

Intellectual Virtue Signalling

Abstract: Discussions of virtue signalling to date have focused exclusively on the signalling of the moral virtues. This article focuses on *intellectual* virtue signalling: the status-seeking advertising of supposed intellectual virtues. Intellectual virtue signalling takes distinctive forms. It is also far more likely to be harmful than moral virtue signalling, because it distracts attention from genuine expertise and gives contrarian opinions an undue prominence in public debate. The article provides a heuristic by which to identify possible instances of intellectual virtue signalling. When people with no relevant expertise rapidly move to offer their opinions on a wide range of topics as soon as these topics become fashionable or newsworthy, and especially when these opinions are contrarian, we should suspect them of intellectual virtue signalling.

Discussions about virtue signalling (VS) are oddly unbalanced.¹ While theories of virtues and vices were first and most extensively developed by philosophers concerned with *moral* dispositions, there has recently been a flowering of work on *intellectual* virtues and vices (see, for example (Cassam 2018; Kidd, Cassam, and Battaly 2020; Roberts and Wood 2007; Zagzebski 1996)). So far, however, discussions of VS have focused exclusively on *moral* talk and virtues.² I aim to correct this omission. I'll argue that intellectual VS has characteristics that moral VS tends to lack. I'll also argue that the understanding the former promises to shed light on the latter. With an account of intellectual VS in hand, we'll be in a position to see that paradigm instances of moral VS are a much less serious problem than their critics allege. Rather, it's talk that closely mimics substantive and genuinely valuable contributions to public discourse that are more likely to be VS of a pernicious kind.

There's a lively debate about which dispositions and character traits might count as intellectual virtues (see, e.g., Daukas, 2011; Frost-Arnold, 2014). I don't aim to contribute to this debate, nor even to take a stand on whether intellectual VS signals genuine virtues (some of the traits I will mention are much more plausible candidates for genuine intellectual virtues than others). Of course, those engaged in signalling need not concern themselves with whether they're signalling genuine virtues: they aim to signal characteristics that other people will value, whether or not they're really virtues.

I. Virtue Signalling: A Spotter's Guide.

Moral VS consists of moral talk ostensibly about some state of affairs or action that is designed in part to draw attention to the person doing the signalling. Tosi and Warmke (2020) argue that VS is partially motivated by what they call the *recognition desire*: the desire to impress others with the signaller's moral qualities. They recognize that motivations can be complex, and offer a test for when a contribution to moral discourse is sufficiently motivated by the recognition desire to count as VS: would the person who engaged in the moral talk be disappointed if they did not receive any recognition in response?

Tosi and Warmke contrast VS – this “abuse of moral talk” (4) – with genuinely virtuous moral discussion. Rather than drawing attention to the person talking, virtuous moral talk aims “to identify publicly certain moral features of a state of affairs, and sometimes additionally to explain the evaluation of that state or recommend some fitting response” (Tosi and Warmke, 2016, pp. 209-210). Virtuous instances of moral talk draw attention to wrongs that might have been overlooked, and advance our moral understanding. These functions are related, of course: we're rarely in a position to expose abuses that need only to be pointed to or described to be seen to be wrong. When we draw the attention of others to (what we take to be) moral wrongs, we usually do by way of argument: for instance, by describing behaviors that are already familiar and giving *reasons* why that behavior is wrong (think, for example, of the way in which feminists have brought us to recognize the moral harms associated with what might once have been seen as banter or flirtation).

Tosi and Warmke recognize that VS can occur in the course or in the form of offering arguments. That someone offers an argument is no guarantee that they're not engaged in VS. But we should expect a rough correlation between lack of substantive argument and VS, at least in the kinds of contexts in which the accusation is likely to levelled at people (paradigmatically, on social media).³ VS is more likely than virtuous moral talk to manifest in simple and short moral utterances. VS is, after all, signalling, and we should expect signals to be easy to detect and decode (Funkhouser 2017). Just as effective adverts tend to be short and punchy, so effective signals should be concise and clear. Genuine contributions to moral talk are likely too complex to play the signalling role.

Indeed, most of the properties Tosi and Warmke (2020, pp. 44ff) take to characterize VS are much more likely to be seen in simple and/or short utterances. *Piling on* consists in the reiteration of what has already been said. *Ramping up* consists in making the same sort of claim in a more extreme

version. *Excessive outrage* and *appeals to self-evidence* even more obviously don't require or typically manifest in long or complex utterances. Moreover, all the examples Tosi and Warmke give of VS are of brief utterances (admittedly, they may use such examples for the sake of convenience). Only *trumping up* – the identification of supposed moral harms in innocuous actions or utterances – doesn't tend to manifest in short and simple moral talk, though of course it may. Considerations like these suggest we might use a rough and ready length heuristic to identify probable instances of VS: if it's short, it's likely a signal.

I'll have more to say about why moral VS tends to manifest in brief or simple communications later in the paper. Whatever the merits of a length heuristic for the identification of moral VS, I'm going to suggest that when it comes to intellectual VS, this heuristic has things backwards. Intellectual VS tends to occur at length, and in the form of genuine argument. Concision and simplicity are strengths in most advertising, but they tend to be ineffective for intellectual VS. Intellectual VS tends to manifest in the form of genuine and (somewhat) substantive arguments. I'll also argue that once we recognize *that*, and *why*, intellectual VS tends to manifest in the presentation of genuine arguments, we'll have a good reason to rethink the extent to which *moral* VS is pernicious.

II. Hot Takes

A river of commentary flowed in the wake of the Russian invasion of Ukraine. Much of that commentary was by journalists and by academics, of course, but in addition many people on social media added their voices. I had no idea that so many of my friends were so expert on Eastern European politics. Indeed, their expertise extended much further: not only did they know a great deal (more than the media were telling us, some of them assured me) about Putin's motivations and goals, but they also had a deep understanding of Russian military doctrine and of Ukrainian politics. Everything had unfolded, they told me, as they had predicted it would.

These near omniscient commentators represented an extreme (often in their politics, as well as in their claims to expertise). But milder versions of much the same phenomena were common. Lots of people who had never mentioned Russia or Ukraine before, as far as I can recall, who had never lived east of Paris (or, in many cases, east of New York) and who had no relevant background or expertise now shared their opinion on the war, its morality, its progress and what it portended for Europe.

Of course, the invasion was a shocking event that riveted the attention of the world, and it's not at all surprising that many people felt the need to discuss it with others. Nor is it surprising that these discussions should happen on social media: for all its many ills, it's well established as a forum where people can seek and share intelligent discussion. There's nothing wrong with any of that. But the people I have in mind, the apparently near omniscient and the more modest alike, did not (merely) take to social media to express horror at the loss of life or at a full-scale invasion of a nation that seemed to resemble their own. Nor did they (only) share links to media stories they'd found interesting or illuminating, or to analyses by experts. Rather, they wrote posts and threads that put their own view of the situation. Often, their posts rejected more mainstream views; in all cases, they were apparently original. They might provide evidence – in the form of links to other sources – but the evidence was offered as support for *their* view. Even when the view was entirely mainstream, the implicit message conveyed was not “I have read these stories and this is what they say” but instead “I have looked into this matter carefully, and I can ratify this view”. The people I have in mind offered what I will call their own *take* on the situation.

This wasn't the first time social media had revealed unsuspected depths of expertise in the posts that populated my timeline. The same phenomenon occurred at the start of the coronavirus pandemic. Over the years of the pandemic, there's no doubt that some (few) of the people who posted their takes on COVID-19 came to acquire a degree of genuine expertise. Some began with skills that gave them a head start, after all. But they often didn't wait until they had read multiple epidemiological studies, and analyses of the economic impact of lockdowns, and so on, to post their own takes. They posted them within days, and sometimes mere hours, of the news hitting the headlines.

Is this outpouring of hot takes VS? It's very different, both in form and content, from paradigm instances of moral VS. First, not all of it has a moral content. While often people offer their hot takes on moral and political topics, it's the fashionableness of an issue – roughly, its aptness to feature in headlines – that seems to attract this behavior, rather than its content. Hot takes on COVID began to appear well before it was a politically polarizing topic. More generally, we should expect moral VS to focus not only on moral and political topics, but in particular on those topics on which there are *sides*; especially those that feature in the culture wars. At the time of writing, the war in Ukraine is not apt for such signalling, because most people in the major parties are on the same side with regard to it.⁴ Second, these takes look nothing like the pithy expressions of attitudes we'd expect of moral signallers. Changing one's profile picture to 'BLM' or tweeting *Black Lives*

Matter; that's the kind of moral expression the critics identify with VS. A 17-tweet thread outlining Putin's alleged real motives, challenging the mainstream media's representation of the war and assessing its geopolitical implications seems to look more like virtuous talk than signalling.

It bears emphasising how genuinely substantive much of the talk I have in mind really is. One may learn from the person I will call the 'instant expert'. They may produce fluent paraphrases and summaries of genuine experts, transmitting these experts' knowledge to a wider audience. Some achieve even more: some have the ability to synthesize work, or even to extend it, sufficiently well to deserve credit as the author. The instant expert might have read a dozen news items, half a dozen longer pieces and innumerable tweets, and distilled them into their take. They may even have contributed a genuinely original element: the instant expert is typically intelligent and may in fact be a genuine expert, albeit in some other field. Their intelligence and expertise may allow them to contribute ideas and perspectives that are novel. Despite the genuine value of their contributions, I'm going to argue that these hot takes are paradigms of intellectual VS.

III. The Properties of Intellectual VS

For Tosi and Warmke, VS is behavior motivated, at least in part, by the recognition desire: moral talk counts as VS only if the agent would be disappointed if she received no recognition of her moral respectability or insight. As we've seen, moral VS tends to manifest in quite undemanding forms: the expression of attitudes. The production of a hot take is quite different: it requires a great deal of time, thought and energy. Given how demanding it is, the question of motivation is raised in an especially pointed form: why would someone go to all that effort for no tangible reward?

We shouldn't be too cynical about this kind of talk. Some people feel a strong need to learn about the topics that come to dominate our headlines (perhaps as a way to gain a feeling of control over events). This need might lead them to immerse themselves in reading and watching as much relevant material as they can; having absorbed and synthesized this material, they may want to share it with others. Such behavior surely can be well-motivated. I'm concerned, though, with only a subset of this posting: the take, where the person presents their own considered view. When the take is presented by the instant expert, and presented very rapidly, we may reasonably suspect a less laudable motivation. The person isn't *merely* passing on what they've learned. The take is offered rapidly, by someone with no previous expertise, and presented as the person's own; these facts

should lead us to think it likely that the desire to acquire expertise *for the sake of being able to offer a take* is playing an important role among their motivations.⁵

Of course, the motivation to thrust themselves into the spotlight needn't be the only one the instant expert has. It's quite possible, for example, that some begin with the desire to inform themselves and then share what they've learned (without passing it off as their own take or only offering an opinion after a relatively long period of reflection and learning), and only discovered they are rewarded for offering their own take by accident; in such cases, the desire to inform themselves might remain powerful. Nor do I claim that every person who engages in behavior like this is engaged in VS: for some, the desire to inform may play by far the strongest role. I'll suggest some further indications of VS later in the paper. For the moment, I'll limit myself to noting that the extent to which the posts and tweets of the instant expert are an outgrowth of a desire for knowledge is to some degree belied by their confidence. The certainty with which most takes are expressed doesn't seem to suggest the tentativeness or humility of someone attempting to understand. They present themselves as settled wisdom. These facts suggest that they're better understood as instances of intellectual VS than as attempts at virtuous contributions to public discourse.

If hot takes had the function of signalling intellectual virtues and (thereby) achieving higher status, what properties would we expect them to have? Above, we saw that simplicity and concision are virtues in signals. But the best signals are *credibility enhancing* (Henrich 2009). Credibility enhancing signals are often costly signals: they're credibility enhancing because they give evidence of a willingness to pay the associated costs. Many signals of religious belonging are credibility enhancing because they're costly: think of tithing or fasting (Irons 2001; Sosis and Alcorta 2003). But some signals are credibility enhancing because they're self-ratifying. A stotting gazelle – leaping high in the air in proximity to a predator – *manifests* its athleticism and thereby signals that it would be hard for the predator to run down (FitzGibbon and Fanshawe 1988). It's much easier for intellectual VS to be credibility enhancing than moral VS.

It's *hard* to signal moral virtue with much credibility, at least in the kinds of contexts in which the accusation is most often made. We can manifest moral virtue in ongoing interactions with others, or when they can observe our behavior over a range of contexts, but on social media, or wherever interactions are with strangers, it far harder to manifest the virtues credibly. Moreover, in such contexts talk is usually cheap: signalling virtue to likeminded followers and friends, or when there are many such people as an audience to play to, is cost-free. It is in large part because credibility is

so hard to achieve in moral VS that such signals tend toward the brief and relatively low effort: why strain to achieve a goal that is out of reach? But intellectual VS can manifest some of the intellectual virtues. An intelligent and complex take *shows* what the person who produced it is capable of, rather than just *saying* it. The virtues of simplicity and concision retain their attractions, though, and we'd expect these properties to be traded off against one another. Effective intellectual VS will have some degree of complexity, but signals must be short and simple enough to keep audiences' attention and be comprehensible to them as a signal. The Twitter thread is ideal, especially in the way it breaks up the overall argument into bite-sized pieces.

Intellectual VS will therefore tend to manifest in complex takes. Attention to the goals of the signaller and the context in which he signals allows us to predict other properties and virtues such signals are likely to manifest.⁶ The signaller aims to attract and keep attention: for instance, to increase follower numbers on Twitter. To achieve this goal, they must manifest traits that others will value. The time and attention of their potential audience are both limited, and the competition is often intense. The successful signaller must offer something unique, or – at any rate – difficult for audience members to find elsewhere. While it is conceivable that some individuals might attract attention by signalling intellectual virtues like empathy and humility, in practice these virtues will rarely do well. They're relatively hard to display, and – worse – they don't attract attention (indeed, they may deflect attention, away from the person who displays it and on to those to whom they defer or those they empathise with). Other virtues are much better at attracting attention.

We should expect the successful signaller to exhibit not merely intelligence, but a kind of intelligence that manifests in *quickness of mind*. Attracting and keeping attention requires takes that offer something novel to the audience, and the logical space of views fills rapidly: the successful signaller must move early. He will probably also benefit from displaying *intellectual autonomy*. Reflectors of the mainstream line are fungible; the mainstream by definition has many sources. Those who repeat mainstream views cannot easily compete for attention on the basis of the content they offer. Thus offering a contrarian take gives the person a competitive advantage. This fact offers instant experts a motivation for proclaiming and for demonstrating their intellectual autonomy: the intellectually autonomous person is more likely to offer further unique takes in the future and is therefore worth following (to the extent their takes are valuable). For similar reasons, we should expect a display of *intellectual courage*. The successful displayer will tend to offer a unique take – and therefore dissent from the mainstream view, if there is one – and such dissent must be presented as virtuous.

We should therefore expect more successful instant experts to produce hot takes with a degree of genuine complexity (exhibiting intelligence) and that to some degree take issue with the mainstream view (exhibiting intellectual autonomy and courage). We should expect him to produce these takes rapidly (manifesting quickness of mind). Pity the instant expert! He has a genuinely hard task before him: to produce a complex, and at least somewhat plausible and somewhat distinct, take on a difficult topic, and do it quickly, before the logical space has been staked out. The pressure is the greater in light of the significant first mover advantages that can accrue. Of course, it is possible to compete for eyeballs on other bases: for example, one might defend the mainstream view especially persuasively or eloquently. Nevertheless, we should expect the niche offered by contrarian takes to be hotly contested. When we encounter hot takes that have these features, we are within our rights to suspect intellectual VS.

IV. The harms of VS

Tosi and Warmke argue that (moral) VS is pernicious because it breeds cynicism and increases polarization (Tosi and Warmke 2020). *Cynicism* arises because ordinary people recognize the pervasiveness of VS and come to doubt the sincerity of public moral talk. That, in turn, prevents it from functioning as it should, to allow us to identify moral wrongs and act to correct them. It reduces moral talk to a partisan shadow play. *Polarization* occurs because people are motivated to adopt more and more extreme political positions in order to distinguish themselves from others and thereby display themselves as admirable. Now that we have the example of intellectual VS before us, it's worth revisiting these claims. I'm going to suggest that these are much more persuasive with regard to intellectual VS than moral. In short, moral VS is unlikely to cause either problem because moral VS is so transparent: it can't easily be mistaken for anything other than a mere expression of attitudes.

As we've seen, paradigms of moral VS consist in bumper stickers and their equivalents: brief, simple declarations. The very features that make them apt for signalling – their clarity of content and ease of reception – also make them easy to distinguish from genuinely serious contributions to moral discourse. When I say *I stand with Ukraine*, or *Black Lives Matter*, I'm not pretending to offer an argument, an insight or a novel point of view. I'm merely expressing an attitude. Their simplicity and the concision of such declarations make them well-suited for the signalling role. But there's no particular reason to think these declarations constitute a pernicious kind of VS.

First, there's no reason to expect them to generate cynicism about serious moral contributions. These signals can't reasonably be mistaken for attempts to contribute to moral discourse, because they don't have the right sort of content. Their simplicity and concision entail that they offer little in the way of argument, and that fact is plain to see. So even if people become cynical about *these declarations*, seeing them as more concerned with displaying the insights or respectability of those who engage in them than with their ostensible objects, there's no reason to think they will become cynical about careful moral argument. Since mere declarations or expressions of attitude rarely play the role that Tosi and Warmke identify with the function of public moral discourse – “to identify publicly certain moral features of a state of affairs, and sometimes additionally to explain the evaluation of that state or recommend some fitting response” (Tosi and Warmke, 2016, pp. 209-210) – cynicism about such declarations risks little harm.

Cynicism about mere expressions of attitude isn't worrying; it doesn't threaten their capacity to play their role. Indeed, there's a case for saying it's not really cynicism at all to see such expressions of attitudes as signals: they just *are* signals. It's only when people begin to suspect that many tokens of apparently serious discussion are really just manifestations of signalling that pernicious cynicism might creep in. For that to occur, signallers are going to have to produce something more substantial than a mere expression of an attitude.⁷ In order for VS to induce cynicism about serious contributions to public discussion, it must *look* enough like such a contribution, and most moral VS doesn't have the right features to risk confusion with such contributions. For VS, moral *or* intellectual, to risk cynicism, it must be substantive. It must have the features that intellectual VS tends to have: complexity and length. It must *be* intellectual VS (or at any rate possess the properties distinctive of intellectual VS), whether or not it is *also* moral VS.

For similar reasons, *mere* expressions of attitude are unlikely to contribute significantly toward polarization. VS is supposed to cause polarization through people attempting to outdo one another in the competition for status (Grubbs et al. 2020). Success in such a competition requires distinguishing oneself from the crowd, not repeating the opinions already expressed; it's this fact that makes *ramping up* attractive. One can signal one's greater moral perceptiveness or concern by claiming to see even more iniquity in the behavior others have condemned. Again, though, *mere* signalling is likely to be ineffective, both in the competition for status (where it will be routinely outdone by signalling accompanied by credibility-enhancing displays) and in convincing others to take one's position seriously. To do well in the competition for status, and to drive polarization, the

signal must be accompanied by substance: it must be, or possess the properties distinctive of, intellectual VS.⁸

Intellectual VS is therefore far more likely to be harmful than mere signalling, and it is only when moral VS *also* has the properties distinctive of intellectual VS that we're likely to see it giving rise to pernicious consequences. Not only is intellectual VS far more likely than mere moral VS to give rise to cynicism and polarization, in addition, intellectual VS has bad consequences, and criticisable features, beyond those to which Tosi and Warmke point.

Whatever their motivations – whatever mix of a desire to thrust themselves into the spotlight and a desire to educate others might be in play – by moving so rapidly to offer their own take, instant experts manifest a disrespect for genuine expertise. Whether he cynically disregards his own lack of qualifications and the need for long apprenticeship in an area, or arrogantly takes himself to be uniquely gifted (or the genuine experts to be especially obtuse or corrupt), his apparent belief he is able to cast light on a topic better than those who have immersed themselves in it for years encourages his audience to follow him in disregarding expertise. To the extent that he succeeds in convincing others that his insights are valuable, he risks crowding out the genuine experts and contributing to the growing problem of a loss of trust in expertise. Moral VS – when it is not also intellectual VS – threatens neither harm.

Of course, there are borderline cases, and perhaps there are kinds of apparent expertise that doesn't deserve to be respected. Perhaps it is reasonable to believe that on some questions there is no genuine expertise. Alternatively, perhaps the instant expert reasonably believes that his genuine expertise on some *other* topic might allow him to shed light on this one (think of how someone with expertise in game theory might take themselves to be well positioned to understand military strategy). It's not easy to assess individual cases. Nevertheless, while I'm sceptical that any heuristic will be highly reliable, I think we can often navigate through difficult cases with some degree of confidence.

I noted above that content is a highly unreliable guide to intellectual VS: The skilled instant expert might not merely be able to mimic a genuine expert, he might in fact be able to produce genuinely new content. Confidently distinguishing the signaller from the better motivated contributor, given that both may produce substantive content, will require some knowledge of the person and their history. First, we need to ask whether they have relevant expertise: expertise on topic they're

discussing, or expertise reasonably held to be relevant to it. Possession of such expertise might indicate that a love of learning or a desire to educate plays a more dominant role in his motivations than the motivations of those engaged in signalling. We should be generous here, especially in light of the fact that expertise from unexpected quarters has proven valuable on many questions. But we should be more suspicious if the person has a *history* of hot takes: if they've thrust themselves into many discussions on many topics, as soon as each topic becomes fashionable (and therefore apt for exploitation as way to attract attention). Of course, some kinds of expertise *are* relevant to many topics. Economists, for instance, might reasonably take themselves to have something to contribute to discussion of COVID-19 (with regard to the costs and benefits of lockdowns, say), to the Russian invasion (the effects of sanctions), climate change (the costs of transitioning to renewables, versus the economic costs of the current trajectory), and so on. To the extent that they limit themselves to economic questions, wherever they arise, the economist is less suspect. We should balance generosity here in assessing what is within the scope of their possible expertise (where economics ends and politics begins is vague, after all) with increased scepticism arising from a history of inserting themselves into many debates very quickly, and obviously acquiring new knowledge in order to do so.

In brief: To the extent to which someone has a history of commenting *rapidly* and at *length* on difficult but disparate issues, we should suspect them of engaging in intellectual VS. Such suspicion may be *reduced* by their possession of expertise plausibly related to the topic, but *increased* when the person does not limit themselves to commenting on aspects with regard which their expertise is reasonably taken to be relevant. Suspicion is also increased by one feature of the content of their takes: contrarianism. To the degree that the person typically or always adopts a contrarian line on unrelated issues, they are more likely to be engaged in intellectual VS (it's important that the issues are unrelated: else the contrarian takes might reflect a single contrarian attitude, rather than a need to distinguish themselves from the crowd).

This heuristic won't allow us to be highly confident that any particular person is engaged in intellectual VS, and it will certainly miss very many people who are (just because someone *does* have relevant expertise doesn't entail they're not significantly motivated to comment by a desire for attention). It allows us to be somewhat confident that some specific individuals contribute to public debate to attempt to throw the spotlight on themselves, by advertising their intellectual virtues: their intelligence, their insight, their intellectual courage and their intellectual autonomy. Ironically, in doing so they manifest what are plausibly taken to be epistemic *vices*. Most obviously, they

manifest a lack of intellectual humility in disregarding the expertise of others. Perhaps they are often motivated by a genuine desire to teach others, but this motivation itself appears alongside intellectual vice, insofar as it manifests only when combined with the arrogance of believing that the instant expert can provide insights that domain experts can't.

It is worth emphasizing, however, that intellectual VS *per se* may not always be harmful; it is intellectual VS in the environment in which accusations of VS are most common that gives rise to these problems. Tosi and Warmke's recognition heuristic for moral VS is unhelpful because it is likely that almost all genuinely valuable contributions to moral and intellectual debate are sufficiently motivated by the desire for recognition that contributors would be disappointed were they not to be recognized for their virtues. Academia is structured so that contributors are more likely to receive recognition for genuinely valuable contributions than for those that attract attention on Twitter (though academic rewards are very far from perfectly correlated with genuine excellence). When the recognition desire is channelled such that agents must acquire discipline-specific expertise and exhibit care in evidence gathering, in formulating and testing hypotheses, in citing related work, and so on, intellectual VS may perform a valuable function.

Objections and clarifications

1. A critic might concede that there's something objectionable about the behavior of the instant expert, but argue that it's not VS. The behavior is objectionable because it manifests intellectual vices – immodesty, perhaps epistemic insouciance – not virtues. The concept of VS is not needed to explain what's objectionable about the behavior or its motivations.

Reply: There's something to this objection; the concept of VS isn't needed to explain what's objectionable about intellectual VS. It's the facts that it manifests vices and has bad consequences that make it objectionable, not the fact (if it is a fact) that it's VS. But that's characteristic of VS generally: it's not because it's VS that it is (or is held to be) objectionable; it's because it causes harm. Tosi and Warmke themselves (2020) argue that (moral) VS is bad on consequentialist grounds. Nevertheless, the behavior *is* VS: the instant expert seeks to manifest his insight, intelligence and intellectual autonomy to establish himself as someone worth following in the future. The concept of intellectual VS isn't needed to explain what's objectionable about it, but is it needed to explain its motivations.

2. I've conceded that the instant expert may work hard to acquire sufficient knowledge to produce takes that are credit-worthy. I've gone even further, suggesting that their intelligence and their occasional genuine expertise on some other topic may allow them to contribute real insights on the topic they address. Perhaps we're better off, epistemically, for their serial contributions to topic after topic. Even if most contribute nothing of value on most topics, the real possibility of insight is sufficient to ensure we should encourage, or not discourage, their contributions.

Reply: This suggestion shouldn't be construed as an objection to the claim that instant experts are engaged in intellectual VS. It's possible that instant experts are engaged in intellectual VS *and* that we're better off epistemically for their contributions. Indeed, defenders of (moral) VS don't deny that it's common; rather, their debate with Tosi and Warmke centres on its costs, not its existence (Levy 2021; Westra 2021). Perhaps intellectual VS is beneficial but nevertheless ill-motivated. That said, I'm sceptical that intellectual VS is beneficial.

On the credit side, we've conceded that instant experts may produce genuine insights. On the debit side, we should note two problems. First, attention is a limited resource, and by attempting to thrust themselves into the centre of debates, instant experts ensure that it's divided. Some of them are very effective at garnering attention, and thereby ensuring that genuine experts receive less attention than they otherwise would. Second, we've seen that instant experts have an incentive toward contrarian views. Since such views are contrary to the opinions of the genuine experts, they're significantly less likely to be correct across a range of domains.⁹ We should therefore expect instant experts to ensure that false claims get a wider hearing than they deserve. These facts suggest that the overall epistemic costs of their proliferation are greater than the benefits.

3. The charge of intellectual VS proves too much: it entails that journalists and (worse!) philosophers are guilty of objectionable behavior. The behavior of the instant expert is identical to what we expect of the good journalist. The journalist, too, often lacks domain expertise: rather, she's expected to cram about a subject as soon as world attention turns to it. We rely on journalists to inform us about the world; this behavior can't be objectionable. Philosophers, too, often lack domain expertise: most philosophers are not experts on epidemiology, or climate change, or Ukraine, yet many of us (us philosophers, that is) would see nothing wrong with the media asking our colleagues to write op-eds about these topics – indeed, many think they should do so much more often.

Reply: It's worth distinguishing two kinds of journalists. The first conveys information: her role is to report *the situation on the ground in Kyiv*, or *the pressure on the NHS*, or what have you. Her reports convey information she has gathered and vetted, to ensure a sufficient degree of reliability, as well as the opinions of *other people*. We often expect such journalists to strive for 'balance': that is, to report the full range of reasonable opinions on a topic, usually by interviewing people and crediting them with the opinions reported. Of course carrying out this work requires the journalist to acquire some degree of knowledge about the topic she's covering, but does not require her to thrust *herself* into the debate. Her role is to *report*, and she may be relatively self-effacing in carrying it out. To the extent to which she makes herself the story, she may be failing at her job.

The second kind of journalist produces opinion pieces. Her role *is* to produce takes, and she does so in precisely the way in which the instant expert does. Indeed, she just *is* an instant expert. Does the fact that we applaud and reward such behavior indicate that the takes of the Twitter expert should be seen as innocuous or beneficial? I think the inverse inference is more plausible: we should see at least some of the people who produce regular op-eds for our newspapers (or for TV panel shows) as engaged in intellectual VS.

There's a continuum along which op-ed writers can be located. At one end is the person who relies on their institutional role – as a columnist for the *New York Times*, say – to ensure attention. At the other end is the person who actively competes for attention. To the extent to which they engage in such competition, their behavior probably counts as especially high-profile intellectual VS. They aim to maintain their position by offering something unique. They may be less objectionable, on average, than Twitter experts because their institutional role ensures that they can compete for attention without needing to offer contrarian takes. But the pressure to produce takes on topic after topic, combined with the need to stay relevant and capture eyeballs, ensures that they're under pressure to engage in similar behavior.

Op-ed writers (and the like) may nevertheless fulfil a valuable function. Some may have genuine expertise that they reasonably take to bear on a wide range of topics: if they limit themselves (mainly) to the ways in which their expertise sheds light on these topics, they may produce insights. Some may range wider, or lack any special expertise, but serve as conduits for the expertise of others; there's room for elegant reporters and aggregators of expert opinion. But we probably would be better off without a stream of hot takes by people who lack relevant expertise. Op-ed space

would be better reserved for a much wider range of individuals, writing about topics they're well qualified to discuss (either on the basis of their expertise or on the basis of their experience).

Finally, let's consider the philosopher. Sadly or not, philosophers don't get many invitations to write op-eds. On the rare occasions we do, it's usually because we have some specific research interest that qualifies us to write on a particular topic. But we tend to think that we could turn our hand to most any topic. And some of us do express opinions – say on blogs – on a wide variety of topics. Is this behavior intellectual VS?

The philosopher is someone with genuine expertise, and like others with expertise might be qualified to comment on a wide range of topics in virtue of that expertise. Just as the expertise of the economist equips her to comment on lockdowns and sanctions and the minimum wage, so the philosopher might be qualified to comment on the same issues in virtue of her grasp of (say) theories of justice or domain-general expertise in argument. However, just as the economist should tend (at least) to restrict herself to the *economic* issues, so the philosopher should focus where her expertise is relevant. Philosophical expertise might be domain-general, but it doesn't follow that philosophers are well equipped to, say, comment on Russian military strategy. We may be well placed to identify mistakes in the arguments of others, but we're less well placed to generate new arguments of our own, except when we have domain-specific expertise. We should not expect each other to do better than non-philosophers in the production of takes on a wide variety of topics.

Conclusion

It's notoriously hard to be confident about other people's motivations (and perhaps even our own), but it's likely that VS is far from rare. However, most VS is probably harmless. There's no likelihood that anyone will mistake brief expressions of support for a view as anything more than a vote for a position. It's not going to get confused for a serious contribution to moral debate. It's unlikely to induce cynicism, and probably contributes little to polarization.

We might have thought that substantive arguments are less suspect when it comes to the charge of VS. That's almost backwards: it's only when substance is offered that VS risks cynicism, polarization and deflection from better justified views. When *moral* VS runs these risks, it is also, or at least also possesses the characteristics distinctive of, *intellectual* VS. The signaller may work hard to produce good simulacra of well-motivated contributions, and sometimes may even produce

genuine insights. Nevertheless, intellectual VS crowds out better contributions, and – worse – distorts debates by giving contrarian views much more attention than they deserve. The competition for attention may be an important driver for the recent prevalence of contrarian views on everything from climate change to COVID: the genuine skills of the successful instant expert at attracting attention combined with their bias toward contrarian views is a recipe for distorting debates.

We should be *less* cynical about the bumper-stickers and lawn signs of our neighbors or our Facebook friends, and *more* cynical about the genuine and often genuinely intelligent arguments of the instant experts. They both might count as VS, but only one does real harm.

Macquarie University and University of Oxford

ACKNOWLEDGEMENTS

This paper greatly benefited from feedback from audiences at the University of Melbourne, Macquarie University and the 2022 conference of the Gesellschaft für Analytische Philosophie. I am especially grateful for comments from Alex Gillett, Christopher Kavanagh, Matthew Browne and two anonymous reviewers for the *American Philosophical Quarterly*.

NOTES

REFERENCES

- Cassam, Quassim. 2018. *Vices of the Mind: From the Intellectual to the Political* (Oxford: Oxford University Press).
- Daukas, Nancy. 2011. “Altogether Now: A Virtue-Theoretic Approach to Pluralism in Feminist Epistemology,” in *Feminist Epistemology and Philosophy of Science: Power in Knowledge*, ed. Heidi Grasswick (Dordrecht: Springer), pp. 45–64.
- FitzGibbon, C. D., and J. H. Fanshawe. 1988. “Stotting in Thomson’s Gazelles: An Honest Signal of Condition.” *Behavioral Ecology and Sociobiology* 23 (2), pp. 69–74.
- Frost-Arnold, Karen. 2014. “Imposters, Tricksters, and Trustworthiness as an Epistemic Virtue,” *Hypatia* 29 (4), pp. 790–807.
- Funkhouser, Eric. 2017. “Beliefs as Signals: A New Function for Belief,” *Philosophical Psychology* 30 (6), pp. 809–31.
- Grubbs, Joshua B., Brandon Warmke, Justin Tosi, and A. Shanti James. 2020. “Moral Grandstanding and Political Polarization: A Multi-Study Consideration,” *Journal of Research in Personality* 88, 104009.
- Grubbs, Joshua B., Brandon Warmke, Justin Tosi, A. Shanti James, and W. Keith Campbell. 2019. “Moral Grandstanding in Public Discourse: Status-Seeking Motives as a Potential Explanatory Mechanism in Predicting Conflict,” *PloS One* 14 (10), e0223749.
- Henrich, Joseph. 2009. “The Evolution of Costly Displays, Cooperation and Religion: Credibility Enhancing Displays and Their Implications for Cultural Evolution,” *Evolution and Human Behavior* 30 (4), pp. 244–60.
- Hill, Scott, and Renaud-Philippe Garner. 2021. “Virtue Signalling and the Condorcet Jury Theorem,” *Synthese* 199 (5–6), pp. 14821–41.

- Irons, William. 2001. "Religion as a Hard-to-Fake Sign of Commitment," in *Evolution and the Capacity for Commitment*, ed. Randolph Nesse (New York: Russell Sage Foundation), pp. 292-309.
- Jones, Karen. 1999. "Second-Hand Moral Knowledge," *Journal of Philosophy* 96 (2), pp. 55-78
- Kidd, Ian James, Quassim Cassam, and Heather Battaly. 2020. *Vice Epistemology* (London: Routledge).
- Levy, Neil. 2021. "Virtue Signalling Is Virtuous," *Synthese* 198 (10), pp. 9545–62.
- Marie, Antoine, and Michael Bang Petersen. 2022. "Moralization of Rationality Can Stimulate, but Intellectual Humility Inhibits, Sharing of Hostile Conspiratorial Rumors," OSF Preprints. <https://doi.org/10.31219/osf.io/k7u68>.
- Roberts, Robert C., and W. Jay Wood. 2007. *Intellectual Virtues: An Essay in Regulative Epistemology* (Oxford: Oxford University Press).
- Sosis, Richard, and Candace Alcorta. 2003. "Signaling, Solidarity, and the Sacred: The Evolution of Religious Behavior," *Evolutionary Anthropology: Issues, News, and Reviews* 12 (6), pp. 264–74.
- Tetlock, Philip E., and Dan Gardner. 2016. *Superforecasting: The Art and Science of Prediction* (New York: Random House).
- Tosi, Justin, and Brandon Warmke. 2016. "Moral Grandstanding," *Philosophy & Public Affairs* 44 (3), pp. 197–217.
- . 2020. *Grandstanding: The Use and Abuse of Moral Talk* (New York: Oxford University Press).
- Vickers, Peter. 2023. *Identifying Future-Proof Science* (Oxford: Oxford University Press).
- Westra, Evan. 2021. "Virtue Signaling and Moral Progress," *Philosophy and Public Affairs* 49 (2), pp. 156–78.
- Williams, Daniel. 2022. "The Marketplace of Rationalizations," *Economics & Philosophy*, March, 1–25. <https://doi.org/10.1017/S0266267121000389>.
- Wylie, Alison. 2003. "Why Standpoint Matters", in *Science and Other Cultures: Issues in Philosophies of Science and Technology*, ed. Robert Figueroa and Sandra G. Harding (New York: Routledge), pp. 26–48..
- Zagzebski, Linda Trinkaus. 1996. *Virtues of the Mind: An Inquiry Into the Nature of Virtue and the Ethical Foundations of Knowledge* (Cambridge: Cambridge University Press).

1 Tosi and Warmke (2020) prefer the term ‘moral grandstanding’ to VS. Their principal argument for this view is that the former is a better defined construct (Grubbs et al. 2019). That can’t be right, however: if those who prefer the term VS to ‘moral grandstanding’ find their operationalization useful, they can simply borrow it. I think VS is preferable, on the grounds that it preserves continuity with the public discussion of VS that Tosi and Warmke themselves acknowledge as fuel for their analysis. Note, too, that philosophical discussions of the phenomenon have adopted VS as the appropriate label (Hill and Garner 2021; Levy 2021; Westra 2021).

2 Two political scientists come closest to discussing intellectual VS. Marie and Petersen (2022) find that sharing of fake news is predicted by what they call the moralization of rationality – that is, the disposition to hold that it is a moral virtue to be rigorously rational – and that the effect largely disappears when ‘moral grandstanding’ is controlled for. That suggests that for many people, the moralization of rationality is driven by status seeking, with signalling as a means to this end. Marie and Petersen do not identify a distinct kind of epistemic VS; rather, they see the VS they focus on as moral VS, with rationality falling within the scope of moral concern.

3 In face-to-face interactions, signals of moral virtue may be very different, as a reviewer for this journal pointed out. I’ll say more about why moral VS on social media tends to manifest in brief communications in the third section of this paper.

4 Of course there are sides in a war, but there is little signalling value *when it comes to the audience who might read my tweets* in signalling my opposition to Russia.

5 Williams (2022) has recently probed the motivations of those who produce narratives about current affairs congenial to political partisans. While of course financial incentives play the dominant role for many, he suggests that a drive for attention and prestige explains much of this behavior. We should expect the same mix of motivations to explain intellectual VS.

6 I use the masculine pronoun for the signaller deliberately; my impression is that men are significantly overrepresented among instant experts.

7 Above, we noted that trumping up, alone of the characteristics Tosi and Warmke identify, won’t tend to appear in the form of mere expressions of attitude. Trumping up alone risks cynicism, in part, because it appears in the context of genuine argument, with a complexity that allows it to mimic serious moral discourse. But accusations of trumping up are likely too controversial to play a useful role in debates over pernicious VS, because such accusations often reflect substantive disagreements over whether something really is morally innocuous. There’s a structural reason for this. If trumping up is too transparent, it won’t induce cynicism. The trumped-up charge can’t be *too* unbelievable (“eating fruit is racist!”). But the more plausible the claim is, the more likely it is intended seriously. Examples of claims that are both trumped up *and* apt to induce cynicism about serious contributions to moral discourse are hard to find, because the two properties are in tension with one another.

8 There is one important exception to this rule. Feminists (in particular) have emphasised the need for moral deference in some contexts. Some people – women, people of color, and other members of marginalized groups – may possess a sensitivity to certain facts that makes them more reliable at discerning them than even other intelligent and well-intentioned individuals. Importantly, they may not always be able to offer dispositive arguments for their conclusions (Jones 1999). Instead, they might rightly call upon us to trust their insights. That entails that there are spheres in which there is room for a confusion between mere expression of attitudes and serious contributions to moral discourse: when the person who makes a declaration can plausibly claim to belong to a group that possesses such special insight into certain moral properties *and* when the grounds of these insights plausibly can’t be fully conveyed. Jones provides plausible examples. Outside this fairly narrow sphere, serious contributions to moral discourse are backed up by serious arguments. Note that saying this is a fairly narrow sphere is not taking a stand on the limits of standpoint epistemology: the epistemic insight that certain groups possess, according to standpoint epistemology, need not entail an inability to provide grounds for their judgments to those outside these groups. Wylie (2003) emphasises how standpoints are the product of reflection and provide tools for the articulation of arguments.

9 Expert judgment does not outperform lay judgment across all domains. Some domains are not expertise conducive, because feedback isn’t sufficiently fast and reliable and/or because causal mechanisms can’t be isolated sufficiently well (Tetlock and Gardner 2016). Some domains are expertise conducive but reliable expertise is nevertheless lacking because epistemic networks are not structured appropriately or lack appropriate diversity (Vickers 2023). But expert judgment is significantly more reliable than lay judgment in very many domains.