

Do Conspiracies Tend to Fail?

Philosophical Reflections on a Poorly Supported Academic Meme

This work was first published in *Episteme*, 2022, available online at:
<https://doi.org/10.1017/epi.2021.55>

Kurtis Hagen

Abstract Critics of conspiracy theories often charge that such theories are implausible because conspiracies of the kind they allege tend to fail. Thus, according to these critics, conspiracy theories that have been around for a while would have been, in all likelihood, already exposed if they had been real. So, they reason, they probably are not. In this article, I maintain that the arguments in support of this view are unconvincing. I do so by examining a list of four sources recently cited in support of the claim that conspiracies tend to fail. I pay special attention to two of these sources, an article by Brian Keeley, and, especially, an article by David Grimes, which is perhaps the single “best” article in support of the idea that conspiracies tend to fail. That is, it offers the most explicit and elaborate attempt to establish this view. Further, that article has garnered significant (uncritical) attention in the mainstream press. I argue that Grimes’s argument does not succeed, that the common assertion that conspiracies tend to fail remains poorly supported, and that there are good reasons to think that at least some types of conspiracies do not tend to fail.

Keywords Conspiracies • conspiracy theory • epistemology • Grimes • Tuskegee study

1. Introduction

In this article, I argue that the idea that conspiracies tend to fail, which is often offered as justification for dismissing conspiracy theories, is not well grounded and does not provide a strong reason for skepticism about such theories. This adds to the literature supporting the view that each conspiracy theory ought to be judged on its own particular evidentiary merits or faults, not merely by appeal to generalities.

The notion that conspiracies tend to fail appears to play a significant role in influencing “sophisticated” people’s opinions about conspiracy theories. Though this notion has been often explicitly asserted, it is even more often implied. The dismissive remark, “someone would have talked,” for example, is fairly ubiquitous. And the perceived implication plays a major role in supporting the view that conspiracy theories are not to be taken seriously. However, there is little reason, *prima facie*, to think that conspiracies do tend to fail. We don’t have strong grounds to believe, for example, that most petty conspiratorial crimes are foiled, or that covert operations are either exposed or backfire. After all, it is difficult to know how many succeed. If we are to be confident that conspiracy theories can be dismissed because conspiracies tend to fail, we need a substantial argument showing that the relevant kinds of conspiracies actually do fail. But such an argument, I maintain, appears to be lacking.

I begin by reviewing and augmenting David Coady’s reasons for skepticism regarding the claim that conspiracies tend to fail.¹ I then turn to a particular recent assertion that conspiracies “tend to fail” (Douglas et al. 2019: 4), seemingly supported by four sources. I discuss two of them at length, namely Keeley (1999) and Grimes (2016). I focus particularly on showing that

¹ Coady presents similar versions of his argument in Coady 2012: 115-120, and 2018: 179-181.

one of those studies, namely Grimes (2016), is seriously flawed in multiple ways, and thus fails to provide a substantial reason to dismiss conspiracy theories, even fairly large ones. My critique of Grimes adds substantially to a more limited, and a slightly flawed, critique by M Dentith (2019), and puts it squarely in the context of the general claim that “conspiracies tend to fail.”

2. Do Conspiracies Tend to Fail?

A number of people, apparently going back as far as Machiavelli (see Coady 2018: 179), have claimed that conspiracies tend to fail. David Coady also cites Karl Popper and Daniel Pipes as making essentially the same claim. In their own way, Cass Sunstein and Adrian Vermeule (2009) make the claim as well, at least with respect to “open societies.”² Responding to this general idea, Coady writes, “The argument that conspiracies tend to fail because they always or usually end up being exposed is mistaken in two ways. First, there is no reason to believe the premise is true. Second, the conclusion does not follow from the premise” (Coady 2012: 117; cf. 2018: 180). Supporting the first point—that there is no reason to believe that conspiracies are usually exposed—Coady points out that no examples of secrets that have been kept successfully

² Sunstein and Vermeule put the point this way: “Conspiracy theories that posit machinations by government officials typically overestimate the competence and discretion of officials and bureaucracies, who are assumed to be able to make and carry out sophisticated secret plans, despite abundant evidence that in open societies government action does not usually remain secret for very long” (Sunstein and Vermeule 2009: 208-209). See Hagen 2010: 158-159, for my response.

to this day can be offered, *no matter how numerous they are*.³ So, we cannot determine the fraction of exposed conspiracies relative to the total number of conspiracies (including both those revealed and those not revealed), because we have no way of knowing the number of those not revealed. Further, based on what we do know, Coady rhetorically asks, “Does ‘the U.S. government regularly engage in conspiratorial and clandestine operations’? No one familiar with US history could think otherwise” (Coady 2012: 120). This consideration actually gives us a reason to think that certain kinds of conspiracies do not tend to fail, at least in the sense of their specific details becoming widely known.

Supporting Coady’s second point—that the tendency of conspiracies to *be exposed* does not imply that they have a tendency *to fail*, Coady points out that temporary and imperfect secrecy is sufficient for the success of many conspiracies. So, just because a conspiracy is exposed, that does not mean that the conspiracy failed. The conspiracy only needs to be sufficiently secret long enough for the purposes of the conspirators to be achieved or their interests furthered. The orchestrated propaganda supporting the Iraq war could be considered an example. At least some fairly mainstream sources accept this as having been a genuine conspiracy, but this acceptance came too late to have any substantial impact.

However, it could be responded that, for a conspiracy, being exposed is *a kind of failure*. And it is indeed *the* kind of failure that is relevant when addressing the plausibility of conspiracy theories that have been around for a considerable time. However, one may reasonably wonder

³ I’ve made this point before as well: “As for secrets that have been kept so well that they have never been revealed in any way (at least not yet), I admit I can give no examples. I hope the reason I cannot do so is obvious, and the implication clear” (Hagen 2010: 159).

whether different conspiracies requiring different degrees of secrecy have corresponding differences in design, execution, and coverup efforts. And so, one cannot draw a firm conclusion about conspiracies that would seem to require long-lasting secrecy—perhaps because of their degree of “toxicity” (see Basham 2018)—from facts about the revelation of conspiracies that did not require such long-lasting secrecy to achieve their goal and avoid serious blowback. Further, the *sort* of secrecy that a conspiracy requires should be considered. Generally, it is not necessary that *nobody* besides the conspirators *knows* about it. It may be enough that the majority of the most relevant population does not *believe* it, or even just feels powerless to do anything about it. This may require nothing more than plausible deniability.

3. Sources Supposedly Supporting the Idea that Conspiracies Tend to Fail

The above considerations notwithstanding, the claim that conspiracies tend to fail continues to be made in ways that seem to cast doubt on the plausibility of conspiracy theories. For example, in a recent survey of the social science literature on conspiracy theories, called “Understanding Conspiracy Theories,” Karen Douglas and her colleagues remark, “Conspiracies such as the Watergate scandal do happen, but because of the difficulties inherent in executing plans and keeping people quiet, they tend to fail” (Douglas et al. 2019: 4).⁴ Douglas et al. cite four sources: Dai and Handley-Schachler (2015), Grimes (2016), Keeley (1999), and Popper (1972), giving the impression that this is a fact that has been established in the literature.

⁴ It should be noted that the authors of this article include Karen M. Douglas, Joseph E. Uscinski, and Robbie M. Sutton, who are among the most well-published social scientists writing about conspiracy theories.

However, none of these sources provide convincing evidence in support of this claim, and only Grimes (2016) really tries to do this. Dai and Handley-Schachler (2015) and Popper (1972), first of all, are not really germane. Dai and Handley-Schachler (2015) is specifically about accounting fraud, and it does *not* suggest that conspiracies to commit such fraud tend to fail.⁵ And Popper’s argument is about “the conspiracy theory of society,” which imagines that conspirators are controlling everything. It argues that this must be wrong because “*nothing ever comes off exactly as intended*” (Popper 1972; see Coady 2006: 130). Popper’s argument does not provide strong reasons to think that less radical conspiracies tend to fail.⁶ As for Keeley’s article, though it does not provide a substantial argument in support of the notion that conspiracies tend to fail, it is nevertheless worth discussing in this context, and thus is the subject of the following section. Of the four, only Grimes really provides an argument that, if successful, would support Douglas et al.’s claim. However, as I will argue below, it is not successful.

Before addressing Keeley and Grimes, it is worth noting that, after Douglas et al. claim that conspiracy theories tend to fail, the very next sentence asserts: “When conspiracies fail—or are otherwise exposed—the appropriate experts deem them as having actually occurred”

⁵ Dai and Handley-Schachler do write, “[T]he existence of a conspiracy makes detection by management or other agencies more probable because of the chance that one of the conspirators or someone who has been approached to join the conspiracy will report the fraud” (2015: 1). However, this only suggests that the existence of a conspiracy provides an additional means of detecting fraud, besides just examining the books. It does not follow that conspiracies tend to fail, and the authors never suggest that it does. Indeed, they seem to imply quite the opposite.

⁶ See Pigden (1995) for a relevant critique of Popper.

(Douglas et al. 2019: 4). In support of this claim, Douglas et al. cite Levy (2007). But that article, which does seem to assume something like this, does not actually provide any *evidence for it*. In fact, unless the claim is understood as a tautology—where “exposure” is defined in terms of expert acceptance—the idea that “the appropriate experts” will endorse the claims of whistleblowers appears to be empirically false. For leaks, confessions, and such, pointing to a conspiracy do *not* reliably lead to experts and mainstream sources deeming that a conspiracy actually occurred, as discussed below (both in the next section and section 5.4). Determining whether or not experts *should* affirm such leaked allegations of conspiracy, in any particular case, is often a nontrivial task.

4. Keeley 1999, “Of Conspiracy Theories”

Keeley’s essay does not, strictly speaking, argue that conspiracies tend to fail. Rather, Keeley seems to suggest that we are not warranted in believing “mature” conspiracy theories, namely, those that have been around a while and have been investigated. Keeley seems to think that if the national media does not report damning evidence, then it is probable that no such evidence exists. For otherwise we would have to increase the scope of the conspiracy, implicating the press and others, to the point that we couldn’t have confidence in anything. While I think there are problems with Keeley’s analysis, the relevant point here is that what Keeley is arguing is *different* from the idea that conspiracies tend to fail. If Douglas et al. want to use his article as evidence, they should *explain* how Keeley’s position justifies theirs.

Keeley does make some incidental remarks about the likelihood of explosive secrets being successfully kept. Perhaps this is what Douglas et al. had in mind: Regarding conspiracy theories about the Oklahoma City bombing, Keeley writes, “[I]t is impossible to believe that not

a single member of the BATF stationed in Oklahoma would be moved by guilt, self-interest, or some other motivation to reveal that agency's role in the tragedy, if not to the press, then to a lover or a family member" (Keeley 1999: 124). While seemingly reasonable, this is really an expression of personal incredulity stated as though it is an objective fact. By citing Keeley in the way that they did, Douglas et al. give the appearance of objective grounding to an idea that is, it now appears, grounded merely in incredulity that may not be warranted. Indeed, as far as this quotation goes, the idea that conspiracy theories tend to fail seems more a premise than a conclusion. Keeley does offer some support for that premise, however, which I'll get to in a moment. But first notice some problems with Keeley's inference. He seems to be arguing: (1) If the BATF was involved, someone would have talked. (2) Nobody did. (3) Therefore, the BATF was not involved. However, not only is the first premise questionable, the second one is as well. We have *no idea* how many secrets of the kind in question are told to friends, lovers, and relatives. In general, we just don't have any access to that type of information. If a secret is passed in this way, it may just stop there. Or, it may spread a bit among a few other intimates. Or, it may spread a bit further and become a rumor within certain circles. Most people would have no knowledge of these minor "leaks." If the rumors spread even further, we might hear of it. But even if we did, it would just be a rumor to us.

Further, in general, it is *not true* that such rumors reliably bring down conspiracies, nor is it true that such rumors do not exist in relation to mature conspiracy theories. Some of these are more than mere rumors; they are descriptions of ear-witnesses of confessions or at least seemingly self-incriminating statements. Examples regarding the JFK case include the following: E. Howard Hunt, a CIA operative and convicted Watergate conspirator, made a deathbed "confession" regarding the JFK assassination, though he admits only to being a "bench warmer"

on the periphery of the plot. In addition, while apparently intoxicated, David Morales, a senior CIA official who was reportedly involved in assassination plots against Castro and in other controversial covert CIA operations, is reported to have made a seemingly incriminating exclamation about the JFK assassination. Specifically, regarding JFK, Morales said, “We took care of that son of a bitch, didn’t we?”⁷ Also, regarding the assassination of Martin Luther King, a man named Lloyd Jowers confessed to playing a small role, and he implicated government agencies. A jury agreed in a civil case brought by the King family, but the government just pooh-poohed the whole thing, and the media paid little attention. I offer these examples not as evidence of government involvement in these assassinations. Here they serve as evidence that confessions and self-incriminatory statements *are not enough* to turn a conspiracy theory into an acknowledged fact. For it is often hard to know how much stock one should put in such statements. After all, in addition to the examples I have mentioned, there are many more clearly dubious “confessions” or self-incriminatory statements. And there is no shortage of interesting borderline cases.⁸ (The fact that leaks, or what appear to be leaks, do not reliably lead to the

⁷ Brief summaries with links to more information regarding the Hunt and Morales “confessions,” as well as several others, can be found on the Mary Ferrell Foundation website, on a page called “Confessions.” <https://www.maryferrell.org/pages/Confessions.html>

⁸ Former investigator for the House Select Committee on Assassinations, Gaeton Fonzi, has written a book (1993) centering on the case of Antonio Veciana, who links Lee Harvey Oswald to the CIA’s David Atlee Phillips. See Hagen 2020: 431, for a short description of this and similar cases, along with remarks regarding their relevance to the issue here. See Hancock (2006), *Someone Would Have Talked*, for more examples regarding the JFK assassination. In

acceptance of conspiracy theories is also relevant to the “demonstration” offered by Grimes, and will be discussed more below.)

Now back to Keeley’s reasoning. As a way of supporting the claim quoted above, Keeley notes, “Government agencies, even those as regulated and controlled as the military and intelligence agencies, are plagued with leaks and rumors” (Keeley 1999: 124). However, it is not clear how true this is, or what the implications are. First, there is an important difference between a leak and a rumor. Rumors are generally given little weight, and even more clearly fail to bring down conspiracies. Indeed, rumors are not that different from conspiracy theories themselves.⁹ Second, it is not clear that *intelligence agencies* are plagued by leaks, constantly foiling their covert operations and outing their operatives. That actually seems fairly rare, given that such operations are presumably continuously ongoing. And so, thinking about the military and intelligence agencies, again, suggests that certain types of conspiracies do not tend to fail in the relevant sense.

Keeley also asserts, “To propose that an explosive secret could be closeted for any length of time simply reveals a lack of understanding of the nature of modern bureaucracies” (Keeley 1999: 124). That seems to be a common view. But not every informed and sophisticated person shares it. Fletcher Prouty, author of *The Secret Team*, for example, does not seem to agree. And Prouty had relevant firsthand knowledge of how plans were compartmentalized to help maintain secrecy and assure that the explosiveness of any particular leak would be limited. Prouty was the

particular, for relevant details about Phillips, including a quasi-confession, see pp. 113-115 and 179-183.

⁹ Coady (2018) argues that neither rumors nor conspiracy theories deserve their bad reputation.

liaison between the military and the CIA in the early 1960s. He also suggests that some secret-keepers are *professional* secret-keepers. And *their* reliability as secret-keepers shouldn't be conflated with that of loose-lipped political operatives.

In addition, making a point that will also be made by Dentith against Grimes (discussed below), Lee Basham remarks:

Keeley overlooks the implications of the remarkably *hierarchical* nature of our civilization's institutions of commerce and control. ... Far from supposing that leaks inevitably expose the "explosive secrets" of conspiracy, we might reason that the more fully developed and high placed a conspiracy is, the more experienced and able its practitioners at controlling information and either co-opting, discrediting or eliminating those who go astray or otherwise encounter the truth. ... [W]e do know this: the existence of "openly secretive" governmental and corporate institutions is the norm in contemporary civilization. Despite the occasional leaks they appear to have been quite successful in their control of extremely disturbing information. This is all the greater an achievement given their status as widely recognized agencies of secrets and secret projects. We can only boggle at the difficulties of reliably revealing the aims and means of competent organizations that systematically hide their very existence. (Basham 2001: 272)

As for the press, Basham has emphasized that some stories are too “toxic” to report (2018: 281-288).¹⁰ In that context, even if a few individual reporters believed that a story deserved to be told, they may easily be thwarted by their superiors. Knowing that bucking the system in such a case would likely be career limiting, few would attempt it. And those that try, ostracized by the mainstream, would typically have only marginal influence anyway. Since this is predictable, there is little incentive to buck the system. Or, at least, such a dynamic is plausible and should not be summarily dismissed.¹¹ (One might think of Gary Webb in this context, see Schou 2009.)

Admittedly, Basham’s considerations are far from decisive. The adequacy of his reasoning and evidence, in comparison to Keeley’s, could be debated at length. But the important point here is precisely that, contrary to the impression given by Douglas et al., the degree to which the truth about a conspiracy may be effectively repressed, even if it is not fully concealed, *is not a settled matter*.

¹⁰ Sometimes toxic stories are reported, but only marginally, and then left to die. Project Censored publishes an annual collection of essays discussing what they consider to be the most underreported stories of the year.

¹¹ While potential counterexamples might be offered, such as partisan news outlets not shying away from undermining the very legitimacy of a president, one may wonder if such stories are truly toxic *from the perspective of those news outlets*. It is issues on which there is broad bipartisan agreement that the risk of toxic truths preventing investigation and revelation are greatest.

Even Keeley himself gives us a reason to think that the kinds of conspiracies in question might *not* tend to fail. He writes: “[T]he conspiracy theorist is working in a domain where *the investigated actively seeks to hamper the investigation...* As evidenced by any number of twentieth century, U.S. government-sponsored activities (take your pick), we have reason to believe that there exist forces with both motive and capacity to carry out effective disinformation campaigns” (Keeley 1999: 120-121, emphasis in original).

Taking all this into consideration, we can’t say with justified confidence that conspiracies of the kind in question will tend to fail on account of the inability of people to keep quiet. And so, we can’t dismiss conspiracy theories on the assumption that, if true, they would have been revealed. Even Keeley admits, ultimately, that “The best we can do is track the evaluation of given theories over time and come to some consensus as to when belief in the theory entails more skepticism than we can stomach” (1999: 126).

One final note about the citation of Keeley (1999): Keeley’s essay was certainly a groundbreaking one, but it has been much criticized in the philosophical literature on conspiracy theories (see Basham 2001, 2018; Clarke 2002: 140-143; Coady 2003; and Rääkkä 2009: 188-198).¹² Indeed, it is perhaps the most contested article on this topic. So, citing it generically, as if it represented uncontested scholarship, is questionable in itself. Further, the specific notion at issue here, that conspiracies tend to fail because maintaining adequate secrecy is nearly

¹² Except for Basham (2018), these articles, and Keeley’s as well, were all republished in Coady’s *Conspiracy Theories: The Philosophical Debate* (2006). The reprint of Basham (2001), in Coady (2006), directly follows Keeley’s article, which it focuses its critique upon.

impossible, has been explicitly challenged in the philosophical literature.¹³ But Douglas et al. give no indication that there is any controversy regarding either Keeley's arguments in general or this issue in particular. In sum, the citation of Keeley (1999) gives a contested claim that is not even unequivocally supported by that article an undue air of authority.

5. Grimes 2016, "On the Viability of Conspiratorial Beliefs"

Finally, we are left with Grimes (2016), which I will discuss at greater length because, unlike the rest, it does at least *seem* to provide significant support for Douglas et al.'s claim, at least if that claim is assumed to refer to fairly large conspiracies. Indeed, Grimes offers what is probably the "best" (i.e., most *seemingly* authoritative) case for the claim that conspiracy theories will tend to fail in a way that is relevant to our assessment of conspiracy theories. Grimes purports to have provided a "simple mathematical demonstration of the untenability" of certain conspiracy theories (2016: 14).

Grimes's model is built on the assumptions that conspirators are generally dedicated to the concealment of the conspiracy and that the conspiracy would be exposed by any leak from one of the conspirators (2016: 3). He uses the following three exposed conspiracies to estimate the probability, over time, of such exposure, relative to the size of the conspiracy: (1) The NSA's PRISM surveillance program exposed by Edward Snowden. (2) The Tuskegee syphilis experiment, exposed by Peter Buxtun, which involved the US Public Health Service deceiving

¹³ Coady offers two versions of his critique of this idea (see 2012: 115-120, and 2018: 179-181). Dentith addresses the issue, in a critique of Grimes (2016), in the context of whether or not conspiracies posited by conspiracy theories are "too big." See Dentith 2019: 2255-2258.

African-American men about their disease, and observing them but essentially not treating them. And, (3) the FBI forensic scandal, in which Supervisory Special Agent Frederic Whitehurst uncovered and exposed scientific misconduct in the FBI crime lab. Grimes uses his results to estimate the likelihood of certain specific conspiracy theories as well as to justify the general claim that large conspiracies are likely to quickly fail and so mature theories that imply large conspiracies are untenable.

In Grimes's own words, he developed "a simple mathematical model for conspiracies involving multiple actors with time, which yields failure probability for any given conspiracy" (Grimes 2016: 1). He uses this to argue that "large conspiracies (≥ 1000 agents) quickly become untenable and prone to failure" (2016: 1). Grimes acknowledges that "historical examples of exposed conspiracies do exist." Indeed, it is the three above-mentioned "known scandals" that provide a basis for his model's parameters. He attempts to use his model to estimate the likelihood of "some commonly-held conspiratorial beliefs," and to thereby demonstrate that they are highly unlikely. Specifically, he uses his model to evaluate the following ideas: "the moon-landings were faked, climate-change is a hoax, vaccination is dangerous and that a cure for cancer is being suppressed by vested interests" (2016: 1). He finds, "Simulations of these claims predict that intrinsic failure would be imminent *even with the most generous estimates* for the secret-keeping ability of active participants" (2016: 1, emphasis added).

Grimes's article garnered significant attention. It is described in the *Internet Encyclopedia of Philosophy* as follows: "Grimes (2016) has conducted simulations showing that large conspiracies with 1000 agents or more are unlikely to succeed due to problems with

maintaining secrecy” (Pauly, no date).¹⁴ In addition, Grimes’s study was covered favorably in the press. An article in *Mother Jones* magazine suggests that Grimes succeeded in showing certain conspiracy theories to be “really dumb” (McDonnell 2016). BBC likewise covered the study uncritically (Berezow 2016), as did PBS (Barajas 2016)¹⁵ and *The Washington Post* (Ohlheiser 2016), as well as *LiveScience*, *ScienceDaily*, and *Wired*. The *Wired* article prominently proclaims, “The equation shows *most* conspiracy theories would have unravelled by now”

¹⁴ While most mentions of Grimes (2016) that I have seen are uncritically supportive, Pauly’s is rather balanced. He mentions Dentith’s critique, and also briefly contrasts Grimes’s view of conspiracies with Basham’s hierarchical view.

¹⁵ I often encounter the idea that we should trust journalists to “fact check” conspiracy theorists. I’m not convinced this is wise. Why should people who crank out stories on a variety of topics, about which they often know very little, be regarded as particularly reliable? For example, the cited PBS article on Grimes’s study describes the FBI scandal as: “an FBI scandal involving Dr. Frederic Whitehurst, whose misleading forensic tests led to wrongful detainment and execution of several innocent people” (Barajas 2016). I can’t be sure whether this is a genuine misunderstanding rather than just bad writing, but it certainly gives the wrong impression. It makes Whitehurst sound like a villain, rather than a hero. But that’s backwards. Whitehurst discovered the fraud and exposed it; he didn’t commit it. While I don’t think Barajas meant to mislead, his presentation of the facts suggests ignorance regarding the case. That is in addition to being completely credulous about the notion that it would take 400,000 people to fake the moon landing, though Grimes’s article gives no substantial reason for believing this.

(Reynolds 2017, emphasis added)—which it did not even claim to do, as it only addressed purported conspiracies that are large and have been around for some time.

Intuitively, it seems that some *versions* of prominent conspiracy theories posit conspiracies so large and ambitious that it seems very likely that they, if true, would have been exposed. Grimes aims to provide “a clear rationale for clarifying the outlandish from the reasonable” (Grimes 2016: 2). He develops a model for the likelihood of exposure of a conspiracy based on real conspiracies that have been exposed, and as mentioned above, he applies his model to conspiracy theories involving the moon landing, climate change, cancer cures, and vaccines. However, *in reality*, not all versions of such theories imply the same scope of conspirators or have the same susceptibility to exposure.¹⁶ And though Grimes refers to “commonly-held conspiratorial beliefs” (2016: 1), it is not clear that the versions that are

¹⁶ Given Grimes’s assumptions, however, some different versions of conspiracy theories do seem to have the same scope. For example, Grimes draws the following distinction: “[T]hose sceptical of the scientific consensus on anthropogenic climate change may take either a ‘hard’ position that climate-change is not occurring or a ‘soft’ position that it may be occurring but isn’t anthropogenic” (Grimes 2016: 9). Curiously, he then writes, “For this investigation, we’ll define climate change conspiracy as those taking a hard position for simplicity.” But why should it be any simpler to apply his analysis only to the “hard” position? After all, given his assumptions, the “hard” and “soft” positions would seem to involve the same numbers, indeed the very same people and organizations. Perhaps it seemed best for Grimes to only claim to prove what was most obvious already. It might look bad if his theory proved too much.

genuinely commonly held really have the characteristics that make them susceptible to the kind of failure Grimes focuses on.

In any case, if Grimes's method is valid, it should work for all conspiracy theories that have numbers in the right ranges. And further, Grimes seems to assume that, for example, all conspiracy theories alleging some sort of cover-up of vaccination harms would have to be large enough in a relevant way to be essentially ruled out as too likely to have already been exposed (see section 5.5 below). If he has not shown that, the door remains open to the possibility that there are versions of these types of conspiracies that might not be reasonably expected to have failed. Grimes clearly does not want that. The whole point of his project is to shut that door. This is fairly clear from the article itself, but it is unambiguously clear from Grimes's later comments about the article, and it is corroborated by the way the article has been read by journalists seemingly eager to promote Grimes's findings. And this is, in any case, the interesting question: Has he shown that we may safely dismiss conspiracy theories of the general types that he describes? I argue that he has not.

5.1 General problems with Grimes's analysis

Grimes's argument has serious flaws. For one, as discussed above, although we know of some conspiracies that were exposed, we don't know how many were not exposed. Grimes seems to think that we can take a few cases that were exposed and use those to judge how long it takes for exposure. However, he is working from a potentially biased sample. We also need to know how many real conspiracies have never been discovered or proven. But, of course, we cannot know that. While it might make *some* sense to make approximations based on the ranges we find in existing cases, Grimes is dealing with a *very* small sample (three revealed

conspiracies) which differ considerably from each other. They are as follows: (1) the Tuskegee syphilis experiment, in which African American men with syphilis were studied without being significantly treated, while the fact that they had syphilis was concealed from them; (2) the FBI forensic scandal, which involved the conviction of many people based (at least in part) on misleading testimony and bogus forensics, revealed by Frederic Whitehurst; and (3) the NSA PRISM project, involving the U.S. National Security Agency's massive data collection, which was brought to light by Edward Snowden.

Indeed, the three examples do not even seem to be of the same *type*, such that modeling for that type might be thought appropriate. (1) The Tuskegee Experiment was completely different from the other two, particularly in involving the *publicly available* documentation of what was going on, published in scientific journals. This made what was going on undeniable once sufficient light was shined upon it. In this important respect, it does not resemble most purported conspiracies. (2) The FBI scandal, compared to the other two examples, seems somewhat pedestrian, involving ongoing corrupt practices at a relatively low level, rather than a particular elaborate plot run from on high. In this way, it also differs from many conspiracy theories. And, while I don't mean to minimize the severity of the retaliation suffered by Whitehurst for his revelation of the fraud that was occurring, it is hardly comparable to Edward Snowden's case. (3) Snowden's revelations about the NSA were truly extraordinary, requiring great sacrifices and even greater risks on the part of Snowden to reveal them. These revelations also required incredibly meticulous execution. So, it is hard to see how Snowden's case can be used to determine what should be *expected* to happen.

5.2 Grimes's model is problematic

Dentith has argued that the examples on which Grimes’s mathematical analysis is based do not actually fit his conceptual model of how conspiracies fail. Grimes’s model is ostensibly based on examples of conspiracies that were exposed by leaks from insiders.¹⁷ And yet the revelations do *not* come from conspirators in any of the three examples that form the basis for his model. Regarding the FBI Forensic Scandal, Frederic Whitehurst revealed fraud in his laboratory that he had not participated in. And, in contrast to how Grimes portrays the revelation of the NSA surveillance programs, “Snowden’s narrative is that of being a *whistleblower*: an outsider—rather than a conspirator—who discovered the existence of the conspiracy” (Dentith 2019: 2257). Similarly, regarding the Tuskegee experiment, Dentith explains:

Information about the experiment was openly published in medical journals; the cover-up, so to speak, was that the patients were not told about the experiment. Once again, it was outsiders who then revealed the existence of the conspiracy. For example, Peter Buxton—who is often considered as *the* whistleblower in this case—came to know about the experiments because of his job with the United States Public Health Service. He was not a conspirator but, rather, a worried public official who leaked the information to the press because his worries were not taken seriously by management. (2019: 2257)

Dentith concludes that Grimes’s examples “fail to capture the very thing he wants to measure” (2019: 2257).

¹⁷ Grimes writes, “We concern ourselves only with potential intrinsic exposure of the conspiracy and do not consider for now the possibility that external agents may reveal the operation” (2016: 17-18).

While there is a flaw in Dentith's critique, the general thrust can nevertheless be sustained. Specifically, Grimes could legitimately respond that the conspiracies that he chose for his model, which occurred within the NSA, NHS, and FBI, were exposed by people who were in those institutions. So, the conspiracies *were* exposed by *insiders*, even if they were not *conspirators*. While other real conspiracies were not exposed by insiders in this sense (e.g., COINTELPRO), in Grimes's examples they were. And it would be fair for Grimes to point out that people within such institutions may be more likely than outsiders to become aware of conspiracies going on within those institutions. Still, the number of individual employees that are actually *positioned to become aware* of a conspiracy within an institution is likely to vary widely, depending both on the nature of the institution and the nature of the conspiracy. Presumably, fewer people were in a position to know or find out about the scope of NSA spying than were in a position to know or find out about the Tuskegee experiment, about which scientific articles were being published. And this scope of potential access has little if anything to do with the size of the institution as a whole. And yet Grimes's model is based on the size of these large institutions, not the size of the actual conspiracy, nor the number of people with a heightened potential of finding out. One can hardly blame Grimes for not basing his calculations on the latter category, even though it is much more relevant than the size of the whole institution. For it seems hopelessly vague and immeasurable. But that just means that the kind of study that he wants to do may not be feasible.

5.3 Grimes's estimates are not reasonable

Relatedly, Grimes's *application* of his model is also problematic. When he evaluates "commonly-held conspiratorial beliefs," he is not sufficiently generous regarding the size of the conspiracy they entail. As Dentith again explains:

A conspiracy can look big, yet only a small number of people involved in it might know its full extent or aim. Some members of the conspiracy will be lackeys, goons or even unwitting conspirators. Not everyone in the NSA need necessarily know that the data they are collecting and processing has been illegally obtained, and FBI agents who were using forensic evidence to secure convictions may not have been informed by senior personnel that the kind of evidence they were relying upon was of dubious merit. It is even possible to be involved in a conspiracy without realising you are conspiring.

(Dentith 2019: 2258).¹⁸

Ignoring all such nuances, Grimes simplistically assumes that the appropriate number of conspirators to assign to the "Moon-landing Hoax" conspiracy theory, for example, is the "Peak NASA employment" in 1965, that is, 411,000 people. BBC dutifully reports this without question, under the title, "Math Study Shows Conspiracies 'Prone to Unravelling'" (Berezow 2016). Grimes offers no good reason to think that all NASA employees would be in on it, or aware that it was a hoax. Indeed, if such a conspiracy had been attempted, it may have been managed so as to keep those "in the know" as limited as possible. If the *right people* were involved, there may have been some creative way to do this by leveraging NASA's institutional

¹⁸ Although the examples Dentith uses here are from Grimes's paradigmatic cases, rather than from the cases to which Grimes applies his model, *the point* Dentith is making is also applicable to the latter cases, as my discussion shows.

structure. It is not intuitively clear why the number of people truly in a position to leak the truth would have been greater than, say, a few hundred.¹⁹ So, it seems quite possible that Grimes's estimate is off by three orders of magnitude. Grimes has done nothing to show that such an intuitive estimation is wrong.

Regarding conspiracies involving climate change, vaccine science, and suppressed cancer cures, Grimes considers the size of the conspiracy to be the size of the relevant scientific community. Consider the justification that Grimes offers for the idea that a “scientific conspiracy” would have to be “vast”:

[E]ven if a small devious cohort of rouge [*sic*] scientists falsified data for climate change or attempted to cover-up vaccine information, examination by other scientists would fatally undermine the nascent conspiracy. To circumvent this, the vast majority of scientists in a field would have to mutually conspire—a circumstance the model predicts is exceptionally unlikely to be viable. (Grimes 2016: 11)

Grimes seems to be suggesting that if there were a conspiracy involving scientific fraud, even if it was perpetrated by only a small number of scientists, it would have been exposed by other

¹⁹ On the other hand, as Ted Goertzel points out about the moon landing, “Soviet astronomers and others around the world ... tracked the event” (Goertzel 2010: 494). This observation has considerable weight not so much because the numbers are large but because some of these people would have a strong interest in exposing the hoax, if it was one. Grimes's analysis, being based just on the numbers, ignores such highly relevant considerations as *who* might be in a position to expose the hoax, what their biases and motives might be, and how strong their evidence would be.

scientists, unless they too became party to the conspiracy—and at that point the theory becomes implausible. But it is far from clear that this picture is accurate. If, for example, a group of scientists at Merck stumble upon a cure to cancer, and management decides to suppress it, why would *everyone* at Merck have to know about that, or be in a position to discover it? Yet Grimes seems to assume that even everyone at Novartis, Pfizer, Roche, Sanofi, Johnson and Johnson, GlaxoSmithKline, and AstraZeneca would also be in a position to blow the whistle (see Grimes 2016: 8). Do theories about suppressed cancer cures necessarily posit or imply this? It is not clear that they do. (See section 5.5, below, for a similar real-life example involving vaccination science.) Grimes seems to overestimate the relevant scrutiny that exists and ignores the possible influence of pervasive bias. If a small group of scientists fudge some numbers, there may well be no indications of that fudging that peer reviewers, or other non-involved scientists, would easily be able to detect. And if the peer reviewers and other fellow scientists found the results both pleasing and well within their expectations—fitting the current narrative—they may not be particularly motivated to scrutinize the results.

It might be argued that a fraudulent study or two would not be sufficient to change the discourse on a topic like vaccine safety or climate change. But that is not entirely clear either, at least not in every case. For example, in an evaluation of the science regarding the relationship between vaccination and pervasive developmental disorder (which includes autism), the number of studies that the Institute of Medicine relied upon was quite small. Starting with 32 somewhat relevant studies, they whittled that down to only four studies that they judged “might help with a study of the [vaccine] schedule” (Institute of Medicine 2013: 86). They found that the scientific evidence “does not suggest a causal association” between autism and the overall vaccination schedule, but they describe this as being based on evidence that is “limited both in quantity and

in quality” (p. 88). So, if just a couple experiments were carefully manipulated, and those specific individual frauds were not discovered, that may plausibly have a significant influence on the ongoing consensus, especially if the fraud was convenient to powerful interests.

Regarding the claim that climate change is “a hoax,” there are a range of positions that might fall into that category. However, for many who promote these conspiracy theories, “hoax” may be hyperbole. Let’s try to imagine one possible version of this: Suppose a few people at the top of organizations with ample financial resources had interests in generating concern about carbon emissions, and so created an incentive structure that subtly encouraged scientific thought leaders to support an overly pessimistic view of the climate hazard, which in turn led to a consensus that was not fully rational (perhaps by being based on biased, though not fraudulent, studies). How many people would be in a position to expose this situation? It is not even clear exactly how this situation *could* be convincingly revealed, even if it were true. And, although many climate change conspiracy theories are explicitly more radically conspiratorial than this, it is not clear that they really *need* to deviate very far from this to preserve their central thesis. For Grimes’s argument to be genuinely informative—to do more than show that some versions of such theories overstate their case—it would have to be applicable to theories resembling this. But it is not at all clear that it is. Grimes’s analysis assumes that if a conspiracy theory were true, it could easily be shown to be true by any random scientist who uncovered fraud. But that doesn’t seem to be correct. Even if some element of actual fraud was involved and was discovered, that would probably not suffice to invalidate the consensus about climate change, nor would it show that the *central thesis* of the conspiracy theory was true. The deeper conspiracy would not have thereby been revealed. For many potential cases, Grimes’s assumptions about what it takes to reveal a conspiracy do not hold.

Incidentally, largely independent conspiracies of different kinds could be loosely held together thematically by a confluence of interests or ideology—petty science fraud on one level, and more subtle influence on incentive structures, or propaganda campaigns, on another level. One level could not rat out the other if they did not actually conspire together. But is this a plausible model for how the world really works? Given that petty scientific fraud occurs, and that organizations do intentionally influence incentive structures (largely aiming at the good, one may assume), the question of whether the world is like this is really one of degree, and thus is a very difficult empirical question.

In general, contrary to Grimes’s central assumption, the size of an institution in which a conspiracy is alleged to mainly occur does not even roughly indicate the size of the conspiracy itself, nor the circle of people with knowledge of it. In this regard, it may be useful to consider that institutions are *institutions-within-institutions*, and so one needs some substantial reason for focusing on a particular institutional scope. For example, NASA is an institution within the US government. Grimes’s estimates for the size of a “Moon-landing Hoax” conspiracy would have been silly if he had determined its size based on the size of the entire federal government. Grimes’s estimates are less silly, but still not appropriately nuanced.

5.4 Grimes’s critical assumption is unsafe: leaks are not enough

Grimes’s argument depends upon a dubious assumption. Namely, Grimes explicitly *assumes* that “a leak of information from any conspirator is sufficient to expose the conspiracy and render it redundant” (Grimes 2016: 3). Not only is this a mere assumption, it seems to be a false one. There are, in fact, no shortage of leaks, confessions, and quasi-confessions in controversial cases that have *not* succeeded in overturning official accounts. As mentioned

above, E. Howard Hunt, David Morales, and several others have made self-incriminating statements regarding the JFK assassination, and Loyd Jowers admitted involvement in Martin Luther King's murder. In the latter case, a jury even found that this was true, and that government agencies were in on the conspiracy. But that has not led to the existence of a conspiracy being accepted by the establishment. Of course, there may be good reasons to *doubt* those confessions, or quasi-confessions. But to merely assume that the reason existing admissions in controversial cases were not accepted by the establishment is because they were *false* confessions seems to beg the central question. Recall that even Deep Throat's revelations to Woodward and Bernstein were not *in themselves* enough to bring down Nixon. If it hadn't been for the Whitehouse tape recordings, many of the crimes involved in Watergate may well never have become widely acknowledged.

Regarding NSA spying, one of Grimes's own examples, Snowden was *not* the first person to reveal the kind of thing that was going on. In 2006, AT&T technician Mark Klein alleged that AT&T was cooperating with the NSA to allow the latter access to vast amounts of data (Wolf 2007). This was reported, and Congress got involved, but ultimately nothing much happened. In addition, NSA employees Thomas Drake and Bill Binney revealed relevant information. Drake alerted Congress and resorted to leaking information to a newspaper reporter only after his efforts to use official channels failed to be effective (Welna 2014). He was then prosecuted; one might even say "persecuted" (see Hertsgaard 2016). Binney also faced heavy handed efforts to keep him quiet (Welna 2014). Despite some coverage of their situations, their revelations did not turn the massive NSA spying operation into an acknowledged reality. Other

than conspiracy theorists, no one, including the mainstream media, paid much attention.²⁰ After all, Drake and Binney could be muzzled well enough, and otherwise written off as cranks. What made Snowden different is that he provided *proof*. But that is not at all an easy thing to do, as Snowden's harrowing saga dramatizes. This is, as mentioned above, a reason to think that extrapolating from this case would be of dubious validity.

Regarding the Tuskegee Experiment, which figures prominently in Grimes's analysis (and is discussed more below), even Peter Buxtun, who is credited with exposing the study in 1972, tried repeatedly to expose it before finally finding success. And further, according to Susan Reverby, a group of African American professionals that included Public Health Service employee Bill Jenkins also tried to bring attention to the study but were unsuccessful: "The group wrote an editorial denouncing the Study's racism and ethics and sent it to the *New York Times* and the *Washington Post*, and then waited for something to happen. Nothing did" (Reverby 2009: 83).

Whether or not a purported conspiracy becomes acknowledged as a genuine conspiracy is not a simple matter of how much time passes before someone talks. It is rather a complex function involving a number of variables, including how carefully compartmentalized the conspiratorial activity is, the vulnerability of the whistleblowers to discreditation, the availability of evidence sufficient to prove the most damning claims (less than that and the story may well die), the plausibility of any semi-innocuous fallback position or "limited hangout," and also the degree of desperation that exists to quash the story, which may depend on its toxicity and on

²⁰ As *The Guardian* belatedly admits, "warnings about the dangers of the NSA's surveillance programme were largely ignored" (Hertsgaard 2016).

what financial, political, or ideological interests may be at stake. Grimes’s study does not in any way address any of these potentially confounding factors.

5.5 Grimes’s problematic characterization of his study

One of the test cases Grimes considers is labeled “Vaccination conspiracy.” However, Grimes seems to conflate “conspiratorial beliefs about vaccination” with “ill-founded beliefs” and “dubious information” that is critical of vaccination (Grimes 2016: 3). This has the effect of suggesting that his analysis of the likelihood of conspiracies to fail has some relevance to critics of vaccination that make no conspiratorial accusations. Many people who raise concerns about vaccination, whether they are well-founded or not, emphasize the published science and argue on that basis for a more cautious approach to vaccination, or else point to specific flaws and limitations in studies supportive of vaccine safety. Contrary to what he suggests, Grimes’s analysis is not actually relevant to such challenges to scientific orthodoxy.²¹ A scientific consensus can be wrong without any conspiracy to maintain an incorrect view.

While Grimes’s article does not specifically mention Andrew Wakefield,²² his category of “Vaccination conspiracy” does include the idea that “there is a link between autism and the

²¹ For a discussion of the use of the term “conspiracy theorist” to discredit people who are critical of vaccination, see Martin (2020), which covers Grimes (2016) specifically on p. 412.

²² Wakefield is commonly regarded as “discredited.” However, it should be noted that the process that is often cited to support his discreditation has itself been discredited. Wakefield and his colleague John Walker-Smith were both stripped of their licenses as a result of a proceeding conducted by the UK’s General Medical Council—the *same* proceeding. Walker-Smith appealed

MMR vaccine,” which is an idea that is often traced back to Wakefield, who was the first author of a retracted *Lancet* article that hinted at a possible connection between autism and MMR vaccination. Wakefield has more recently argued that a conspiracy occurred involving CDC scientists, which covered up an association between autism and the timing of MMR vaccination in African American boys.

A revealing radio “debate” was broadcast between Grimes and Wakefield, during which Grimes suggests that his article had established that the conspiracy Wakefield alleged would have involved twenty-five to fifty thousand people. He says, “a CDC/WHO conspiracy would need at least fifty—twenty-five to fifty—thousand active collaborators.”²³ Now you have to insinuate that twenty-five to fifty thousand doctors, researchers, medics, people you meet, people who treat you, are actively conspiring. . . .”²⁴ To this Wakefield responds, “It took five people to conspire to produce fraudulent data in this case.” That is quite a discrepancy. When challenged to justify his numbers, Grimes stated, “The number of people involved in the CDC/WHO was in

the decision and won, with the judge issuing a sharp rebuke of the GMC proceedings. *The Independent* reports: “The judge criticised the disciplinary panel’s ‘inadequate and superficial reasoning and, in a number of instances, a wrong conclusion’” (Aston 2012). That fact does not exonerate Wakefield, but it does raise questions about whether he was treated fairly. (See also Elisha et al. 2021.)

²³ For the record, Grimes’s paper lists the size of a combined CDC and WHO vaccine conspiracy at twenty-two thousand, though that discrepancy is not important for our purposes.

²⁴ See “96fm Vaxxed Debate with Andy Wakefield.” *Fiona O’leary*. YouTube.

<https://www.youtube.com/watch?v=sVzMcZzqC5E> The most relevant portion starts at 00:19:00.

a paper I did earlier on this year when I worked out how many people would have to be involved.” This clearly refers to Grimes 2016. But just as clearly, contrary to Grimes’s claim, his paper did nothing to *establish* the size of the conspiracy alleged. It merely *assumed* (rather unrealistically) that it must encompass the whole of both institutions.

To be fair, Grimes can most plausibly be read as merely suggesting that members of the scientific community would have uncovered and exposed the fraud, had it occurred, not that they would have been involved in the underlying conspiracy itself. Grimes explains that he “assume[s] all scientists involved would have [to] be aware of an active cover-up, and that a small group of odious actors would be unable to deceive the scientific community for long timescales” (Grimes 2016: 8). But it is at best unclear that this assumption would hold in a case like this. If five scientists produced a fraudulent study that purported to show that there was no connection between the MMR and autism—a result that nearly everyone in the scientific community both expected and desired—it is not obvious that this would have been easily uncovered by uninvolved scientists. The number of people in position to detect the deception might be quite small, and those people may not be particularly suspicious or eager to find fault.

Interestingly, the purported conspiracy that Wakefield and Grimes were arguing about involved an insider, CDC scientist William Thompson, who leaked information in a way that seems to fit Grimes’s model. Thompson was even involved in the “conspiracy,” not just a member of the same institution who stumbled upon it. However, his revelations did *not* result in the “conspiracy theory” being accepted as true. This seems to further confirm that it is a mistake to assume that leaks decisively expose conspiracies. Revelations can simply be walked back, “contextualized,” and/or largely ignored by the mainstream media. When that happens, it is hard to know what to believe. In this case, in a series of recorded telephone conversations, Thompson

made several quasi-confessions. He then gave a statement in which he said the collaborators on the study “intentionally withheld controversial findings” and met to get rid of documents.²⁵ However, another formal statement released by Thompson seems to moderate his assertions enough to allow an interpretation in which the appropriateness or inappropriateness of what had occurred may be viewed as a matter of reasonable scientific disagreement (see Thompson 2014). Regardless of the underlying truth in this case, the point here is that the genie can sometimes be put back in the bottle. Even if someone “talks,” we still don’t know if they should be believed, or how they should be interpreted. These quite reasonable doubts can be seized upon to nullify, or disarm, the revelations. Whether or not that is really *proper* may be hard to adjudicate.

These difficulties are, perhaps, typical of what we can expect when there are revelations about most conspiracy theories. Grimes’s examples of revealed conspiracies, in contrast, may be quite exceptional. Given that information about the Tuskegee experiment had been published along the way, once the light was shown on it, it was impossible to deny. But most controversial conspiracy theories cannot be assumed to be analogous in this respect. And Snowden’s feats, by which he was able to provide dramatic evidence, are difficult, costly, and dangerous to replicate. As for the FBI scandal, let’s not forget that Whitehurst first took the matter to his superiors, who did nothing. Then, after he took the matter to the Department of Justice, he was fired, and his credibility was attacked. It was ten years later that he was finally vindicated—and that was not a forgone conclusion. The Whistleblower Network News calls him “America’s first successful FBI

²⁵ Thompson’s statements to this effect were read into the Congressional record (see Willingham 2015). For more background, see Barry (2015).

whistleblower” (WNN Staff 2017). Again, this is not exactly a case that can be assumed to be representative with respect to its ultimate success in becoming an acknowledged fact.

5.6 Grimes’s moral error

Grimes’s calculations are explicitly premised on the notion that the Tuskegee experiment was not unethical *from its inception* (in 1932). Grimes writes, “The [Tuskegee] study became unethical in the mid 1940s, when penicillin was shown to effectively cure the ailment and yet was not given to the infected men” (Grimes 2016: 6). This suggests that, prior to penicillin, it was not unethical for doctors to deceive these African American men, to conceal the nature of their disease, and to treat them as *subjects* rather than as *patients*, performing experiments on them without their informed consent, rather than providing the standard of care current at the time (ineffective and possibly harmful though it may have been). Let’s not forget that not only were these sharecroppers *not informed* that they had a terrible and deadly *transmissible* disease but were *actively blocked* from finding out. And yet, based on the notion that this was somehow *not* unethical, Grimes considers the length of time during which the conspiracy was effectively concealed to be twenty-five years, rather than forty. Twenty-five years being the “Time calculated from unethical experiment duration—1947 to 1972” (Grimes 2016: 7).

In terms of invalidating Grimes’s conclusions, this is not nearly the most significant consideration. But it provides a clear way of showing that his study is seriously flawed. While the study *seems* to be based on objective numbers, this critical number is skewed by an ethical judgement that is at least dubious. And yet there is no indication that the journal editor, peer reviewers of the article, Grimes’s social science colleagues who cite the article, or the media outlets that promote its findings raised any objection to this. That is worth repeating: both

scholars and journalists appear to have registered no serious objection to the claim that, for the first fifteen years, the Tuskegee Study was not unethical.

Let's be clear about what Grimes is doing by introducing the issue of when the experiment became unethical. Although it is not possible to enter someone else's mind, it seems he is trying to prove a preconceived idea, namely, that belief in mature conspiracy theories is not rational.²⁶ His strategy is to use some examples of conspiracies that were revealed as a basis for determining how long it can be expected to take for a conspiracy to be revealed as a function of its size. And for his purpose, it would be helpful if the sampled conspiracies were shorter rather than longer. So, he came up with a reason for shortening the conspiracy associated with the Tuskegee study—a practice that can be euphemistically called “correcting the data”—by implicitly defining the length of the conspiracy as the length of time the study was unethical, and then treating the first fifteen years of the study as therefore not part of the period of the conspiracy. But whether or not the Tuskegee study was ethical during those early years is not even the right question. The issue is how long a secret can be held by a fairly large number of people, not whether the secret can be considered (however implausibly) as ethical. Further, different people may have different ideas about what is ethical or right. And, by all appearances, many if not most of the people involved in the study continued to believe it was ethically justified *even after* penicillin was found to be an effective treatment.

²⁶ There is nothing inherently wrong with trying to prove a preconceived idea. But we can see in cases like this how it might lead to the introduction of bias into what appears superficially to be objective science and mathematics.

5.7 *The simple proof that Grimes's analysis is invalid*

What is the relevance of suggesting that conspiracies tend to fail, or are unlikely to succeed? Surely it is this: to suggest that conspiracy theories are likely to be false, at least mature ones. But being *likely to be false*, by itself, doesn't imply that a theory should be dismissed. After all, most theories, of all kinds, including scientific theories, are likely to be false. This is where the mathematical analysis comes into play. If Grimes can show the *degree* of unlikeliness is sufficiently high, a dismissive attitude toward (certain) conspiracy theories might be warranted. That is what is at stake here. But has Grimes succeeded? By focusing on the Tuskegee Study (as, in the end, Grimes himself does), we can clearly see why the answer must be "no," at least for those who accept that the Tuskegee Study was unethical and conspiratorial *from the beginning*.

As mentioned above, Grimes claims, "the results of this model suggest that large conspiracies (≥ 1000 agents) quickly become untenable and prone to failure" (Grimes 2016: 1). But the Tuskegee conspiracy lasted about 40 years (from 1932 to 1972). Grimes estimates the number of people in a position to blow the whistle on the study to be 6,700. I've suggested that Grimes's method of estimation is invalid, even a bit silly. Let's nevertheless accept it for the sake of argument. If we do, no valid analysis, mathematical or otherwise, could reasonably support the conclusion that a conspiracy theory involving several thousand conspirators (or less) and which has been around for four decades (or less), can be safely dismissed. Once we have a precedent, conspiracies of similar scope must be acknowledged to have the potential to last a similar length of time before exposure. No amount of fancy math can change that, no matter how inventive or obscure.

There is one caveat to the idea that a clear precedent can be so powerfully revealing. If something significant has *changed* since the precedent-setting case, it may no longer have the

same implications. However, in this case, the main thing that has changed is that potential conspirators may have learned that they must be more careful, and thus hold their secrets more closely—perhaps by not publishing what they were doing in the scientific literature the whole time! So, although there may have been many people in some way involved in or aware of the Tuskegee experiment, if something similar were attempted today it may well be designed to be handled more discretely. At least, that is a possibility that Grimes has done nothing to refute.

Grimes's argument thus fails. This is proven straightforwardly with Grimes's own example and his own assumptions, correcting only his moral error. Indeed, his corrected analysis, if his assumptions had any merit, would give us reasons to believe that large conspiracies actually could succeed for long periods.

6. Conclusion

The idea that conspiracy theories are unlikely because conspiracies “tend to fail” is commonly asserted in academic work. But the arguments that have been given for it are not convincing. Douglas et al. (2019), for example, which is authored by influential figures in the relevant subfield, cited four sources seemingly to support that idea. Yet only one of those sources provided (seemingly) significant support for that claim. And that one, Grimes (2016), did not stand up to scrutiny. More specifically, Grimes purports to offer a mathematical demonstration showing that theories that have been around for a long time that posit large conspiracies are highly likely to have been exposed, if they had been true. But his argument has crippling flaws. It is based on cases that are unlikely to be representative. It depends on a method of approximating the size of conspiracies that is highly unreliable, potentially off by orders of magnitude. It wrongly assumes that leaks by conspirators necessarily lead to the failure of the conspiracy in the

sense of mainstream acknowledgment that it was a conspiracy. It ignores the fact that we have no good way of estimating the number of real conspiracies that have never been revealed. And it fudges the length of the Tuskegee Experiment, which is *the* conspiracy upon which Grimes's findings ultimately depend. Further, even if Grimes's argument had been valid, it would not show that conspiracies of a more modest size "tend to fail," though Douglas et al.'s claim was unqualified.

Building on the arguments of Coady (2012, 2018), I maintain that the evidence that conspiracies tend to fail remains weak. Indeed, not only is it *possible* that lots of conspiracies have remained concealed from us, we have reasons to think that many conspiracies, such as covert operations, are *actually* successfully kept from us, even though we don't know the details. And we know that some conspiracies remained secret enough for a fairly long time. We also have reasons to think that even a secret that is leaked may not result in a failed conspiracy, as it may be difficult to determine if the leak is actually true.

In my own view, the fact that a larger conspiracy has more potential leakers than a smaller conspiracy does indeed give us *a reason* to think that the larger one would be more prone to being exposed and thereby widely accepted as true, *all else being equal*. But, generally speaking, not all else is equal, and size is not *all* that matters. The particular structure of the conspiracy also matters—How well compartmentalized is it? The institutional structures involved matter—How easily can people be made to contribute to the conspiracy without even knowing it? The individuals involved matter—How "professional" are they with regard to their conspiratorial activities? And circumstantial considerations matter—How plausibly deniable is the conspiracy if a leak occurs? These are just some of the other considerations, besides size, that influence a conspiracy's risk of failure. Such considerations must be assessed on a case-by-case

basis, along with the relevant evidence, in order to judge the plausibility of any particular conspiracy theory, even a mature one that posits a large conspiracy.²⁷

References

- Aston, J. 2012. “MMR Doctor John Walker-Smith Wins High Court Appeal.” *The Independent* (March 7, 2012). <https://www.independent.co.uk/life-style/health-and-families/health-news/mmr-doctor-john-walker-smith-wins-high-court-appeal-7543114.html>
- Barajas, J. 2016. “How Many People does it Take to Keep a Conspiracy Alive?” *PBS News Hour* (Feb 15, 2016). <https://www.pbs.org/newshour/science/math-formula-charts-the-lifespan-of-hoaxes>
- Barry, K. 2015. *Vaccine Whistleblower: Exposing Autism Research Fraud at the CDC*. New York: Skyhorse Publishing.
- Basham, L. 2001. “Living with the Conspiracy.” *The Philosophical Forum* 32 (3), 265-280. (Reprinted in Coady 2006)
- Basham, L. 2018. “Joining the Conspiracy.” *Argumenta* 3 (2) (“Issue 6”), 271-290.
- Berezow, A. 2016. “Math Study Shows Conspiracies ‘Prone to Unravelling’.” BBC (January 26, 2016). <https://www.bbc.com/news/science-environment-35411684>
- Clarke, S. 2002. “Conspiracy Theories and Conspiracy Theorizing.” *Philosophy of the Social Sciences* 32, 131–150. (Reprinted in Coady 2006)

²⁷ I thank Brian Martin and Brian Keeley for helpful feedback on earlier versions of this article.

- Coady, D. 2003. "Conspiracy Theories and Official Stories." *International Journal of Applied Philosophy* 17 (2), 197-209. (Reprinted in Coady 2006)
- Coady, D. (ed.) 2006. *Conspiracy Theories: The Philosophical Debate*. Burlington, VT: Ashgate Publishing.
- Coady, D. 2012. *What to Believe Now: Applying Epistemology to Contemporary Issues*. Malden, MA: Wiley-Blackwell.
- Coady, D. 2018. "Anti-Rumor Campaigns and Conspiracy-Bating as Propaganda." In Dentith 2018 (171-187).
- Dai, Y., & Handley-Schachler, M. 2015. A Fundamental Weakness in Auditing: The Need for a Conspiracy Theory. *Procedia Economics and Finance* 28, 1-6.
[https://doi.org/10.1016/S2212-5671\(15\)01074-6](https://doi.org/10.1016/S2212-5671(15)01074-6)
- Dentith, M R. X. 2019. "Conspiracies on the Basis of the Evidence," *Synthese* 196, 2243-2261.
<https://doi.org/10.1007/s11229-017-1532-7>
- Douglas, K. M., Uscinski J. E., Sutton R. M., Cichocka A., Nefes T., Chee C. S., and Deravi, F. 2019. "Understanding Conspiracy Theories." *Advances in Political Psychology* 40 (1).
 doi: 10.1111/pops.12568
- Elisha, E., Guetzkow, J., Shir-Raz, Y., and Ronel, N. 2021). "Retraction of Scientific Papers: The Case of Vaccine Research." *Critical Public Health*.
 doi: [10.1080/09581596.2021.1878109](https://doi.org/10.1080/09581596.2021.1878109)
- Fonzi, G. 1993. *The Last Investigation*. New York: Skyhorse.
- Goertzel, T. 2010. "Conspiracy Theories in Science." *EMBO Reports* 11 (7), 493-499. doi: 10.1038/embor. 2010.84

- Grimes, D. R. 2016. "On the Viability of Conspiratorial Beliefs." *PLOS ONE* 11 (3).
<https://doi.org/10.1371/journal.pone.0147905>
- Hagen, K. 2010. "Is Infiltration of 'Extremist Groups' Justified?" *The International Journal of Applied Philosophy* 24 (2), 153-168.
- Hagen, K. 2020. "Should Academics Debunk Conspiracy Theories?" *Social Epistemology* 34 (5), 423-439.
- Hancock, L. 2006. *Someone Would Have Talked: The Assassination of President John F. Kennedy and the Conspiracy to Mislead History*. Southlake, TX: JFK Lancer Productions & Publications.
- Hertsgaard, M. 2016. "How the Pentagon Punished NSA Whistleblowers." *The Guardian* (May 22). <https://www.theguardian.com/us-news/2016/may/22/how-pentagon-punished-nsa-whistleblowers>
- Institute of Medicine 2013. *The Childhood Immunization Schedule and Safety: Stakeholder Concerns, Scientific Evidence, and Future Studies*. Washington, DC: The National Academies Press. <https://doi.org/10.17226/13563>.
- Keeley, B. L. 1999. Of Conspiracy Theories. *Journal of Philosophy* 96: 109–126.
<https://doi.org/10.2139/ssrn.1084585>
- Levy, N. 2007. "Radically Socialized Knowledge and Conspiracy Theories." *Episteme* 4 (2), 181-192.
- Martin, B. 2020. "Dealing with Conspiracy Theory Attributions," *Social Epistemology* 34 (5), 409-422. doi: 10.1080/02691728.2020.1748744

- McDonnell, T. 2016. “This Chart Shows Why Your Conspiracy Theory Is Really Dumb.”
Mother Jones (Feb. 1, 2016). <https://www.motherjones.com/politics/2016/02/you-are-wrong/>
- Ohlheiser, A. 2016. “Why the Internet’s biggest conspiracy theories don’t make mathematical sense” *The Washington Post*. (Jan. 28, 2016) <https://www.washingtonpost.com/news/the-intersect/wp/2016/01/28/the-internets-favorite-conspiracies-would-involve-too-many-people-to-stay-secret-says-science/>
- Pauly, M. (no date). “Conspiracy Theories.” *Internet Encyclopedia of Philosophy*.
<https://iep.utm.edu/conspira/> (last accessed 3-19-21).
- Pigden, C. 1995. Popper Revisited, or What is Wrong with Conspiracy Theories? *Philosophy of the Social Sciences*, 25, 3–34. <https://doi.org/10.1177/004839319502500101>
- Popper, K. 1972. *Conjectures and Refutations* (4th edition). London, UK: Routledge Kegan Paul.
(The relevant pages, i.e., pp. 123-125, are reprinted in Coady 2006: 13-15.)
- Räikkä, J. 2009. “On Political Conspiracy Theories.” *The Journal of Political Philosophy* 17.2: 185-201.
- Reynolds, E. 2017. “The Moon Landing is Not a Conspiracy Theory – and this Equation Explains Why.” *Wired* (May 2, 2017). <https://www.wired.co.uk/article/conspiracy-theories-equation>
- Reverby, S M. 2009. *Examining Tuskegee: The Infamous Syphilis Study and its Legacy*. Chapel Hill, NC: University of North Carolina Press.
- Schou, N. 2009. *Kill the Messenger: How the CIA's Crack-Cocaine Controversy Destroyed Journalist Gary Webb*. New York: Bold Type Books.

- Thompson, W. W. 2014. "Statement of William W. Thompson, Ph.D., Regarding the 2004 Article Examining the Possibility of a Relationship Between MMR Vaccine and Autism." <https://legislature.vermont.gov/Documents/2016/WorkGroups/House%20Health%20Care/Bills/H.98/Witness%20Testimony/H.98~Jennifer%20Stella~William%20Thompson%20Statement~5-6-2015.pdf>
- Welna, D. 2014. "Before Snowden: The Whistleblowers Who Tried To Lift The Veil." *NPR, Morning Edition* (July 22). <https://www.npr.org/2014/07/22/333741495/before-snowden-the-whistleblowers-who-tried-to-lift-the-veil>
- Willingham, E. 2015. "A Congressman, A CDC Whistleblower [*sic*] and an Autism Tempest in a Trashcan." *Forbes* (Aug. 6, 2015). <https://www.forbes.com/sites/emilywillingham/2015/08/06/a-congressman-a-cdc-whistleblower-and-an-autism-tempest-in-a-trashcan/?sh=417978d65396>
- Wolf, Z. B. 2007. "Big Brother Spying on Americans' Internet Data?" *ABC News*. (Nov. 7, 2007) <https://abcnews.go.com/Politics/story?id=3833172&page=1>
- WNN Staff. 2017. "Dr. Whitehurst and the FBI Lab Scandal." *Whistleblower News Network* (Nov. 21, 2017). <https://whistleblowersblog.org/2017/11/articles/intelligence-community-whistleblowers/dr-whitehurst-and-the-fbi-lab-scandal/>