

## **Culpable Ignorance and Suppressed Disagreement in the Context of Evidence-Based Policy: The Case of South African AIDS Denialism**

### **Preface**

It is relatively uncontroversial to suggest that policy makers should defer to the experts on issues of scientific fact when developing policy. However, it is less clear what policy makers should do in situations where the scientific community is divided, or when it is difficult to tell from the lay-perspective what the mainstream scientific view is. One approach in these situations might be to consult one's epistemic peers – those who have the same or similar evidence as the policy maker, and who are equally good at reasoning – on which experts to trust. However, consulting one's epistemic peers might itself not be very straightforward. Today I am going to talk about some problems with drawing on epistemic peers in the context of policy development.

This talk comes out of work I have recently done in attempting to assess former South African President, Thabo Mbeki's, moral responsibility for the AIDS denialist policies he pursued in South Africa in the early 2000s.

From 28 October 1999 to 26 September 2000 Mbeki endorsed the position of 'denialist' AIDS scientists – a marginal group who oppose the claim that HIV causes AIDS, and believe that ARVs (the medication used to treat the disease) are toxic. He went on to develop policy on the basis of this endorsement. In particular, he prevented ARVs from being distributed via the public health system. This policy persisted until 2004, with tragic consequences. Best estimates indicate that it resulted in 171, 000 avoidable new infections and 343,000 deaths over the 1999 – 2002 period (Nattrass, 2008, p. 157). As a result, members of South African civil society have called for an assessment to be made as to whether Mbeki can be held personally morally responsible for the harms that resulted from this policy (see, for example, Geffen (2009)). I take this call from civil society seriously.

In order to make the moral assessment in this case (and in policy cases like this one), it becomes important to consider issues of responsibility and ignorance in the policy making process – this will be done in the process of this talk. I will attempt to do this by linking the issues of culpable ignorance in the ethics literature to topics of peer disagreement in the social epistemology literature, particularly in instances where peer disagreement is complicated (like it was in the Mbeki case), and I to do this with an eye on the context of policy development.

## **Introduction**

The puzzle to be addressed in this talk is as follows. Expressed disagreements place an obligation on the agent to pay attention to the dissenting interlocutor, or risk culpable ignorance for which they might later be found blameworthy. Silence, on the other hand, is typically taken as assent. However, in cases of suppressed disagreement, the silenced interlocutor has information that could save the agent from ignorance in scenarios where that ignorance might lead to harmful action, and silence does not actually indicate assent. The problem is further complicated because the agent might not be aware of the fact that a silenced interlocutor has information that could prevent ignorance, and consequent harmful action. This talk will provide an account of the obligations on agents in cases of suppressed disagreement. This is particularly relevant to the Mbeki case, because members of his political party had information that could have prevented his ignorance, but their disagreement was suppressed.

This talk will be structured as follows. I begin by explaining the relationship between moral responsibility, culpable ignorance and disagreement, and I introduce the problem of suppressed disagreements. I then look at the standard philosophical response to examples of suppressed disagreement, which has been to distinguish between cases of actual disagreement (disagreement that someone has actually expressed) and merely possible disagreement (disagreement that has not been expressed by anybody, but is possible in the modal sense – there is a possible world in which an interlocutor disagrees). However, as has been correctly pointed out in the literature, giving epistemic weight to merely possible disagreement leads to strange consequences and should be avoided. Instead, I propose distinguishing between expressed disagreements, and suppressed actual disagreements (instances in which someone does actually disagree with the agent, but has been unable to express that disagreement due to suppression). I will argue that in certain circumstances there will be markers that suppression has occurred (such as silence when one would expect there to be debate). Depending on the context, and the role of the agent in that context, the agent might have an obligation to be on the lookout for the markers of suppression. I also argue that the agent will be additionally blameworthy if they created the conditions under which suppression occurred. I conclude by applying the lessons from the rest of the talk to the Mbeki case; assessing whether he actually suppressed disagreement, and the implications for his blameworthiness.

## **Moral Responsibility, Culpable Ignorance and Disagreement**

When you undertake an action with harmful consequences you might be found blameworthy for it (Smith, 1983, p. 543). Being blameworthy, on at least one popular account, means that we can legitimately adopt certain negative reactive attitudes toward you – such as anger, resentment, and disappointment (Strawson, 1962/2008) –

and that it would be appropriate for you to adopt negative reactive attitudes toward yourself – such as regret, guilt, and remorse (Williams, 1981 ; Fischer & Ravizza, 1998, pp. 5-7; Levy, 2005, p. 2). Blameworthiness is distinct from merely being accountable for your actions. Accountability tracks whether the harmful action can be causally attributed to the agent, while blameworthiness tracks the appropriateness of adopting negative reactive attitudes towards them. If you put arsenic in my tea and it kills me, then you are accountable. However, ignorance frequently excuses in moral cases (Rosen, 2004, p. 298; Smith, 1983; Zimmerman, 1997; Strawson, 2008, p. 3; Alvarez & Littlejohn, Forthcoming; Aristotle, NE1113b24-16). For instance, if the sugar in your sugar bowl underwent a freak chemical transformation that created the arsenic, and you did not know that this had occurred when you put it into my tea, then you would not be blameworthy for my death. However, ignorance only excuses if it is itself blameless. If you had been storing arsenic in your sugar bowl and had just forgotten about it when preparing my tea, then your ignorance *is* blameworthy (because it is the result of negligence) and so it does not provide an excuse for your action. It does not get you off the hook morally.

One can avoid culpable ignorance by satisfying one's 'procedural epistemic obligations'. This is the requirement that one should take due care when forming beliefs that inform actions with potentially harmful consequences (Rosen, 2004, p. 301). In Rosen's words:

Now among the required precautions against negligent harm are certain *epistemic* precautions. As you move through the world you are required to take certain steps to inform yourself about matters that might bear on the permissibility of your conduct. You are obliged to keep your eyes on the road while driving, to seek advice before launching a war and to think seriously about the advice you're given; to see to it that dangerous substances are clearly labelled, and so on. These obligations are your procedural epistemic obligations... As I understand them, these procedural obligations are always obligations to *do* (or to refrain from doing) certain things: to ask certain questions, to take careful notes, to stop and think, to focus one's attention in certain directions, etc. The procedural obligation is not itself an obligation to know or believe this or that. It is an obligation to take steps to ensure that when the time comes to act, one will know what one ought to know (2004, p. 301).

One way to take care when forming beliefs is to pay attention when an epistemic peer disagrees. For instance, imagine that you are a high-ranking member of the army trying to decide whether to bomb a site and you have a thermo-imaging map to help you make the decision. You notice that there is a warm spot to the side of the map and conclude that the heat pattern is consistent with there being a weapons manufacturing plant in that location. Your colleague, who has the same training as you and is looking at the same map, disagrees and suggests that the heat pattern is also consistent with the site being a hospital. You should pay attention to what your disagreeing colleague has to say.

Otherwise, you risk being culpably ignorant and morally blameworthy if things go very badly wrong (in this case, if it turns out that it was really a hospital and not a weapons factory).

Now consider a completely different case. A tyrannical dictator kills off any would-be dissenters before they have a chance to disagree with him. On some particular issue there would have been considerable disagreement, but because everyone who would have disagreed is now dead, there is nobody to give expression to that disagreement (Kelly, 2005, pp. 181-182). We can imagine that the dictator undertakes some harmful action (completely independently of his initial killing of the would-be dissidents) on the basis of a belief that went uncontested. However, as noted above, ignorance can provide an excuse for seemingly blameworthy action – perhaps he did not know that the action would have the harmful consequences that it did. In which case, we need to assess whether the ignorance itself was blameworthy, and the dictator might argue in response that he satisfied his procedural epistemic obligations by looking out for disagreeing interlocutors, but that no disagreement was forthcoming (because everyone who would have disagreed is dead). Would we be convinced by his excuse from ignorance? It seems unlikely that we would (or that we should).

Cases of suppressed disagreement, like this, are puzzling. Expressed disagreement from one's epistemic peers creates an obligation to pay attention to that disagreement (Christensen, 2007 ; Feldman, 2006 ; Elga, 2007). Silence, on the other hand, is often taken as agreement with what has been asserted (Goldberg, Forthcoming; Goldberg, 2010; Goldberg, 2010a). But taking silence as agreement in cases of suppressed disagreement would be a troubling outcome. The implication would be that if we were able to silence potentially disagreeing interlocutors before they have a chance to actually express disagreement that would make our ignorance less culpable and our resulting actions less blameworthy. Or as Lammenranta phrases the problem: "[t]his would mean that we could gain knowledge by killing our opponents" (Lammenranta, 2011, p. 211) – presumably, he means so long as we manage to kill off our opponents *before* they have a chance to express their disagreement. This is clearly the wrong outcome, so we need to think more carefully about how to formulate the requirement to take account of epistemic peers.

The problem of suppressed disagreement is an issue in the Mbeki case. Part of what seems to have gone wrong in this case is that he failed to be appropriately receptive to disagreement with his view. If correct, this would make his ignorance culpable and his resulting actions blameworthy. The most obvious form of disagreement that he should have paid attention to came from the scientific community in the form of the Durban Declaration – a petition signed by over 5,000 scientists affirming the mainstream view on HIV and AIDS (Durban Declaration, 2000). But, on a charitable reading of the Mbeki case, he was confused about who the experts were, thinking that the denialists had a genuine claim to scientific expertise, and that they had been unfairly marginalised by

the scientific community. The petition would have done nothing more than restate one position in the AIDS debate as far as Mbeki was concerned. Thus, when he looked to his epistemic superiors for guidance on what to do (which he did when he assembled a Presidential Advisory Panel) he got what looked like contradictory advice. The debate thus shifts a level, and instead of being directly about the issues of scientific fact, it becomes a debate about which scientists to trust. As Coady (2006) argues, being able to assess who the experts are is a task which itself requires expertise (p. 71). In assessing which scientists to trust, Mbeki should have turned to his epistemic peers for guidance.

But why pay attention to epistemic peers in what is essentially an expert debate? Suppose that I am a philosopher and I don't know anything about cars. My epistemic peers are all philosophers too and they also don't know anything about cars. I think I need a new car battery. My epistemic peers disagree, but my mechanic agrees with me. Disagreement from my epistemic peers in this case is irrelevant. The reason why disagreement from my epistemic peers is irrelevant is because there is a clear epistemic superior (someone with more/better evidence, or better reasoning skills than me) to whom I ought to defer.

Imagine a slightly different case. I am a philosopher who knows nothing about cars. This is the first time that something has gone wrong with my car. I think it needs a new battery, but because my car has always run well, I don't have the details of a reliable mechanic that I can contact for advice. In this case it seems like a completely appropriate epistemic strategy for me to ask around the philosophy department tearoom to see if anyone has the number for a good mechanic. When there is no obvious epistemic superior to defer to, then consulting one's epistemic peers on which experts to trust is a good strategy.

In the Mbeki case, his epistemic peers on the question of which scientists to trust would plausibly have been other members of his political party (the ANC) – it is reasonable to assume that they had similar levels of reasoning ability to Mbeki and access to similar evidence, thus making them his epistemic peers on this topic. However, when we look at state of disagreement from within the party, there is eerie silence from within the ANC caucus. Commentators on the case argue that the silence was the result of members of the party fearing Mbeki, even though they did actually disagree with him (Steinberg, 2017, forthcoming, p.5; Feinstein, 2009). But why single out members of his political party as the interlocutors to pay attention to, when there was substantial and actually expressed disagreement from other groups?

Let us briefly have a look at the state of disagreement with Mbeki's view, in order focus in on why the state of debate within the ANC was so important, and thus why the issue of suppressed disagreement is relevant for understanding this case. As already noted, there was disagreement from the scientific community in the form of the Durban Declaration, but from Mbeki's perspective this was just a restatement of one side of a

contested debate, with the Declaration merely endorsing the more popular view. There was substantial disagreement from the international press with his view (see Herbst (2005) and Johnson (2007) for examples of this), but Mbeki might plausibly have dismissed the press as being his epistemic inferiors in this case – they did not have access to the same evidence he had in the form of the expert testimony he received from the Presidential Panel (the Panel’s meetings were closed to the press, except for the opening and closing events (Cherry, 2009)). There was also disagreement from civil society groups, most notably the Treatment Action Campaign (TAC) (Geffen, 2010; Natrass, 2007 ), but again he might have dismissed them as his epistemic inferiors, because they did not have access to the evidence that he had.<sup>1</sup>

Mbeki also believed that the international community and local civil society groups were biased, and that their views on AIDS were motivated by racism (Fassin & Schneider, 2003, p. 496). The concern about racism being attached to beliefs about AIDS was not without historical precedent. The Truth and Reconciliation Commission in the early days of South African democracy uncovered that experiments had been conducted by the apartheid state in an effort to create biological weapons to target the black majority (with particular emphasis on sexually transmitted diseases, and with the aim of causing sterility). When that project failed, HIV positive prostitutes were allegedly planted on the mines, with the intention that this would facilitate the transmission of the virus throughout the migrant labour paths associated with the industry. Senior members of the apartheid state also expressed approval of HIV, hoping that it would “eliminate” the black majority (Fassin & Schneider, 2003, p. 496; Van der Vliet, 2001, p. 156).

While it seems like a stretch to link the apartheid state’s racist attitudes about AIDS to the views of the international community and to civil society groups, it is clear that Mbeki felt that these groups were biased in similar ways.<sup>2</sup> In the *Castro Hlongwane* document (an anonymous AIDS denialist monograph that was circulated in South African parliament in 2002, and for which Mbeki finally claimed authorship in 2016) it is argued that there are large financial gains for the “global north” if HIV causes AIDS and if there is an AIDS crisis in southern Africa, not just for the pharmaceutical industry (and their associated governments) who would make large profits from the sale of drugs; but also for the global health “industry” of NGOs and international organisations, who receive grant money and salaries on the back of the African AIDS crisis. From Mbeki’s perspective, pharmaceutical companies, their governments, international health organisations, NGOs and the press from global north were all untrustworthy

---

<sup>1</sup> I am sceptical of characterising civil society groups in this way, given the compelling case that Epstein gives for how AIDS-orientated civil society groups gained substantial expertise in the United States (Epstein, 1996), and the evidence that a very similar process took place in the TAC (Geffen, 2010; Stephen, 2009 p.174). But for the sake of a charitable interpretation, I will take this to have been the case.

<sup>2</sup> Fassin (2007) provides an extensive account of how apartheid era racism surrounding AIDS later became connected to suspicion surrounding AIDS during the denialism year in his book, *When Bodies Remember*.

because of their vested interests in the existence of the disease, and it being treatable via antiretroviral therapy. Further, these financial incentives had a racialised component. The following section taken from the *Castro Hlongwane* document shows how Mbeki saw the connection between the financial incentives attached to the African AIDS epidemic and the racist beliefs about African sexuality that helped to support the financial bias:

[T]he conviction has taken firm hold that sub-Saharan Africa will surely be wiped out by an HIV/AIDS pandemic unless, most important of all, we must access anti-retroviral drugs. This urgent and insistent call is made by some of the friends of the Africans, who are intent that the Africans must be saved from a plague worse than the Black Death of many centuries ago. For their part, the Africans believe this story, as told by their friends. They too shout the message that – yes, indeed, we are as you say we are! Yes, we are sex-crazy! Yes, we are diseased! Yes, we spread the deadly HI [sic] Virus through our uncontrolled heterosexual sex! In this regard, yes we are different from the US and Western Europe! Yes, we, the men, abuse women and the girl-child with gay abandon! Yes, among us rape is endemic because of our culture! Yes, we do believe that sleeping with young virgins will cure us of AIDS! Yes, as a result of all this, we are threatened with destruction by the HIV/AIDS pandemic! Yes, what we need, and cannot afford, because we are poor, are condoms and anti-retroviral drugs! Help! (Anonymous, 2002)

He was similarly distrustful of local civil society groups. In a public lecture at the University of Fort Hare in 2001 he made the following statement (clearly about the TAC):

And thus does it happen that others who consider themselves to be our leaders take to the streets carrying their placards, to demand that because we are germ carriers, and human beings of a lower order that cannot subject its passions to reason, we must perforce adopt strange opinions, to save a deprived and diseased people from perishing from self-inflicted disease (Mbeki, 2001).

It is clear from the above statement that Mbeki took the TAC's stance on AIDS to be motivated by racist views of African sexuality. Steinberg points out that this was not helped by the fact that most of the senior leadership of the TAC were white, thus making the clash between a black government and a white-led civil society group seem especially liable to interpretations of racism (Steinberg, 2017, forthcoming).

While there was substantial disagreement from a variety of groups with Mbeki at the time, we can see that he might dismiss them as his epistemic inferiors due to their not having access to as much evidence as he had, and he suspected that racist and financial biases motivated their disagreement. Goldman (2001) argues that if you have evidence that a purported expert is biased, then it is legitimate to downgrade your trust in their testimony (p. 93). Presumably the same considerations would also apply to those further down the epistemic food chain – we should place less trust in *anyone* who we

have reason to suspect of being biased. So while Mbeki might have thought there was reason to doubt various forms of disagreement with his view, he could not dismiss other senior members of the ANC as easily. They would have had access to the same (or very similar) evidence, and there was also no reason for him to suspect them of being racially biased – they were his “comrades” in the struggle against apartheid, and they were now leading the newly democratic South Africa. Members of Mbeki’s political party thus occupy a unique position in the South African AIDS debate in the late 1990s and early 2000s, in that they are the one group of individuals who could have disagreed with Mbeki with unquestionable moral legitimacy (from Mbeki’s perspective). However, singling out Mbeki’s fellow ANC members as the ones to pay attention to makes things complicated, because disagreement from within the party was suppressed.

This looks remarkably similar to Kelly’s tyrannical dictator, in that both are cases of suppressed disagreement. The following section will outline the philosophical reactions to Kelly’s dictator case to see if they provide any useful advice on how to handle the Mbeki case, given the striking similarity between the two cases.

### **Is there a distinction between actual and merely possible disagreement?**

In this section I provide an outline of the debate that has arisen in response to Kelly’s example of the tyrannical dictator – in particular, the question of whether ‘merely possible disagreement’ should be afforded the same epistemic status as ‘actual disagreement’. This is relevant to the broader project of this talk because if it is the case that possible disagreement is as epistemically weighty as actual disagreement, then this would solve the challenge posed by the cases of the tyrannical dictator and Mbeki. In both cases, if merely possible disagreement were as epistemically important as actual disagreement, then they should have been on the lookout for possible disagreements with their view; they were not, and so their ignorance is culpable and they are blameworthy. However, as will be seen, this route out of the problem is not available.

Kelly takes his example of the tyrant to indicate that there is no distinction between actually articulated disagreement and merely possible disagreement. What matters for determining whether a disagreement should threaten one’s existing beliefs, he argues, is the strength of the arguments and the evidence put forward in support of that disagreement – regardless of whether the disagreement is actual or merely possible. As such, the disagreement itself is irrelevant and one should focus on the arguments and evidence that could be levelled against a particular belief, not the disagreements themselves (Kelly, 2005, pp. 181-182). Kelly:

Whether we find the possibility of disagreement intellectually threatening, I suggest, will and should ultimately depend on our considered judgments about *how rational* the merely possible dissenters might be in so dissenting. And our assessment of whether rational dissent is possible with respect to some question



(or our assessment of the extent to which such dissent might be rational) will depend in turn on our assessment of the strength of the evidence and arguments that might be put forward on behalf of such dissent. But if this is correct, then the extent to which merely possible dissent should be seen as intellectually threatening effectively reduces to questions about the strength of the reasons that might be put forward on behalf of such dissent... The role of disagreement, whether possible or actual, ultimately proves superfluous or inessential with respect to the case of scepticism (Kelly, 2005, pp. 181-182).

The way that Kelly sets up the argument leaves us with something of a dilemma. Kelly uses the case of the tyrannical dictator to argue that there is no difference between merely possible disagreement and actual disagreement, which in turn he uses to cast doubt on testimony as a form of evidence – it is not the testimony that matters, it is the underlying arguments and evidence that matter. If we side with Kelly, then it is unclear why an agent would have any obligation to pay attention to disagreeing interlocutors, because testimony (in general) is not a form of evidence, and disagreement (a particular type of testimony) does not matter. At first glance, this leaves us at a loss for how to deal with a large portion of cases that seem like they should be cases of culpable ignorance. Kelly might argue that in these cases the agent should have thought about whether someone would have good reasons for dissent, but he places a lot of faith in the abilities of ordinary epistemic agents to anticipate all the possible reasons for dissent and treat them with appropriate weight. If we want to resist Kelly (and hang on to the position that neglecting peer disagreement is a route to culpable ignorance), then it seems as though we need to reject the claim that actual disagreement and merely possible disagreement are on par. But then we lack the resources to deal with the tyrannical dictator case. Even though Kelly's position is plausible – it is the arguments and the evidence that matter, not the testimony itself<sup>3</sup> – Kelly asks too much of reasoners that they anticipate all counter-examples – testimony continues to have an important role to play. In this section I will focus on the other horn of the dilemma. I argue that actual disagreement and merely possible disagreement should not be put on a par, but I will also show that this does not actually strip us of the intellectual resources to deal with the tyrannical dictator and Mbeki cases.

In the rest of this section I will consider two arguments that have been put forward in the literature in response to Kelly's argument regarding the tyrannical dictator case: 1) accepting the position that actual disagreement and merely possible disagreement are as epistemically weighty as each other results in unacceptable scepticism (the Sceptical Argument); 2) actual disagreements provide a signal that an error may have occurred and merely possible disagreements cannot do this (the Signalling Argument), and so Kelly is wrong to put them on a par.

---

<sup>3</sup> Goldberg (2006) also argues that testimony should not be taken as a form of evidence.

The Sceptical Argument tells us that if actual disagreement and merely possible disagreement were on par, this would put us in an untenable philosophical position, because for any belief that one might have, there is always a *possible* disagreeing interlocutor. This would require that we suspend (or substantially revise, or diminish our confidence in) all of our beliefs, resulting in widespread scepticism (Kornblith, 2010, p. 34; Carey, 2011, pp. 374-377). It should be noted that Kelly's position does not commit him to the sceptical conclusion. As explained above, Kelly argues that disagreement is irrelevant, which means that no one is required to alter their beliefs in the face of actual disagreement or merely possible disagreement, so he avoids the sceptical conclusion. But regardless of what Kelly himself is actually committed to, it still seems worth considering the implications of taking merely possible disagreement as seriously as actual disagreement, because if this were the case then this would provide a solution to the suppressed disagreement problem.

One way of taking merely possible disagreement seriously, without falling into the trap of scepticism, would be to consider possible disagreements only in nearby possible worlds. This would amount to engaging in some abstraction from the actual world, but not considering disagreement from *all* possible worlds. This would evade the problem that there is always *some* possible world in which an interlocutor disagrees. This approach is appealing, and coheres with some common sense views about what one ought to do when considering potentially harmful action. Dryzek and Niemeyer (2008) advocate for something similar when they argue that deliberative democracy requires that all relevant viewpoints be considered when making a decision, not just those that are actually represented:

Rationality may even benefit from the presence of a vantage point to which nobody subscribes; such was presumably the rationale for the use of a "Devil's Advocate" when evaluating cases for sainthood in the Catholic Church (Dryzek & Niemeyer, 2008, p. 482)

Considering disagreement from nearby possible worlds amounts to taking into account viewpoints that may not actually be represented at the decision making table, but which bear on the issue under consideration, and this looks like very sensible advice. For instance, we can imagine that a committee needs to make a decision that will affect the homeless community in the city. A proposed policy is being discussed, but there are no homeless people in the room to express their disagreement with a proposed policy, and the homeless community has no one to represent their views to the committee. In these circumstances, it would be good for the committee members to imagine what disagreements the homeless community might have with the proposal, even though there are no actual homeless people in the room to give expression to their viewpoint. This is the same as (or very close to) considering a nearby possible world in which there is a homeless person in the room to give expression to the dissenting view. While this would be a good thing to do, we now seem to have strayed from the realm of disagreements.

What is going on when you consider disagreements from nearby possible worlds? Given that you are acting as your own disagreeing epistemic peer in this scenario – you have the same reasoning ability and the same evidence as you actually have – considering disagreements from nearby possible worlds amounts to carefully thinking through the evidence that you already have. It therefore looks like taking account of disagreements from nearby possible worlds just means that you should think really hard before engaging in actions with potentially harmful consequences, and this was already covered by the advice provided by Rosen earlier in this talk, when his stance on procedural epistemic obligations was described. He advises that satisfying one’s procedural epistemic obligations requires that one ought “to stop and think, to focus one’s attention in certain directions...” before engaging in an activity that could have harmful consequences (Rosen, 2004, p. 301), and so advising that one should consider disagreements from nearby possible worlds does not seem to add anything useful. Considering only nearby possible worlds does not provide us with a route out of the Sceptical Argument, because we have changed the case too much – we are no longer really dealing with actual/possible peer disagreement.

Another way of dealing with the Sceptical Argument would be to bite the bullet and accept the sceptical conclusion. Ballantyn (2015 ) does something similar when he suggests that we should be epistemically modest about almost all of our beliefs, due to existing, but un-possessed, evidence that runs counter to our views. Ballantyne’s argument goes like this. We live in an era of information glut. As such, it is very likely that for any belief that you hold, someone has made a good argument with good evidence against that view. Even if you don’t yet know the content of the dissenting argument, you know that it is likely to exist and that it is likely that at least some of the dissenting arguments are compelling. He suggests that all you need to do is walk through your nearest university library to see that for many of your most interesting beliefs there are a wide variety of books that have been written on the topic, many of which go against your view. This problem is augmented if you do a Google Scholar search. His suggestion is that the appropriate reaction to this realisation is that you should be epistemically modest about nearly all of your beliefs, and he thus accepts a weaker version of the sceptical conclusion.

Ballantyne’s argument still does not help us very much with the cases that have motivated this talk – that of the tyrannical dictator, or with the Mbeki case – because his hypothetical agent is in a very different epistemic position to the kinds of agents under consideration. In Ballantyne’s cases, the disagreements have been actually expressed (people have written books and articles on the relevant topics, those books have been published and put in the library, the articles have been catalogued and put on Google Scholar), and the agent knows that they have been expressed (they have browsed the university library, and they have done a Google Scholar search); she just doesn’t yet know what the content of the disagreement is. This is different from the Dictator and

Mbeki cases, in that disagreement from epistemic peers did not have a chance to be actually expressed in either of these cases. It is also unclear whether the Dictator and Mbeki are even aware that there is a disagreement that has been suppressed. So while Ballantyne offers us a plausible way of accepting something akin to the Sceptical Argument – that we should be epistemically modest about most of our beliefs – he does not offer us a way out of the Sceptical Argument.

Ultimately, the problem with the Sceptical Argument is that it proves too much; it means that almost all people are at fault for having normal epistemic convictions. Claiming that most people should suspend almost all of their beliefs (as the Sceptical Argument requires), or risk culpability, does not seem plausible.

Another argument for why actual disagreement should not be put on an epistemic par with merely possible disagreement is that actual disagreements give us information that merely possible disagreements do not – this is the Signalling Argument (Tersman, 2013 ; Carey, 2011 ). To see how this argument works it is helpful to take a step back and note that the reason conciliationists – those who think that we are required to revise or reconsider our beliefs when we discover that an epistemic peer disagrees with our view – believe we should take disagreements seriously is due to a kind of inference to the best explanation. Having a disagreement with someone who is as competent at reasoning as I am and has access to the same evidence that I do indicates that one of us has made an error, but we cannot tell from the mere fact of disagreement which one of us has erred. Sidgwick makes this point as follows:

For if I find any of my judgements, intuitive or inferential, in direct conflict with a judgement of some other mind, there must be some error somewhere: and I have no more reason to suspect error in the other mind than in my own... (Sidgwick, 1981/1907, p. 342).

Christensen's (2007) famous dinner bill case makes this particularly clear.<sup>4</sup> Christensen asks us to imagine a scenario in which friends go out for dinner, an activity that they frequently participate in. At the end of the meal, two friends figure out how to split the bill. They are equally good at the kind of arithmetic that is required to perform the task and they both have the same evidence (the bill), but they come up with two slightly different amounts at the end of their calculations. It is clear that one of them has made an error, but it is not clear just from their disagreement which one of them it is (Christensen, 2007 , p. 194). Their disagreement has provided them with evidence of a mistake. That there is some modal world in which someone disagrees with a particular belief does not give the agent any new evidence (Carey, 2011 , p. 378). Merely possible disagreement does not serve the signalling function that actually expressed disagreement does, and so it should not be given consideration in the way that we give actual disagreements consideration.

---

<sup>4</sup> Both Kornblith (2010) and Carey (2011 ) use this example to make a similar point to what I make here.

I find both the Sceptical Argument and the Signalling Argument convincing. It does seem that failing to distinguish between the epistemic significance of actual and merely possible disagreement results in a scepticism so widespread that it cannot be acceptable. Further, giving merely possible disagreement epistemic status fails to recognise that the reason we give actual disagreements epistemic weight is because they provide us with evidence of potential errors.

However, accepting that merely possible disagreement should not be afforded the same epistemic status as actual agreement, or any epistemic status at all, does not mean that we are stuck with the conclusion that the tyrannical dictator and Mbeki are off the hook because no disagreement was actually forthcoming in their cases. In order to make sense of these cases we need to distinguish between “merely possible disagreement” (there is some possible world in which there is disagreement with my view) and “suppressed actual disagreement” (there is actual disagreement in this world, but that disagreement has not been expressed, because of suppression). While the former is not epistemically salient, the latter should be. Further, both the case of the tyrannical dictator and Mbeki are cases of suppressed disagreement, not of merely possible disagreement.

### **Actual, but suppressed, disagreement**

There is still something puzzling about how instances of suppressed actual disagreement might be epistemically significant. In both of the cases that have motivated this talk – the tyrannical dictator and the Mbeki case – we can imagine that the agents did not know that disagreement had been suppressed. The tyrannical dictator might have killed off the potential dissidents without knowing that they would later go on to disagree with him (he had them killed for reasons other than their dissent). Mbeki might not have known that the members of his party feared him.<sup>5</sup> In both of these instances, from the agent’s perspective it might just look like nobody has disagreed and silence might (wrongfully) be interpreted as assent – this is clearly a more widespread problem in policy making. How then could suppressed disagreement be epistemically relevant if the agents do not even know that it is happening? Phrased slightly differently, suppressed disagreement does not give them any new information, much like the case of merely possible disagreement, and so it is unclear how suppressed disagreement can have any epistemic relevance at all.

To get a grip on this issue, it will be helpful to better understand what a reasonable response to silence is. Imagine a case in which an agent makes an assertion, and is greeted only with stony silence in response – what would a reasonable interpretation of

---

<sup>5</sup> In fact, this seems like a plausible reading of Mbeki’s own assessment of what happened, given the content of his recent letters to the public, in which he attempts to defend himself against accusations that he had been ‘aloof’ and ‘overly sensitive toward criticism’ during his time as president (Mbeki, 2016).

that silence be? Are they getting any information from the silence? Should they be getting any information from it? Answering these questions will help us understand what the obligations are on agents in scenarios of suppressed disagreement. In particular, if it turns out that the agent should be getting information from the silence (which I will argue that they should, in certain circumstances), then this distinguishes suppressed disagreement from merely possible disagreement, because the agent is actually getting information from the (silent) response, while they would get no new information cases of merely possible disagreement.

In what follows, I address only instances of literal silence. But there is clearly a range of related phenomena in contexts of suppression. In particular, there will be cases where those who are suppressed still make utterances, but in which they have been effectively silenced. For instance, when West and Langton (1999) talk about pornography “silencing” women, they do not mean that women are literally no longer permitted to speak. They mean that women are no longer properly heard when they speak (to use the appropriate technical language, their speech acts do not achieve “uptake”) (McGill, 2013, p. 206). We can also imagine cases in which suppressed individuals explicitly assert assent, but only as a result of their suppression— when Kim Jong-un makes an assertion into a crowd in North Korea everyone will cheer regardless of whether they actually agree with what has been asserted.<sup>6</sup> The members of the North Korean crowd have been effectively silenced, even though they are not literally silent. I will not be dealing with these more complicated cases here, but hopefully some of the lessons from the following discussion will also help to shed light on these other kinds of cases. For now, I will just address the question of what an agent should reasonably take from literal silence in response to their assertion.

Goldberg argues that silence in response to an assertion is typically taken to indicate assent (forthcoming), and he later goes on to argue that one has a normative entitlement to interpret silence as assent under normal circumstances, because typically if someone disagrees with your assertion they will speak up (unpublished manuscript). By contrast, Tanesini (forthcoming) argues that a speaker has no such entitlement. On Tanesini’s view, when a speaker makes an assertion they invite listeners to respond, but the listener is under no obligation to do so, even if they do not accept the content of the assertion. Because the hearer is under no obligation to offer a response, even if they disagree, the speaker is not entitled to take silence to be acceptance (Tanesini, forthcoming, pp. 7-8). For example, I might give a talk in the Philosophy Department, and none of the audience members will be obliged to express their disagreement with me during the question and answer period, even if they do actually disagree. If nobody raises their hand when it comes time for discussion, that does not mean that I am entitled to believe that everyone in the audience agrees with

---

<sup>6</sup> Thanks to Susanne Burri and Wlodek Rabinowicz, who both independently suggested this example to me.

everything I have said – a more plausible explanation might be that I was so boring that nobody can muster the energy to disagree with me by the time the opportunity becomes available.

Also in contrast to Goldberg, Beatty and Moore (2010) suggest that we should find complete consensus with some position to be suspicious. Complete consensus, they argue, might indicate that some coercion was involved. In explaining a comment by Elster (1986/1997), in which he argues that he would feel more comfortable passing a law that had a persistent minority opposed to it than one that had unanimous support, they say:

... [W]hat had worried him was the possibility that a unanimous decision might be due to some sort of conformism – perhaps resulting from intimidation, the suppression of alternative viewpoints, or self-censorship – rather than from a proper deliberation of the alternatives (Beatty & Moore, 2010, p. 198).

They go on to say:

...[N]o one expects unanimity in politics; not in light of culture, class, gender, and other differences. If a diverse voting body were to report unanimous agreement on an issue or candidate, one might well wonder if all the parties had freely spoken their minds (Beatty & Moore, 2010, p. 199).

In summary, when an agent asserts something and receives silence as a response, it seems that this might indicate a range of things. Under some circumstances it might indicate assent, but in others it may just be that the listeners were too bored to respond, or that the disagreement had been suppressed. Listeners can have a range of reactions to an assertion and silence might mask some of those reactions. Or in some cases the silence might be an expression of the reaction – such as might be the case when the agent is too bored to bother responding.

In the second quote from Beatty and Moore above, they point out that silence is particularly suspicious in political environments, where we would *expect* there to be a wide variety of differing opinions and debates. When an agent makes an assertion about something in a political context and receives only silence in response, it seems like something strange is going on.

So far in this talk I have been treating the tyrannical dictator and Mbeki as though they were ordinary epistemic agents assessing their evidence in ordinary ways. In some ways, this is obviously the correct approach – an agent will only ever have the capacities of an ordinary epistemic agent (unless they are some kind of genius), regardless of the role that they occupy. However, it does seem that agents might have very specific epistemic obligations, depending on their social positioning – we expect medical doctors to keep abreast of recent developments in treatments, but we do not have

similar expectations of epistemic agents in other roles. Goldberg (2015) refers to this as an agent's 'role-specific obligations' (p. 3). How one ought to react to silence in response to an assertion might vary depending on context and the role that the agent is playing in that context. It is plausible that bearers of political office and policy makers have role-specific obligations to pay attention when their assertions are greeted only with silence and to find this a bit suspicious.

There seem to be at least two reasons why those occupying political office might have role-specific obligations to pay attention to the reactions to their assertions and be alert to the possibility of suppression. The first (as correctly pointed out by Beatty and Moore) is that one would expect debate in a political context, given the wide variety of perspectives and opinions that are expressed in political environments, and silence might indicate coercion or suppression. The other is that bearers of political office might have power that could result in others being cautious about expressing contrarian viewpoints, and they should be alert to this possibility.

This then solves the problem of 'suppressed disagreement' associated with culpable ignorance, without falling into the sceptical conclusion. This is done by distinguishing suppressed disagreement from merely possible disagreement, in that the former sometimes carries information, while the latter never does. In particular, silence as a response sometimes counts as information or evidence of suppressed disagreement – this is true in cases where the normal response would not be silence, such as in political contexts.

This then raises the question of how the silence came about in the first place. So far, the emphasis has been on the agent's responses to expressed disagreement and to silence. Nothing has been said about the blameworthiness of the subject relative to the circumstances that resulted in the disagreement being suppressed. This will be addressed in the following section.

### **The circumstances of silence**

In assessing whether an agent is blameworthy, and exactly what they are blameworthy for, it is important not just to have an account of what an appropriate reaction to silence might be, but also of how blameworthiness might be connected to the circumstances in which that silence arose. There are at least four possible scenarios here, and depending on which scenario the agent finds herself in, this will have an impact on whether she is blameworthy and what she is blameworthy for. They are (in order of what I suspect is increasing culpability):

- 1) The agent did not suppress the disagreement, and she could not reasonably have been expected to know about the existence of the disagreement, or that the



disagreement had been suppressed (e.g. silence in apparently normal conditions).

- 2) The agent did not suppress the disagreement, but there were cues that should have alerted her to the suppression (e.g. there was silence in a context where one would have ordinarily anticipated debate). That is, there were aspects of the context that were relevant to the assessment.
- 3) The agent suppressed the disagreement, but she did not know that she had suppressed the disagreement (perhaps she is oblivious to how frightening her peers find her), but there were cues that should have alerted her to the suppression (e.g. there was silence in a context where one would have ordinarily expected there to be debate).
- 4) The agent suppressed the disagreement and she knew that she had suppressed the disagreement (e.g. she is an Idi Amin style dictator who kills those who disagree with her, but it might not be as extreme as in this example suggests – intentional side-lining in politics would be sufficient).<sup>7</sup>

I take it the agent in scenario 1) is obviously not culpable for her ignorance, while the agent in scenario 4) is obviously culpable. The agents in scenarios 2) and 3) are more complicated

The agent in 2) will not be culpable for the circumstances of the suppression, but will still be culpable for her ignorance if she failed to pick up on the silence that was masking the suppressed disagreement (assuming the silence was unusual in that context). The agent in 3) will be responsible for her ignorance if she failed to pick up on the cues. Additionally, she may also be blameworthy for the circumstances of the suppression (Tanesini, forthcoming). For instance, perhaps the agent is arrogant, and it was her arrogance that suppressed the disagreement. Under these circumstances she might also be blameworthy for the epistemic injustice inflicted on those who were silenced (Fricker, 2007). This is not a small harm to inflict on someone. Miranda Fricker (2012) describes it thus:

The intrinsic wrong of testimonial injustice is the epistemic insult: the subject is undermined in their capacity as a knower, and so as a rational being. The insult goes deep. If we accept that our rationality is part of the essence of human beings' distinctive value, then to be perceived and treated as lesser in one's capacity as a knower is to be perceived and treated as a lesser human being. (p. 294)

Additionally, it should be noted that there might be aspects of the very structure of scenarios in which silence is the result of a power imbalance that may make it less likely

---

<sup>7</sup> Thanks to Jonathan Birch for suggesting these four scenarios to me.

for the agent to notice that they are suppressing disagreement. Again I draw on Fricker (2012):

If *you* are the one doing the crushing [suppressing the disagreement] ...then not only are you not in a position to know what it is like to be crushed, but also – and this is a separate point – your *general* picture of the social world in which such crushings take place will be an unhelpfully partial perspective, the perspective of the powerful. (p. 288)

It might just be an unfortunate fact about the social world that those who are most likely to unwittingly suppress disagreement with their views – the arrogant, the powerful, etc. – are those who are the least likely to notice that they are doing so. But this won't get them off the hook morally.

The case that Fricker uses to motivate what is at stake in scenarios of unwitting suppression is that of a police investigation into the stabbing and killing of a black teenager in London in the early 1990s by a group of white teenagers. The victim's friend witnessed the attack, but the police acted in ways that effectively suppressed his testimony. They did not help to calm him down after he had just witnessed the murder of a friend so that he would be in an appropriate state to testify, they did not ask him for his testimony, when he offered his testimony they did not take it seriously, and they did not take him along on searches of the local area even though he would have been able to identify the perpetrators of the crime if he had seen them. Fricker describes him as having been 'pre-emptively silenced' (2012, p. 293) (much like the potential dissidents in the tyrannical dictator case and the party members in Mbeki's case). As a result of pre-emptively silencing his testimony, the police were unable to gather enough evidence to convict anybody for the crime. In the subsequent report that assessed what went wrong in the investigation, it was concluded that institutional racism had played a role in the suppression of the key witness's testimony – the police officers on the scene assumed that the black teenager was part of the trouble and not a valuable source of information about what had happened. The report also noted that the policemen involved might not have been aware of their racism – it was just part of the culture of being a member of the British police force in the early 1990s (Fricker, 2012, pp. 291-300). It would be very strange to conclude that the police officers are not morally responsible for silencing the key witness's testimony, and ultimately bungling the case, because they were unaware of their racism, even though they might still be less culpable than they might otherwise have been due to their position in a social context where racism was not questioned.

This is slightly puzzling. Normally we only hold agents morally responsible for things that are under their control. That assumption has underpinned this whole talk – ignorance can provide a plausible excuse for otherwise blameworthy action because it can show that the agent was not fully in control of that action; had they known

differently they would have (hopefully) acted differently. But how then can we hold someone morally responsible for something that was ultimately the result of something that they may not even have been aware of – being arrogant or racist, or having some other character flaw that resulted in suppressing interlocutors who would have had valuable information to offer had they been allowed to offer it?

Given the similarities between the cases of the unknowingly racist police officer and the unwittingly arrogant interlocutor, a useful first place to look for help is to the literature on moral responsibility and implicit bias. The implicit bias literature is potentially useful because part of what characterises implicit biases is that the agent is unaware that they have them.

A range of views on culpability is represented in the implicit bias literature. On the one end of the spectrum are those who bite the bullet, and argue that we cannot hold individuals morally responsible for their implicit biases, because they lack awareness. Saul comments: “A person should not be blamed for an implicit bias of which they are completely unaware” (2013, p. 55). Similarly, Zimmerman (1997) argues that in cases where an individual is unaware of the wrongness of their racist action they cannot be held morally responsible for it (Zimmerman, 1997, pp. 425-426).<sup>8</sup>

Holroyd (2012) argues against Saul’s position, by drawing on Doris’s (2002) work in behavioural science and social psychology. Doris argues that we lack stable characters, and our actions are more the result of context and habit than anything else; most (perhaps all) of our actions are automatic. Holroyd uses Doris’s empirical findings to argue that if the vast majority of our actions are automatic, then we are not aware of the underlying motivations for almost all of our actions, and so awareness is too demanding a requirement for moral responsibility, and would require that we abandon attributions of moral responsibility in almost all cases (Holroyd, 2012, pp. 293-294).

Abandoning the awareness requirement strikes me as quite a blow to the way we usually think about moral responsibility, and too controversial a way to try and maintain attributions of blameworthiness in cases of arrogant interlocutors. I am going to approach the awareness requirement from a slightly different angle. The reason we are interested in awareness is because we are ultimately concerned with whether the agent is fully in control of her action, and it is control that allows for attributions of

---

<sup>8</sup> There is some difference between Saul’s view and Zimmerman’s. Saul is concerned with cases in which the individual is unaware that they hold implicit biases, even though they might be aware of the wrongfulness of holding such biases, while Zimmerman is concerned with cases in which the individual is unaware of the wrongness of the bias, regardless of whether or not they are aware of holding the bias. Saul’s concerns are closer to the concerns that I have in this talk – that is, how do we assess culpability in cases where the individual is unaware of aspects of their beliefs or character that result in their engaging in wrongful action?

moral responsibility. Ultimately awareness of one's bad "character"<sup>9</sup> and of one's implicit biases is irrelevant if one has control over the processes that form one's character and one's implicit associations.

The natural first place to look for guidance on responsibility and character is Aristotle. Aristotle argues that individuals can be held morally responsible for their bad character traits (and therefore the harms that flow from those traits) because individuals have control of their characters. In his discussion of what kinds of ignorance will successfully excuse in cases of wrongful action, he argues that if ignorance is the result of inattentiveness it will not provide a successful excuse, even if the agent is just an inattentive sort of person (that is, if inattentiveness is part of his character), because the agent has control over whether or not they are an inattentive sort of person. He states it thus:

But presumably he is the sort of person who is inattentive. Still, he is himself responsible for becoming this sort of person, because he has lived carelessly. Similarly, an individual is responsible for being unjust because he cheated, and for being intemperate, because he has passed his time in drinking and the like; for each type of activity produces the corresponding sort of person. This is clear from those who train for any contest or action, since they continually practice the appropriate activities. [Only] a totally insensible person would not know that a given type of activity is the source of the corresponding state; [Hence] if someone does what he knows will make him unjust, he is willingly unjust (Aristotle, NE 114a3-13).

If it is the case that agents have control over whether or not they have arrogant characters (even Doris (2002) believes that habituation is important and something over which we have control, and Holroyd (2012) argues that we have long-range control over our implicit associations), then it is irrelevant whether the agent is aware that they are being arrogant and crushing disagreement in any particular case, because they could have avoided being an arrogant sort of person in the first place. Being unaware that one was arrogant and suppressing disagreement therefore won't get the agent off the hook morally – that is, it won't get one off the hook for actually suppressing disagreement, as well as for not noticing a case of suppressed disagreement given reasonable cues.

We now have an account of what a reasonable response to silence might be, and we have some idea of the agent's varying culpability relative to their role and the context that produced the silence.

---

<sup>9</sup> Character is in inverted commas here because Doris contests whether there is such a thing as character at all, but even he agrees that our language is imbued with talk of character (Doris, 2002, p.15). In the interests of not straying too far from my main argument here, I am going to keep talking about character.

## **Did Mbeki really suppress disagreement?**

So far it has been argued that suppressed disagreement cannot count as an excuse from ignorance in many cases. There will be some instances in which the agent did not suppress the disagreement and could not have known that there was a disagreement that had been suppressed (leaving them blameless), such as in case 1 discussed in the previous section. But in many circumstances, even if the agent had not been responsible for the suppression, they should have noticed its markers – such as silence where one would have expected debate (case 2). In other situations, like case 3, the agent will also be responsible (either knowingly or unknowingly) for having suppressed the disagreement, in which case they will be responsible both for their ignorance and for the epistemic injustice associated with silencing their interlocutors.

Where does Mbeki fall? Up until this point it has been assumed that Mbeki did suppress disagreement within the ANC, but there is some controversy surrounding this claim. In this section I will defend the claim that Mbeki suppressed disagreement, and thus that his ignorance about HIV and AIDS was culpable, and that he was morally responsible for the consequences of his AIDS denialism.

On the one end of the spectrum, it is often argued that Mbeki obviously suppressed disagreement, particularly on the issues of HIV and AIDS, by ousting those who openly disagreed with him. In particular, individuals who support this view cite the removal of Madlala-Routledge from her position as Deputy Minister of Health, due to her outspoken support for mainstream AIDS science as evidence of this (Keeton, 2009; Natrass, 2007). However, this only happened in 2007 – seven years after the main epistemic action concerning AIDS in South Africa had already occurred (that is, seven years after Mbeki had considered the evidence and formed his beliefs concerning HIV and AIDS). While Madlala-Routledge's removal might count as evidence of the persistence of AIDS denialism in South African politics, it seems strange to suggest that this was the kind of disagreement that Mbeki should have been on the lookout for when he was forming his beliefs about HIV and AIDS in 1999. It should also be noted that Madlala-Routledge was the only cabinet member who was removed from her position by Mbeki during his presidency (Vale & Barrett, 2009). So it seems implausible that Mbeki was suppressing disagreement in quite such a heavy-handed way as actively ousting dissenting interlocutors.

On the other end of the spectrum is the position that Mbeki was not suppressing disagreement at all within the ANC. Rather, it might be suggested that Mbeki was just one amongst many within the ANC who held denialist views about HIV and AIDS, and that he was merely continuing a long tradition of suspicion about AIDS within the party. On this view, he did not suppress disagreement within the ANC, because there was no disagreement to suppress. Steinberg (2011) gestures toward this view when he states

in an article on AIDS in South Africa that: “Mbeki was not necessarily the outlier he is often said to be” (Sternberg, 2011). It is clear that some forms of AIDS scepticism within the ANC do substantially pre-date Mbeki. In 1988 Jabulani Nxumalo (aka ‘Mzala’) – a celebrated intellectual within the ANC in the 1980s – wrote:

[T]hat the theory of African origin of AIDS was ‘yet another justification for... racist prejudice’, that Africans were being deliberately misdiagnosed, and that there was evidence that HIV might have been invented in the ‘laboratories of the military-industrial complex’ of the West? (Gevisser, 2007, paraphrasing and quoting Mzala)

Similarly, Jonny Steinberg reports an incident in Johannesburg, in August 1988, where a group of sociologists presented early research on HIV and the South African mining sector. They predicted that a combination of migrant labour paths and the broken families that had become closely associated with the mining sector would result in a substantial increase in HIV infection rates. At this point, HIV prevalence in South Africa was only 0.1 %. Cyril Ramaphosa, a leading member of the ANC (it is often claimed that Ramaphosa was Mandela’s preferred candidate to succeed him as president), was in the audience. It is rumoured that after the presentation Ramaphosa contacted the head of the research team to request that the paper not be published, “complaining that the research presumed black men to be promiscuous and was thus tinged with racism” and that he successfully prevented the research findings from being published within South Africa (Steinberg, 2011).

There was clearly scepticism surrounding HIV and AIDS within the ANC long before Mbeki entered the debate. It should also be noted that Mbeki himself endorsed mainstream AIDS science until 1998 (that is, he accepted that HIV causes AIDS and endorsed the position that an anti-retroviral programme should be made available) (Chikane, 2013, p. 258).

However, both Mzala and Ramaphosa’s scepticism is of a very different sort to Mbeki’s. Unlike Mbeki, neither question the connection between HIV and AIDS. Mzala explicitly endorsed the existence of HIV, but believed a conspiracy theory in which HIV was created in a laboratory. Again, Ramaphosa does not question the existence of HIV, or the causal relationship between HIV and AIDS. Rather, he argued that the sociologists’ predictions of exponential increases in infection rates in South Africa were based on racist beliefs about African sexuality, which is consistent with mainstream AIDS science. Neither of these positions should thus be taken as precursors to Mbeki’s view, or as offering support to it, given how different they are from Mbeki’s own stance.

The two extreme positions – that Mbeki openly ousted those he disagreed with, or that he was just continuing a long history of ANC scepticism about AIDS – both seem implausible. But it does seem obvious that Mbeki was suppressing disagreement in

more subtle ways. In a more recent piece by Steinberg (2017, forthcoming) he describes the debate about AIDS within the party as having been “muted”:

The politics of the controversy were painful and difficult. While it is clear that there was a great deal of unease about Mbeki’s position both in the ruling African National Congress (ANC) and in the health department, criticism of Mbeki from within the ranks of the country’s former liberation movement was muted (Steinberg, 2017 , p. 5).

Steinberg’s commentary gains support from Feinstein, who was a member of the ANC at the time, and a Member of Parliament. He describes the climate of debate at the time as follows:

During the Mandela years the caucus room had resonated with sharp debate and discussion, passionate argument and profound polemic, the discourse that has characterised the ANC and the internal resistance movement, a broad church all of whose congregants felt able to speak their mind and argue their view. On 28 September 2000, two years after Thabo Mbeki had assumed the leadership of the ANC, the caucus reflected a more disciplined, choreographed and constrained party, a party fearful of its leader, conscious of his power to make or break careers, conscious of his demand for loyalty, for conformity of thinking. (Feinstein, 2009, p. 111)

It does seem that Mbeki suppressed the debate about HIV and AIDS within the ANC in the early 2000s. It would be overstepping my mark to theorise too much on how the mechanism of suppression operated, so I will keep my speculation on this issue brief.

One plausible account might be that his arrogance suppressed disagreement. Tanesini (forthcoming) provides an explanation for how this might occur. She argues that when an interlocutor’s contributions are ignored (or in more severe cases, belittled) by an arrogant agent, this demoralises the interlocutor, and eventually she might give up on trying to make any contribution at all (McGill (2013) makes the same argument). The quote at the start of this chapter suggests that this is what was happening in the Mbeki case – Mbeki had ignored disagreement on various historically contentious issues (such as negotiating with representatives from the apartheid state, and abandoning communism once democracy had been achieved), and had luckily been right on previous issues despite neglecting peer disagreement with his view (by all appearances, moral luck had operated in his favour even though he acted irresponsibly). In the case of AIDS, he assumed that he would come out on the right side yet again, and so refused to entertain alternative perspectives on this issue. This reading of the case gains further credibility in light of another quote from Mbeki’s biographer, Gevisser:

One of Mbeki’s great weaknesses, even some of his most loyal comrades have told me, is that he often does not trust others to filter data for him: he likes to

hear it himself directly from the source, and to make his own decisions (Gevisser, 2007 , p. 734).

An alternative reading of the case would be that Mbeki did not suppress disagreement himself, but that the culture of the ANC at the time was one of deference to authority. Those who make this argument appeal to the fact that the ANC was a militarised resistance group under apartheid, where deference to authority was a necessary part of the organisation's survival (as is the case with any militarised group) (Butler, 2005). If this the correct description of what happened, then Mbeki would not be responsible for suppressing the disagreement, but he might still be responsible for not noticing that the stony silence within the ANC caucus was masking disagreement with his view – he would thus be responsible for one thing fewer than under the alternative description. However, it seems unlikely that this is the correct description of events. In the quote from Andrew Feinstein above, he reminisces about the lively culture of debate within the ANC under the Mandela presidency, and explicitly points out the contrast to the Mbeki presidency, where debate was not permitted within the ANC caucus. If it were the case that a culture of deference to authority had been carried over from the ANC's years in exile, we would expect to see this in the Mandela presidency too, which we don't. It thus seems more likely that Mbeki suppressed disagreement within the ANC.

It is thus plausible that Mbeki's arrogance silenced peer disagreement with his view. Regardless of how Mbeki silenced disagreement within the party, Steinberg and Feinstein offer a compelling case that debate *was* suppressed during the Mbeki years. Further, it is clear that political contexts are such that one would expect debate, and so the silence should have been troubling. He is thus culpable for his ignorance and blameworthy for his resulting action.

## **Conclusion**

In this talk it has been argued that in instances in which it isn't clear who the scientific experts are, or when the scientific community is genuinely divided, epistemic peers can provide a valuable resource to the policy maker grappling with what to do. However, relying on epistemic peers is not always straightforward, particularly in the context of policy development. One of the things that the Mbeki case makes clear is that even when epistemic peers have important information that is relevant to a serious policy question, they might not articulate the information that they have – perhaps because they fear the consequences of disagreeing with a policy maker who occupies a position of power. That is, their view might be suppressed. One of the role specific obligations that might apply to occupying political office, or being a policy maker, is that one is required to be on the lookout for markers of suppressed viewpoints – in particular, one should be wary of silence at the table where policy is being developed.



## Bibliography

- Anonymous. (2002, March). Castro Hlongwane, caravans, cats, geese, foot & mouth and statistics: HIV/AIDS and the struggle for the humanisation of the African.
- Ballantyne, N. (2015 ). The Significance of Unpossessed Evidence. *The Philosophical Quarterly* , 1-21.
- Beatty, J., & Moore, A. (2010). Should we aim for consensus? . *Episteme*, 198-214.
- Butler, A. (2005). South Africa's HIV/AIDS Policy, 1994 - 2004: How can it be explained? . *African Affairs*, 591-614.
- Carey, B. (2011 ). Possible disagreements and defeat . *Philosophical Studies* , 371-381.
- Cherry, M. (2009). The President's panel. In K. Cullinan, & A. Thom, *The Virus, Vitamins and Vegetables: The South African HIV/AIDS Mystery* (pp. 16-35). Auckland Park: Jacan.
- Christensen, D. (2007 ). Epistemology of Disagreement: The Good News . *The Philosophical Review* , 187-217.
- Coady, D. (2006 ). When Experts Disagree. *Episteme*, 68-79.
- Doris, J. (2002). *Lack of Character: Personality and Moral Behavior*. New York: Cambridge University Press.
- Dryzek, J., & Niemeyer, S. (2008). Discursive Representation. *American Political Science Review*, 481-493.
- Elga, A. (2007). Reflection and Disagreement. *Nous*, 478-502.
- Elster, J. (1986/1997). The Market and the Forum: Three Varieties of Political Theory . In J. Bohman, & W. Rehg, *Deliberative Democracy* (pp. 3-34). Cambridge, MA: MIT Press .
- Fassin, D., & Schneider, H. (2003). The Politics of AIDS in South Africa: Beyond the Controversies. *BMJ*, 495-498.
- Feldman, R. (2006 ). Epistemological Puzzles about Disagreement. In S. Hetherington, *Epistemology Futures* (p. 223). Oxford: Oxford University Press.
- Fischer, J., & Ravizza, M. (1998). *Responsibility and Control: A Theory of Moral Responsibility*. Cambridge: Cambridge University Press.
- Fricker, M. (2007). *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford: Oxford University Press .
- Fricker, M. (2012). Silence and Institutional Prejudice . In S. Crasnow, & A. Superson, *Out of the Shadows: Analytical Feminist Contributions to Traditional PhilosophyI* (pp. 287- 306). Oxford: Oxford University Press.
- Geffen, N. (2010). *Debunking Delusions: The Inside Story of the Treatment Action Campaign*. Auckland Park: Jacana .
- Gevisser, M. (2007 ). *Thabo Mbeki: The Dream Deferred*. Johannesburg & Cape Town: Jonathan Ball Publishers .
- Goldberg, S. (2010). The Epistemology of Silence. In A. Haddock, A. Millar, & D. Pritchard, *The Epistemology of Silence* (pp. 243- 261). Oxford: Oxford University Press.
- Goldberg, S. (2010a). *Relying on Others: An Essay in Epistemology*. Oxford: Oxford University Press.
- Goldberg, S. (2015 ). Should have known. *Synthese*, 1-32.
- Goldberg, S. (Forthcoming). Arrogance, Silence and Silencing. *Proceedings of the Aristotelian Society*.
- Goldman, A. (2001). Experts: Which ones should you trust? *Philosophy and Phenomenological Research*, 85-110.
- Holroyd, J. (2012). Responsibility for Implicit Bias. *Journal of Social Philosophy*, 274-306.

- Keeton, C. (2009). 'Speaking truth to power'. In K. Cullinan, & A. Thom, *The Virus, Vitamins, and Vegetables* (pp. 186-197). Auckland Park: Jacana .
- Kelly, T. (2005). The Epistemic Significance of Disagreement. In *Oxford Studies in Epistemology, 1* (pp. 167-196).
- Kornblith, H. (2010). Belief in the Face of Controversy. In T. Warfield, & R. Feldman, *Disagreement* (pp. 29-52). New York: Oxford University Press .
- Lammenranta, M. (2011). Skepticism and Disagreement. In D. Machuca, *Pyrrhonism in Ancient, Modern, and Contemporary Philosophy* (pp. 203-216). Dordrecht: Springer.
- Langton, R., & West, C. (1999). Scorekeeping in a Pornographic Language Game. *Australasian Journal of Philosophy*, 303-319.
- Levy, N. (2005). The Good, the Bad and the Blameworthy. *Ethics and Social Philosophy*, 1-16.
- Mbeki, T. (2001, October 12). Address by President Thabo Mbeki at the Inaugural ZK Matthews Memorial Lecture. University of Fort Hare.
- McGill, J. (2013). The Silencing of Women. In K. Hutchinson, & F. Jenkins, *Women in Philosophy: What needs to change?* (pp. 39-60). Oxford: Oxford University Press.
- Nattrass, N. (2007 ). *Mortal Combat: AIDS Denialism and the Struggle for Antiretrovirals in South Africa* . Scottsville : University of KwaZulu-Natal Press .
- Nattrass, N. (2008). AIDS and the Scientific Governance of Medicine in Post-Apartheid South Africa . *African Affairs*, 157-176.
- Rosen, G. (2004). Skepticism about Moral Responsibility . *Philosophical Perspectives* , 295- 313 .
- Saul, J. (2013). Implicit Bias, Stereotype Threat, and Women in Philosophy. In K. Hutchinson, & F. Jenkins, *Women in Philosophy: What needs to change?* (pp. 39-60). Oxford: Oxford University Press.
- Sidgwick, H. (1981/1907). *The Methods of Ethics* (Hackett first edition, unabridged republication of the seventh (1907) Macmillan edition ed.). Cambridge: Hackett.
- Smith, H. (1983). Culpable Ignorance. *The Philosophical Review* , 543-571.
- Steinberg, J. (2017 , January ). A Taste for Medicine: Re-examining the early years of anti-retroviral treatment in South Africa. *African Affairs* .
- Tanesini, A. (forthcoming). "Calm Down Dear": Intellectual Arrogance, Silencing and Ignorance. *Aristotelian Society Supplementary Volume*.
- Tersman, F. (2013 ). Moral Disagreement: Actual vs. Possible. In D. Machuca, *Disagreement and Skepticism* (pp. 90-108 ). New York: Routledge.
- Vale, P., & Barrett, G. (2009). The curious career of an African modernizer: South Africa's Thabo Mbeki. *Contemporary Politics*, 445-460.
- Van der Vliet, V. (2001). AIDS: Losing "The New Struggle". *Daedalus*, 151-184.
- Williams, B. (1981 ). *Moral Luck: Philosophical Papers 1973-1980*. Cambridge : Cambridge University Press .
- Zimmerman, M. (1997). Moral Responsibility and Ignorance. *Ethics*, 410-426.