be reinforcing those very social divisions and further undermining epistemic trust.

## VI. REFERENCES

Aaronson, S. (2005). The complexity of agreement. *Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing*, 634–643. ACM.

Aumann, R. J. (1976). Agreeing to disagree. *The Annals of Statistics*, 1236–1239.

Aumann, R. J. (1998). Common priors: A reply to Gul. *Econometrica*, 929–938.

Bacharach, M. (1985). Some extensions of a claim of Aumann in an axiomatic model of knowledge. *Journal of Economic Theory*, *37*(1), 167–190.

Bartle, J., & Bellucci, P. (2014). *Political Parties and Partisanship: Social Identity and Individual Attitudes*. Routledge.

Billig, M., & Tajfel, H. (1973). Social categorization and similarity in intergroup behaviour. *European Journal of Social Psychology*, *3*(1), 27–52.

Bohnet, I., Frey, B. s, & Huck, S. (2001). More Order with Less Law: On Contract Enforcement, Trust and Crowding. *American Political Science Review*, *95*, 131–144.

Bonanno, G., & Nehring, K. (1997). Agreeing to disagree: A survey. *Document Prepared of an Invited Lecture at the Workshop on Bounded Rationality and Economic Modelling*.

Bosworth, W. (2016). An interpretation of political argument , *European Journal of Political Theory*, https://doi.org/10.1177/1474885116659842

Bosworth, W., & Taylor, B. R. (2020) The impossibility of a Bayesian liberal? Unpublished manuscript

Brennan, G., & Buchanan, J. M. (1985). *The reason of rules: Constitutional political economy*. Cambridge: Cambridge University Press.

Brennan, G., & Hamlin, A. (1998). Expressive voting and electoral equilibrium. *Public Choice*, *95*(1), 149–175.

Brennan, G., & Lomasky, L. E. (1993). *Democracy and decision: The pure theory of electoral preference*. Cambridge: Cambridge University Press.

Buchanan, J. M. (1975). *The limits of liberty: Between anarchy and Leviathan*. Chicago: University of Chicago Press.

Buchanan, J. M., & Tullock, G. (1962). *The calculus of consent: Logical foundations of constitutional democracy*. Ann Arbor: University of Michigan Press.

Caplan, B. (2007). *The myth of the rational voter*. Princeton University Press.

Chalmers, D. J. (2011). Verbal disputes. *Philosophical Review, 120*(4), 515–566.

Christenson, D. P., & Weisberg, H. F. (2019). Bad characters or just more polarization? The rise of extremely negative feelings for presidential candidates. *Electoral Studies*. https://doi.org/10.1016/j.electstud.2019.03.008

Coase, R. H. (1960). The problem of social cost. *Journal of Law and Economics, 3*(1), 1–44.

Cowen, T., & Hanson, R. (2004). *Are disagreements honest?* Retrieved from http://www.mercatus.org/uploadedFiles/Mercatus/Publications/Are%20Disagreements%20Honest%20-%20WP.pdf

Downs, A. (1957). *An economic theory of democracy*. Harper.

Fiorina, M. (2017). *Unstable Majorities: Polarization, Party Sorting, and Political Stalemate*. Hoover Press.

Gambetta, D. (1990). *Trust: Making and breaking cooperative relations*.

Geanakoplos, J. D., & Polemarchakis, H. M. (1982). We can't disagree forever. *Journal of Economic Theory, 28*(1), 192–200. https://doi.org/10.1016/0022-0531(82)90099-0

Glazer, A. (2008). Voting to anger and to please others. *Public Choice, 134*(3), 247–254.

Green, D., Palmquist, B., & Schickler, E. (2002). *Partisan Hearts and Minds: Political Parties and the Social Identities of Voters*. New Haven, CT: Yale University Press.

Greene, S. (2004). Social Identity Theory and Party Identification*. *Social Science Quarterly, 85*(1), 136–153. https://doi.org/10.1111/j.0038-4941.2004.08501010.x

Hamlin, A., & Jennings, C. (2004). Group Formation and Political Conflict: Instrumental and Expressive Approaches. *Public Choice, 118*(3/4), 413–435. https://doi.org/10.1023/B:PUCH.0000019908.98720.1e

Hamlin, A., & Jennings, C. (2011). Expressive Political Behavior: Foundations, Scope and Implications. *British Journal of Political Science, 41*, 645–670.

Hanson, R. (2002). Disagreement is unpredictable. *Economics Letters*, *77*(3), 365–369. https://doi.org/10.1016/S0165-1765(02)00149-0

Hanson, R. (2003). For Bayesian Wannabes, Are Disagreements Not About Information? *Theory and Decision*, *54*(2), 105–123. https://doi.org/10.1023/A:1026251404228

Hanson, R. (2006). Uncommon priors require origin disputes. *Theory and Decision*, *61*(4), 319–328.

Harsanyi, J. C. (1955). Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility. *The Journal of Political Economy*, *63*(4), 309–321.

Hillman, A. (2010). Expressive behavior in economics and politics. *European Journal of Political Economy*, *26*(4), 403–418. https://doi.org/16/j.ejpoleco.2010.06.004

Hindmoor, A. M., & Taylor, B. R. (2015). *Rational choice* (2nd ed.). Basingstoke: Palgrave Macmillan.

Iyengar, S., Lelkes, Y., Levendusky, M., Malhotra, N., & Westwood, S. J. (2019). The Origins and Consequences of Affective Polarization in the United States. *Annual Review of Political Science*, *22*(1), 129–146. https://doi.org/10.1146/annurev-polisci-051117-073034

Iyengar, S., Sood, G., & Lelkes, Y. (2012). Affect, not ideology: A social identity perspective on polarization. *Public Opinion Quarterly*, *76*(3), 405–431.

Iyengar, S., & Westwood, S. J. (2015). Fear and loathing across party lines: New evidence on group polarization. *American Journal of Political Science*, *59*(3), 690–707.

Kukathas, C. (2003). *The liberal archipelago: A theory of diversity and freedom*. Oxford: Oxford University Press.

Layman, G. C., Carsey, T. M., & Horowitz, J. M. (2006). Party polarization in American politics: Characteristics, Causes, and Consequences. *Annual Review of Political Science*, *9*(1), 83–110.

Lindenberg, S. (2000). It Takes Both Trust and Lack of Mistrust: The Workings of Cooperation and Relational Signaling in Contractual Relationships. *Journal of Management and Governance*, *4*(1), 11–33. https://doi.org/10.1023/A:1009985720365

Monderer, D., & Samet, D. (1989). Approximating common knowledge with common beliefs. *Games and Economic Behavior*, *1*(2), 170–190. https://doi.org/10.1016/0899-8256(89)90017-1

O'Callaghan, D., Greene, D., Conway, M., Carthy, J., & Cunningham, P. (2015). Down the (white) rabbit hole: The extreme right and online recommender systems. *Social Science Computer Review*, *33*(4), 459–478.

Rawls, J. (1971). *A theory of justice*. Cambridge, MA: Harvard University Press.

Rawls, J. (2005). *Political liberalism*. New York: Columbia University Press.

Rozner, G. (2018). Inside the intellectual dark web. *Institute of Public Affairs Review: A Quarterly Review of Politics and Public Affairs, The*, *70*(3), 6.

Samet, D. (2010). Agreeing to disagree: The non-probabilistic case. *Games and Economic Behavior*, *69*(1), 169–174.

Samuels, D. J., & Zucco, C. (2018). *Partisans, Anti-Partisans and Non-Partisans: Voting Behavior in Brazil*. Cambridge University Press.

Sunstein, C. R. (2015). Partyism. *University of Chicago Legal Forum*, *2015*, 1–27.

Sunstein, C. R. (2018). *# Republic: Divided democracy in the age of social media*. Princeton University Press.

Tajfel, H. (2010). *Social Identity and Intergroup Relations*. Cambridge University Press.

Taylor, B. R. (2015). Strategic and expressive voting. *Constitutional Political Economy*, *26*(2), 159–170.

Taylor, B. R. (2020). The psychological foundations of rational ignorance: Biased heuristics and decision costs. *Constitutional Political Economy*, *31*(1), 70–88.

Taylor, B. R., & Crampton, E. (2010). *Anarchy, preferences, and robust political economy*. Presented at the Australasian Public Choice Society Conference, Christchurch.

Torres, P. (2017, July 29). How did "new atheism" slide so far toward the alt-right? Retrieved July 8, 2019, from Salon website: https://www.salon.com/2017/07/29/from-the-enlightenment-to-the-dark-ages-how-new-atheism-slid-into-the-alt-right/

Wittman, D. (1989). Why democracies produce efficient results. *The Journal of Political Economy*, *97*(6), 1395–1424.

Wittman, D. (1995). *The myth of democratic failure: Why political institutions are efficient*. Chicago: University of Chicago press.