




ARTICLE

Bullshit in Politics Pays

Adam F. Gibbons 

Department of Philosophy, Lingnan University, Hong Kong SAR
Email: adamgibbons@ln.edu.hk

(Received 13 July 2021; revised 12 November 2022; accepted 8 December 2022)

Abstract

Politics is full of people who don't care about the facts. Still, while not caring about the facts, they are often concerned to present themselves as caring about them. Politics, in other words, is full of bullshitters. But why? In this paper I develop an incentives-based analysis of bullshit in politics, arguing that it is often a rational response to the incentives facing different groups of agents. In a slogan: bullshit in politics pays, sometimes literally. After first outlining an account of bullshit, I discuss the incentives driving three different groups of agents to bullshit: politicians, the media, and voters. I then examine several existing proposals to combat bullshit in politics, arguing that each will fail because they ignore the relevant underlying incentives. I conclude somewhat pessimistically that a certain amount of bullshit in politics is inevitable.

Keywords: Bullshit; fact-checking; fake news; incentives; political ignorance; political irrationality; social media reform

Chaos, rage, headlines front page,
Break out the spotlight, bullshit centre stage. (Pase Rock)

Introduction

Politics is full of people who don't care about the facts. They vote in elections, run for office, report the news, write opinion pieces, share clickbait, and more. Rather than caring about the facts, they're content instead to merely present themselves as caring. But why do politicians play fast and loose with the facts? Why are voters so often vocal on matters they know so little about? Why do people create (and consume) fake news? In short, why is there so much *bullshit* in politics?

In this paper I develop an incentives-based analysis of bullshit in politics, arguing that it is often a rational response to the incentives facing different groups of political agents. In a slogan: bullshit in politics pays, sometimes literally. This has important implications, both for how to understand bullshit in politics as an empirical phenomenon and how to address it as a practical problem. As we'll see, existing interventions to reduce the amount of bullshit in politics, as well as similar proposals to combat the spread of related phenomena such as fake news, fail to recognize the extent to which

© The Author(s), 2023. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

it is a product of widespread incentives.¹ This failure renders the success of such interventions highly unlikely. Indeed, these interventions might even exacerbate the problems they're supposed to fix.

I begin in section 1 by first outlining what I mean by 'bullshit'. Although there are many different things we pick out in ordinary language with this expression, I restrict my attention to a phenomenon wherein people communicate without regard for the truth of what they communicate.² Next, in section 2, I discuss the sort of incentives that commonly drive three different groups of political agents to bullshit, focusing on politicians, the media, and voters. After highlighting the many ways in which it can be *individually* rewarding to bullshit, I explain how bullshit in politics can be *collectively* harmful. Among other things, the more agents in politics bullshit, the more likely the spread of harmful misinformation. This sets the stage for section 3, where I examine several existing proposals for how to combat bullshit in politics: reliance upon personal responsibility, media literacy training, fact-checking organizations, social media reform, and governmental regulation. Each is likely to fail, I argue, since they ignore the underlying incentives that drive agents to bullshit in the first place. I conclude in section 4 on a pessimistic note. The incentives to produce bullshit emerge whenever large groups of agents with conflicting aims and ambitions, and each with finite cognitive and attentional resources, interact with one another. Accordingly, a certain amount of bullshit in politics is inevitable.

1. What is Bullshit?

In ordinary language, people use the expression 'bullshit' in many ways. For instance, speakers often use 'bullshit' *dismissively*. We *call bullshit* when we want to dismiss or reject what another says (Bergstrom and West 2020). We're also happy to call habitual liars bullshitters *because* of their habitual lies, thus viewing bullshitters as a kind of liar. In this paper, however, bullshit and lying are understood as distinct (though not mutually exclusive) phenomena.³ Lies characteristically involve an agent uttering something they believe to be false in order to get their addressee to believe the content of the utterance.⁴ Bullshit, though, involves agents communicating without regard for the truth of

¹There are exceptions to this general trend in the literature, with some recognizing the extent to which incentives drive agents in politics to behave as they do. However, while these analyses nicely capture many of the relevant incentives, they suffer from some important drawbacks. First, some do not canvass practical interventions to reduce the spread of the relevant phenomena whatsoever (Mathiesen and Fallis 2017). Second, those that do discuss practical interventions fail to consistently apply the underlying incentives-based analyses to the practical interventions they propose (Tullock 1972; Schauer and Zeckhauser 2009). As a result, they fail to consider the possibility that these interventions, which rely on the good faith and reliable epistemic conduct of agents who implement them, are themselves compromised because they create incentives for agents to lie about lying, bullshit about bullshit, and so on. This possibility is discussed at length in section 3.

²Throughout the paper I use locutions like 'regard for the truth' and 'regard for the facts' interchangeably.

³This account thus contrasts with Frankfurt's (2005 [1986]) classic account outlined in *On Bullshit* (which understood bullshit and lying as different and *mutually exclusive* phenomena), but it meshes well with Frankfurt's later work according to which lying and bullshitting can overlap (Frankfurt 2002). For the purposes of this paper, I remain agnostic about the extent to which these phenomena can overlap, though see Carson (2010: 46–64) and Stokke (2018: 162–70) for related discussion.

⁴Some philosophers argue that the intent to deceive is inessential for lying. One can, they claim, engage in *bald-faced lying* while knowing that one won't deceive one's audience. For relevant discussion, see Carson (2006), Sorensen (2007), Saul (2012: 8–10), and Mahon (2015).

what they communicate. So understood, lying and bullshitting can overlap – someone can tell lies *while* being fundamentally indifferent towards the truth – but they are nonetheless distinct and often pull apart. One can lie without bullshitting, and one can bullshit without lying.

To illustrate the difference, consider two competing politicians running for office, Jessie and James. Both are willing and able to use underhanded tactics to gain advantage over the other. Predictably, both want to diminish the standing of their opponent in the eyes of the electorate. To do this, Jessie fabricates vicious rumors about James' personal life, all while believing them to be false. When she tweets about his conduct to her millions of followers, her tweets contain lies. (Indeed, her statements are lies even if, *pace* her expectations, the rumors turn out to be true.)⁵ James, meanwhile, does not fabricate rumors about Jessie, opting instead to opportunistically share rumors about her that were already circulating online. Importantly, he knows neither their origin nor whether they are true. In fact, he does not care to know. The truth of these rumors is of secondary importance to their usefulness as a political weapon. In short, James is a bullshitter.⁶

Several additional features of this account of bullshit are worth clarifying. First, when thinking about bullshit, the focus is on how an agent's mental state relates to what they communicate rather than the truth or falsity of what they communicate. Bullshitters can sometimes say true things. But when they get things right, they do so accidentally, for the truth is not something that concerns them. In other words, something (whether an utterance, written statement, news report, and so on) counts as bullshit because of the attitude(s) of the agent(s) who produced it and not because it is false, misleading, or otherwise dubious (although much of what bullshitters communicate will in fact be false, misleading, or otherwise dubious).

Second, when using “communicate without regard for the truth of what one communicates”, I intend to pick out two different ways of relating to the truth. On the one hand, one might not care whatsoever about the truth of what one communicates. On the other hand, one might *doubt* the truth of what one communicates and yet still choose to communicate it (Mukerji 2018: 934). Additionally, the comprehensiveness of one's indifference towards the truth can vary. For instance, some agents might not care about the truth *in general*, while others might not care about the truth of a specific claim (or series of claims).⁷ In this paper, the focus is on agents in politics who lack regard for the truth of specific claims (or series of claims), whether this means they do not care whatsoever about the truth of such claims or they doubt such claims while communicating them nonetheless.⁸ Such agents may subsequently end up caring

⁵Some may resist this claim. Although many philosophers think one can lie even while uttering a truth, one might think that such utterances instead constitute *attempted lies* (Cappelen and Dever 2019: 42). For defenses of the claim that one can lie while uttering truths, see Faulkner (2007) and Stokke (2018). Since this issue is unimportant for the purposes of this paper, I set it aside. For more on lying in politics, see Tullock (1972: Ch. 9), Bok (1978: Ch. 12), Weissberg (2004), Osborne (2005), Jay (2010), and Mearshamer (2011).

⁶It is worth noting that given her goal to discredit James, Jessie would arguably share rumours about him even if she had not fabricated them herself and neither knew nor cared about their accuracy. If so, she may be most fundamentally a bullshitter (Frankfurt 2002: 341). Still, as it stands, her case involves lying while James' involves bullshitting. Thanks to an anonymous referee for helpful discussion on this point.

⁷Also consider cases of the sort discussed by Stokke wherein one has a general concern for the truth but no concern for any particular truths (Stokke 2018: 140–1).

⁸It is plausible that very few agents completely disregard the facts. More likely, agents care very much about some facts, care somewhat about others, not at all about still others, and so on. For instance,

about the truth of what they communicate in the sense that they go on to hope, for whatever reason, that they get things right. But at the time of utterance (or of writing, etc.), bullshitters lack regard for the truth of what they communicate insofar as they either possess no attitude whatsoever towards the truth or they do not care enough to make any effort to dispel whatever doubts they may feel.

Third, while other accounts of bullshit focus on *saying* (or *uttering*) things without regard for their truth, I focus on the much broader phenomenon of *communicating* without regard for the truth of what one communicates. This is largely because political bullshit is *multimodal*. It can be spoken, written, signed, tweeted, and more. As the example involving James shows, the bullshit one communicates often has its origin elsewhere, and needn't involve the bullshitter *saying* anything. One can instead *transmit* a message of some kind first produced by another agent. Indeed, what is transmitted by a bullshitter may be something both true and first produced by someone who cared about the facts. To return to the above example, James is still bullshitting even if the incriminating stories about Jessie are true, having been first uncovered by a whistle-blower greatly concerned with their veracity.

Of course, bullshit so understood doesn't exhaust the range of phenomena that may aptly be described as 'bullshit'.⁹ This paper offers no definition in terms of severally necessary and jointly sufficient conditions, and other accounts aiming for a more comprehensive picture of bullshitting in all its varieties may pick out other phenomena. But the frequency with which such bullshit appears in politics (as well as its possible harms) warrants special attention. In politics, with so much at stake so often, a disregard for the facts can have serious consequences. Why, then, do so many people in politics have such a casual relationship with the truth?

The explanation defended in this paper is that such behaviors are often an *instrumentally rational* response to the incentives political agents face.¹⁰ Instrumentally rational agents adopt means that, given their beliefs, are suitable to their ends (Kolodny and Brunero 2020). The more suitable the means, the more rational they are, even if the ends strike us as outlandish or outright immoral. For example, a fact-checker who cares about properly discharging her duties acts rationally when she engages in painstaking research on the relevant subject, but so too does the fact-checker who, because she is too deeply involved in party politics to be unbiased, selectively fact-checks in ways flattering to her preferred party. While one warrants praise and the other blame, both act rationally in the intended sense.

Going forward, I will write often of the *costs* and *benefits* of bullshit. This should always be read as elliptical for costs and benefits that are relative to an agent's preferences. Benefits are outcomes that are broadly consistent with one's preferences, while costs will hinder the satisfaction of one's preferences. Since agents frequently possess different preferences, what constitutes a cost for one agent might constitute a benefit

James may not care whether the incriminating stories about Jessie are true, but he presumably cares a great deal about whether propagating these stories is sufficiently damaging to Jessie's campaign. Facts about whether her campaign will indeed be damaged, then, are important to him (largely because it is in his interest to know such facts).

⁹Other phenomena have been discussed by philosophers under the heading of 'bullshit'. For example, see Cohen (2002), Ivankovic (2016), Sarajlic (2019), and Heffer (2021). See especially Fallis and Stokke (2017), who argue that bullshitting is compatible with *not* lacking a regard for the facts.

¹⁰Cappelen and Dever gesture in this direction (2019: 55). He correctly notes that people sometimes bullshit to achieve something. This paper deepens and refines this suggestion, while also showing how the rationality of bullshit in politics complicates efforts to reduce its prevalence.

for another. Continuing our previous example, the biased fact-checker faces costs when forced to fact-check impartially. Relatedly, she benefits to the extent that she is free to act in a biased manner. In contrast, the dutiful fact-checker faces no such costs.

Seen this way, the decision to bullshit is just one strategy among many that political agents can choose in order to minimize their costs (or maximize their benefits), with political agents often deliberately bullshitting as a means to their varied ends.

2. Rational Bullshit

In this section I explore the incentives driving three groups of political agents to bullshit: politicians, the media, and voters. Though the details differ in each case, a common pattern emerges – namely, agents typically have a plurality of ends with which truth-directed inquiry competes. Given these competing ends, the costs of caring about the truth frequently outweigh the benefits. Moreover, merely presenting oneself as caring about the truth of what one communicates is significantly less costly than sincerely caring. Accordingly, a certain degree of bullshit in politics is rational.

2.1. Bullshitting Politicians

Let's begin with politicians. The claim that politicians are occasional bullshitters is not a hard sell, and cases of bullshitting politicians are likely what springs to mind for most when thinking about political bullshit. Still, it is worth asking why politicians are often such flagrant purveyors of bullshit.

In general, politicians must decide between competing courses of action when deliberating about how to pursue their ends. If they are thinking strategically about how to act, they will assess the overall costs and benefits of these competing courses of action to determine which will, in expectation, provide the most benefits and the least costs. Some campaign strategies, for instance, may offer greater expected benefits than others. For another, entering some coalitions rather than others may yield greater benefits. Decisions regarding how to communicate are no different. Communicating in a clear, honest, and well-informed manner might bring benefits that are inferior to communicative strategies that do not require one to be clear, honest, or well-informed. Accordingly, a politician deciding whether or not to care about the truth of what they communicate must determine the overall costs and benefits from either course of action.

Consider, then, the costs and benefits of *not* bullshitting (that is, the costs and benefits of caring about whether one communicates the truth). The most significant and common costs are related to the acquisition of information itself. Simply put, actively attempting to get things right is in many cases an extraordinarily demanding endeavor, bringing with it steep *procedural costs*. In politics one often grapples with complex issues where the facts are difficult to ascertain, and where significant time and effort must be devoted in pursuit of the truth. In short, exerting such effort is costly, and these costs push politicians in the direction of caring less about the facts than they would otherwise. Moreover, such efforts have high opportunity costs, for many politicians will have aims not best served by undertaking them. Time spent devoted to truth-directed inquiry is time that could be spent on other, more valuable things (Hardin 2009: 2–3).

Naturally, communicating the truth often either yields no benefits or actively harms politicians. Telling the electorate something that they don't want to hear – for instance,

that their preferred economic policies are in fact bad for the economy – is a risky choice for self-interested politicians seeking (re)election, even if what they say is true. Of course, politicians looking to avoid such electoral costs are also incentivized to engage in more straightforward lying; simply tell the electorate what they want to hear, even if you believe otherwise. But one can tell the electorate what it wants to hear without even knowing or caring about the truth. Such disregard for the truth is the hallmark of bullshit, not lying. In short, then, both caring about and communicating the truth can be costly.

The costs of bullshitting are notably different. At first glance, one might think that bullshit has very few procedural costs, taking no more effort to engage in than the effort required to communicate in general. But this is inaccurate. For example, some bullshit is produced with the provocation and excitement of voters in mind (Jacquemet 2020). Consider again the case involving Jessie and James. James may attempt to rile up his supporters as much as possible by packaging the relevant rumors in a way that angers them, motivates them to continue opposing Jessie, and the like. But for politicians like James to know what animates their base, they must first expend some effort. Nonetheless, the procedural costs of determining what animates one's base are in many cases lower than the costs associated with doggedly pursuing the many facts relevant to policymaking and the like. The latter requires one to seriously engage with multiple, complex policy-relevant fields (or to delegate such engagement to subordinates), while the former will sometimes require only a series of well-executed public opinion polls, or to possess some general knowledge of the sorts of key issues that are reliably emphasized by voters (such as the economy, healthcare, or education). Even *strategic* bullshit designed to gain favor with the electorate, then, frequently has lower procedural costs than not bullshitting.

Politicians may suffer reputational harms if they are known bullshitters. It's plausible to think that such behavior will frustrate some of the electorate, and this frustration will in many cases translate to electoral setbacks for the relevant politicians. With that said, voters will likely level charges of bullshitting in a partisan manner, consistent with the level of partisan bias and hostility we find in politics more generally.¹¹ Voters themselves will often bullshit about whether politicians are bullshitting, not caring about whether the politicians they support are actual bullshitters, or about whether the politicians they dislike are engaged in honest pursuit of the truth. This suggests that the reputational harms associated with being labelled a bullshitter arise to a large extent regardless of what one does. In other words, one's reputation may be harmed even if one is falsely believed to be a bullshitter, and such reputational harms stem in part from simple partisan bias.

Additionally, the extent to which one suffers reputational harms because of one's bullshitting may vary with one's ability to bullshit in a convincing manner. Some politicians may lack the skill to bullshit in ways that accurately mimic the behavior of agents sincerely interested in the facts, while others may be better able to bullshit in non-obvious ways. Reputational harms can thus be more readily reduced by members of

¹¹Indirect evidence for this claim comes from Robert Michael and Brooke Breaux, who demonstrate that people's political affiliation influences which sources of news they believe are fake news (Michael and Breaux 2021). If partisan biases influence which sources of news people classify as fake news, it's plausible that the very same biases will influence their beliefs about similar phenomena such as which politicians (and political parties) are guilty of propagandizing, which political figures are liars, who is and isn't a bullshitter, and so on.

the latter group, skilled as they are at presenting themselves as concerned with the truth, while members of the former group need to tread more carefully.

Moreover, since one's base will often not care if one is a bullshitter, politicians can safely benefit from a certain amount of bullshit. Free from a commitment to getting things right, one can speak for all sorts of reasons. Politicians can engage in grandstanding that makes them appear attractive to voters (Tosi and Warmke 2016, 2020).¹² They can save face when asked difficult questions, speaking *as if* they possess expertise on complex matters with which they lack acquaintance entirely, safe in the knowledge that their supporters will either not know or not care about the bullshit.¹³ They can evade questions they would prefer not to answer (Carson 2016: 56–7). They can disparage and denigrate their political opponents without the need to find something worth disparaging. One can bullshit for a host of reasons, and it will often be beneficial and easy to do so. When paired with the high costs and low benefits of not bullshitting, it all too often becomes rational for politicians to bullshit.

2.2. Bullshitting Media

Let's now consider the extent to which the media bullshits rationally, whether traditional mass media connected to large and powerful news networks or smaller online platforms. First, media platforms can sometimes function as conduits for state-manufactured bullshit. Simple cases where the media is directly controlled by the state collapse the distinction between bullshitting politicians and bullshitting media, with the rational bullshit of the state being the rational bullshit of the media. If the media is a mere extension of the state, we can assume that many of the underlying incentives driving media bullshit are identical to those outlined in the previous section. News reports may be designed to maximize the provocation of citizens rather than to accurately present the facts, headlines written to discredit political opposition, propaganda produced to control the citizenry, and the like (Stanley 2015: 125–77).

Different cases arise when we consider more indirect mechanisms by which the media comes to distribute the state's bullshit. In some cases, the media might be coerced to spread bullshit at the behest of the state because of its control over licencing and regulation of interest to the media (Chomsky and Herman 1988: 13). Capitulating to the state's demands, though often regrettable, is instrumentally rational for agents if it helps them to avoid sufficiently high costs (in this case, the potential loss of relevant licences). In this case, once more the bullshitters are agents of the state since they are the source of the content. Additionally, though, media outlets may also lack regard for the truth of what they spread, being primarily concerned with acceding to the state's demands. If so, the media simply bullshits about state-manufactured bullshit.

¹²This is not to imply that caring about the facts and grandstanding are inconsistent. One can communicate with the intention of elevating one's social standing while at the same time caring about the relevant facts.

¹³Frankfurt articulated this motivation to bullshit nicely when he wrote that “[bullshit] is unavoidable whenever circumstances require someone to talk without knowing what he is talking about. Thus, the production of bullshit is stimulated whenever a person's obligations or opportunities to speak about some topic exceed his knowledge of the facts that are relevant to that topic. This discrepancy is common in public life, where people are frequently impelled – whether by their own propensities or by the demands of others – to speak extensively about matters of which they are to some degree ignorant” (Frankfurt 2005 [1986]: 63). Thanks to an anonymous referee for reminding me of this insightful passage.

Other cases lack this coercive element, with the state and the media voluntarily maintaining close ties while being nominally independent. Paradigmatic instances of such voluntary ties involve members of the media and agents of the state forming mutually beneficial relationships. For example, members of the media may *capture* whatever regulatory bodies are tasked with regulating their conduct, shaping legislation to their ends.¹⁴ But the direction of influence may run in both directions as closer ties are established between the pair. If the price of favorable legislation is the dissemination of bullshit, many will be willing to pay if the benefits of favorable legislation outweigh any relevant costs.

Of course, the media does not distribute political bullshit solely at the behest of the state. There are possibly cases where media platforms share certain content at the urging of corporate sponsors, financial partners, and other private entities while not caring whether the stories are accurate. Such content may be tailored to depict the relevant private organizations in a flattering light, to spread misinformation that politically benefits these private organizations, or for a host of other reasons.¹⁵ Whatever the case, expected financial costs from a failure to acquiesce to the demands of financial partners and sponsors might incentivize self-interested members of the media to treat the facts as of secondary importance. Similarly, and much like earlier cases involving the state and the media, such bullshit may be entirely collaborative, involving parties in mutually beneficial partnerships.

Naturally, the media often bullshits about politics even when facing no pressure from external sources. Financial incentives loom large in explaining the propensity of various media platforms to bullshit in different ways.¹⁶ These financial incentives can exist for media organizations in the absence of any ties to other groups, or even in the presence of pressure from other groups to *not* produce such bullshit. Consider first some incentives related to the differential costs and benefits of consistently aiming for accuracy as against bullshitting. Accurate reporting is a labor-intensive, costly endeavor that will not always yield compensatory benefits since potential customers will not always pay to read the boring truth. Eye-grabbing clickbait, on the other hand, is comparatively easy to generate, even if it's wildly inaccurate. If producing clickbait without regard for its truth is what sells, and if media outlets are sometimes concerned with maximizing profits, media outlets will sometimes produce bullshit in the form of clickbait.¹⁷

Closely related concerns revolve around *fake news*. Some of what people call fake news is bullshit – stories produced without regard for their truth that also happen to be substantially false.¹⁸ Some of this fake news is produced by instrumentally rational agents seeking to maximize profits (Rini 2017: 45). Much like 'bullshit', though, 'fake news' is used in many ways by different speakers. Some fake news involves lying,

¹⁴On capture more generally, see Stigler (1971), Posner (1974), Guerrero (2014), and Lindsey and Teles (2017).

¹⁵Some scholars distinguish between *misinformation* (roughly, the communication of false or otherwise misleading information) and *disinformation* (roughly, the *intentional* communication of false or otherwise misleading information) (Benkler *et al.* 2018: 6; Brown 2021: 3). Since both can constitute bullshit, I set this distinction aside in what follows.

¹⁶See McBrayer (2021: 24–39) for insightful discussion of this point.

¹⁷Sometimes, of course, such clickbait is generated with intentional deception in mind. Some of these cases will involve straightforward lying rather than the sort of bullshit that concerns us here.

¹⁸Although see Gelfert (2018) for an account of fake news where bullshit is *not* a defining characteristic. Interestingly, Mukerji (2018) claims that fake news is *always* bullshit.

some of it involves non-bullshitting, non-lying misleading, and so on.¹⁹ Still, whenever fake news constitutes bullshit, and whenever such fake news is fake *political* news, we have a potential source of rational political bullshit.

Lastly, we shouldn't forget that members of the media, like everybody else, have their own political biases. Though in more traditional settings there may be journalistic norms prescribing a certain degree of impartiality, for many the temptation to bullshit on matters they feel strongly about will be too much.²⁰ Needless to say, the bewildering variety of online blogs and websites from which many get their daily news may lack even those partly effective norms recommending impartiality that are (sometimes) internalized by members of mainstream media sources. As far as bullshitting about politics is concerned, the result is predictable. It can be psychologically taxing for partisan individuals to sincerely engage with ideologically incongruent viewpoints. Correspondingly, it is much easier and less costly to engage in biased reporting, expressing largely congruent viewpoints.

Some citizens will react negatively to such reporting, and this may harm the reputations of the relevant media outlets. But as Gordon Tullock once wrote, "inaccuracy in your information is only important if the person whom you are trying to persuade knows or is likely to find out that you are in error" (Tullock 1972: 128). So long as there are sufficiently many who will reward the media even when they bullshit (whether by tuning in, paying for subscriptions, or what have you), bullshitting will remain rational. The question, then, is whether media outlets can expect their audience to be habitually ignorant in the requisite manner. If so, the reputational harms of having their bullshit discovered will often be outweighed by the expected benefits of continuing to bullshit. As we shall see in the following section, there are indeed reasons to think that audiences will often remain ignorant of the media's bullshit.

2.3. Bullshitting Voters

It is a commonplace among political economists that citizens in democracies are often *rationaly ignorant* (Downs 1957). Becoming well-informed is just too hard while offering few rewards, and so the costs of becoming politically well-informed typically far outweigh the benefits.²¹ Consequently, most citizens are ignorant of even basic political matters. The underlying rational choice analysis also generalizes to *political* irrationality – it is often instrumentally rational to be epistemically irrational about politics (Caplan 2007). Again, epistemic rationality is costly, requiring serious time and effort. When its costs outweigh its benefits, we find widespread irrationality in politics.

¹⁹With that said, there is some degree of uniformity in how academic philosophers think of fake news. As Jaster and Lanius (2021) demonstrate, most philosophical accounts of fake news agree that it contains false content that is spread intentionally. However, they also acknowledge that some fake news may be propagated by agents indifferent to the truth of what they propagate. In other words, at least some fake news is bullshit. For additional accounts of fake news, see Levy (2017) and McIntyre (2018). See also Habgood-Coote (2019) for critical discussion of this literature.

²⁰I remain silent in this paper on whether members of the media ever discharge their duties in a fully impartial manner. I claim only that *even if* there are norms prescribing impartiality, these norms will often be outweighed by opposing incentives to bullshit (or lie, or engage in deception, etc.).

²¹There is a vast literature documenting the extent of voter ignorance. For helpful overviews, see Caplan (2007), Somin (2013), Achen and Bartels (2016), and Brennan (2016). For more on rational ignorance, see Somin (2006).

How does this bear on the prevalence of bullshit in political discourse? Many citizens are both rationally ignorant and rationally irrational, but they are also regularly pressed into social settings that incentivize making one's voice heard on political matters. In short, they face pressure to speak while not knowing or caring about the facts (Frankfurt 2005 [1986]: 63; Cappelen and Dever 2019: 55).

Why might citizens be incentivized to speak without knowing or caring about the facts? First, political discussions provide frequent opportunities to *grandstand*. Speaking in ways that boost one's level of esteem within one's social group incentivizes saying things that earn the approval of the relevant group, even if one doesn't know or care whether what one says is true (Tosi and Warmke 2016; Simler and Hanson 2018: Ch. 16). An agent embedded in an online community characterized in part by, say, adherence to the claim that the Earth is flat, is motivated to publicly disavow scientific orthodoxy on this matter, even if the agent is unfamiliar with any of the relevant scientific evidence. The public disavowal acts as a signal to their in-group about the propriety of their beliefs, boosts their credibility and status, and helps to avoid penalties associated with believing the "wrong" things (e.g., scientific orthodoxy regarding the ellipsoidal shape of the Earth).²²

Secondly, and relatedly, political discussions provide ample opportunity to dismiss, discredit, and disparage one's political opponents. One can earn the esteem of one's peers by casting aspersions on the right people (members of opposing parties, ideologically opposed peers, and the like).²³ In short, one earns esteem by attempting to lessen the esteem of those perceived to be in opposition. Independently of such grandstanding motives, there are plausibly benefits from the mere expression of disapproval of politically opposed individuals alone. Sometimes one just wants to *express* how one feels; it can feel good to express anger, venting one's negative attitudes towards others. These are occasionally substantial psychological benefits, and it is important to note that one can attain these benefits without a serious regard for the truth of what one communicates.

Thirdly, and setting aside such tribalistic motives, there are more general pressures to appear well-informed about important political matters (Petrocelli 2018). One can *lose face* by appearing uninformed, and so one may speak in ways that mimic being well-informed without incurring the costs required to become genuinely well-informed.²⁴ More generally, one might fear losing face even if admitting one's political ignorance would not in fact earn the disapproval of one's peers. Consider a case where an individual consistently engages in preference falsification to avoid social sanctions from their community – presenting themselves as *knowing* that prevailing ideologies are sound while in fact being wracked by doubt regarding their veracity – all while, unbeknownst to them, their doubts are shared by the wider community and would be praised rather than punished.²⁵ In such cases, one's false beliefs about the social costs of public political ignorance incentivize face-saving bullshit, since bullshitting allows one to avoid the expected social costs associated with having one's ignorance exposed.

²²On the importance of political identities in political psychology, see Achen and Bartels (2016) and Mason (2018). For a discussion of the role played by signaling in motivations to share fake news, see Bergamaschi Ganapini (2021).

²³Relatedly, similar partisan motivations seemingly lie behind much sharing of fake news on Twitter (Osmundsen *et al.* 2020; Oyserman and Dawson 2021).

²⁴On face-saving interaction more generally, see Goffman (1967, 1969).

²⁵See Kuran (1995), who first introduced the concept of preference falsification.

We should take care to avoid depicting citizens in an overly unflattering manner. First, although it is hard to be precise on this matter, not all citizens are bullshitters. Some are intellectually and morally serious people who care deeply about the truth, at least some of the time. Second, much bullshit results from merely wanting to avoid the costs of acquiring political information, rather than the largely partisan and reputational motives discussed above (Pennycook and Rand 2018, 2021; McBrayer 2021: 40).²⁶ The former type of bullshit is arguably more excusable than partisan bullshit, being a natural consequence of the time-consuming and effortful nature of the acquisition and evaluation of political information. Still, when the costs of acquiring political information is paired with the sort of partisan social incentives discussed above, the result is a considerable amount of bullshitting on the part of citizens. We don't all bullshit all the time, but we do it enough that politics is seemingly filled with it. As we'll see below, this has important implications.

2.4. *The Collective Harm of Bullshit in Politics*

One might be tempted to dismiss bullshit as a mere annoyance rather than something seriously harmful. Given the language used to characterize the phenomenon, one might even view discussions of it (such as this one) as flippant or somehow tongue-in-cheek. But one shouldn't be misled by the language used. Bullshit in politics is in fact a tremendous source of harm. The harmfulness of bullshit becomes more obvious when one recalls that we are first and foremost concerned with political agents lacking regard for the truth. When one lacks regard for the truth of relatively unimportant matters in one's private life, the consequences are trivial. Indeed, even when faced with important decisions in one's private life, one frequently bears the costs of any poor decision-making alone. But in politics where we must interact with others and make decisions collectively, others bear the costs of our bullshit. Agents pursuing local gains by rationally bullshitting contribute to widespread collective harms. In effect, the problem of rational bullshit in politics constitutes a tragedy of the epistemic commons (Hardin 1969; Joshi 2021).

Consider first the bullshit of politicians. Speaking without regard for the facts is an easy way to spread inaccurate information, especially when done – as is frequently the case – in highly visible ways in public fora. Bullshitters will sometimes stumble across the truth, but they'll often get things wrong too. When we, as voters, believe what bullshitting politicians say, we often form false beliefs. Specifically, when we believe bullshitting politicians who get the facts wrong, their mistakes become ours. When falsehoods circulate too readily, the epistemic health of democracies is threatened.²⁷ We can't vote in optimal ways if consistently wrong about the facts, whether about the character of political candidates, the content of their policies, or about politically relevant facts more generally.

²⁶Similarly, much bullshit likely arises because many citizens are *epistemically insouciant* – that is, they are simply habitually indifferent to whether their beliefs are supported by available evidence (Cassam 2018). Such insouciance is not a deliberate choice. As Cassam (2019: 84) writes, “[one] doesn't choose to be excessively casual and nonchalant towards the challenge of finding answers to complex questions. One just is.” This unreflective insouciance is different from the more deliberate bullshit often produced by citizens and other political actors. Thanks to an anonymous referee for highlighting the connection between bullshit and epistemic insouciance.

²⁷See Blake-Turner (2020) and Brown (2021) for further discussion of this point.

In general, bullshit in politics inhibits our ability to make reliably accurate decisions on politically important matters. As we have seen, when there exists an incentive to produce fake news or bullshit clickbait, organizations will gladly supply them. But when the media is saturated with bullshit and fake news, we again become increasingly more likely to form false beliefs on important issues. When we're continuously reasoning from false beliefs, we're going to make bad decisions. Moreover, as the volume of fake news and bullshit becomes noticed by consumers of online media, we become increasingly distrustful of sources of news in general (Jaster and Lanius 2021). Indeed, fostering mutual distrust among already polarized voters might be the primary aim of some creators and distributors of fake news; and in conditions of heightened distrust, we might not even trust generally accurate and reliable sources of information. We ignore those who mostly get things right, or who at least make a sincere effort to get things right, unfairly dismissing them as bullshitters. If we don't become outright skeptics, we at least become *epistemic cynics*, viewing everybody as out to sell us their self-serving bullshit.²⁸

Of course, bullshitting citizens play a central role in this process. While some people are driven to *create* fake news because of financial incentives, far more are driven to consume and *spread* fake news. When we encounter an inaccurate story online, the first impulse of many of us is to share it – provided, naturally, that it vindicates our pre-conceptions. By sharing it, we can signal to others about our political beliefs, express our affiliation, and so on. Sharing content with the click of a button is far less costly than taking the time to fact-check its content. Even when we don't share fake news, the article may be skimmed, with its content subsequently being taken on board. In this way, misinformation and bullshit spread throughout online communities. As before though, the more bullshit we share online, the more diminished the quality of information we have access to. We partly create the conditions for epistemically defective political environments, for it is easier to do so than it is to create epistemically healthy ones.

Although the focus of this section has been squarely on potential epistemic harms of bullshit, one should bear in mind that these epistemic harms often lead to serious *non-epistemic* harms. When voters are not acquainted with the facts, they may vote for candidates they wouldn't support if they were fully informed. Additionally, candidates and political representatives more generally are incentivized to respond to the preferences of an electorate constituted by frequent bullshitters. Pairing such an electorate with self-serving politicians (who, let's recall, are also themselves frequent bullshitters) is a recipe for improperly designed laws and policies that can adversely impact the lives of millions of people: economic policies that impoverish us, public health policies that make us sick or unsafe, geopolitical policies that put our lives in danger, and more.

3. What Not to Do About Bullshit in Politics

All things considered, bullshitters and their bullshit cause much harm in politics. While some agents secure localized benefits for themselves through their bullshit, it is hard to see how we could collectively, in aggregate, gain more than we lose through such

²⁸Cf. Pritchard (2021: 63). With that said, while epistemic cynicism may cause us to unfairly dismiss some people as bullshitters, it might also cause us to dismiss people we would otherwise uncritically – and unwisely – view as engaged in a sincere pursuit of the facts. Thus, epistemic cynicism has some potential upshots in addition to the costs outlined above.

behavior. Naturally enough, then, some may wish to take action to safeguard against this bullshit. If bullshit is so harmful, isn't it worthwhile to try to reduce its prevalence? And if so, how might we go about this?

First, one might think that if the severity of the problems created by political bullshit were more adequately communicated to the public, one could rely on people to make more of an effort to do what they can to prevent its spread. In short, people could be relied upon to exercise *personal responsibility*, thus tackling the problem from the bottom up. For instance, Christopher Blake-Turner argues that we should stress the importance of holding others accountable whenever they share or create fake news (Blake-Turner 2020: 13). Perhaps by reproaching each other, we can effectively disincentivize the creation and transmission of fake news and other forms of bullshit.²⁹

Second, we could train people to be better able to spot political bullshit. More broadly, we can strive to inculcate in as many people as possible various epistemic virtues that render it more likely they make a sincere attempt to ascertain the facts, more likely that they can successfully navigate epistemic environments saturated with misinformation, and so on (Pritchard 2021). Such efforts would involve vigorous attempts to improve overall media literacy among the electorate so that citizens possess the tools they need to identify and manage misinformation created, disseminated, and shared by bullshitters (Holcombe 2017; Orlando 2017; Rini 2019; McBrayer 2021: 175–6).

Third, we could rely on fact-checkers to flag content that contains falsehoods, misleading claims, and other epistemically unhelpful statements (Amazeen 2013, 2015; Rini 2017). Perhaps if such content is consistently flagged, people would be less likely to uncritically accept it and, importantly, less likely to share it. Fact-checkers, of course, are not in the business of identifying bullshit *per se*; their quarry is a much broader category. Still, bullshit in politics is harmful to the extent that it enables the spread of misinformation, and if fact-checkers can reliably identify such misinformation (some of which has its roots in bullshitters), then they can help with this problem nonetheless.

Fourth, and closely related to the previous proposal, we could reform social media in such a way that the spread of harmful misinformation becomes less likely (Woolley and Howard 2019: 243–4). For example, social media platforms could highlight content that has been flagged by independent fact-checkers, providing to its users a warning that they are about to read and/or share potential misinformation (Rini 2017: 56–7). Alternatively, social media platforms could conduct their own fact-checking operations. Social media platforms could also assign reputation scores to users who repeatedly share flagged content (Rini 2017: 57–8). Such reputation scores would in principle act as a signal to other users that the relevant agents are epistemically unreliable, being prone to sharing fake news and political bullshit more generally. By undermining the credibility of habitual bullshitters, we make it more likely that they're ignored by their peers.

Fifth, and lastly, we could go as far as to empower lawmakers to enact legislation that attempts to curb the spread of misinformation generally (Sim 2019; Brown 2021; Fritts and Cabrera 2022a). Legislation could be targeted at the sources of bullshit themselves, in much the same way that we already prevent certain industries from advertising their products in deliberately misleading ways (O'Connor and Weatherall 2019). Alternatively, we could proceed indirectly, incentivizing social media platforms to take greater care to curate their content and keep it as free as possible from bullshit (Rini 2019). Instead of approaching the problem of political bullshit from the bottom

²⁹In a similar vein, Schauer and Zeckhauser (2009) argue that informal reputational mechanisms that raise the expected costs of paltering may effectively reduce its prevalence.

up, we work from the top down, entrusting our political representatives and lawmakers more broadly with the task of preserving and improving the epistemic health of our democracies.

Although an overview of this length naturally overlooks various details regarding these five broad strategies, it suffices to highlight one major complication that harms the prospects of each – namely, each ignores to varying degrees the extent to which the production and transmission of political bullshit is incentivized. Unfortunately, though, we overlook the relevant incentives at our peril. While some of the proposals would be relatively harmless (even if ineffective), others create additional incentives to bullshit in new, potentially more harmful ways. Let's consider how each fails to account for the underlying incentives.

First, reliance upon personal responsibility overlooks the fact that, without the appropriate incentives, it's hard to see what could motivate people to hold each other accountable in the right way. One could grant that *if* we each held each other accountable fairly and diligently, then the creation and transmission of political bullshit might be curbed. If we each became far more epistemically vigilant, and if we each became more receptive to reprobation from those who criticize our epistemic conduct, then personal responsibility would constitute a plausible, bottom-up approach to the problem. But if political agents are instead incentivized to be epistemically negligent and even outright biased, such an approach will be unsuccessful. Of course, we have already seen that agents in politics can often secure rewards by bullshitting. These same rewards are present in contexts where we could hold others accountable. Accurately holding others accountable will be an occasionally difficult, time-consuming effort.³⁰ Additionally, an honest assessment of our peers might force us to hold our ideologically likeminded allies to account for their epistemic vices. Biased political agents looking to avoid such costs, whether voters or members of the media or politicians, will rationally bullshit about the bullshit of others.

Second, improved media literacy and training may be ineffective in the absence of appropriate incentives to use such training in a consistent, unbiased manner. On the one hand, such training may occur in relatively cloistered environments free of the sort of pressures and incentives that political agents face in quotidian political settings. For instance, even if people are educated about the identity of consistently egregious producers of fake news and characteristic hallmarks of bullshitting politicians, they may fail to apply what they have learned in cases when they are faced with political bullshit that meshes well with their political convictions, or when it is shared by a trusted friend, and so on (Martens 2010). On the other hand, those who provide the relevant training may do so in a biased fashion, influenced by powerful incentives to bullshit about the degree to which various political factions are prone to bullshitting, prone to producing and sharing fake news, and more. Were we to implement media literacy programs on a large scale, those who organize and implement them will suddenly find themselves in a position where they can potentially influence the character of the political informational landscape in a serious way. While there would no doubt be some who discharge their duties in an impartial manner, some will be unable to resist the temptation to abuse their position.

Third, the use of fact-checkers may be ineffective for several reasons. Most obviously, we need unbiased fact-checkers who are willing and able to set aside their partisan

³⁰This is already recognized in the wider literature in political philosophy. For example, see Guerrero's discussion of a lack of meaningful accountability in contemporary electoral democracies (Guerrero 2014).

commitments and discharge their duties in a consistent and fair manner. But fact-checkers will often find it difficult to refrain from disproportionately scrutinizing political figures they dislike while subjecting claims made by favored figures to less intense scrutiny. Fact-checkers may not even be aware that they are guilty of focusing differentially on different figures, instead being driven by largely unconscious biases. Still, whether their biases are conscious or not, it can be difficult for people to question others with whom they largely agree. Correspondingly, it is much easier and more rewarding to critique one's political opponents. These twin pressures, as we have seen, incentivize a biased mode of political cognition that is characteristic of the habitual bullshitter. Additionally, the effectiveness of fact-checkers relies upon a citizenry willing to accept the verdicts of such fact-checkers. Even if fact-checkers are generally reliable and unbiased, citizens may disregard their findings when they are incongruent with their political beliefs (Walter *et al.* 2019; Fritts and Cabrera 2022b). Naturally, this sort of knee-jerk dismissal would be even more predictable were instances of genuine bias among fact-checkers to be documented, with bullshitting citizens erroneously tarring all fact-checkers with the same brush.

Similar concerns plague the suggestion that social media reform could meaningfully halt the spread of political bullshit. First, there is no guarantee that social media platforms who conduct their own fact-checking will do so without bullshitting about which political figures or media outlets are themselves guilty of bullshitting. Second, if they forego conducting their own fact-checking and choose to rely on independent fact-checkers, this simply raises the sort of problems that could arise with independent fact-checkers. Third, even if social media platforms hire agents willing and able to conduct these tasks without bullshitting, users of social media may bullshit about the relevant mechanisms, whether fact-checking verdicts, reputation scores, or both. One can even imagine social media users wearing their negative reputation scores as badges of honor, earned by bravely standing up to the biased verdicts of overly powerful and influential social media organizations.

Lastly, we should proceed cautiously before ceding power to political representatives and legislators to decide who is bullshitting, who is guilty of manufacturing and spreading fake news, and the like.³¹ Such proposals expose a curious short-sightedness and naivete among many – specifically, a naivete regarding the degree to which we can trust politicians with obvious incentives to abuse their position for political gain to fairly and accurately hold creators and sharers of political bullshit accountable. As we saw in section 2.1, politicians are incentivized in various ways to bullshit. With the power to punish those they deem guilty of bullshitting, we are in effect giving an even greater amount of institutional power to people who are likely to bullshit about when political opponents are bullshitting, bullshit about when stories critical of them are fake news, and so on. Such an outcome may be even worse than the status quo where, at the very least, few people are being punished by agents backed by the power of the state for their alleged bullshit.

In short, by failing to address the underlying incentives, these proposals are likely to be ineffective. Even worse, some of them will afford new, dangerous ways for bullshitters to bullshit. Nonetheless, one might maintain that all that matters from a practical vantage point is whether any given intervention would produce, on balance, a positive impact. And whether an intervention would do so or not is clearly an empirical question. But as we have seen, there is little reason to think that these interventions will be

³¹Cf. McBrayer (2021: 178).

successful. More worryingly, we have some strong reasons to think that not only will these interventions fail, at least some of these interventions will create even worse problems than those they are supposed to fix. Absent strong countervailing evidence, then, we should err on the side of caution before implementing some of these more radical proposals, and proponents of such interventions should shoulder an evidentiary burden which they have yet to discharge.

4. The Inevitability of Bullshit in Politics

That so many people could care so little about the facts is in part a consequence of powerful incentives that reward them for bullshitting. By failing to consider these incentives, existing proposals for reducing the harm of political bullshit are unlikely to succeed. Worse, they may simply make things worse by providing new avenues for bullshitters to bullshit. However, one might think that recognizing the fact that perverse incentives breed bullshitters can help us do better. Perhaps if we can figure out how to create *different* incentives more conducive to an epistemically well-functioning polity, we can reduce the prevalence of political bullshit, or at least mitigate its harms. This, no doubt, will be a difficult task. But it won't be an impossible one.

How might this work? Plausibly, successful anti-bullshit interventions will focus on making it costly for agents to disregard the facts, reward those who make a serious effort to determine the facts, or both. In other words, successful interventions will create incentives directly conflicting with existing incentives that, *inter alia*, make it costly to acquire political information and rewarding to bullshit. For instance, an ambitious strategy would be to strive to inculcate in children from a young age an urge to pursue the facts even when doing so is difficult or time-consuming. Perhaps this could be achieved either by portraying truth-directed inquiry as *intrinsically* rewarding or by stressing the *instrumental benefits* of such inquiry (or both). In either case, children raised to find truth-directed inquiry rewarding will, one hopes, be less likely to consider serious attempts to figure out the truth as costs to be avoided. Put another way, children raised in such a manner are less likely to be epistemically insouciant (Cassam 2018). In effect, by socializing children to internalize the benefits of pursuing the truth, we could in principle raise a generation of people who, by disposition, are strongly averse to bullshit in politics.³²

To take another less ambitious proposal, we could ensure that fact-checking organizations are visibly non-partisan, ideally being endorsed, supported, or otherwise associated with salient figures from a range of political parties, interest groups, and more. Partisan agents may be more willing to relinquish their beliefs if figures they trust are the ones either suggesting such beliefs be relinquished or publicly backing those who do make such suggestions (Margolin *et al.* 2017). When fact-checking processes are associated with trustworthy and ideologically likeminded people, changing one's mind is less likely to be seen as a cost to avoid in order to save face.

Relatedly, if important, charismatic public figures are seen to reliably pursue the facts, change their minds in light of the facts even if it is inconvenient for them, or otherwise strive to be epistemically virtuous, a healthy regard for the facts may come to be seen as something to emulate (Mohseni and Williams 2019).³³ If careful, diligent

³²Compare this with Gordon Tullock's remarks about the desirability of moral education for children in order to inculcate in them an aversion to lying (Tullock 1972: 137).

³³See also O'Connor and Weatherall (2019: 178–9) for relevant discussion.

deliberation were seen less as a cost and more as something that popular, well-respected people regularly engage in, people could begin to secure social rewards for behaving similarly. We could exploit our tendency to conform to widespread practices and norms by making an effort to popularize figures who are known to reliably resist the temptation to bullshit. In the same way that in current political settings one can secure social rewards by toeing the party line, one could in principle secure similar rewards by conforming to increasingly widespread and popular practices of virtuous epistemic conduct.

However, it is much easier to outline such proposals than it is to successfully implement them, let alone implement them in ways that would achieve our ambitious aims, for the very same incentives that hamper existing attempts to reduce the amount of bullshit in politics reappear. Will those who inculcate in our children the relevant attitudes do so fairly and without bullshitting? Will the internalized rewards for properly regarding the facts outweigh the potential rewards of bullshit? How can we incentivize the formation of truly non-partisan fact-checking organizations? What incentives are there for popular public figures to avoid bullshit, and what incentives are there for the public to follow their lead? Do these incentives outweigh conflicting incentives to bullshit? If not, even those interventions that attempt to create the right sort of incentives will likely fail. In a sense, we would need to *already* possess sufficiently strong beneficial incentives for the appropriate interventions to succeed.

One might object that there must be *some* interventions that would have a beneficial impact. Even if existing interventions fail, why assume that *none* can succeed? Additionally, proponents of such interventions are likely aware that eliminating bullshit from politics is not feasible. Still, merely reducing its prevalence or mitigating its impact is much more realistic. If so, we shouldn't conclude that no intervention could succeed where others have failed.

The claim that there are *possible* interventions that could fare better than existing proposals should be granted. With that said, we need to be realistic about what anti-bullshit interventions can achieve. The incentives that drive agents to bullshit inevitably arise in the sort of conditions that characterize large political communities. In politics, people will always have a plurality of ends, only some of which are best served by serious truth-directed inquiry. For agents like us with finite cognitive and attentional resources, the acquisition and evaluation of political information will always have its costs.³⁴ Meanwhile, presenting oneself as knowledgeable will often yield benefits. Right away, such facts already incentivize a certain amount of bullshit. Politics

³⁴An anonymous referee wonders whether such costs, rather than incentivizing agents to bullshit, instead give them reason to use heuristics that reduce the costs of information acquisition. This suggestion mirrors existing work in democratic theory seeking to downplay the significance of widespread voter ignorance, with heuristics purportedly helping voters to reduce the costs of acquiring political information that would otherwise render political ignorance rational (Popkin 1991; Wittman 1995). Heuristics can indeed reduce error in some cases (Sinhababu 2016). However, though I cannot fully defend this claim here, I think we have reason to worry that agents will frequently bullshit nonetheless. First, agents may bullshit about the reliability of various heuristics, claiming to use them for the purposes of reliably forming true beliefs while primarily choosing them for reasons unrelated to their purported connection to the facts. For instance, some heuristics may be chosen primarily for their entertainment value, while others are chosen primarily because they are congruent with preexisting convictions (Somin 2013: 91). Second, many of the heuristics adverted to in the wider literature (such as political parties and mass media) involve the use of agents who, if the arguments in this paper are sound, are habitual bullshitters. Accordingly, even if these heuristics are *sometimes* reliable, the extent to which they rely on bullshitters compromises this reliability.

always involves agents with competing aims and ambitions, and factions inexorably arise, whether we like it or not. When factions arise and the agents constituting them develop and intensify their partisan commitments, the various social incentives discussed earlier emerge, and such incentives drive the production of much bullshit in politics. Whenever there is a profit to be made from pandering to one's audience, and wherever costs can be cut by shirking one's duties to report the facts accurately, there will be bullshit in the media.

It's plausible that there is no escape from such conditions in politics. Consequently, while we may yet be able to reduce the amount of bullshit in politics through clever design or happy accident, it seems that we'll never be rid of it entirely. In politics, bullshit will always pay.³⁵

References

- Achen C.H. and Bartels L.** (2016). *Democracy for Realists: Why Elections do not Produce Responsive Government*. Princeton, NJ: Princeton University Press.
- Amazeen M.A.** (2013). 'Making a Difference: A Critical Assessment of Fact-Checking in 2012.' *New America Foundation Media Policy Initiative*, research paper.
- Amazeen M.A.** (2015). 'Revisiting the Epistemology of Fact-Checking.' *Critical Review: A Journal of Politics and Society* 27(1), 1–22.
- Benkler Y., Faris R. and Roberts H.** (2018). *Network Propaganda: Manipulation, Disinformation, and Radicalization in American Politics*. Oxford: Oxford University Press.
- Bergamaschi Ganapini M.** (2021). 'The Signaling Function of Sharing Fake News Stories.' *Mind & Language*. <https://doi.org/10.1111/mila.12373>.
- Bergstrom C.T. and West J.D.** (2020). *Calling Bullshit: The Art of Skepticism in a Data-Driven World*. New York, NY: Random House.
- Blake-Turner C.** (2020). 'Fake News, Relevant Alternatives, and the Degradation of our Epistemic Environment.' *Inquiry*. <https://doi.org/10.1080/0020174X.2020.1725623>.
- Bok S.** (1978). *Lying: Moral Choice in Public and Private Life*. New York, NY: Vintage Books.
- Brennan J.** (2016). *Against Democracy*. Princeton, NJ: Princeton University Press.
- Brown E.** (2021). 'Regulating the Spread of Online Misinformation.' In M. Hannon and J. De Ridder (eds), *The Routledge Handbook of Political Epistemology*. New York, NY: Routledge.
- Caplan B.** (2007). *The Myth of the Rational Voter*. Princeton, NJ: Princeton University Press.
- Cappelen H. and Dever J.** (2019). *Bad Language*. Oxford: Oxford University Press.
- Carson T.** (2006). 'The Definition of Lying.' *Noûs* 40(2), 284–306.
- Carson T.** (2010). *Lying and Deception: Theory and Practice*. Oxford: Oxford University Press.
- Carson T.** (2016). 'Frankfurt and Cohen on Bullshit, Bullshitting, Deception, Lying, and Concern with the Truth of What One Says.' *Pragmatics and Cognition* 23(1), 53–67.
- Cassam Q.** (2018). 'Epistemic Insouciance.' *Journal of Philosophical Research* 43, 1–20.
- Cassam Q.** (2019). *Vices of the Mind: From the Intellectual to the Political*. Oxford: Oxford University Press.
- Chomsky N. and Herman E.** (1988). *Manufacturing Consent: The Political Economy of the Mass Media*. New York, NY: Pantheon.
- Cohen G.A.** (2002). 'Deeper into Bullshit.' In S. Buss and L. Overton (eds), *Contours of Agency: Themes from the Philosophy of Harry Frankfurt*. Cambridge, MA: MIT Press.
- Downs A.** (1957). *An Economic Theory of Democracy*. New York, NY: Harper and Row.
- Fallis D. and Stokke A.** (2017). 'Bullshitting, Lying, and Indifference Toward Truth.' *Ergo: An Open Access Journal of Philosophy* 4(10), 277–309.
- Faulkner P.** (2007). 'What is Wrong with Lying?' *Philosophy and Phenomenological Research* 75(3), 535–57.

³⁵Thanks to Patrick Brooks, Elisabeth Camp, Andy Egan, Alex Guerrero, Philip Pettit, Oliver Traldi, audiences at Dublin and Genoa, and three anonymous referees for their helpful feedback on several versions of this paper.

- Frankfurt H.** (2002). 'Reply to G.A. Cohen.' In S. Buss and L. Overton (eds), *Contours of Agency: Themes from the Philosophy of Harry Frankfurt*. Cambridge, MA: MIT Press.
- Frankfurt H.** (2005 [1986]). *On Bullshit*. Princeton, NJ: Princeton University Press.
- Fritts M. and Cabrera F.** (2022a). 'Fake News and Epistemic Vice: Combating a Uniquely Noxious Market.' *Journal of the American Philosophical Association* 1–22. doi: <https://doi.org/10.1017/apa.2021.11>.
- Fritts M. and Cabrera F.** (2022b). 'Online Misinformation and 'Phantom Patterns': Epistemic Exploitation in the Era of Big Data.' *Southern Journal of Philosophy* 60(1), 57–87.
- Gelfert A.** (2018). 'Fake News: A Definition.' *Informal Logic* 38(2), 84–117.
- Goffman E.** (1967). *Interaction Ritual: Essays on Face-to-Face Behavior*. New York, NY: Anchor Books.
- Goffman E.** (1969). *Strategic Interaction*. Philadelphia, PA: University of Pennsylvania Press.
- Guerrero A.** (2014). 'Against Elections: The Lottocratic Alternative.' *Philosophy and Public Affairs* 42, 135–78.
- Habgood-Coote J.** (2019). 'Stop Talking About Fake News!' *Inquiry: An Interdisciplinary Journal of Philosophy* 62(9–10), 1033–65.
- Hardin G.** (1969). 'The Tragedy of the Commons.' *Science* 162, 1243–68.
- Hardin R.** (2009). *How Do You Know? The Economics of Ordinary Knowledge*. Princeton, NJ: Princeton University Press.
- Heffer C.** (2021). 'Bullshit and Dogmatism: A Discourse Analytical Perspective.' In A. Tanesini and M.P. Lynch (eds), *Polarisation, Arrogance, and Dogmatism: Philosophical Perspectives*. New York, NY: Routledge.
- Holcombe M.** (2017). 'Reading, Writing, Fighting Fake News.' *CNN*, 29 March.
- Ivankovic V.** (2016). 'Steering Clear of Bullshit? The Problem of Obscurantism.' *Philosophia* 44, 531–46.
- Jacquemet M.** (2020). '45 as a Bullshit Artist: Straining for Charisma.' In J. McIntosh and N. Mendoza-Denton (eds), *Language in the Trump Era: Scandals and Emergencies*. Cambridge: Cambridge University Press.
- Jaster R. and Lanius D.** (2021). 'Speaking of Fake News: Definitions and Dimensions.' In S. Bernecker, A. Flowerree and T. Grundmann (eds), *The Epistemology of Fake News*. Oxford: Oxford University Press.
- Jay M.** (2010). *The Virtues of Mendacity: On Lying in Politics*. Charlottesville, VA: University of Virginia Press.
- Joshi H.** (2021). *Why It's Ok to Speak Your Mind*. New York, NY: Routledge.
- Kolodny N. and Brunero B.** (2020). 'Instrumental Rationality.' In E.N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2020 Edition). <https://plato.stanford.edu/archives/spr2020/entries/rationality-instrumental/>.
- Kuran T.** (1995). *Private Truths, Public Lies*. Cambridge, MA: Harvard University Press.
- Levy N.** (2017). 'The Bad News About Fake News.' *Social Epistemology Review and Reply Collective* 6(8), 20–36.
- Lindsey B. and Teles S.M.** (2017). *The Captured Economy: How the Powerful Enrich Themselves, Slow Down Growth, and Increase Inequality*. Oxford: Oxford University Press.
- Mahon J.E.** (2015). 'The Definition of Lying and Deception.' In E.N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/archives/win2016/entries/lying-definition/>.
- Margolin D.B., Hannak A., and Weber I.** (2017). 'Political Fact-Checking on Twitter: When Do Corrections Have an Effect?' *Political Communication* 35(2), 196–219.
- Martens H.** (2010). 'Evaluating Media Literacy Education: Concepts, Theories, and Future Directions.' *Journal of Media Literacy Education* 2, 1–22.
- Mason L.** (2018). *Uncivil Agreement: How Politics Becomes Our Identity*. Chicago, IL: University of Chicago Press.
- Mathiesen K. and Fallis D.** (2017). 'The Greatest Liar Has His Believers: The Social Epistemology of Political Lying.' In E. Crookston, D. Killoren and J. Trerise (eds), *Ethics in Politics*. New York, NY: Routledge.
- McBrayer J.P.** (2021). *Beyond Fake News: Finding the Truth in a World of Misinformation*. London: Routledge.
- McIntyre L.** (2018). *Post-Truth*. Cambridge, MA: MIT Press.
- Mearsheimer J.J.** (2011). *Why Leaders Lie: The Truth About Lying in International Politics*. Oxford: Oxford University Press.

- Michael R.B. and Breaux B.O.** (2021). 'The Relationship Between Political Affiliation and Beliefs About Sources of Fake News.' *Cognitive Research: Principles and Implications* 6(6). <https://doi.org/10.1186/s41235-021-00278-1>.
- Mohseni A. and Williams C.R.** (2019). 'Truth and Conformity on Networks.' *Erkenntnis*. <https://doi.org/10.1007/s10670-019-00167-6>.
- Mukerji N.** (2018). 'What is Fake News?' *Ergo: An Open Access Journal of Philosophy* 5, 923–46.
- Osborne P.** (2005). *The Rise of Political Lying*. London: Free Press.
- O'Connor C. and Weatherall B.** (2019). *The Misinformation Age: How False Beliefs Spread*. New Haven, CT: Yale University Press.
- Orlando J.** (2017). 'How to Help Kids Navigate Fake News and Misinformation Online.' *ABC News*, 26 June.
- Osmundsen M., Bor A., Wahlstrup P., Bechmann A. and Petersen M.B.** (2020). 'Partisan Polarization is the Primary Psychological Motivation Behind "Fake News" Sharing on Twitter.' *PsyArXiv*. <https://psy-arxiv.com/v455bk/>.
- Oyserman D. and Dawson A.** (2021). 'Your Fake News, Our Facts: Identity-based Motivation Shapes What We Believe, Share, and Accept.' In R. Greifeneder, M.E. Jaffé, E.J. Newman and N. Schwarz (eds), *The Psychology of Fake News: Accepting, Sharing, and Correcting Misinformation*. New York, NY: Routledge.
- Pennycook G. and Rand D.G.** (2018). 'Lazy, not Biased: Susceptibility to Partisan Fake News is Better Explained by Lack of Reasoning than by Motivated Reasoning.' *Cognition* 188, 29–50.
- Pennycook G. and Rand D.G.** (2021). 'The Psychology of Fake News.' *Trends in Cognitive Sciences*. <https://doi.org/10.1016/j.tics.2021.02.007>.
- Petrocelli J.V.** (2018). 'Antecedents of Bullshitting.' *Journal of Experimental Social Psychology* 76, 249–58.
- Popkin S.L.** (1991). *The Reasoning Voter: Communication and Persuasion in Presidential Campaigns*. Chicago, IL: University of Chicago Press.
- Posner R** (1974). 'Theories of Economic Regulation.' *Bell Journal of Economics and Management Science* 5, 335–58.
- Pritchard D.** (2021). 'Good News, Bad News, Fake News.' In S. Bernecker, A. Flowerree and T. Grundmann (eds), *The Epistemology of Fake News*. Oxford: Oxford University Press.
- Rini R.** (2017). 'Fake News and Partisan Epistemology.' *Kennedy Institute of Ethics Journal* 27(2), 43–64.
- Rini R.** (2019). 'Social Media Disinformation and the Security Threat to Democratic Legitimacy.' In *Disinformation and Digital Democracies in the 21st Century*. Publication of the NATO Association of Canada.
- Sarajlic E.** (2019). 'Bullshit, Truth, and Reason.' *Philosophia* 47, 865–79.
- Saul J.** (2012). *Lying, Misleading, and What is Said*. Oxford: Oxford University Press.
- Schauer F. and Zeckhauser R.** (2009). 'Paltering.' In B. Harrington (ed.), *Deception: From Ancient Empires to Internet Dating*. Stanford, CA: Stanford University Press.
- Sim S.** (2019). *Post-Truth, Scepticism, and Power*. London: Palgrave Macmillan.
- Simler K. and Hanson R.** (2018). *The Elephant in the Brain: Hidden Motives in Everyday Life*. Oxford: Oxford University Press.
- Sinhababu N.** (2016). 'In Defense of Partisanship.' In D. Killoren, E. Crookston and J. Trerise (eds), *Ethics in Politics: New Papers on the Rights and Obligations of Political Agents*. New York, NY: Routledge.
- Somin I.** (2006). 'Knowledge About Ignorance: New Directions in the Study of Political Information.' *Critical Review* 18(1–3), 413–58.
- Somin I.** (2013). *Democracy and Political Ignorance: Why Smaller Government is Smarter*. Stanford, CA: Stanford University Press.
- Sorensen R.** (2007). 'Bald-Faced Lies! Lying Without the Intent to Deceive.' *Pacific Philosophical Quarterly* 88(2), 251–64.
- Stanley J.** (2015). *How Propaganda Works*. Princeton, NJ: Princeton University Press.
- Stigler GJ** (1971). 'The Economic Theory of Regulation.' *Bell Journal of Economics and Management Science* 2, 3–21.
- Stokke A** (2018). *Lying and Insincerity*. Oxford: Oxford University Press.
- Tosi J. and Warmke B.** (2016). 'Moral Grandstanding.' *Philosophy and Public Affairs* 44(3), 197–217.
- Tosi J. and Warmke B.** (2020). *Grandstanding: The Use and Abuse of Moral Talk*. Oxford: Oxford University Press.
- Tullock G.** (1972). *Toward a Mathematics of Politics*. Ann Arbor, MI: University of Michigan Press.

- Walter N., Cohen J., Holbert R.L. and Morag Y.** (2019). 'Fact-Checking: A Meta-Analysis of What Works and For Whom.' *Political Communication* 37(3), 350–75.
- Weissberg R.** (2004). 'Mr. Pinocchio Goes to Washington: Lying in Politics.' *Social Philosophy and Policy* 21(1), 167–201.
- Wittman D.** (1995). *The Myth of Democratic Failure*. Chicago, IL: University of Chicago Press.
- Woolley S.C. and Howard P.N.** (2019). 'Conclusion: Political Parties, Politicians, and Computational Propaganda.' In S.C. Woolley and P.N. Howard (eds), *Computational Propaganda: Political Parties, Politicians, and Political Manipulation on Social Media*, pp. 241–8. Oxford: Oxford University Press.

Adam F. Gibbons is an Assistant Professor of Philosophy at Lingnan University, where he also works with the Hong Kong Catastrophic Risk Centre. He works primarily in social and political philosophy, especially as they intersect with epistemology, political economy, and political science. Most recently, his research has focused on issues related to institutional design, political epistemology, and language in politics.