

Word Count (including notes): 8,729

Samuel Director

Florida Atlantic University

sdirector@fau.edu

Abstract

One of the lessons of the COVID-19 pandemic is that the lay public relies immensely on the knowledge of public health officials. At every phase of the pandemic, the testimony of public health officials has been crucial for guiding public policy and individual behavior. The reason is simple: public health officials know a lot more than you and I do about public health. As lay people, we rely on experts. This seems straightforward. But, the current pandemic has shown that public health officials seem undecided as to what, precisely, their role is; are they providing the public information as it presents itself, or are they informing the public in a way that produces a desired or optimal outcome? In this paper, I answer the following question: what are public health officials morally obligated to tell the public? As I see it, these are the main options: public health officials should (1) tell the full truth, regardless of outcome, or they should (2) tell partial truths or lies that are aimed to promote a socially optimal outcome. My answer to this question is that public health officials are only allowed to lie under very narrow and rare conditions.

Public Health Officials Should *Almost* Always Tell the Truth

One of the lessons of the COVID-19 pandemic is that the lay public relies immensely on the knowledge of public health officials. At every phase of the pandemic, the testimony of public health officials has been crucial for guiding public policy and individual behavior. The reason is simple: public health officials know a lot more than you and I do about public health. As lay people, we rely on experts. This seems straightforward. But, the current pandemic has shown that public health officials seem undecided as to what, precisely, their role is; are they providing the public information as it presents itself, or are they informing the public in a way that produces a desired or optimal outcome? In this paper, I answer the following question: what are public health officials morally obligated to tell the public? As I see it, these are the main options: public health officials should (1) tell the full truth, regardless of outcome, or they should (2) tell partial truths or lies that are aimed to promote a socially optimal outcome. My answer to this question is that public health officials are only allowed to lie under very narrow and rare conditions.

In section I, I outline my general theory, which is that liberal states have a duty of honesty and transparency to their citizens. This duty is strong but defeasible within narrow specifications. I outline and defend a set of necessary conditions for when a liberal government may permissibly lie to its people to produce a desired outcome.¹ In section II, I argue that, if my theory is correct, then permissible public health lies will be very rare. I will illustrate this with examples from the history of public health communication (esp. the Philadelphia Liberty Loan Parade of 1918 and early communication about masks during the COVID-19 pandemic).

A few clarifications are in order. I won't be addressing what the content of public health communication should be. There are difficult questions about how to best communicate information to lay people, how complex or simple to make it, etc. I am only concerned with whether that information should be, to the best of officials' knowledge, true. Furthermore, I am not suggesting that public health officials have an easy task. With constantly changing crises, there are pragmatic and informational constraints on public health officials. But, within those constraints, the behavior of public health officials should be governed by principles of morality. It is those general principles that I want to discuss. Finally, I should note that although this article focuses on the United States, the analysis generalizes beyond the American context.

II. The liberal presumption against government lies

Liberal democracies owe their citizens transparency.ⁱⁱ Lying to the public is incompatible with this.ⁱⁱⁱ Thus, there appears to be a straightforward liberal argument against governments and government agents lying, in public health or elsewhere. I won't devote much time to explaining why liberal governments owe transparency to their citizens, as I take that to be uncontroversial.^{iv} Such a duty could be based on a plurality of grounds; for example, liberal legitimacy requires the consent of the governed, deception undermines consent, thus legitimate governments may not lie to the governed.

While liberal states have a duty of honesty, there are limits to this duty. If telling the truth would lead to nuclear war or millions of deaths, lying would be permissible. Even extreme deontologists agree with this.^v Given this, liberals should endorse the view that governments have a strong, but *defeasible*, duty to not lie to the public. Here, I offer a set of very narrow conditions which set the parameters for when governments or government agents may permissibly lie to the public.

To set these parameters, I start from our general, interpersonal duties regarding lying. I outline an account of our obligations regarding lying to our friends and then apply this to government agents. This move, as I will explain, sets my view apart from extant defenses of transparency. I later address an objection to my move from obligations between friends to obligations between the state and citizens. There are a variety of views about the ethics of lying. I will not settle that debate here. Instead, I endorse what I take to be the common-sense, middle ground view, namely that we have a *prima facie* duty not to lie, but this duty can be defeated.^{vi} This view avoids Kant's absolute prohibition on lying (regardless of how good the effects would be) and the act-utilitarian's absolute permission to lie whenever it achieves better effects than not doing so. My view most closely resembles W.D. Ross's position on lying.^{vii}

So, lies are sometimes permissible, but under what conditions? I offer a set of conditions that maps out the intuitive, middle ground conception of lying mentioned above. Lying to my friend seems to be justified *only if* doing so would (1) prevent a sufficiently bad harm, (2) lying is net beneficial, (3) I'm confident that lying to my friend would achieve the goal of avoiding the harm, and (4) lying is non-trivially more effective at preventing the harm than methods other than lying. This last condition requires clarification. Suppose that lying and some other method that doesn't involve lying (i.e. changing the subject) are equally effective at preventing the harm; the *prima facie* wrongness of lying tips the scale against choosing the lie in such a case. For that reason, lying is only permitted if it is more effective than the methods of avoiding the harm which don't involve lying (and which are not themselves wrong for some other reason).

On this test, white lies to one's friend would be allowed if they avoid sufficiently bad effects, and occasional lies about serious matters would also be justified. For example, suppose my friend makes me a meal that does not taste good and asks if I liked it. According to my view, if lying to him would avoid hurting his feelings in a non-trivial way, would produce net benefits, and is more effective than merely avoiding the question, then I am permitted to lie. However, if it doesn't matter all that much to my friend if I like his food, or if I can easily change the subject (without otherwise deceiving him), I am not permitted to lie.

So far, I take these conditions to be mapping a common-sense view (compatible with moderate deontology and rule-utilitarianism) of when lying to a friend is permissible. Furthermore, these conditions are meant to express a set of necessary conditions for lying, not necessary and sufficient conditions. In that regard, the claim is not overly ambitious. I take this view to be a plausible way of capturing the common-sense view that lying is *prima facie* wrong but permissible in some cases where duties of beneficence trump duties not to lie.

We can now apply these conditions to governments and government officials. There is no reason to think that, once someone dons the cap of being a government agent, they are released from their general obligations regarding lying.^{viii} In fact, government officials' obligations about lying may be even stronger than their interpersonal obligations about lying to our friends. Regarding the first two conditions (how serious the harm must be and how high the net benefits are), the standard should go up for government agents. If we are lying to a friend, the harm to be avoided by lying must be important but not catastrophic. With governments it seems to me that the harm avoided/benefit gained by lying must be incredibly serious if the lie is to be justified. Why raise the bar this way? Interpersonal relationships are very different from citizen to state relationships. There is far more power inequality between states and individuals, there is less trust, there is no special obligation of friendship, etc. For these reasons, I argue that government lies must meet a higher standard of harm prevention than interpersonal lies. Regarding condition 3, it seems that the epistemic bar for governments is higher than it is for friends. We generally have good reason to trust our friends and to think that they know us and our interests well enough to occasionally lie to us. This is not the case with governments. Even when trust in government is high, we clearly trust our friends more than government officials who we have never met, who do not know our interests and desires, etc. If someone is a stranger to me, they must meet a much higher bar for lying to me than my friend needs to. For this reason, if the state lying is to be justified, then the state must have a high level of confidence that its lie will achieve its goal. Additionally, for any state actions that have higher stakes (i.e. conviction for 1st degree murder), the epistemic bar goes up. The same would apply to lying in a high stakes case.^{ix} Lastly, in the context of interpersonal relationships, I argued (in condition 4) that the lie must be non-trivially more effective at avoiding the harm than other non-lying methods. In the context of governments, this bar also goes up. Given the power differentials, the fact that lying undermines government legitimacy, and the potential for establishing a dangerous precedent, for a government to be justified in lying to the people, the lie must be *substantially* more effective at avoiding the harm than non-lying methods. With this clarification, we can set out the necessary conditions under which governments may lie to their people:

Lies from the state to the public can be justified only if the state/official has a justified, high degree of confidence that:

1. the potential harm to the public is very high.
2. the net benefit of the lie is sufficiently high.
3. lying would avoid the harm.
4. lying is substantially more effective at preventing the harm than morally permissible methods other than lying

Some government lies pass this test. Suppose that covert defense operations are legitimate. Transparency about such operations would undermine their effectiveness. If a terrorist cell sees on C-SPAN that the CIA is coming after it tomorrow, this operation will fail. If the other conditions are also satisfied by a covert operation, then officials may permissibly lie about it. To borrow an example from Edyvane, it may be that politicians who don't believe nuclear retaliation is permissible are justified in publicly saying they would engage in nuclear retaliation, otherwise the deterrent effect of mutually assured destruction wouldn't obtain.^x As an example on the other end, when James Clapper lied about the NSA's surveillance program, he committed a grave wrong. Clapper (then Director of National Intelligence) was asked before congress (under oath) "does the NSA collect any data at all on millions or hundreds of millions of Americans?", to which he replied "no, sir...not wittingly."^{xi} This lie failed all four of the conditions, because the NSA's surveillance program did not prevent serious harms, it was not net beneficial, and so forth.

My argument is importantly different than extant defenses of transparency from government officials. For example, Kelsall's view is "grounded in neither general moral precepts nor instrumental reasons," but is

instead “grounded in considerations of the relationships of trust between laypersons and epistemic authorities.”^{xii} While my argument makes use of trust, it is not based on trust. My argument is based on our general obligations regarding lying, applied to the case of public health officials. Although we reach a similar conclusion, Kelsall’s view has an entirely different basis than mine. The same is true of Edyvane’s defense of transparency. Edyvane says, “I...briefly consider efforts to address this tension [between transparency and deception] by appeal to moral principle,” but “I shall not pursue that line of argument, partly because I do not think it especially promising.”^{xiii} Instead, he bases his argument on “what sort of moral character we should look for in our politicians.”^{xiv}

Before applying my theory, a few objections need to be addressed.

Some (i.e. Thompson) have argued for the view that governments may engage in first-order deception but only if they engage in second-order honesty. In other words, governments should be honest about the fact that they engage in deception sometimes, even if they don’t tell citizens what they are lied to about. Bok defends a related view, that government deception is legitimate if people consent to it in advance.^{xv} Although these views can be pulled apart, they make for a stronger view together, namely that the state can only engage in deception if it has the ex-ante consent of the public to engage in occasional beneficial deception; the only way to get that ex-ante consent would be to engage in second-order honesty about the first-order deception. While I feel the force of this view, I do not think that ex-ante consent is a necessary condition that should be added to my theory. While many people would agree with the higher-order claim that governments may lie under certain conditions and would consent to be governed under such a regime, many would not. So, if the consent of the governed is a necessary condition in my theory, it would immediately rule out any public health lies. The defender of this view could say that it’s not unanimous consent that is required, just some kind of majoritarian consent via a democratic procedure of some kind. While I’m very sympathetic to this addition, it would have the implication that a liberal state which has lost substantial trust with its citizens cannot tell a beneficial public health lie in an extreme emergency (as the people would not likely vote to give such a government higher-order permission to deceive them). While such states are blameworthy for losing the trust of their people, I don’t think this should stop them from acting to benefit their citizens in a dire emergency. Finally, ex-ante consent is not a standard we apply to the interpersonal ethics of lying. We generally think that we are allowed to lie to our friends to benefit them under some conditions, even if they would not consent ex-ante to being deceived by us.^{xvi}

One might object to condition 3 on the grounds that it doesn’t consider level of risk. For example, if lie A has a 10% chance of avoiding a mild negative effect, while lie B has a 10% chance of avoiding an extinction level, catastrophic event, it seems like B is permissible, while A is not. In short, the same degree of confidence may permit or rule out a lie, depending on the expected harm to be avoided, with the standard going down for catastrophic potential harms. I’m inclined to agree with this point.^{xvii} The reader can add a caveat to condition 3 stating that in potentially *catastrophic* events, the epistemic bar lowers.

One might object on the grounds that public health officials are not elected officials and thus are not bound in the same way to be honest. I think that public health officials in a liberal democracy are bound to the political norms of liberal democracy because they are (often) serving in roles as government officials, appointed by a democratically elected representative. Second, public health officials in general (in government, outside of government, or outside of liberal democracies) should be honest based on their pre-existing duties regarding honesty to other human beings. In that manner, my argument doesn’t have to rely on public health officials standing in a certain relationship with the government.^{xviii}

One could object that to develop an account of when the government may lie to citizens, it doesn’t make sense to begin with an account of when we may lie to our friends. Instead, we should start with our obligations regarding lying to strangers (which are likely not as strong) and then generalize from this to

the obligations of governments. While I agree that the government is not like a friend, the government does have an existing special relationship with citizens, unlike my non-existent relationship with strangers. The government has a duty of beneficence and transparency toward citizens, and so do friends to each other. I realize that these are not perfectly analogous relationships, but for the purposes of my argument, the citizen-state relationship is more analogous to the friend-friend relationship than it is to the stranger-stranger relationship.

An important objection/clarification needs to be considered here concerning the difference between lying and misleading. Owing to Jennifer Saul, there is a strain in the literature on the ethics of lying which argues that lying and misleading are not morally different. Briefly, misleading is defined as “a form of deception that occurs without lying.”^{xix} While I won’t restate Saul’s argument here, I think that it does not pose a threat to my view. If it turns out (as Saul and others have convincingly argued) that lying and misleading are morally on a par, I’d be happy to revise my view to cover cases of public health officials misleading the public as well. My view would then be restated as saying that state officials are allowed to lie or mislead only under the conditions outlined previously. All I would be saying is that the conditions under which the state may mislead citizens are the same as when the state may lie to citizens. Given that many accounts of the wrongness of lying are based on the fact that lying invites a false trust, if misleading also invites a false trust, then my argument (in the next section) that lying undermines trust in public health would fit neatly with misleading also being wrong for the same reason as lying.^{xx}

The worry for my view may enter in my discussion of whether the agent can avoid the question rather than lie. One might worry that this amounts to misleading. In many situations, I’m not convinced that avoiding a question is misleading. If I pretend not to hear my friend’s question, then I am misleading, because I intend for my friend to believe the false proposition “they didn’t hear me.” But, I can also imagine cases where avoiding a question need not amount to misleading. In the end, if avoiding the question ends up being a form of misleading, and if misleading is on a par with lying, then my view would also rule out avoiding the question, which would be consistent with condition 4 of my theory. Condition 4 would just have to be understood to say that misleading is not on the list of morally permissible methods other than lying.

If, on the other hand, Saul and others are wrong (meaning that lying and misleading are not on a par), I can accommodate this as well. If, as Strudler argues, lying invites more trust “than merely implicating something” invites, it would not be true that lying and misleading are equally wrong.^{xxi} Should this turn out to be true, I could change the theory back to the original formulation about lying. I do not intend to come off wishy washy here; I only mean to show that whichever answer turns out to be correct in the debate about whether lying and misleading are morally on a par, my view can accommodate it.^{xxii} Either version is still a philosophically interesting thesis about government transparency.

II. Why Permissible Public Health Lies are Rare

I take it that the theory defended in the previous section is plausible, as it is based in a common-sense view of the ethics of lying, and the move from interpersonal obligations to state obligations has been defended. At least for the purposes of this paper, I’ll treat my theory as sufficiently motivated for the next step. In this section, I argue that if the theory I’ve defended is correct, permissible public health lies will be *exceedingly rare*. I’ll go through each condition and argue that, in general, it is difficult for public health officials to satisfy it. Along the way, I’ll use examples from the history of public health testimony to illustrate.

Once again, my view is that lies from the state to the public can be justified only if the state/official has a justified, high degree of confidence that:

1. the potential harm to the public is very high.
2. the net benefit of the lie is sufficiently high.
3. lying would avoid the harm.
4. lying is substantially more effective at preventing the harm than morally permissible methods other than lying

Beginning with condition 1, are public health officials generally able to have a justified belief about the degree of the threat posed to the public by a health crisis? In general, it seems like the answer is ‘yes.’ Although there are certainly problems with methods like pandemic modelling, public health officials do seem equipped to assess the potential severity of a given threat.^{xxiii} So, this condition will often be satisfied. For example, in the case of COVID-19, although there was disagreement about the reliability of the models and the veracity of the fatality predictions, there was good reason to believe that the threat posed was substantial.

Conditions 2-3 is where the trouble begins. As I’ve argued, to justify lying to the public, the official in question must have sufficient justification to claim that the net benefit of telling a lie is sufficiently high and that telling the lie would avoid the harm in question. I’ll argue here that, in general, this is a very hard bar to meet. More precisely, it is unlikely that public health officials will have sufficient justification that telling a lie will be overall beneficial, because (1) there are easily predictable long-term harms of telling lies and (2) the likelihood of achieving the purported benefits of lying is reliably low. This leaves us with a presumption that, in general, lying in public health settings *probably* has a higher chance of causing harm than of achieving its desired benefits. I’ll show this by going through both (1) and (2) in the coming paragraphs.

There are easily foreseen long-term harms of lying in public health. Let’s stipulate for the sake of argument that there will be some likely short-term benefit from telling a lie. We can reasonably expect the aforementioned harms, so the potential benefits of the lie must overcome these. What are these long-term harms that accompany all public health lies? Lying in public health communication is likely to (A) lower public trust in public health officials, which makes compliance worse and (B) establish potentially dangerous legal and institutional precedents. Meanwhile, the benefits are generally uncertain. They require a kind of coordination that seems very unlikely, and they only make sense if we attribute a kind of expertise to public health experts that they don’t have. When we add this together, we see that public health lies generally have likely long-term harms without a strong likelihood of achieving their goals.

Here, I’ll go through each of likely harms and benefits, starting with trust.

In general, all lies in public health contexts have the potential to undermine public trust in public health. Once trust in public health is undermined, people comply at lower rates, and the health outcomes are worse.^{xxiv} The explanation is simple: if people feel lied to or feel like the guidance flip-flops for no reason, they begin to distrust the source of the information. For example, in 2021 officials in Colorado said that a highway was closed due to a flash flood. There was no flood; but officials wanted to reduce traffic on that road. If this becomes known, Coloradans will be less likely to follow future flood warnings, because they may see them as lies.^{xxv} Of course, in the short term, the public may be willing to follow flip-flopping guidance, but this can only last for so long. Eventually, if people feel that they are being manipulated, they begin to distrust the source of that information and stop following public health guidance altogether. As Peterson et al. note: “polarizing or sensationalized media coverage involving cases where public health officials have modified or reversed recommendations, such as evolving mammography guidelines, may lead to confusion and disintegration of trust.”^{xxvi} Ultimately, giving shifting advice lowers trust in public health officials, which lowers compliance, which raises infections, which leads to more deaths.^{xxvii} The claims of the previous paragraph have support from the literature on trust in public health. At a definitional level, dishonesty and trust work against each other. The empirical evidence bears this out as

well. As Peters et al. report in survey study, they found “strong support” for the hypothesis that “perceptions of trust and credibility are dependent on three factors: perceptions of knowledge and expertise; perceptions of openness and honesty; and perceptions of concern and care.”^{xxviii} In other words, perceptions of secrecy and dishonesty undermine trust. Once trust is undermined, public health goals are difficult to achieve. As Peterson et al. note:

In public health emergencies...individuals who report high trust in government health agencies respond more quickly and are more likely to comply with the health recommendations provided by the agencies...Trust in government health agencies is also associated with routine health outcomes, such as seasonal and pediatric vaccination uptake, medical adherence, and fewer emergency room visits...In all, the research suggests that maintaining public trust is essential for public health researchers, communicators, and practitioners to effectively disseminate health information and perform necessary public health functions in society.^{xxix}

One might object that there isn't a necessary connection between honesty and trust. My argument relies on the claim that public officials being honest fosters trust and that officials being dishonest leads to distrust and lower compliance with public health guidance. But, if this link doesn't obtain reliably, a key premise in my argument fails. Why think that trust and honesty need to go together? Consider relationships where we trust a person but do not always expect honesty from them in every situation. For example, during a difficult time in my life, if I ask a friend or my spouse, “will it get better?”, I don't always want the truth. I may want to be told a comforting lie, but this is compatible with broadly trusting my friend or spouse. Perhaps the same is true of public health officials; during crises, they may lie to us for a greater good, consistent with them being trustworthy. Indeed, a similar argument is advanced by John, who argues that maintaining trust requires lack of transparency.^{xxx}

While I agree that honesty does not *necessarily* lead to trust, there is compelling reason to doubt this objection. First, the kinds of relationships where we trust people or institutions but do not always expect them to tell us the truth differ importantly from public health officials. When I want my spouse to lie to me for my own good, it's because that lie will hopefully help me get through a difficult time and because the truth might be debilitating in that moment.^{xxxi} However, it would be strange if someone wanted a similar relationship with public officials during health emergencies. For example, I do not want my doctor to lie to me (unless it is for my extreme benefit), nor do I want public health officials to do so (unless it is extremely beneficial).^{xxxii} When it comes to having the information necessary to make safe decisions during an emergency, I suspect nobody would want anything but the truth. So, although there are trusting relationships in which honesty is not a requirement, this is not the kind of relationship we stand in with public health experts.

Furthermore, while dishonesty does not always lead to distrust, there is evidence to show that it does so more often than not, suggesting that honesty is a best practice for public health officials. In a recent book, Vallier synthesizes the literature on social and political trust. His findings support this view. As he says regarding lying in particular, “social trust...carries a generic expectation that members of society will comply with moral rules generally, especially moral rules essential to social order, like prohibitions on harm and *lying*.”^{xxxiii} Additionally, when officials are seen as deviating from norms, trust decreases: “in general, observed compliance with political norms by officials, especially in the civil service, will tend to increase political trust of some varieties, as will observed compliance with norms calling for citizen involvement.”^{xxxiv} Thus, while we can conceive of cases where government agents lying does not lead to distrust, this would buck the trend, showing that public health officials lying will reliably connect with distrust of public health institutions, and lower compliance with health guidance.^{xxxv}

An objector could argue that successful long-term deception actually increases trust in public health. Only if the deception is found out would there be any effect on trust. However, I am extremely doubtful that

long-term deception could be successful. If the public were to find out that they had been lied to on a long-term basis, the resulting loss of trust would be huge. Although not in the language of trust, John does in fact make the aforementioned objection:

You might worry that hiding the workings of a system is deceitful...However, the relevant wrong, while regrettable, may be mitigated by the fact that transparency would not lead to understanding, but to greater confusion. Imagine that a physician is an excellent diagnostician. Her patient believes that this is because she has a great intuitive genius. In fact, the physician is excellent at using complex disease databases. It might be regrettable, but still all-things-considered, justifiable for the physician to hide her actual practices from the patient. Of course, ‘playing along’ with false views of proper practice may compound those views, hence undermining the long-term goal of building a more realistic view of practice. While these are important concerns, it is unclear that the potential long-term benefits of transparency must outweigh the short-term costs.^{xxxvi}

I believe that the same response, about the fragility of lying, applies here as well.

Not only do all public health lies likely undermine trust, they also establish a potentially dangerous legal and institutional precedent that will cause long-term harms. Whenever we discuss general principles intended to guide public policy, we must be concerned about their institutional implications. How will typical people in positions of power use these principles? On average, will they be used for good or bad? Of course, general moral principles are such that we don’t test their success based on whether people will likely use them for good or bad. But, in public policy, we should only endorse principles that we would be comfortable with our ideological enemies using on us. While I have no doubt that many public health officials have the goal of saving lives when they lie, the same principle (lie to the public to achieve an optimal outcome) would permit bad actors to use deception for immoral ends. This is vividly illustrated by the Woodrow Wilson Administration’s handling of the Spanish Flu. Many government officials lied to the public during the 1918 pandemic. Historian John Barry documents how, in 1918, a parade (The Liberty Loan Parade) with hundreds of thousands of people was planned in Philadelphia. The goal of the parade was to raise money for the war effort. According to Barry, public health officials warned against holding this event, but the public health official in charge (Dr. Wilmer Krusen) declined and went forward with the parade.^{xxxvii} When Krusen was warned about the growth of influenza in the city, he and a colleague “feared that taking any such steps [such as quarantines and limiting public events] might cause panic and interfere with the war effort. Keeping the public calm was their goal.”^{xxxviii} Barry goes on: “none of the anxiety of the moment was reported in any of the city’s five daily papers, and if any reporter questioned...the wisdom of the parade’s proceeding, no mention of it appeared in print.”^{xxxix} Many people died from this parade spreading the flu.^{xl} This example represents the broader way in which the Wilson Administration controlled the narrative surrounding the Spanish Flu. As Barry puts it, “to keep morale up during the war, the government lied...national public health leaders said things like, ‘This is ordinary influenza by another name.’ They tried to minimize it.”^{xli} It seems wrong that the Wilson Administration did this, and it seems wrong that some public health officials went along with it, although many probably had little choice. The lesson for our purposes is that if public health officials are allowed to lie without strict conditions imposed, they may use those lies for ends that we disagree with. In general, lying is a precarious tool of public policy that I would only trust in the hands of noble and virtuous people. But, people are people, so we should not build our government as if only angels will run it. We need to be aware of the long-term harms that are likely introduced, via dangerous precedent, by lying in a way that does not meet very strict conditions. Of course, it may be that the benefit of lying in a certain case is so high that it justifies incurring the long-term potential harm of this dangerous precedent. I mean only to say that the dangerous precedent is a very likely long-term harm of any public health lie.

To recap, all public health lies have a chance of leading to the long-term harms of decreased trust in public health and dangerous precedent. But, my theory is open to the possibility that if the benefits are

good enough or the risk high enough, telling a lie could be justified even in the face of these predictable long-term harms. However, we should be generally doubtful that public health officials can be sufficiently justified in claiming that the good effects of a particular lie will outweigh the likely harms of that lie. This means that, on balance, we should not be confident in the efficacy of public health lies and should adopt the view that permissible lies are very rare.

We should be doubtful that public health officials can be sufficiently confident in claiming that the good effects of any one lie will outweigh the easily foreseeable harms of that lie, because of (1) coordination costs and (2) the limits of expertise.

First, in a global crisis, a policy of lying requires all (or most) government health agencies to state the same lie. We should be doubtful of policies whose success depends on multiple countries and multiple NGOs to all maintain the same lie. Such policies are unlikely to be successful, because they require a kind of coordination and consensus that is unlikely. Although on a smaller scale, domestic health lies would face a similar coordination challenge in any country with multiple health authorities at different levels of governance, as the lie would be less effective if these officials did not all repeat the same lie.

Second, experts are importantly limited in their ability to make accurate predictions about the likelihood of good outcomes that may come from telling public health lies. Take the example of early mask guidance during the COVID-19 pandemic in the spring of 2020. In short, there is evidence to suggest that members of the public health community knew that masks were effective at preventing the spread of COVID-19 but told the public not to buy them so that they could preserve the, then, scarce supply of masks for frontline workers. If this is true, this means that public health officials lied to the public when they told them both that they should not buy masks and that masks were not necessary or effective.^{xlii} Although I lack space to outline a full case for the claim that the public health community did, indeed, lie about this issue, I think there is compelling evidence for this. In case the reader does not agree with my assessment, we can take the example as a conditional, i.e. “if the public health community lied about masks, then...” To make the calculations necessary to determine if lying to the public about masks was a good idea, public health officials needed to justifiably believe a lot. They needed to justifiably believe that the supply of masks would be depleted if the public knew the truth, that there wouldn’t have been enough masks for hospital staff, and they needed to have a justified belief about the long-term costs and benefits of lying. This is a lot to predict. We should be skeptical of predictions that require the aforementioned information for several reasons. Public health officials are experts, but they are not experts in every domain necessary to make these calculations.^{xliii} To assess the tradeoffs and moral costs and benefits of telling a lie to the public, we would need to consult a diverse group of doctors, public policy experts, economists, and moral philosophers. To my knowledge, there was no such group that made the decision to lie about masks. Without those other voices, public health officials are not experts regarding the question at hand, because the question involves more than epidemiology.^{xliii} An objector could respond that if such a diverse group of deliberators were to do the math and think that lying was justified, wouldn’t it then overcome this objection? In principle, I agree. As I’ve stated, if a government lie could meet the conditions of my test, I believe it would be permissible. However, that is not what happened in the case of lies about masks. While I agree that a diverse group of experts would be better suited to assess if a lie is net beneficial, when we are discussing society level policies (especially during a crisis), there is good reason to be skeptical of how often such a group could reach the requisite level of confidence in their predictions. Society is a complex system; when we add the additional complexity of a pandemic with many unknown variables, this calls into question our ability to make predictions. At best, such a group of experts could make a semi-educated guess as to whether their lie would be worth it. Barring extreme threats, public officials must meet a higher bar than this to justify lying to the public. The general lesson I’m trying to illustrate is that even with a very simple and straightforward lie (“tell the public that masks don’t work, even though we know they do, so that we can preserve the supply of masks”) it is far from obvious that

the appropriate group of experts could have sufficient confidence that lying would achieve its harm avoidance goal.

As a further illustration, take the use of risk compensation methods in public health. The basic idea of risk compensation is that giving the public safety measures may have a negative effect, because it will make people less risk averse. Mantzari and Rubin describe risk compensation in this way:

The central idea is that people have a target level of risk they are comfortable with and they adjust their behaviour to maintain that risk level...It [risk compensation] famously has its origins in opposition to driver safety regulations, with an early proponent arguing that regulations were “at best useless and at worst counterproductive.” This was based on the assumption that people offset any gain in safety with an increase in risky behaviour.^{xlv}

This was used as a justification for not telling the public to wear masks. As Mantzari et al. report:

Concerns have been raised...that wearing face coverings might lead people to forgo other protective behaviours...Early in the pandemic, the World Health Organization warned that wearing medical masks when not indicated can “create a false sense of security that can lead to neglecting other essential measures such as hand hygiene practices.”^{xlvi}

I bring up risk compensation as an example of a tool that involves lies to the public with the goal of protecting them. Mantzari et al. state that “there is...no compelling evidence that such risk compensation exists at a population level.”^{xlvii} They discuss four examples where officials have used risk compensation (“wearing of bike or ski helmets (purportedly leading to riskier cycling or skiing), circumcision to prevent HIV infection, HIV pre-exposure prophylaxis, and HPV vaccination (all purportedly leading to increased unprotected sexual activity”), and they find that “the results of the most recent systematic reviews for each intervention provide no evidence for an increase in any outcome deemed to reflect risk compensation.”^{xlviii} I bring this up as an example of how intelligent, socially-concerned officials can get it wrong when they interact with an entire human population as if it were an equation with simple inputs and outputs. The lesson of risk compensation is that it’s difficult to use lies to make desirable policy outcomes.

Overall, any public health lie has easily foreseeable harms and tenuous, hard to predict benefits. While there will certainly be some cases where the balance tips in favor of telling a lie, even with these predictable harms, the default presumption seems to strongly fall in favor of public health lies not being justified.

Before wrapping up, I’ll address condition four of my theory, that lying is only justified if it is substantially more effective at preventing the harm than morally permissible methods other than lying. If this is a necessary condition for the permissibility of public health lies, this will make such lies even harder, in general, to justify. The reason is simple: there will almost always be methods other than lying which are likely just as effective at avoiding the harm in question. Take the aforementioned COVID example. Instead of lying, the government could have simply mandated production of additional masks, told the public the truth, and told the public to make their own masks with cloth. Early promotion of mask wearing would have increased demand for masks, which would have greatly incentivized private firms to make masks in bulk very quickly, solving the supply problem. Furthermore, the state could have simply mandated that masks only be sold to medical providers until the shortage was relieved. One might object it’s odd for me to say that the government using coercion is morally superior to it lying. I’m not convinced that this is always the case. If it’s the case that lying has long term harms, and if the coercion would be short-lived (which it would have been), then it seems that using coercion to promote public health is better than lying to do so. So long as the coercion from the state in the name of public health is

very circumscribed, I see no problem. For example, we could endorse Director and Frieman's recent proposal for the limitation of civil liberties in a pandemic, according to which the state may restrict a civil liberty in the name of public health if this restriction fits four conditions: (1) the coercion must not outlast the crisis, (2) the coercion must only be used to pursue public health goals, (3) the net welfare effects of the coercion must be positive, and (4) it must be the case that there are no equally successful, non-coercive measures which would solve the problem. To my mind, the state temporarily mandating that all sellers of masks cease selling them to non-medical workers, and then ending this order once the shortage has been solved, passes this test.

Conclusion:

Overall, I've defended a 4-part theory of when the state may lie in the pursuit of public health. I've also argued that 3 of those 4 conditions are, generally, very unlikely to be met by public health officials. This leaves us with the view that government lies in the name of public health are permissible but extremely rare. In other words, public health officials should *almost* always tell the truth.

-
- ⁱ I view an expert as "someone who is recognized as having a special competence to draw inferences from evidence within a certain domain. That competence typically derives from access to a large body of evidence and from socialisation into specialised ways of perceiving and reasoning about evidence of that kind" (Ward 263). I won't discuss when expert testimony produces justified beliefs. My concern is just with the intent to deceive by experts. For a discussion of these other questions, see Ward (2016). My view of transparency in public health is in concert with, although grounded in a different approach, Grill and Hansson (2005).
- ⁱⁱ See Kogelmann (2021) for an historical overview of political philosophers' views on transparency. Kogelmann traces the liberal importance placed on transparency back to Bentham and Mill, reaching a contemporary articulation in Williams and Waldron. Philosophers give various rationales for government transparency. Kogelmann (2022, chapter 5) presents a helpful summary. For Rawls, transparency is a requirement of publicity. Williams defends several arguments for transparency: transparency is necessary to prevent tyranny, the nature of democratic representation requires transparency, and that legitimacy requires truthfulness. Waldron (2016) makes an accountability argument for transparency (see Kogelmann, 2022, 130).
- ⁱⁱⁱ I will be assuming the standard view of the nature of lying. As Fallis says, "the 'standard philosophical definition of lying' is the following: 'you lie to X if and only if: you say that P to X. You believe that p is false. By saying that p, you intend to deceive X'" (1-2). This definition is not without controversy. See Carson (2006) for a discussion of this debate. I do, however, think that my view does not rely on this definition of lying as a set of necessary and sufficient conditions. Even if the definition I mention is not a successful analysis of lying, we can agree that it at least specifies sufficient conditions for lying. That's all I need.
- ^{iv} Some philosophers reject transparency as a requirement of government and/or science communication. As John says, "Given an independently plausible account of how non-experts learn from experts, transparency, openness, sincerity and honesty are not basic norms for ethical scientific communication" (75). I admit that, for the purpose of argument, I am assuming that citizens generally want their governments to trust them and that they generally would only trust noble people to lie to them. We especially trust public health experts, in their capacity as experts, to be honest.
- ^v See Nozick (*Anarchy*, 30), Hooker (*Ideal*, 98-99, Hooker is not a deontologist), and surprisingly Rawls (*A Theory*, 40).
- ^{vi} The goal of this paper to apply a plausible view of the ethics of lying to public officials. As such, I do not spend time arguing that my account of the ethics of lying is superior to all existing theories, only that my account is reasonable and common-sensical. This fits with Carson's recent argument that the "range of reasonable views [about lying] excludes all views that permit lying in cases in which both act-utilitarianism and Ross's theory prohibit it" (145). My view counts a reasonable on Carson's account.
- ^{vii} For Ross's view of lying, see *The Right and the Good* chapter 2. In brief, Ross's view is that we have a prima facie duty to not lie, which may be defeated by our other prima facie duties if the consequences are sufficiently high. Whether this obtains in a given situation will depend on which of our prima facie duties becomes our all-things-considered duty in that situation. As Carson notes, Ross does not spell out a precise threshold for this in the case of lying (Carson, 2005, 144).

-
- viii See Brennan (2019) for a defense of the view that state actors are not free from normal moral obligations. There are exceptions to this rule. The government may be permitted to imprison me for breaking the law, while my friend cannot. This is because we have delegated this authority to the government to avoid the inefficiency of private enforcement of law. When it comes to lying, I don't see any power citizens have delegated to government agents which would free them from their general duties. In fact, people may care much more about honesty from their government than from their friends.
- ix I borrow this point from Winsberg, Brennan, and Surprenant (2020).
- x Edyvane, "The Ethics," 313.
- xi Wyden (2013).
- xii Kelsall, "The Trust," 288.
- xiii Edyvane "The Ethics," 310-311.
- xiv Edyvane, "The Ethics," 310-311.
- xv Bok 177.
- xvi As an anonymous referee suggested to me, this meta-consent approach is vulnerable to the arguments regarding hypocrisy in Tillyris (2016). The meta-consent approach requires that citizens know, ex-post, that they were deceived or know, in general, that deception may occur in certain circumstances. Tillyris argues that hypocrisy is a necessary vice for democratic politicians, but hypocrisy seems to require that politicians not disclose their deceptions ex-post. To address this defense of hypocrisy would be beyond the scope of this paper.
- xvii The intuition in this example relies on the idea of pragmatic encroachment on epistemic justification, the phenomenon when differing practical stakes alter the epistemic standard of justification. See Fantl and McGrath (2007).
- xviii An anonymous referee made me aware of this point.
- xix Strudler, "The Distinctive," 173. I am grateful to an associated editor of this journal for raising this point.
- xx For example, Faulkner accounts for the wrongness of lying in this way. Regarding lying and misleading both inviting a false trust, Strudler says, "when a person makes a misleading statement, through conversational implicature, he invites a trust that he intends to breach, just as a person invites a trust that he intend to breach when he lies. Lying and misleading seem equally to aim at breaching trust" (174). To be clear, Strudler does not agree with this argument.
- xxi Ibid, 175.
- xxii As an anonymous referee pointed out to me, my view possibly has the implication that other deceptive speech practices (bullshitting, spin, pandering, etc.) would be ruled out as well. Although that is beyond the scope of this paper, I am in principle fine with that implication. For an extended discussion of spin, see Manson (2012).
- xxiii For a discussion of pandemic modelling, see Northcott (2022).
- xxiv As Wynia (2007) argues, "lack of trust will breed both fear and non-adherence" (6).
- xxv Lopez, "Did CDOT."
- xxvi Peterson et al., "Trust in," 979.
- xxvii As Wynia (2006) cautions, "effective risk communication depends heavily on the trustworthiness of the information source" (3).
- xxviii Peters et al., "The determinants," 44 and 53. This study was in the context of environmental risk communication.
- xxix Peterson et al., "Trust in," 978.
- xxx John, "Epistemic Trust," 81-82.
- xxxi As a sad example of this, trust in government can be very high under authoritarian regimes (Vallier, 75, 2020).
- xxxii As an anonymous referee pointed out to me, it may be much easier for doctors to pass this extreme benefit threshold, given that the stakes are often high and given that there is less scrutiny on them. I'm not convinced that this is true; however, even if it were true, what matters for my purposes is that the principle (health workers should only lie to us if it's to our extreme benefit) is still true, even if it's easier for doctors to meet the extreme benefit threshold than it is for public health officials.
- xxxiii Vallier, *Trust*, 49-50. Emphasis added.
- xxxiv Ibid, 75. One might object that public health officials may be governed by a different set of norms that elected officials. See my earlier response to this claim in section I.
- xxxv An anonymous referee pointed out to me that if lies undermine trust, then (in the pursuit of trust) public health officials should lie about the fact that they sometimes lie. I share this concern and think that some kind of legal requirement for ex-post honesty (i.e. telling the public they were lied to once the emergency is over) could address this.
- xxxvi John, "Epistemic Trust," 82.

-
- xxxvii Barry, *The Great*, 208. The official in question was the director of the Philadelphia Department of Public Health and Charities, Dr. Wilmer Krusen. Although he was a doctor and served in an official public health role, Barry notes that “he lacked background in, commitment to, or understanding of public health issues” (2004, 199). But, his failure to cancel the parade was not born of ignorance; Barry notes that Krusen was aware of reports that influenza “was already roaring full speed through the Navy Yard [in Philadelphia]” (2004, 208).
- xxxviii Ibid, 203.
- xxxix Ibid, 209.
- xl Ibid, 220.
- xli Tapper, “What the author,” 2020.
- xlii Kogelmann briefly addresses this example (*Secret Government*, 114-115). I am not attributing malice to public health officials. They likely acted out of genuine concern for public health.
- xliii Ballantyne has helpfully labelled this as “epistemic trespassing,” which occurs when a person who has “competence or expertise to make judgments in one field” moves “to another field where they lack competence and pass judgment nevertheless” (367).
- xliv Kelsall makes a similar point (“The Trust,” 299).
- xlvi Mantzari et al., “Is Risk,” 1.
- xlvii Ibid.
- xlviii Ibid, 2.
- xlvi Ibid, 1.

Acknowledgements: Thanks are due to David Boonin, Wyatt McDaniel, Chris Heathwood, Jason Brennan, Chris Freiman, Greg Robson, Iskra Fileva, Lawrence Torcello, Takunda Matose, audience members at the 2023 Eastern APA, audience members at the 2021 RoME. Thanks are due to Emily Director for her continued support.

References

- Ballantyne, Nathan. “Epistemic Trespassing.” *Mind* 128 (2019): 367-396.
- Barry, John. *The Great Influenza: The Epic Story of the Deadliest Plague In History*. New York. Penguin, 2004.
- Bok, Sissela. *Lying: Moral Choice in Public and Private Life*. Quartet, 1980.
- Brennan, Jason. *When All Else Fails: The Ethics of Resistance to State Injustice*. Princeton. Princeton University Press, 2019.
- Carson, Thomas. “The Range of Reasonable Views about the Morality of Lying.” In *Lying: Language, Knowledge, Ethics, and Politics*, edited by Andreas Stokke and Eliot Michaelson, 145-160. Oxford University Press, 2018.
- Carson, Thomas. “The Definition of Lying.” *Nous* 40, no. 2 (2006): 284-306.
- Carson, Thomas. “Ross and Utilitarianism On Promise Keeping and Lying: Self-Evidence and the Data of Ethics.” *Philosophical Issues* 15 (2005): 140-157.
- Clark, Cory, et al. “Predictors of COVID-19 voluntary compliance behaviors: An international investigation.” *Global Transitions* 2 (2020): 76-82.
- Director, Samuel and Christopher Freiman. “Civil Liberties in a Lockdown: The Case of COVID-19.” Forthcoming *Journal of Medicine & Philosophy*. 1-24.
- Edyvane, Derek. “The Ethics of Democratic Deceit.” *Journal of Applied Philosophy* (2014): 310-325.
- Fallis, Don. “Lying and Deception.” *Philosophers’ Imprint* 10, no. 11 (2010): 1-22.
- Fantl, Jeremy and Matthew McGrath. “On Pragmatic Encroachment in Epistemology.” *Philosophy and Phenomenological Research* 75, no. 3 (2007): 558-589.
- Faulkner, Paul. “What Is Wrong with Lying?” *Philosophy and Phenomenological Research* 75, no. 3 (2007): 535-557.
- Grill, Kall, and Sven Ove Hansson. “Epistemic Paternalism In Public Health.” *Journal of Medical Ethics* 31 (2005): 648-653.
- Hooker, Brad. *Ideal Code, Real World*. Oxford: Oxford University Press, 2000.
- John, Stephen. “Epistemic trust and the ethics of science communication: against transparency, openness, sincerity, and honesty.” *Social Epistemology* (2018): 75-87.

-
- Kelsall, Joshua. "The Trust-Based Communicative Obligations of Expert Authorities." *Journal of Applied Philosophy* 38, no. 2 (2021): 288-305.
- Kogelmann, Brian. "Secrecy and transparency in political philosophy." *Philosophy Compass* (2021): 1-10.
- Kogelmann, Brian. *Secret Government: The Pathologies of Publicity*. Cambridge: Cambridge University Press, 2022.
- Lopez, Meghan. "Did CDOT lie? Internal CDOT emails show why employees decided to misinform drivers." *Denver 7 News*, Aug 12, 2021. <https://www.thedenverchannel.com/news/local-news/did-cdot-lie-internal-cdot-emails-show-why-employees-decided-to-misinform-drivers>
- Mahon, James Edwin. "The Definition of Lying and Deception." *The Stanford Encyclopedia of Philosophy* (2016): <https://plato.stanford.edu/entries/lying-definition/#TraDefLyI>
- Manson, Neil C. "Making Sense of Spin." *Journal of Applied Philosophy* 29, no. 3 (2012): 200-213.
- Mantzari, Eleni, G James Rubin, and Theresa Marteau. "Is risk compensation threatening public health in the covid-19 pandemic?" *British Medical Journal* 370, no. 2913 (2020): 1-2.
- Northcott, Robert. "Pandemic Modeling, Good and Bad." *Philosophy of Medicine* 3, no. 1 (2022): 1-20.
- Nozick, Robert. *Anarchy, State, and Utopia*. New York: Basic Books, 1974.
- Peterson, Emily, et al. "Trust in national health information sources in the United States: comparing predictors and levels of trust across three health domains." *Transitional Behavioral Medicine* 10, no. 4 (2020): 978-988.
- Peters, R. G., Covello, V. T., & McCallum, D. B. (1997). The determinants of trust and credibility in environmental risk communication: An empirical study. *Risk Analysis*, 17(1), 43-54.
- Rawls, John. *A Theory of Justice*. Cambridge, MA: Harvard University Press, 1999.
- Ross, William David. *The Right and the Good*, Edited by Philip Stratton-Lake. Oxford: Oxford University Press, 2002.
- Saul, Jennifer. "Just go ahead and lie." *Analysis* 72, no. 1 (2012): 3-9.
- Strudler, Alan. "The Distinctive Wrong in Lying." *Ethical Theory and Moral Practice* 13, no. 2 (2010): 171-179.
- Tapper, Jake. "What the author of 'The Great Influenza: the Story of the Deadliest Pandemic in History' can teach us about the coronavirus." *CNN*, March 27, 2020. <https://www.cnn.com/2020/03/27/politics/interview-john-barry-great-influenza/index.html>
- Tillyris, Demetris. "The virtue of vice: a defence of hypocrisy in democratic politics." *Contemporary Politics* 22, no. 1 (2016): 1-19.
- Thompson, Dennis. "Democratic secrecy." *Political Science Quarterly* 114 (1999): 181-193.
- Tufekci, Zeynep. "5 Pandemic Mistakes We Keep Repeating." *The Atlantic*, Feb 26, 2021. <https://www.theatlantic.com/ideas/archive/2021/02/how-public-health-messaging-backfired/618147/>
- Vallier, Kevin. *Trust In A Polarized Age*. Oxford: Oxford University Press, 2021.
- Waldron, Jeremy. *Political Political Theory: Essays on Institutions*. Cambridge: Harvard, 2016.
- Ward, Tony. "Expert Testimony, Law and Epistemic Authority." *Journal of Applied Philosophy* (2016): 263-277.
- Williams, Bernard. "Truth, Politics, and Self-Deception." *Social Research* 63 (1996): 603-617.
- Williams, Bernard. *Truth and Truthfulness: An Essay in Genealogy*. Princeton: Princeton University Press, 2002.
- Winsberg, Eric, Jason Brennan, and Chris Surprenant. "How Government Leaders Violated Their Epistemic Duties During the SARS-CoV-2 Crisis." *Kennedy Institute of Ethics Journal* 30, no. 3-4 (2020): 215-242.
- Wyden, Ron. "DNI Clapper tells Wyden the NSA does not collect data on millions of Americans." *Youtube*, Mar 12, 2013. <https://www.youtube.com/watch?v=QwiUVUJmGjs>
- Wynia, Matthew. "Public Health, Public Trust, and Lobbying." *The American Journal of Bioethics* 7, no. 6 (2007): 4-7.

Wynia, Matthew, and American Medical Association. "Risk and Trust in Public Health: A Cautionary Tale." *The American Journal of Bioethics* 6, no. 2 (2006): 3-6.