

A modern Pascal's wager for mass electronic surveillance

David Danks

Departments of Philosophy & Psychology, Carnegie Mellon University

Abstract

Debates about the moral permissibility of mass electronic surveillance often turn on whether consequentialist considerations legitimately trump relevant deontological rights and principles. In order to establish such overriding consequences, many proponents of mass surveillance employ a modern analogue of Pascal's wager: they contend that the (negative) consequences of no surveillance are so severe that any probability of such outcomes legitimates the abrogation of the relevant rights. In this paper, I briefly review Pascal's original wager about whether to live a pious life, including two classes of objections that were almost immediately leveled against his argument. I then show that analogues of those objections apply straightforwardly to the modern versions of Pascal's wager. Mass electronic surveillance might be ethically permissible or even obligatory in some circumstances, but the details matter in ways that are systematically ignored or assumed away in a Pascal's wager-type argument. Careful consideration of the complexities of our modern situation is required to decide whether consequentialist considerations override deontological principles in this domain.

A modern Pascal's wager for mass electronic surveillance

Recent disclosures and leaks have revealed that the public, both in the United States and internationally, are subject to some degree of mass electronic surveillance.¹ Much of the resulting debate has focused on exactly who is performing the surveillance (private companies, government agencies, or a partnership between the two), what information is being recorded (message metadata, message content, actual sound or image files, or more), and numerous legal intricacies involving these activities (e.g., exactly what the National Security Agency is legally permitted to collect or analyze). In this paper, I simply take for granted that *some* type of large-scale data collection is occurring, and focus instead on the ethical justification of these activities, without worrying about legal details of exactly who can collect or analyze what communications data.

Public debates about the ethical or moral legitimacy of mass surveillance are typically framed in the language of deontology versus consequentialism. Opponents of such surveillance largely talk in terms of rights that are relatively inviolable, such as the right to privacy of communications, the right to autonomy and self-determination, or the right to be a member of a well-functioning civil society.² Proponents of such surveillance instead focus on talk of consequences, such as the harms that might befall society if nefarious plans were consummated, the relative lack of harms due to mass surveillance on innocent citizens, or the deterrence benefits yielded by public knowledge of the

¹ E.g., Leslie Cauley, "NSA Has Massive Database of Americans' Phone Calls," *USA Today*, May 11, 2006; Barton Gellman and Ashkan Soltani, "NSA Collects Millions of E-Mail Address Books Globally," *The Washington Post*, October 14, 2013.

² Marie-Helen Maras, "The Social Consequences of a Mass Surveillance Measure: What Happens When We Become the 'Others'?", *International Journal of Law, Crime and Justice* 40 (2012): 65–81; Alan Rubel, "The Particularized Judgment Account of Privacy," *Res Publica* 17, no. 3 (August 2011): 275–90; David Wright et al., "Sorting Out Smart Surveillance," *Computer Law & Security Review* 26, no. 4 (July 2010): 343–54.

mass electronic surveillance apparatus.³ This deep mismatch between the rhetoric of the two sides—one using deontological language, the other consequentialist—translates directly into a theoretical challenge in resolving the disagreement. If both sides agreed, for example, that the ethical justification of mass surveillance depended on only its consequences, then we could “just” calculate the different costs and benefits (though that would presumably be exceptionally difficult). Instead, we have two qualitatively different, potentially incommensurable types of criteria that we must somehow reconcile in order to advance the debate beyond mere shouting.

There is relatively widespread agreement that most deontological principles can ultimately be overridden by consequentialist concerns, if the consequences are sufficiently severe. The moral legitimacy of the proverbial choice to “kill one to save a billion” reveals that the maxim not to kill an innocent is not inviolable. Of course, consequentialist considerations do not necessarily make it right to violate the deontological principle, but they may make it permissible (e.g., if we face a moral dilemma in which there is no right choice). Thus, one way to dissolve the impasse over mass surveillance is for the consequentialist to argue that not engaging in large-scale information collection would result in such negative consequences that it is permissible, perhaps even obligatory, to violate the relevant rights and principles. Again, such an argument does not show that it is morally right to engage in mass surveillance, but the consequentialist need only contend that we are trapped in a moral dilemma for which *no* choice will be fully morally acceptable. It is difficult, however, for the proponent to give such an argument, as many of the relevant negative consequences—both their disutility and their probabilities—have thankfully not been observed. One can note activities that have been stopped, such as various terrorist plots, but this information is rarely sufficient to know

³ Jack Goldsmith, “We Need an Invasive NSA,” *New Republic*, October 10, 2013; Katerina Hadjimatheou, “The Relative Moral Risks of Untargeted and Targeted Surveillance,” *Ethical Theory and Moral Practice*, August 2013; Michelle Van Cleave, “What It Takes: in Defense of the NSA,” *World Affairs*, November/December 2013; White House Office of the Press Secretary, “Remarks by the President on Review of Signals Intelligence,” January 17, 2014.

what would have probabilistically happened if the activities had not been blocked. Instead, the dominant consequentialist move has been to appeal to an argument that I suggest is structurally similar to Pascal’s famous wager about the existence of God.

Blaise Pascal was a 17th century French mathematician who made significant contributions to physics, geometry, early notions of computation, and probability theory. He also experienced an intense religious experience at age 31, after which he devoted much of his remaining eight years to writings on philosophy and theology, particularly the nature and justification of belief in God. Pascal believed that one could not simply choose to believe in God; such a belief is not, he thought, under direct, voluntary control. One could, however, choose to live a pious life by abstaining from sin, attending church, and so forth, and Pascal believed that such behaviors would lead (with high probability) to a sincere belief in God’s existence. The question thus naturally arises: ought a non-believer choose to pursue a pious life? Pascal argued that the non-believer should do so, using essentially decision-theoretic considerations. At a high level, there were (for Pascal) two possible worlds—God exists or God does not exist—and two possible actions by the decision-maker—lead a pious life, or lead a sinful one. He then argued that the outcomes for the four different world-action possibilities had a specific structure, which we can represent as a payoff or utility matrix:

	God exists	God does not exist
Pious life	$+\infty - C = +\infty$	$L - C$
Sinful life	$-\infty$	L

In this table, L represents the (finite) utility that one experiences over the course of a sinful life if God does not exist, and C represents the cost of foregoing certain earthly pleasures while living a pious life. The infinite utilities in the “God exists” world arise because one will, in that world, either go to heaven (eternal bliss) or hell (eternal torment) based on whether one comes to believe in God.

Simple decision-theoretic calculations show that, for this payoff matrix, one ought to lead a pious life if there is any positive probability of God’s existence, no matter how small. Arguably, even the most hard-line agnostic would grant, on grounds of epistemic humility, that there is *some* non-zero probability that God exists. Such a concession is, for Pascal, sufficient to complete the argument that one should lead a pious life in order to arrive at a sincere belief in God’s existence.

Present-day consequentialist arguments for mass surveillance have a remarkably similar structure to Pascal’s wager, but with changes to the row and column headings for the payoff matrix:

	Terrorists plan an attack	Terrorists do not plan an attack
Mass surveillance	$+\infty - C = +\infty$	$L - C$
No mass surveillance	$-\infty$	L

In this table, L represents the utility of “ordinary” life and C denotes the costs due to abrogation of people’s rights. This table is obviously an exaggeration, as no one contends that a terrorist attack would truly yield infinite disutility or that stopping such an attack would produce infinite positive utility. Nonetheless, the rhetoric of many consequentialist arguments suggests something like this matrix; we must, for example, do “whatever is necessary” to avoid another terrorist attack.

Moreover, if something like this payoff matrix is correct (perhaps substituting very large numbers for the infinities), then we have a straightforward justification for overriding the deontological principles: as long as the “Terrorists plan an attack” possible world has any non-zero probability, then the expected value of “Mass surveillance” will be so much larger than “No mass surveillance” that we can legitimately abrogate the relevant rights. We do not need to determine the exact probabilities and utilities; the numbers are sufficiently large that we exceed the threshold (whatever it might be) for overriding the rights.

A successful Pascal's wager-type argument would yield the conclusion that the consequentialist requires, but instances of this argument-schema have been subject to serious objections ever since Pascal's original version. I focus here on just two. First, although this type of argument permits some ambiguity in the utility and probability numbers, they still must have the correct, relative orders of magnitude. If the "positive infinity" payoff is, for example, only twice as large as L (and C is a significant fraction of L), then the action in the top row—Pious life or Mass surveillance—will not necessarily be optimal for arbitrarily small probabilities. Pascal's original wager assumed that all pious lives were the same, whether the reason for the piety was religious experience (as for Pascal) or decision-theoretic calculation (as for the audience of his argument). Objectors to Pascal's wager contended instead that the latter motivation would lead to an inauthentic life that did not have the same utilities as the "truly pious" life, both lower gains (if God exists) and higher costs (if God does not exist). In particular, if the positive gains are less-than-infinite (e.g., if God does not completely reward people who believe for the "wrong" reasons), then the payoff utilities matter. In extreme cases, one can even use a Pascal's wager-type argument to conclude that one ought rationally to pursue a *sinful* life!

Essentially the same type of objection can be leveled against the modern Pascal's wagers for mass surveillance. The positive gains from thwarting a terrorist attack are obviously not infinite; the losses from a successful terrorist attack are similarly not negative infinity. The actual numbers lie somewhere in the middle. Moreover, it is quite difficult to assess the costs C , as both the actual impacts and opportunity costs of mass surveillance programs are hard to quantify. Resources—whether governmental or corporate, financial or temporal—spent on mass surveillance are resources that cannot be used for other important projects and endeavors. Without some consideration of those alternate uses, we have little understanding of the "true" utilities in our payoff matrix. Having said that, it is important to be clear that this objection does not show that no Pascal's wager-type

argument is available to the proponent of mass surveillance. It may well turn out that the numbers work out in its favor. However, that conclusion should be reached through careful consideration of the costs and benefits, rather than through rhetoric that implies infinite utility and disutility in one state of the world.

Second, and more importantly, Pascal's wager-style arguments work only if the choice or action—leading a pious life, engaging in mass surveillance, etc.—causally influences the probabilities in the right way. Pascal believed in a benevolent God who rewarded all and only believers; hence, living piously in order to produce belief was the one and only way to reach heaven (and infinite positive utility). As was quickly pointed out, however, this assumption is a substantive one. Suppose instead that God is malevolent (or just a trickster) and so rewards only non-believers or doubters. In that case, heaven (and its infinite positive reward) would await those who live a *sinful* life. If we grant any non-zero probability to the possibility of a malevolent or trickster God (which epistemic humility presumably requires), then Pascal's wager falls apart: the overall expected values of a pious versus sinful life will depend on the specific probability and utility numbers. More generally, Pascal made assumptions about the relationship between one's actions and various outcomes, and those substantive assumptions could well turn out to be wrong. In particular, the table implicitly encodes a rich set of beliefs about the causal impact of one's lifestyle. In fact, it is a substantive assumption to hold that one's beliefs about God matter at all in terms of whether one reaches heaven or hell (if there is a God). A devout Calvinist, for example, would agree that there are infinite positive and negative utilities in the "God exists" column of the payoff matrix, but would contend that one's actions and beliefs are irrelevant to which one receives.

A structurally identical set of worries arises for arguments in favor of mass surveillance. The proposed payoff matrix depends on a large set of assumptions about the causal impact of a mass surveillance program, many of which are even more dubious than the corresponding causal

assumptions in the original Pascal's wager. For example, the surveillance payoff matrix assumes that the creation (and subsequent involuntary disclosure) of mass surveillance programs will reduce the probability of a successful terrorist attack, not increase it. This assumption is possibly correct; the proponents of such programs certainly believe in their efficacy. But it seems similarly possible that the existence of such programs could prompt resentment, fear, or impulsive actions that increase the probability of a terrorist attack. Even if we discount the impact of such programs on others, mass surveillance programs could be causally inefficacious relative to more targeted surveillance programs. In the medical domain, it is widely understood and accepted that population-wide testing can yield a huge number of false positives, and the costs (social and personal) of such false positives can outweigh the benefits of population-wide testing.⁴ The same observation holds for large-scale, population-wide surveillance. In brief, mass surveillance can generate so many false leads that the analysis infrastructure—both computational and personnel—can be overwhelmed. In certain conditions, targeted surveillance can be more efficacious; paradoxically, one can often be more successful precisely by recognizing that one cannot succeed everywhere, and thereby being freed to allocate resources in a more targeted, more optimal manner.⁵ More generally, our actions regarding mass surveillance can have, in totality, a very different causal impact than is assumed in the proponents' payoff matrix, and so the Pascal's wager-type argument dissolves.

As with the first set of concerns, these objections do *not* show that mass surveillance programs are ethically impermissible. Rather, the overall argument in this paper is that the details matter in this

⁴ Richard Hayes, Kalpana Sabapathy, and Sarah Fidler, "Universal Testing and Treatment as an HIV Prevention Strategy: Research Questions and Methods," *Current HIV Research* 9 (2011): 429–45; Bradley G Wagner, James S Kahn, and Sally Blower, "Should We Try to Eliminate HIV Epidemics by Using a 'Test and Treat' Strategy?," *Aids* 24, no. 5 (March 13, 2010): 775–76.

⁵ See also: Marie-Helen Maras, "The Economic Costs and Consequences of Mass Communications Data Retention: Is the Data Retention Directive a Proportionate Measure?," *European Journal of Law and Economics* 33 (2012): 447–72; National Research Council, *Protecting Individual Privacy in the Struggle Against Terrorists: A Framework for Program Assessment*. Washington, DC: The National Academies Press, 2008.

debate. Perhaps mass surveillance programs are causally efficacious in the right way, unlike population-wide disease testing. Perhaps such programs do not trigger exactly the type of behavior that they are supposed to prevent. Perhaps the payoffs exhibit the appropriate relative orders of magnitude to support a Pascal's wager-type argument. All of these conditions might well hold, but that must be established through careful argument and empirical investigation, not implicitly smuggled into the argument through particular ways of characterizing the informal payoff matrix (e.g., by saying that "we cannot allow another 9/11"). There are many cases in which we, as individuals and as a collective society, recognize that consequentialist considerations override deontological rights and principles. Mass electronic surveillance may well turn out to be one such case. However, this question—just like the question of whether to live a pious life in order to prompt belief in God—is too complex to be solved through the blunt instrument of a modern Pascal's wager.