

ARTICLE

Can You Keep a Secret? BS Conspiracy Theories and the Argument from Loose Lips

Ryan Ross 

Johns Hopkins University, Baltimore, Maryland
Email: rross27@jhu.edu

(Received 9 January 2021; revised 6 March 2022; accepted 30 April 2022)

Abstract

According to an argument that I will call *the argument from loose lips*, we can safely reject certain notorious conspiracy theories because they posit conspiracies that would be nearly impossible to keep secret. I distinguish between three versions of this argument: the epistemic argument, the alethic argument, and the statistical argument. I, then, discuss several limitations of the argument from loose lips. The first limitation is that only the statistical argument can be applied to new conspiracy theories. The second limitation is that no version of the argument suffices to rule out the existence of small initial conspiracies that have no need to add further conspirators. The third limitation is that no version of the argument is dialectically efficacious in the context of arguing with the relevant conspiracy theorists because nothing is said to address the alleged evidence that they cite.

Keywords: Conspiracy theories; applied epistemology; social epistemology

1. Introduction

Theories alleging that the CIA conspired with Oswald to assassinate Kennedy, that the 9/11 attacks were an inside job, that the world is secretly run by the Freemasons, etc. are ordinarily labelled as “conspiracy theories.” It is often argued that we can reject such conspiracy theories because the relevant conspiracies would be nearly impossible to keep secret. Call this *the argument from loose lips*. I distinguish between three versions of this argument: the epistemic argument, the alethic argument, and the statistical argument. I, then, discuss several limitations of the argument from loose lips. The first limitation is that only the statistical argument can be applied to new conspiracy theories (e.g., ones about COVID-19). The second limitation is that no version of the argument suffices to rule out the existence of small initial conspiracies that have no need to add further conspirators. The third limitation is that no version of the argument is dialectically efficacious in the context of arguing with the relevant conspiracy theorists because nothing is said to address the alleged evidence that they cite.¹

¹For other discussions of the argument from loose lips, see Basham (2006), Dentith and Orr (2017), Hagen (2018), and Dentith (2019).

© The Author(s), 2022. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution-NonCommercial licence (<http://creativecommons.org/licenses/by-nc/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original article is properly cited. The written permission of Cambridge University Press must be obtained prior to any commercial use.

Keeley and Grimes defend different versions of the argument from loose lips, but they agree on defining “conspiracy theory” as a theory that posits a conspiracy. Keeley says that “A conspiracy theory is a proposed explanation of some historical event (or events) in terms of the significant causal agency of a relatively small group of persons – the conspirators – acting in secret” (1999: 116). Similarly, Grimes cites Sunstein and Vermeule, who say a conspiracy theory is “an effort to explain some event or practice by reference to the machinations of powerful people, who attempt to conceal their role (at least until their aims are accomplished)” (Sunstein and Vermeule 2009: 205; Grimes 2016: 1–2). Importantly, if conspiracy theories are just theories that posit conspiracies, then some conspiracy theories are known to be true.² For example, we know that there was a conspiracy to assassinate Julius Caesar, that Nixon initiated a conspiracy to spy on the Democrats at the Watergate Office Building, and that some people are guilty of conspiracy to commit murder.

Consequently, the target of the argument from loose lips cannot be conspiracy theories in this broad sense; rather, the better interpretation is that the argument is aimed at big-secret (BS) conspiracy theories. *BS conspiracy theories* are theories that posit BS conspiracies. *BS conspiracies* are conspiracies that involve a large group of people and that have not been exposed. This definition raises two questions: What constitutes a big conspiracy? And what constitutes exposure? First, because “big” is vague, we cannot say precisely how big a conspiracy must be to count as a big conspiracy. Consequently, it’s best to focus on conspiracies that are big by any reasonable standard. This is how Keeley proceeds. Meanwhile, this vagueness is not troubling for Grimes, because his argument involves giving numerical estimates of the sizes of the relevant conspiracies. Next, my criterion for exposure is affirmation by a reputable news source: A conspiracy has been exposed only if a reputable news source has affirmed the existence of this conspiracy. Speculation and conjecture do not suffice for exposure. For example, the Watergate Affair was exposed, whereas conspiracy theories about the Freemasons are mere speculation and conjecture.

Keeley and Grimes don’t use the term “BS conspiracy theory.” Nevertheless, their arguments can be understood in terms of BS conspiracy theories. Keeley can be read as arguing that BS conspiracy theories are *prima facie* unwarranted.³ Meanwhile, Grimes can be read as providing a strategy for showing that BS conspiracy theories are probably false.

2. The epistemic argument from loose lips

According to Keeley’s epistemic argument, BS conspiracy theories are *prima facie* unwarranted because there are too many ways for the relevant conspiracies to be exposed. As he puts it,

By supposing that current events are under the control of nefarious agents, conspiracy theories entail that such events are *capable of being controlled* ... [However,] [t]he world as we understand it today is made up of an extremely large number of interacting agents, each with its own imperfect view of the world and its own set of goals ... To propose that an explosive secret could be closeted for any length of time simply reveals a lack of understanding of modern bureaucracies. Like the world itself, they are made up of too many people with too

²For other discussions of the definition of “conspiracy theory,” see Pigden (1995, 2022), Coady (2006), Dentith (2014), Cassam (2019), and Napolitano and Reuter (2021).

³I say *prima facie*, but *pro tanto* or *ceteris paribus* would be more accurate.

many different agendas to be easily controlled. (Keeley 1999: 123–4, italics in original)

To be clear, when Keeley says that an explosive secret could not be closeted, he doesn't mean that this is metaphysically impossible (2006: 111). Rather, he means that it's out of step with how we understand the world today to think that a group of conspirators is capable of keeping an explosive secret closeted (1999: 123–4). The former claim is about metaphysical necessity. Meanwhile, the latter claim is epistemic, not metaphysical: Keeley intends to make a point about the epistemology of BS conspiracy theories rather than about their truth value (1999: 111). Moreover, the latter claim is about capabilities (i.e., what someone has the ability to do), not necessity: Talk of capabilities is weaker, since the fact that I am not capable of doing a backflip doesn't entail that it's metaphysically impossible for me to do a backflip.

Having clarified that Keeley's argument makes an epistemic claim about conspirators' capabilities, we can state it as follows:

- (E1) One is prima facie warranted in believing that there is a BS conspiracy only if one is prima facie warranted in believing that a large group of people is capable of working together in secret.
- (E2) Anyone who knows that a BS conspiracy would face many internal and external threats of exposure is prima facie unwarranted in believing that a large group of people is capable of working together in secret.
- (E3) ∴ Anyone who knows that a BS conspiracy would face many internal and external threats of exposure is prima facie unwarranted in believing any BS conspiracy theories.⁴

The motivation for (E1) is that its negation is absurd. Suppose (E1) is false. Then, one could be prima facie warranted in believing that there is a BS conspiracy without being prima facie warranted in believing that a large group of people is capable of working together in secret. This implies that one could be prima facie warranted in believing that people are doing *x* without being prima facie warranted in believing that people are capable of doing *x*. But no one can be prima facie warranted in believing the former but not the latter. Therefore, (E1) is true.⁵

⁴There are three clarificatory points about my statement of Keeley's argument that are worth making. First, because Keeley defines "conspiracy theories" in terms of small groups, it might seem like a mistake to recast his argument in terms of big conspiracies. But Keeley thinks that the relevant conspiracy theories end up being committed to positing big conspiracies. He says that "an initial claim that a small group of people is conspiring gives way to claims of larger and larger conspiracies" (Keeley 1999: 122). Second, I state the epistemic argument in terms of prima facie warrant, whereas Keeley talks about warrant proper. This slight change is justified by the principle of charity, since (E2) is less plausible without the prima facie rider. Third, it might be thought that Keeley actually gives a different argument: namely, that certain conspiracies theories are unwarranted because believing them forces one into an implausible skepticism about institutions that we should trust (1999: 123). True, this is one of Keeley's arguments, but he also gives an argument from loose lips: This is clearest when says there are "too many people with too many different agendas to be easily controlled" (1999: 124).

⁵This is not to say that one is prima facie warranted in believing that *p* only if one is prima facie warranted in believing every entailment of *p*. Maybe we don't have to be prima facie warranted in believing contextually irrelevant entailments or unobvious entailments. For example, skeptical hypotheses might be irrelevant entailments in normal conversations. And Goldbach's conjecture (if true) is an unobvious entailment of the proposition that $2+2=4$. Nevertheless, the consequent of (E1) is both a relevant and an obvious entailment.

To see why (E2) is plausible, suppose an arbitrary agent knows that any BS conspiracy would face many internal and external threats of exposure. Internal threats of exposure include sloppy conspirators who leave behind evidence, morally outraged conspirators who blow the whistle, conspirators who discuss the conspiracy for a book deal or as part of a plea deal, etc. (Keeley 1999: 124). Meanwhile, external threats of exposure could come from the press, investigative government agencies, observant people with relevant knowledge, etc. (Keeley 1999: 122). But seemingly if one knows that a BS conspiracy could be exposed in all these different ways, then one is at least *prima facie* unwarranted in believing that a large group of people is capable of working together in secret. Since the agent is arbitrary, what goes for this agent goes for everyone. Therefore, anyone who knows that a BS conspiracy would face many internal and external threats of exposure is *prima facie* unwarranted in believing that a large group of people is capable of working together in secret.

3. The alethic argument from loose lips

Similar considerations can be used to motivate the alethic argument. According to this argument, BS conspiracies theories are false, not just unwarranted:

- (A1) If any big conspiracy would have been exposed by now, then every BS conspiracy theory is false.
- (A2) Any big conspiracy would have been exposed by now.
- (A3) ∴ Every BS conspiracy theory is false.

To see that (A1) is true, suppose any big conspiracy would have been exposed by now. Then, every theory that posits a big conspiracy that has *not* been exposed is false. But, by definition, a big conspiracy that has not been exposed just is a BS conspiracy. And, by definition, a theory that posits a BS conspiracy just is a BS conspiracy theory. It follows that every BS conspiracy theory is false. Therefore, if any big conspiracy would have been exposed by now, then every BS conspiracy theory is false.

The motivation for (A2) mirrors the motivation for (E2). If there were any big conspiracies, then there would be a conspiracy consisting of many knowing conspirators who have different worldviews and different agendas. Such a conspiracy would face many threats of exposure inside and out. But seemingly a conspiracy that faced so many threats of exposure would have been exposed by now. Therefore, any big conspiracy would have been exposed by now.

4. The statistical argument from loose lips

Grimes' (2016) statistical argument uses the Poisson probability mass function to calculate the probability that there would be at least one leak sufficient to expose a conspiracy. The Poisson probability mass function is as follows:

$$P(x) = \frac{\lambda^x e^{-\lambda}}{x!} \text{ for } x = 0, 1, 2, \dots$$

Here, x is the number of times the target event occurs, λ is rate at which the target events occur, e is Euler's number ≈ 2.71828 , and $!$ is the factorial function.

To calculate the probability that there will be at least one leak, L , we find the probability that there will be less than one leak (in which case, $x = 0$) and subtract from one:

$$L(x = 0) = 1 - \frac{\lambda^0 e^{-\lambda}}{0!}$$

Since $\lambda^0 = 1$ and $0! = 1$, we can simplify the formula as follows:

$$L(x = 0) = 1 - e^{-\lambda}$$

Grimes says that $\lambda = t\phi$, where t is the number of years the conspiracy has been going on and ϕ is the mean expected number of leaks per year (2016: 4). Then,

$$L(x = 0) = 1 - e^{-t\phi}$$

To calculate the probability of a leak, we need to know what ϕ is. Grimes says that ϕ can be understood in terms of p and $N(t)$, where p is the probability that a given conspirator will leak the conspiracy per year, and $N(t)$ is the number of conspirators at any given time (2016: 4). If the conspiracy consists of two people, then the probability that it will remain secret in a given year is $(1 - p)^2$. If three people, then $(1 - p)^3$. More generally, the probability that a conspiracy will remain secret per year is $(1 - p)^{N(t)}$. And if the probability that a conspiracy will remain secret per year is $(1 - p)^{N(t)}$, then the probability that it will be leaked (i.e., not remain secret) per year is $1 - (1 - p)^{N(t)}$. Hence,

$$\phi = 1 - (1 - p)^{N(t)}$$

Thus, the probability that there will be at least one leak sufficient to expose a conspiracy can be expressed as follows:

$$L(x = 0) = 1 - e^{-t(1 - (1-p)^{N(t)})}$$

Grimes' main question is: How many years and how many conspirators would it take for the probability of a leak to exceed 95%? Put this way, it's natural to rewrite the Poisson probability mass function as a function whose inputs are time and the number of conspirators:

$$L(t; N(t)) = 1 - e^{-t(1 - (1-p)^{N(t)})}$$

Before we put this formula to work, there are two empirical parameters that need to be estimated to solve for L : namely, p and $N(t)$. Grimes estimates p (the probability that a conspirator will leak the conspiracy in a given year) based on considerations relating to the exposure of the NSA PRISM program, the FBI's use of pseudoscience in criminal trials, and the Tuskegee syphilis experiment (2016: 6–7). Of these three examples, the value for p is lowest for people involved with the NSA PRISM program at 4.09×10^{-6} , or approximately one in a quarter-million. Since Grimes wants to be generous to the conspiracy theorist in the way he assigns probabilities, he sets p to 4.09×10^{-6} when calculating the probability that other conspiracies would be leaked (2016: 8).

Next, Grimes considers the value of $N(t)$ (the number of conspirators at any given time) for four conspiracy theories: the moon landing conspiracy theory, the antivaxxer

conspiracy theory, the climate change conspiracy theory, and the cancer cure conspiracy theory. He estimates that:

The moon landing conspiracy would consist of 411,000 people.

The antivaxxer conspiracy would consist of 22,000 people.

The climate change conspiracy would consist of 405,000 people.

The cancer cure conspiracy would consist of 714,000 people. (Grimes 2016: 8)

These estimates are based on data concerning the size of various organizations that would have to be involved if the relevant conspiracy theory were true. For example, the estimate for the number of conspirators who would be involved in faking the moon landing is based on the reported peak NASA employment in 1965, and the estimate for the number of conspirators who would be involved in the antivaxxer conspiracy is the sum of the number of people employed by the CDC (15,000 people) and the WHO (7000 people) (Grimes 2016: 8). Notably, Grimes assumes that everyone employed by these organizations would be a conspirator; he doesn't entertain the possibility that the conspiracies might be limited to small groups within these organizations.

Based on these estimates for p and $N(t)$, Grimes calculates the amount of time it would take for the probability of a leak to exceed 95%. The probability that the faking of the moon landing would be leaked is calculated to exceed 95% after 3.68 years. The probability that the suppression of a cure for cancer would be leaked is calculated to exceed 95% after 3.17 years. Depending on the number of conspirators, the probability that the suppression of a vaccine-autism link would be leaked is calculated to exceed 95% after either 3.15 years or 34.78 years. Last, depending on the number of conspirators, the probability that deceit about climate change would be leaked is calculated to exceed 95% after either 3.7 years or 26.77 years (Grimes 2016: 11).

Strictly speaking, Grimes' calculations are only aimed at these four conspiracy theories. Therefore, his calculations don't constitute a fully general objection to BS conspiracy theories. Nevertheless, Grimes' methodology, if cogent, can be applied to other BS conspiracy theories on a case-by-case basis. Once we select some BS conspiracy theory, we simply calculate the number of years until exposure by keeping the estimation of p at 4.09×10^{-6} and estimating $N(t)$ on empirical grounds.

5. New conspiracy theories

Having clarified the argument from loose lips, I will discuss some of its limitations. The first limitation concerns whether it can be applied to the wave of new conspiracy theories that keep coming up on the news and online. The answer is that it depends on which version of the argument we use. Consideration of new conspiracy theories shows that (E2) of the epistemic argument and (A2) of the alethic argument are unjustified unless they are restricted so that they do not apply to new conspiracy theories.⁶ However, the statistical argument can be used to object to at least some new conspiracy theories assuming that it's otherwise cogent.

⁶The reviewer points out that maybe the epistemic and alethic arguments were never intended to apply to new conspiracies in the first place.

A *new conspiracy* is a conspiracy that has been around for only a short time, and a *new conspiracy theory* is a theory that posits a new conspiracy. For example, the following conspiracy theories apparently posit new conspiracies:

The 2020 election conspiracy theory: There was a conspiracy among Democrats to rig the 2020 US election so that Biden would win.

The microchipped vaccine conspiracy theory: There is a conspiracy to microchip people via the COVID-19 vaccine.

The Capitol riot conspiracy theory: There was a conspiracy among members of Antifa to carry out the riot at the US Capitol Building on January 6, 2021.

New conspiracies confront the epistemic argument with a dilemma. According to (E2), anyone who knows that a BS conspiracy would face many internal and external threats of exposure is *prima facie* unwarranted in believing that a large group of people is capable of working together in secret. But the belief that is supposed to be *prima facie* unwarranted can be interpreted in two ways:

A large group of people is capable of working together in secret *for a long time*.

A large group of people is capable of working together in secret *for a short time*.

If (E2) is about the former, then the epistemic argument doesn't rule out new conspiracy theories, since no new conspiracies have been around for a long time. Meanwhile, if (E2) is about the latter, then (E2) is unjustified: The premise that a BS conspiracy would face many internal and external threats of exposure doesn't support the conclusion that the relevant conspiracy would be exposed in a short time. Internal threats of exposure might take a long time to leak the conspiracy because conspirators would have incentives to keep the conspiracy secret, making them reluctant to leak the conspiracy deliberately and careful to avoid leaking it accidentally. Meanwhile, external threats of exposure might take a long time to cause the conspiracy to be exposed because investigations often take a long time. Thus, since the latter reading of (E2) is unjustified, it's better to restrict it so that it doesn't make a claim about new conspiracies.⁷

New conspiracies also make trouble for the alethic argument. According to (A2), any big conspiracy would have been exposed by now. When (A2) is instantiated with respect to a new big conspiracy, "by now" refers to a short time, since new conspiracies have only been around for a short time. Therefore, (A2) entails that any new big conspiracy would be exposed in a short time. However, the reason given for (A2) is that a big conspiracy would face many internal and external threats of exposure, and this doesn't show that exposure would only take a short time. As with the epistemic argument, this reason doesn't rule out the possibility that the conspirators could be slow to leak

⁷One possible response, which the reviewer mentions, is that what really matters is the thoroughness of the relevant investigation, not the amount of time that has passed. In turn, maybe the relevant belief in (E2) should be this: A large group of people is capable of working together in secret *despite a thorough investigation relating to their activities*. Nevertheless, an argument based on this version of (E2) won't be dialectically efficacious because BS conspiracy theorists will deny that a thorough investigation has been carried out. Moreover, this problem isn't limited to just new conspiracies. For example, people who believe conspiracy theories about Kennedy's assassination, 9/11 being an inside job, or the Freemasons will deny that the official story is supported by a thorough investigation. They think relevant evidence is being hidden. Parallel considerations apply to the alethic argument.

the conspiracy or that an investigation could take a long time. Even if we were to concede that this reason supports the conclusion that any big conspiracy would be exposed *eventually*, it doesn't suggest that exposure will be quick enough to rule out new conspiracy theories. To avoid this problem, (A2) should be reformulated so that it doesn't make a claim about new conspiracies.

Unlike the epistemic and alethic arguments, the statistical argument can be applied to new conspiracies. In a follow-up to his initial article, Grimes argues against various conspiracy theories about COVID-19: e.g., that the COVID-19 pandemic is a hoax, that SARS-CoV-2 was created and released deliberately, and that COVID-19 is caused by 5G radiation. He estimates that the number of people who would need to be involved in these conspiracies is so high that the relevant conspiracies would have been exposed in a matter of weeks. For example, he calculates that if there were a conspiracy among drug companies, public health bodies, and researchers to cover up the fact that the COVID-19 pandemic is a hoax, then there is a 95% chance that this conspiracy would be exposed in 40 weeks (Grimes 2021: 10). Thus, if the statistical argument is otherwise cogent, then it can be used to rule out some new conspiracies. This is not to say that the statistical argument *always* works against new conspiracies, since there could be new conspiracy theories that posit few enough conspirators that the statistical argument cannot rule them out in a timely manner. Rather, my point is that the statistical argument is better equipped to handle new conspiracies, whereas the epistemic and alethic arguments lack a clear basis for maintaining that new conspiracies would have already been exposed.

6. Small conspiracies

Another limitation of the argument from loose lips concerns small conspiracies. Some views that are ordinarily labelled as "conspiracy theories" only need to posit small initial conspiracies that have no need to add further conspirators later on. But the argument from loose lips only applies to conspiracy theories that posit big conspiracies. Therefore, the argument is not general enough to rule out all the views that are ordinarily labelled as "conspiracy theories."

There are several well-known conspiracy theories that do not need to posit big conspiracies. For example,

The JFK conspiracy theory: Lee Harvey Oswald was part of a conspiracy to assassinate John F. Kennedy.

The Oklahoma City bombing conspiracy theory: Some ATF agents were involved in the conspiracy to bomb the Murrah Federal Building. They intended to stop the bombing at the last minute so that they would be portrayed well in the media. However, they somehow failed to stop the bombing and covered up their involvement.

The 9/11 conspiracy theory: There was a conspiracy within the US government to either play an active role in carrying out the 9/11 attacks or deliberately allow the 9/11 attacks to happen.⁸

First, some versions of the JFK conspiracy theory posit a small conspiracy. For example, consider a version involving the CIA:

⁸For further discussion of the Oklahoma City bombing conspiracy theory, see Keeley (1999). For further discussion of the JFK conspiracy theory and the 9/11 conspiracy theory, see Hagen (2018: 27–9).

The CIA attempted to assassinate Fidel Castro without Kennedy's knowledge. When Kennedy found out about this, some high-ranking members of the CIA became concerned that he would disband the CIA. Consequently, they made plans to have Kennedy assassinated by Oswald. These members of the CIA were able to cover up their involvement because Allen Dulles, a former Director of the CIA, was a member of the Warren Commission. Dulles used this position to guide the Commission to the conclusion that Oswald acted alone.⁹

We can divide this conspiracy into two parts: an initial conspiracy to assassinate Kennedy and a later conspiracy to cover up the CIA's involvement. The initial conspiracy need not have involved more than a few people. In fact, the conspiracy would be more likely to succeed if it were kept small. Likewise, the cover-up need not have involved more than a few people. According to this theory, the CIA was able to cover its tracks because Allen Dulles helped guide the Warren Commission to the conclusion that Oswald acted alone. In which case, it would not be necessary to posit a vast conspiracy involving various government agencies or the press.¹⁰

We can also state a version of the Oklahoma City bombing conspiracy theory that posits a small conspiracy. Here is an example:

A small team of ATF agents encouraged McVeigh to bomb the Murrah Federal Building, because they intended to stop the bombing and look like heroes to the public. However, they failed to stop the bombing. At which point, they started covering up their involvement. Their cover-up was convincing enough that other federal agencies were fooled or at least didn't have adequate evidence to implicate the ATF. In either case, other agencies were not in a position to take a public stand against the official story. Moreover, because the cover-up was mostly convincing and no federal agencies contradicted the official story, the press ended up being misled or at least not having much to report that would contradict the official story.

This theory doesn't implicate the press or government agencies other than the ATF. It doesn't even implicate *all* of the ATF – just some of its agents. And as long as the original conspirators within the ATF kept quiet, there is no clear reason why this theory would require positing any further conspirators.

Last, there are versions of the 9/11 conspiracy theory that can avoid positing big conspiracies. Consider the following:

Some members of the US government (e.g., George W. Bush and Dick Cheney) had prior knowledge that the 9/11 attacks were going to be carried out. This knowledge came from intelligence gathered by the CIA. However, rather than using this intelligence to stop the attacks, Bush and his inner circle deliberately allowed the attacks to be carried out. Their motives may have been to profit from the war on terror, to enable the US to take control of Middle Eastern oil, to westernize the Middle East, or a combination of these.

⁹This theory is based on a possibility that Dave Perry says he is unable to fully debunk (Patterson 2018).

¹⁰Some people might argue that this theory is committed to the view that Lyndon B. Johnson was involved, since he selected Allen Dulles to be on the Commission. Even if this is conceded, the addition of one extra person doesn't make this a big conspiracy.

According to this theory, the conspiracy consisted of Bush, some of his most trusted associates, and probably some members of the CIA. If this small group of people could remain quiet and cover up their involvement, there would be no need to involve anyone else.

One objection to my point about small conspiracies is that the views ordinarily labelled as “conspiracy theories” must posit further conspirators to explain why no external investigators have uncovered the conspiracy.¹¹ For example, Keeley writes:

as a conspiracy theory matures, attempt after attempt to falsify a conspiracy theory appears to succeed, and this apparent success must be explained as the nefarious work of the conspirators. As a result of this process, an initial claim that a small group of people is conspiring gives way to claims of larger and larger conspiracies. (Keeley 1999: 122)

According to Keeley, the views that are ordinarily labelled as “conspiracy theories” must explain why the relevant conspiracy has not been exposed, and often the only feasible explanation is that the people who are tasked with investigating the conspiracy are in on it. This has the effect of swelling the estimated number of conspirators. If this pattern applies to conspiracy theories about Kennedy’s assassination, the Oklahoma City bombing, and the 9/11 attacks, then my examples are not problematic for the argument from loose lips.

However, my examples show how this objection can be answered. Instead of constantly expanding the conspiracy, the conspirators could take steps to avoid detection by investigators, circumventing the need to bring the investigators into the conspiracy and allowing the conspiracy to remain small. Allen Dulles could have deliberately misled the Warren Commission. Members of the ATF could have covered up their involvement well enough that no one had sufficient evidence to accuse them. Bush and others could have avoided exposure because there was no way for investigators to prove their involvement provided the conspirators all remained silent.

Another objection is that even if some conspiracies could be kept small, this won’t work for all conspiracies. For example, it’s plausible that the conspiracies posited by the following conspiracy theories could not be kept small:

The climate change conspiracy theory: There is a conspiracy to fake or exaggerate the occurrence of anthropogenic climate change.

The flat Earth conspiracy theory: The Earth is flat, but there is a conspiracy to mislead the public into believing that it is round.

The Holocaust denial conspiracy theory: The Holocaust never happened, but there is a conspiracy to mislead the public into believing that it did.

When we look at these conspiracies on case-by-case basis, it’s hard to see how they could avoid involving a large group of knowing conspirators.

Consider the climate change conspiracy theory. Approximately 150,000 climate scientists have published on climate change from 1991 to 2015 (Powell n.d.). Meanwhile, between 97% and 99.94% of published articles about climate change agree that anthropogenic climate change is occurring (Cook *et al.* 2013; Powell

¹¹As the reviewer mentions.

2016).¹² As Grimes says, it seems that nearly every published climate scientist would have to be part of the climate change conspiracy (2016: 11). After all, it's unlikely that such widespread data manipulation would happen without most of the co-authors being aware. And even if only half of the co-authors on these papers were in on the conspiracy, the conspiracy would still consist of approximately 75,000 climate scientists. But if approximately 75,000 people are part of the conspiracy, then it's big by any standard.

There is also reason to think that the flat Earth conspiracy theory must posit a big conspiracy. This is because there is an extremely broad base of testimony that speaks against the claim that the Earth is flat, which is difficult for flat-Earthers to explain without saying that such testimony comes from conspirators trying to mislead the public.¹³ For example, the flat Earth conspiracy would need to involve a variety of physicists, astronomers, cosmologists, and government agencies (e.g., NASA) across the globe. Moreover, the fact that no government in the world contradicts the view that the Earth is round makes it likely that many governments would have to be involved in the conspiracy, especially the governments that have space programs. I cannot give a precise estimate of the number of conspirators who would have to be involved. Nevertheless, the fact that the flat Earth conspiracy theory is committed to the existence of an international conspiracy among physicists, astronomers, cosmologists, public officials, and government workers suggests that it must posit a big conspiracy.¹⁴

Likewise, there is reason to think that the Holocaust denial conspiracy would need to involve a large group. The USC Shoah Foundation (2021) has recorded over 50,000 interviews in which people (mostly Jewish survivors) describe their experiences of the Holocaust.¹⁵ Moreover, this conspiracy would apparently have to include all four government bodies (those of the US, the UK, France, and the USSR) that worked together to prosecute Nazi war criminals during the Nuremberg trials. Additionally, the Holocaust denial conspiracy theory must posit further conspirators to explain the existence of incriminating Nazi documents, photographs and videos of Nazi atrocities, archeological evidence of mass graves, and statistical data that point to a huge decrease in the Jewish population of Europe. Given the broad range of evidence that would need to be faked, the conspiracy would probably have to be a big one.

Based on these considerations, I grant that some of the views ordinarily labelled as "conspiracy theories" would probably need to posit big conspiracies. Nevertheless, this objection doesn't fully nullify the problem presented by small conspiracies. Assume that conspiracy theories about climate change, the Earth being flat, and the Holocaust must posit big conspiracies. But this doesn't show that conspiracy theories about Kennedy's

¹²There is room for disagreement about these estimates, but not enough to plausibly maintain that this conspiracy is small. Cook *et al.* (2013) say that approximately 97% of climate scientists agree in contrast to Powell's estimate of 99.94%. See Powell (2016) for further references and a defense of his higher estimate. Meanwhile, Grimes estimates that 405,000 people would have to be involved in the conspiracy.

¹³As the reviewer notes.

¹⁴Notably, this theory apparently requires that Russia and the United States have been *working together* for decades to lie about the Earth being round. Another apparent implication is that Republicans and Democrats have agreed to lie about the Earth being round despite all their other disagreements. Though, BS conspiracy theorists might reply that such disputes are staged and only meant to distract the public.

¹⁵The USC Shoah Foundation interviews are paywalled, but other interviews with Holocaust survivors are available for free on the Yad Vashem YouTube channel at <https://www.youtube.com/user/YadVashem/videos>.

assassination, the Oklahoma City bombing, and the 9/11 attacks would have to posit big conspiracies. Therefore, consideration of small conspiracies serves its stated purpose: It shows that the argument from loose lips fails to rule out some important conspiracy theories that it was meant to rule out.

Small conspiracies are problematic for all three versions of the argument from loose lips. Due to the vagueness of “big,” the epistemic argument and the alethic argument only work when the relevant conspiracy is obviously big by any reasonable standard. But the aforementioned conspiracy theories about Kennedy’s assassination, the Oklahoma City bombing, and the 9/11 attacks do not posit obviously big conspiracies. Meanwhile, the statistical argument only works if our calculations show that there is a high probability that a given conspiracy would have been exposed by now. The problem is that each of the conspiracies I described would involve very few people. Even if we assume that 100 people would be involved in each conspiracy, the probability that they would be leaked doesn’t reach 95% until after approximately 7327 years according to Grimes’ methodology.¹⁶ Moreover, even if we waited 100 years after Kennedy’s death, Grimes’ methodology says that there is only about a 4% chance that a conspiracy of 100 people would be exposed.¹⁷ Since the Oklahoma City bombing conspiracy theory and the 9/11 conspiracy theory posit newer conspiracies, Grimes’ methodology entails that the probability of their being leaked in the same amount of time is even lower (assuming they did not involve more than 100 conspirators).

7. Alleged evidence for BS conspiracy theories

A third limitation of the argument from loose lips concerns its dialectical efficacy (i.e., whether its premises provide BS conspiracy theorists with a non-question-begging reason to change their minds). The argument from loose lips doesn’t address the alleged evidence cited in support of BS conspiracy theories. But the argument isn’t dialectically efficacious if it doesn’t address this alleged evidence. Therefore, the argument is not dialectically efficacious.

7.1. Alleged evidence and the epistemic argument

BS conspiracy theorists can grant everything the epistemic argument says while maintaining their favorite conspiracy theories. This is because the epistemic argument only purports to show that believing BS conspiracy theories is *prima facie* unwarranted (rather than unwarranted proper) for people who know that any BS conspiracy would face many internal and external threats of exposure. But the BS conspiracy theorist can respond,

“Yes, the fact that any BS conspiracy would face many internal and external threats of exposure provides *prima facie* warrant for believing that my BS conspiracy theory is false. But there is additional evidence that supports believing my conspiracy theory despite these threats of exposure. And we cannot conclude that my BS conspiracy theory is false until it’s shown that your evidence outweighs my evidence.”

That is, the epistemic argument isn’t dialectically efficacious because it doesn’t address the alleged evidence that is cited in favor of BS conspiracy theories.

¹⁶Using Grimes’ formula, $L(t; N(t)) = 1 - e^{-t(1 - (1-p)^{N(t)})}$, solve for t when $L = .95$, $N(t) = 100$, and $p = 4.09 \times 10^{-6}$. The result is that $t \approx 7327$.

¹⁷ $1 - e^{-100(1 - (1 - 0.0000409)^{100})} \approx 0.040$.

To fix this problem, one would have to add a premise to the epistemic argument that addresses this alleged evidence. Such an addition would have to go roughly as follows:

- (E4) Any alleged evidence for any BS conspiracy theory is, at best, outweighed by the evidence against that BS conspiracy theory.

Plausibly, if (E3) and (E4) are true, then anyone who knows that a BS conspiracy would face many internal and external threats of exposure is unwarranted (and not just *prima facie*) in rejecting any BS conspiracy theory. (Interestingly, this revised argument ends up drawing an even more severe conclusion against the warrant for BS conspiracy theories.)

However, asserting (E4) without argument begs the question against the BS conspiracy theorist. This is because we should expect that BS conspiracy theorists will not just concede that the (alleged) evidence they cite is outweighed by the evidence that we cite. To avoid begging the question, we must support (E4) by defending the view that our evidence is better than the alleged evidence favored by BS conspiracy theorists. But this is something the epistemic argument is not equipped to do.

Notably, once we acknowledge the importance of arguing for (E4), the loose-lips part of the argument becomes less important. If one's interlocutor bases their belief on testimony, then the debate about whether the BS conspiracy theory is true will involve a debate about whose sources of information are better. Meanwhile, if one's interlocutor bases their belief on alleged evidence they have gathered, then the debate will be about the best way to evaluate the alleged evidence. In either case, the debate will no longer be primarily about how big the conspiracy is or how long a secret could be kept; instead, the debate will be about which sources to trust or which pieces of evidence to rely on.

At this point, it's easy to get impatient. Those who reject views that are ordinarily labelled as "conspiracy theories" will claim to already know that so-called "mainstream" sources of information are reliable even if baldly asserting this begs the question in some contexts.¹⁸ Consequently, it's tempting to think that we can simply ignore those who don't trust mainstream sources of information on the grounds that it's not our problem if some people believe a bunch of nonsense. But the problem is that it *is* our problem. For example, conspiracy theories about climate change may lead to preventable damage to the environment. Conspiracy theories about the Holocaust may help the spread of antisemitism. Conspiracy theories about vaccines (in particular, vaccines for MMR and COVID-19) continue to cause preventable deaths and infections. Conspiracy theories about the 2020 US presidential election caused a riot at the US Capitol where people died.¹⁹

7.2. Alleged evidence and the alethic argument

Now, consider the alethic argument. The problem is (A2): that any big conspiracy would have been exposed by now. BS conspiracy theorists can object to (A2) by noting that conspirators would have robust incentives for keeping their conspiracy from

¹⁸As Kelly says, "Even if one genuinely knows that *p*, there might nevertheless be contexts in which it would be inappropriate to cite *p* as evidence" (2008: 73). For example, I can know something based on the premise that I heard it on the news even if giving an argument based on this premise would beg the question against someone who incorrectly believes that my source is fake news.

¹⁹Consequently, I would question Coady's claim that conspiracy theorists usually only harm themselves (2007: 196). Cassam (2019: 63–91) also makes this point against Coady.

exposure. But if the conspirators have robust incentives for keeping their conspiracy from exposure, then they are unlikely to do anything that would cause their conspiracy to be exposed. And if the conspirators are unlikely to do anything that would cause their conspiracy to be exposed, then it's unclear that every big conspiracy would have been exposed by now.

What incentives? If a conspirator were to cause their conspiracy to be exposed, they would stop profiting from the conspiracy, they would be forced to pay a large sum of money for damages, they would ruin their reputations, they would have trouble finding work again, and they would risk facing harsh prison sentences (Hagen 2018: 27–8). Additionally, conspirators would likely pressure would-be whistleblowers to stay quiet via blackmail, threats against their life or family, or the threat of being scapegoated (Basham 2006). Thus, if there were a BS conspiracy, the conspirators would probably have many strong incentives to keep the conspiracy from exposure.

To defend the alethic argument from this objection, one must argue that such incentives wouldn't make a difference. That is, one must argue that

Incentive hypothesis: Any big conspiracy would have been exposed by now even if the conspirators had robust incentives for keeping the conspiracy from exposure.

To avoid begging the question, some argument for the incentive hypothesis is needed. Because the incentive hypothesis is an empirical claim about human behavior, support for it must either come from an empirical study or from our common-sense knowledge of what people are like. But the people who reject BS conspiracy theories on the grounds that such conspiracies would have been exposed by now don't cite any empirical studies.²⁰ Therefore, they must say that common sense is enough to support the incentive hypothesis.

There are different views about what kind of epistemic standing a claim must have to count as commonsensical. I will distinguish between the strict view and the relaxed view. The strict view is that common sense claims have an exceedingly high degree of justification. For example, according to G.E. Moore (1925: 193), common sense is restricted to what we know with certainty. Similarly, Kelly discusses a conception of common sense according to which common-sense claims are ones whose justification is “invulnerable to being undermined *by means of philosophical argument*” (2005: 180, italics in original). Meanwhile, the relaxed view is that a claim's being part of common sense gives it a positive epistemic status that is nevertheless defeasible.²¹ If none of these conceptions of common sense can be used to show that BS conspiracy theorists should accept the incentive hypothesis, then it's unlikely that any other conception of common sense will succeed in doing so.

Let's start with Moore's view according to which common sense is restricted to claims that we know with certainty (1925: 193). Plausible examples of claims that I know with certainty are that I have hands, that there are tables, that other people

²⁰Even if one did cite empirical studies as evidence for the incentive hypothesis, one would still have to show that this evidence isn't outweighed by the alleged evidence that BS conspiracy theorists cite. Moreover, there are currently no empirical studies I know of that help support the incentive hypothesis, and it would be infeasible to conduct such studies. See the appendix.

²¹Another view is that common-sense claims have no special epistemic standing at all, because common sense is a miscellaneous collection of assumptions, biases, and beliefs held by many people often without any good reason. If this is what common sense is, then it doesn't support the incentive hypothesis or any other claim.

have minds, that I know a lot, and that other people know a lot (Kelly 2008). For present purposes of trying to support the incentive hypothesis, it's notable that some counterfactuals could be added to this list (e.g., that if I were to drop a pen right now, it would fall).

Why not add the incentive hypothesis to this list of trivialities? One problem is that it's overly confident to describe the incentive hypothesis as something that we "know with certainty" rather than describing it as "reasonable" or "probable." A further problem is that even if I know with certainty that the incentive hypothesis is true, there are BS conspiracy theorists who do not know it (because they do not believe it). So, how can I convince them? It won't be dialectically efficacious to argue: "I know it with certainty; therefore, it's true." The BS conspiracy theorist will be justified in asking how I know. Either my knowledge of the incentive hypothesis is non-inferential or inferential. If my knowledge is non-inferential, then I will be at a loss to persuade my interlocutor, because the grounds that support non-inferential knowledge (e.g., rational intuition, experience, and memory) are not things one can simply transfer to one's interlocutor. Meanwhile, if my knowledge is inferential, then I should be able to cite the premises that I base my belief on. However, it's doubtful that there are premises known by me with certainty that can be used to defend the incentive hypothesis against someone who already doubts it.

Meanwhile, it's a moot point whether the incentive hypothesis is commonsensical in the sense of being immune to refutation by philosophical arguments. BS conspiracy theorists don't question the incentive hypothesis on *philosophical* grounds; rather, they question it because of alleged anomalies, testimony from their preferred sources of information, or worries that relevant evidence is being hidden. Therefore, BS conspiracy theorists can grant that the incentive hypothesis is commonsensical in this sense and then use non-philosophical considerations to argue that the incentive hypothesis is a part of common sense that should be discarded.

Next, consider the relaxed view of common sense where being commonsensical involves having a positive, but defeasible, epistemic standing. This defeasibility is noteworthy because BS conspiracy theorists think they have good evidence for their BS conspiracy theories. Therefore, they can respond to the relaxed appeal to common sense, saying

"Yes, common sense supports the incentive hypothesis. But there is strong evidence for my favorite BS conspiracy theory. Because common sense is defeasible, we need to balance these competing considerations."

To answer this objection, the relaxed appeal to common sense needs to be supplemented with the premise that any alleged evidence for any BS conspiracy theory is, at best, outweighed by the common-sense case against it. But, like (E4), this begs the question.

So far, I have been talking in terms of common sense. If we don't want to call it "common sense," we could, instead, appeal to our "ordinary cognitive capacity to handle counterfactual conditionals" (Williamson 2007: 136). We often justifiably believe that counterfactuals are true without conducting an empirical study (e.g., we can justifiably believe that throwing a baseball at a window would break the window). Suppose we believe the incentive hypothesis based on this capacity to handle counterfactuals. Nevertheless, this capacity is "manifestly fallible" (Williamson 2007: 155). Therefore, the BS conspiracy theorist can still object that we must weigh competing considerations: This time, the BS conspiracy theorist's alleged evidence is to be weighed against our

confidence that our capacity to handle counterfactuals has led us to a correct judgment about the incentive hypothesis.

7.3. Alleged evidence and the statistical argument

The statistical argument fares no better when it comes to addressing alleged evidence cited by BS conspiracy theorists. The BS conspiracy theorist can concede that the improbability of keeping a big conspiracy secret is a drawback of their view while maintaining that the total evidence supports believing their favorite BS conspiracy theory. For example, consider the statistical argument against the climate change conspiracy theory. This argument implies that the probability that the climate change conspiracy theory is true given that it posits 405,000 conspirators is approximately 0.05 – in symbols, $p(H|E) \approx 0.05$ (Grimes 2016: 11). However, BS conspiracy theorists can concede this conclusion while maintaining that it's irrelevant because it ignores evidence. Let E^* be the alleged evidence that a BS conspiracy theorist cites to support H . For all the statistical argument shows, it could be that $p(H|E)$ is low while $p(H|E \& E^*)$ is high. In which case, the BS conspiracy theorist should think that the fact that $p(H|E)$ is low is irrelevant because it fails to conditionalize on the total evidence. To respond, one must argue that $p(H|E \& E^*)$ is low. This would require addressing the alleged evidence that BS conspiracy theorists cite in favor of their views.²² But this is something that the statistical argument doesn't attempt to do.

It might be objected that the probability of a BS conspiracy being kept secret is so extremely low that we don't even need to consider additional evidence. But consider an analogy. If I see that my lottery numbers have been drawn on TV, then (maybe after doublechecking) I know that I have won the lottery. The fact that winning the lottery is antecedently improbable shouldn't convince me that I have lost. In this case, my evidence from watching TV outweighs the probabilistic consideration that winning the lottery is improbable. Sometimes, unlikely things happen, and we can be justified in believing accordingly. Importantly, the probability of winning the lottery is much lower than the probabilities that Grimes calculates for the exposure of big conspiracies. For example, Grimes argues that there is a one in twenty chance that the climate change conspiracy theory would remain secret for over 26 years, assuming all climate scientists and all members of relevant scientific organizations are conspirators (2016: 11). Meanwhile, the odds of winning the lottery are one in several million. Therefore, if additional evidence can outweigh the antecedent improbability of winning the lottery, then a fortiori we shouldn't assume that additional evidence cannot outweigh probabilistic considerations when it comes to BS conspiracy theories.

Granted, probabilistic considerations can sometimes rationally convince us that qualitative considerations are incorrect. For example, the low probability of a putative childhood memory can rationally convince me that it was just a dream. Similarly, maybe the low probability of a BS conspiracy being kept secret is enough to undermine the alleged evidence that BS conspiracy theorists cite. But, again, defending this point would require engaging with this alleged evidence, something the statistical argument doesn't do.

²²This parallels how some theists respond to the probabilistic version of the problem of evil. The non-theist objects that the conditional probability that God exists given that evil exists is low. The theist concedes the point but says that there is additional evidence to consider and that the total evidence supports the existence of God. For discussion, see Tooley (2019: §6.1).

8. Conclusion

According to the argument from loose lips, we can reject BS conspiracy theories because they posit conspiracies that would be nearly impossible to keep secret. I noted that there are at least three versions of this argument: the epistemic argument, the alethic argument, and the statistical argument. I, then, discussed three limitations of the argument from loose lips. First, only the statistical argument can be applied to new conspiracy theories. Second, no version of the argument is strong enough to rule out the existence of small initial conspiracies that have no need to add further conspirators. Third, no version of the argument is dialectically efficacious because none of them address the alleged evidence cited in favor of BS conspiracy theories. Nevertheless, this is not to say that we should lend more credence to the views that are ordinarily labelled as “conspiracy theories” or give up on trying to convince the people who believe them. For all I have said, there are other arguments that might succeed where the argument from loose lips fails.²³

References

- Basham L. (2006). ‘Malevolent Global Conspiracy.’ In D. Coady (ed.), *Conspiracy Theories: The Philosophical Debate*, pp. 83–106. London: Routledge.
- Cassam Q. (2019). *Conspiracy Theories*. Cambridge: Polity Press.
- Coady D. (ed.) (2006). ‘An Introduction to the Philosophical Debate about Conspiracy Theories.’ In *Conspiracy Theories: The Philosophical Debate*, pp. 1–12. London: Routledge.
- Coady D. (2007). ‘Are Conspiracy Theorists Irrational?’ *Episteme* 4(2), 193–204.
- Cook J., Nuccitelli D., Green S.A., Richardson M., Winkler B., Painting R., Way R., Jacobs P. and Skuce A. (2013). ‘Quantifying the Consensus on Anthropogenic Global Warming in the Scientific Literature.’ *Environmental Research Letters* 8(2), 1–7.
- Dentith M.R.X. (2014). *The Philosophy of Conspiracy Theories*. New York, NY: Palgrave Macmillan.
- Dentith M.R.X. (2019). ‘Conspiracy Theories on the Basis of the Evidence.’ *Synthese* 196(6), 2243–61.
- Dentith M.R.X. and Orr M. (2017). ‘Secrecy and Conspiracy.’ *Episteme* 15(4), 433–50.
- Grimes D. (2016). ‘On the Viability of Conspiratorial Beliefs.’ *PLoS ONE* 11(3), e0147905. <https://doi.org/10.1371/journal.pone.0147905>.
- Grimes D. (2021). ‘Medical Disinformation and the Unenviable Nature of COVID-19 Conspiracy Theories.’ *PLoS ONE* 16(3), e0245900. <https://doi.org/10.1371/journal.pone.0245900>.
- Hagen K. (2018). ‘Conspiracy Theories and the Paranoid Style: Do Conspiracy Theories Posit Implausibly Vast and Evil Conspiracies?’ *Social Epistemology* 32(1), 24–40.
- Keeley B. (1999). ‘Of Conspiracy Theories.’ *Journal of Philosophy* 96(3), 109–26.
- Keeley B. (2006). ‘Nobody Expects the Spanish *Inquisition!* More Thoughts on Conspiracy Theory.’ In D. Coady (ed.), *Conspiracy Theories: The Philosophical Debate*, pp. 107–14. London: Routledge.
- Kelly T. (2005). ‘Moorean Facts and Belief Revision, or Can the Skeptic Win?’ *Philosophical Perspectives* 19(1), 179–209.
- Kelly T. (2008). ‘Common Sense as Evidence: Against Revisionary Ontology and Skepticism.’ *Midwest Studies in Philosophy* 32(1), 53–78.
- Moore G.E. (1925). ‘A Defence of Common Sense.’ In J.H. Muirhead (ed.), *Contemporary British Philosophy* (2nd series), pp. 193–223. London: George Allen and Unwin.
- Napolitano M.G. and Reuter K. (2021). ‘What is a Conspiracy Theory?’ *Erkenntnis*. <https://doi.org/10.1007/s10670-021-00441-6>.
- Patterson T. (2018). ‘One JFK Conspiracy Theory that Could be True.’ *CNN*, 21 March. <https://www.cnn.com/2018/03/21/us/jfk-assassination-conspiracy-theories-debunked/index.html>.
- Pigden C. (1995). ‘Popper Revisited, or What Is Wrong with Conspiracy Theories?’ *Philosophy of the Social Sciences* 25(1), 3–34.

²³Thanks, Cara Cummings, Chris Arledge, and the anonymous reviewer.

- Pigden C.** (2022). 'Conspiracy Theories and the Conventional Wisdom Revisited.' In O. Loukola and L. Donskis (eds), *Secrets and Conspiracies*, pp. 126–56. Leiden: Brill.
- Powell J.** (2016). 'The Consensus on Anthropogenic Climate Change Matters.' *Bulletin of Science, Technology & Society* 36(3), 157–63.
- Powell J.** (n.d.). 'Consensus.' Jamespowell.org. <http://www.jamespowell.org/Consensus/>.
- Sunstein C. and Vermeule A.** (2009). 'Conspiracy Theories: Causes and Cures.' *Journal of Political Philosophy* 17(2), 202–27.
- Tooley M.** (2019). 'The Problem of Evil.' In E.N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2021 Edition). <https://plato.stanford.edu/archives/win2021/entries/evil/>.
- USC Shoah Foundation** (2021). 'VHA Subject Guide.' <https://libguides.usc.edu/vha/overview>.
- Williamson T.** (2007). *The Philosophy of Philosophy*. Malden, MA: Blackwell.
- Yad Vashem** (n.d.). YouTube. <https://www.youtube.com/user/yadvashem>.

Appendix. Additional problems for the statistical argument

There are two additional problems with the statistical argument that do not apply to the epistemic or alethic argument. First, the data used are not a good basis on which to estimate p , the probability of a conspirator keeping the secret per year. Second, the statistical analysis is problematic because the Poisson probability mass function doesn't apply to the exposure of conspiracies in the way the statistical argument requires.

Estimating p

When Grimes estimates p , he is not trying to infer from a representative sample; rather, he is trying to estimate a lower bound that is both charitable and empirically informed. His methodology is to calculate p for three cases (the NSA PRISM program, the Tuskegee syphilis experiment, and the FBI forensic scandal) and choose whichever value for p is most charitable to BS conspiracy theorists (Grimes 2016: 6). The lower the estimate of p , the more charitable. Based on his calculations, p is lowest (at approximately one in a quarter-million) in the case of the NSA PRISM program. Consequently, Grimes estimates that p is one in a quarter-million for any arbitrary conspirator (2016: 9–11).

However, there is little reason to think that the value of p for members of BS conspiracies would be no lower than the value of p for people involved in the NSA PRISM program. Grimes would have to make the following inference:

The value of p for people involved in the NSA PRISM program is approximately one in a quarter-million.

Therefore, the value of p for any member of any BS conspiracy would be no lower than approximately one in a quarter-million.

This argument is best construed as an a fortiori argument whose suppressed premise is that the NSA PRISM program was the pinnacle of discretion. More precisely, the suppressed premise is that:

If the value of p for people involved in the NSA PRISM program is one in a quarter-million, then the value of p would not be lower than one in a quarter-million for any other conspirator.

But this suppressed premise is dubious. We cannot rule out the possibility that BS conspiracies would involve harsher incentives to enforce silence or that they would be much more diligent about who to inform about the conspiracy. If so, we cannot rule out the possibility that members of BS conspiracies would have significantly lower values for p than one in a quarter-million. In which case, we are not justified in believing the suppressed premise.

So far, I have objected that Grimes extrapolates too much from a single example: the NSA PRISM program. Would it have helped if he had used a variety of examples to build up a representative sample? No. If Grimes had used this methodology, then he would be using a large sample of exposed conspiracies to estimate the value of p for members of BS conspiracies. The problem is that this would amount to using information about exposed conspiracies to draw conclusions about unexposed conspiracies. But it may be that

unexposed conspiracies have characteristics that enable them to remain unexposed, whereas exposed conspiracies lack these characteristics. Again, unexposed conspiracies might utilize harsher incentives than exposed conspiracies, or unexposed conspiracies might be more diligent about who to inform than exposed conspiracies.

If we are barred from relying on information about exposed conspiracies, is it even possible to know that we have a sample that is representative of BS conspiracies? A definitional issue with the term “BS conspiracy” is this: It’s possible to know that a BS conspiracy is part of one’s sample only if it’s possible for a BS conspiracy to be exposed. But, by definition, a BS conspiracy is an unexposed conspiracy.²⁴ Therefore, it’s impossible to know that a BS conspiracy is part of one’s sample. But then how can one reasonably expect one’s sample to be representative of BS conspiracies?

To get around this definitional problem, we can define a new term: *BS-adjacent conspiracy*, which is a big conspiracy that has been exposed and whose rate of exposure can reasonably be expected to approximate the rate of exposure of BS conspiracies. In this terminology, what I have argued so far is that there is no reason to think that the NSA’s conspiracy was a BS-adjacent conspiracy. This raises the question: Which conspiracies, if any, are BS-adjacent conspiracies? Currently, there are no known examples of a BS-adjacent conspiracy. For example, there is little reason to think that the rate of exposure of BS conspiracies is approximated by the conspiracies that Grimes discusses. Therefore, we currently have nothing we can use to estimate p for members of BS conspiracies.

How do we know that it’s reasonable to expect that the rate of exposure of a given conspiracy approximates that of BS conspiracies? The best methodology would be to wait for several views that are ordinarily labelled as “conspiracy theories” to be proven correct, confirm that they involve a large group of people, and use these as data points. But even then, we would need to wait for suitably many of these conspiracies to be exposed in order to be sure that we have a representative sample. Currently, we are in no position to carry out this task because we are still waiting for such conspiracies to be exposed. And anyone who is already convinced that BS conspiracy theories are all false should be equally convinced that we will never be able to carry out this task.

An alternative method for estimating p would be to answer the question: What is the probability that a highly discreet person whose incentives militate against revealing a secret of extreme importance would nevertheless reveal that secret (whether intentionally or unintentionally) in a way that would cause the secret to be widely known? The answer to this question would have to come from empirical evidence gathered by social scientists. However, it would be infeasible for social scientists to carry out the relevant kind of study. The first obstacle is that the participants in the study would need to be representative of the people who would be members of a BS conspiracy. Consequently, the study would not be able to use the usual sample of undergraduate students in introductory classes. Instead, the participants would need to be selected based on having a disposition for dishonesty, lacking a disposition for feeling guilt, etc., since these are characteristics that we should expect members of BS conspiracies to have. This would make it difficult to find participants for the study. The second obstacle is that participants would need to face harsh consequences for revealing some secret. However, it’s unclear how a study could replicate the life-altering consequences of leaking a BS conspiracy while satisfying the ethical constraints needed for IRB approval.

Thus, there is apparently no feasible way of estimating p with the information currently available or with information we are likely to have any time soon. In the meantime, we will be at a loss to assign a probability to p . Importantly, this is not the BS conspiracy theorist’s problem; rather, this is a problem for anyone who wants to run a statistical argument against BS conspiracy theories.

Applying the Poisson probability mass function

Another problem for the statistical argument is that it hinges on applying the Poisson probability mass function to draw conclusions about how long BS conspiracies would be likely to remain secret. However, the Poisson probability mass function cannot be applied in the way the argument requires. The Poisson probability mass function only applies to independent events (i.e., events that do not increase the probability of one another). Therefore, to apply the Poisson probability mass function to BS conspiracy

²⁴A conspiracy might be a BS conspiracy today. But if it gets exposed tomorrow, then, by definition, it won’t be a BS conspiracy anymore.

theories, Grimes must assume that a given conspirator's keeping the secret is an independent event. There are two reasons why this assumption is incorrect.

The first reason is that a conspirator's keeping the secret in a given year is not independent of their keeping the secret in previous years. The best indicator of a person's future behavior is their past behavior. Applied to conspiracies, this general point entails that the longer a conspirator keeps the secret, the more likely it is that they will continue keeping the secret in the years to come. If so, then a conspirator's keeping the secret is not independent from year to year. But if these events are not independent, then Grimes cannot use the Poisson probability mass function to calculate the number of years it would take for a BS conspiracy to be leaked.

The second reason is that a conspirator's keeping the secret in a given year is not independent of whether other conspirators keep the secret in that year. Given the hierarchical organization that BS conspiracies would likely have, some conspirators can threaten other conspirators to keep them quiet. This makes the conspirators who are being threatened more likely to keep the secret. Therefore, the fact that the conspirators who are in a position to make threats are keeping the secret makes it more likely that other conspirators will keep the secret. In other words, whether the other conspirators will keep the secret is not independent of whether the conspirators who can make threats keep the secret. Again, we reach the conclusion that the Poisson probability mass function cannot be applied to the present issue.

In defense of the claim that a conspirator's keeping the secret is an independent event, Grimes writes:

We initially assume that for a given conspiracy, conspirators are in general dedicated for the most part to the concealment of their activity. We further assume that a leak of information from any conspirator is sufficient to expose the conspiracy and render it redundant – such leaks might be intentional (in the form of whistle-blowing or defection) or accidental (mistaken release of information). 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. Thus, it follows that the act of a conspiracy being exposed is a relatively rare and independent event. (Grimes 2016: 3–4)

However, the premises of this argument do not support the independence of the event in question. The premises of Grimes' argument are the assumption of dedication, the assumption of sufficient exposure, and the assumption of intrinsicity, respectively. The assumption of dedication supports the conclusion that the target event is relatively rare, but it's irrelevant to independence. Meanwhile, the assumption of sufficient exposure is irrelevant to both rareness and independence, as is the assumption of intrinsicity.²⁵

Can the independence of the conspirators' keeping the secret be treated as an idealizing assumption? Sometimes, idealizing assumptions helpfully allow us to overlook minor complexities that don't affect the big picture. Other times, idealizing assumptions are unhelpful distortions that prevent us from making accurate predictions. The problem is that Grimes provides no reason to think that the relevant idealization would be a helpful idealization rather than an unhelpful one.

Ryan Ross is a PhD student at Johns Hopkins University. He is currently working on the epistemology of conspiracy theories. He has also published on permissivism. Additionally, he is interested in ancient philosophy.

²⁵Additionally, the assumption of sufficient exposure is false. If the leak is a drunken confession to a stranger in a bar, the confession is unlikely to be believed and, therefore, unlikely to expose the conspiracy. This person's testimony would not be credible without evidence (e.g., documents, emails, recordings, or witnesses).