

Betting Against Pascal's Wager

GREGORY MOUGIN AND ELLIOTT SOBER*
University of Wisconsin

Abstract: Only one traditional objection to Pascal's wager is telling: Pascal assumes a particular theology, but without justification. We produce two new objections that go deeper. We show that even if Pascal's theology is assumed to be probable, Pascal's argument does not go through. In addition, we describe a wager that Pascal never considered, which leads away from Pascal's conclusion. We then consider the impact of these considerations on other prudential arguments concerning what one should believe, and on the more general question of when and why belief formation ought to be based solely on the evidence.

1. Introduction

Pascal's wager is a curiosity in its own right, but it also has a wider significance. Within philosophy of religion, it aims to provide a reason to believe that God exists even if all the evidence says that he does not. And within epistemology, it threatens the ideal of objectivity, which demands that we base our beliefs on the evidence and on the evidence alone.

The wager¹ contains two elements. First, there is the assumption of a *theology*. A "theology," as we use the term, does not entail that God exists; rather it provides a set of conditional claims that describe what God would be like if he did exist and what the world would be like if he did not. Pascal assumes that if God exists, he sends theists to heaven and atheists to hell. And if God does not exist, after death there is only dust and ashes. Pascal's theology allows him to describe the *payoffs* that accrue to the theist and to the atheist. Heaven is of great (perhaps infinite) value and hell is comparably awful. Pascal also mentions some this-worldly consequences of theism and atheism. The theist is modestly inconvenienced by the duties of religious observance, whereas the atheist avoids these negligible costs.

In light of the payoffs attaching to theism and atheism,² what is an agent to do? Pascal argues that theism is a prudent wager, even if the probability of God's

existing, based on the evidence alone, is small. It now is customary to represent Pascal's reasoning in a table like the following:

		States of the World	
		God exists	God does not exist
Actions	Theism	a (heaven)	b (modest loss)
	Atheism	c (hell)	d (modest gain)

If p is the probability that God exists, then the *expected utility* $U()$ of theism is $pa + (1 - p)b$ and that of atheism is $pc + (1 - p)d$. We now can deduce that

$$(1) \quad U(\text{Theism}) > U(\text{Atheism}) \text{ if and only if } p(a - c) > (1 - p)(d - b).$$

Pascal argues that $(a - c)$ is *gigantic*, whereas $(d - b)$ is *modest*.³ If so, theism can be the better bet even when p is quite small. What is more, if $(a - c)$ is *infinite* and $(d - b)$ is *finite*, then theism is the better act as long as p isn't zero.

It makes a difference whether the payoffs are finite. If they are, then the Pascalian must show that the evidence confers on the proposition that God exists a probability greater than some threshold that is greater than zero. But if heaven and hell involve infinite payoffs, a stronger conclusion can be drawn: *Grant that there is some chance, however small, that God exists, and prudential considerations lead straight to theism.*⁴

2. Three Traditional Objections

Although W. K. Clifford does not single out Pascal by name, the views he develops in "The Ethics of Belief" (1877) challenge the basic idea of letting prudence, rather than the evidence, influence what one believes.⁵ According to Clifford, "it is wrong always, everywhere, and for anyone, to believe anything upon insufficient evidence (p. 186)." Here is his argument for this conclusion:

...if I let myself believe anything on insufficient evidence, there may be no great harm done by the mere belief; it may be true after all, or I may never have occasion to exhibit it in outward acts. But I cannot help doing this great wrong towards Man, that I make myself credulous. The danger to society is not merely that it should believe wrong things, though that is great enough, but that it should become credulous and lose the habit of testing things and inquiring into them; for then it must sink back into savagery (pp. 185-186).

We do not think that Pascal's prudentialist approach to the question of theism can be dismissed in this way. First of all, we doubt that objectivity is an irreducible end in itself; we suspect that objectivity makes sense only to the extent that it can be expected to have good consequences for the agent.⁶ If so, it would be odd

to criticize Pascal's wager on the grounds that it is a prudential argument when one's alternative epistemology is itself justified by a prudential argument. The issue should be decided by seeing whether Pascal's wager is a *good* prudential argument, not by simply observing that it is, indeed, a prudential argument.

We note, in this connection, that Clifford's argument for objectivity is prudential. And as such, it is none too good. The risk that society will "sink back into savagery" if someone buys Pascal's argument is so small as to be not worth considering. Clifford seems intent on scaring the reader by exaggerating the risks of taking even a single step on a slippery slope.

A second standard objection to Pascal's argument is that theism cannot be treated as a problem of prudential action, because believing a proposition is not an action one can simply decide to perform. Pascal himself anticipates this challenge. He says, with just a trace of irony, that if you accept the logic of the argument, but still find it impossible to believe, then you should start living among religious people. By doing so, habits of faith will gradually take hold. Although one cannot decide to believe, one can decide to live one's life so as to make it probable that one will become a believer.

We agree with the drift of Pascal's reply, although it is questionable how contagious religious beliefs are to those already inoculated with skepticism. In any event, we think that Pascal's principal conclusion is that theism is *prudent*. Whether and how theism may be *attained* is another matter.⁷

The last traditional objection to Pascal's wager that we want to consider is that Pascal offers no justification for his theological assumptions. These assumptions are conditional in form; they describe what happens to theists and atheists if God exists, and what happens to them if he does not. We find this objection quite compelling. If God sends people to heaven and hell for reasons having nothing to do with whether they are theists, then Pascal's analysis will be mistaken.⁸

3. A New Objection to Pascal's Wager

Let us consider this last objection with more care. The objection cannot be that Pascal has not shown his theological assumptions to be *absolutely certain*. This would be an unreasonably high standard to impose. Rather, what we ask for, and what Pascal does not provide, is that his theology be shown to be *reasonable*.

Understood in this way, the objection raises the following question. If Pascal's theology were reasonable, would it follow that theism is more prudent than atheism? We now will show that the answer is *no*. In particular, we will show that

- (2) If Pascal's theology ("P-theology") is highly probable but less than certain, then it does not follow from the payoffs he describes that theism is more prudent than atheism.

Our new objection to the wager stems from a simple point: If $\text{Pr}(\text{P-theology}) < 1$, then the expected utilities of theism and atheism will depend crucially on

what *other possible theologies* say. As it turns out, there are many decision problems involving theism and atheism in which atheism has a higher utility than theism even though the agent is virtually certain that P-theology is true.

The Pascalian may be tempted to respond to this point by saying that if $\Pr(\text{P-theology}) \approx 1$, then alternative theologies will be so improbable that the consequences of theism and atheism under the assumption that an alternative theology is true will have negligible impact on the expected utilities, thus leaving unchanged Pascal's conclusion that $U(\text{Theism}) > U(\text{Atheism})$.

Here is a simple example that shows that this quick response is mistaken:

	G&P	–G&P	G&X	–G&X
theism	heaven	rot	hell	hell
atheism	hell	rot	heaven	heaven

Here G = "God exists," P = Pascal's theology, and X is a theology that says that atheists go to heaven and theists go to hell, regardless of whether God exists or not. The above table describes the consequences of theism and atheism by focusing on what happens after one dies. Further specifying the consequences before one dies, as Pascal troubles to do, will not affect the outcome of the analysis.

X-theology posits an after-life that is not supervised by a deity. After they die, splendid things happen to atheists and horrible things happen to theists. We call these splendid and horrible eventualities "heaven" and "hell" because that is how good and bad they are. If the reader thinks that heaven and hell cannot exist if God does not, then these outcomes should be renamed. All we need for our point is that X-theology says that atheists fare as well (from an experiential point of view, say) as P-theology says that theists do when God exists.

To compute the expected utilities of theism and atheism, the agent must assign probabilities to the four possible states that the world might occupy. The source for these probabilities is the agent's current degrees of belief, where these are calculated solely on the basis of the agent's evidence. The agent we are considering assumes that these four states are exhaustive.

One obvious consequence of the decision table is that

$$\text{If } \Pr(\text{P}) = 1 \text{ and } \Pr(\text{G}) > 0, \text{ then } U(\text{theism}) > U(\text{atheism}).$$

If P-theology is absolutely certain and there is some chance, however small, that God exists, then Pascal's conclusion is correct. However, we now need to consider the case in which $\Pr(\text{P})$ is high, though less than unity.

Let us begin by supposing that heaven is infinitely good and hell is infinitely bad. If so, theism and atheism are equally prudent, as long as $\Pr(\text{G}\&\text{P})$ and $\Pr(\text{X})$ are each nonzero. This is because 1/10th of infinity is the same as 1/1000 of

infinity.⁹ With the assumption of infinite payoffs, we can mimic Pascal's conclusion as follows: grant that there is some chance, however small, that X-theology is true, and prudential considerations lead straight to the conclusion that it doesn't matter whether you are a theist or an atheist.

Let us now explore the supposition that heaven and hell entail enormous, though finite, pay-offs.¹⁰ Given this assumption, it is a consequence of the decision table that

- (3) If $\Pr(P) < 1$, then $U(\text{theism}) > U(\text{atheism})$ if and only if $\Pr(G\&P) > \Pr(X)$.

The point of importance here is that $\Pr(G\&P)$ can be less than $\Pr(X)$ without violating the assumption that $\Pr(P)$ is quite high.

Pascal wants to address someone who thinks that $\Pr(G)$ is small. This is the case of interest, since if someone's evidence already makes $\Pr(G)$ high, there is no need for a prudential argument to justify theism. But notice that if $\Pr(G)$ is small, the same must be true of $\Pr(G\&P)$, even if $\Pr(P)$ is big. It follows that there are many cases in which $\Pr(P)$ is high, $\Pr(X)$ is low, and $\Pr(G\&P) < \Pr(X)$; merely let $1.0 > \Pr(P) \gg 0.5 \gg \Pr(X) > \Pr(G)$.¹¹ So even if P-theology were a quite reasonable theology (in the sense of being highly *probable*), it would not follow that theism is more prudent than atheism.

Of course, there are many more theologies than the two we have described.¹² We have considered P-theology, which Pascal uses to argue that theism is prudent, and X-theology, which can be used to construct a prudential argument for atheism. There are many more theologies of each type. Our suspicion is that widening the range of theologies considered will not displace the conclusion reached here: *The assumption that Pascal makes—that P-theology is evidentially well-supported—does not suffice for theism to have a higher expected utility than atheism.*

There is a general lesson underlying our argument for this point. Typically, agents need to make background assumptions if they are to assign payoffs to actions performed in different states of the world. For example, consider the mundane problem of deciding whether to carry an umbrella. Either there will be rain or there will not. We then consider what the consequences will be in each of four situations. In particular, we help ourselves to such assumptions as (A)—*that we will get wet if we are outside in the rain without an umbrella.*

We make such assumptions, but precisely how certain are we that they are true? A moment's reflection reveals that we are not absolutely certain of these assumptions. At best, we regard them as highly probable. However, this lack of total certainty does not undercut the legitimacy of our analysis; it would be overly fastidious to demand that we take into account the improbable eventuality that (A) is false. The reason is that *a more precise representation of our epistemic situation would not materially affect the outcome of deliberation.* Whether it

makes sense to take an umbrella will not be affected by whether we assume that (A) is certain, or only that its probability is 0.99.

The problem we have just described in connection with Pascal's wager is that his analysis is not similarly robust;¹³ his argument succeeds if his theology is assigned a probability of unity, but founders when a lesser value is specified.

4. A Wager that Pascal Never Considered: The Second New Objection

We think that the problems with Pascal's wager go deeper than the one described in the previous section. Suppose that an atheist has good reason to be very confident in the truth of P-theology, and assigns probabilities to the other possible theologies in such a way that $U(\text{theism}) > U(\text{atheism})$. Pascal supposes that considerations of prudence now oblige the agent to become a theist. However, what is to prevent the atheist from reasoning in the following way? Instead of accepting P-theology on evidential grounds and evaluating G on prudential grounds, why not accept $\neg G$ on evidential grounds and evaluate the prudential value of continuing to believe P-theology? Pascal asks us to assume a theology and to decide whether to shift from atheism to theism. The new problem is to assume atheism and to decide whether to shift from P-theology to some other theology.

Pascal cannot reject this problem out of hand. After all, if convictions about the existence of God may be modified for prudential reasons, why can't theologies? Let us now trace out the consequences of this new problem.

If an atheist is to consider alternatives to P-theology, what options should be contemplated? We will focus on X-theology, which, recall, says that atheists go to heaven and theists go to hell, regardless of whether God exists. For an atheist, the payoffs associated with believing P-theology (B(P)) and believing X-theology (B(X)) are as follows:

	P-theology is true	X-theology is true
B(P)	rot and feel bad	heaven and feel bad
B(X)	rot and feel good	heaven and feel good

Atheists who accept P-theology are apt to suffer from anxiety. This is because atheists sometimes contemplate what will happen to them if they are mistaken and God in fact exists. How much better for one's peace of mind to see one's atheistic error, if such it be, go unpunished!¹⁴ This is why, from a psychological point of view, it is better for an atheist to embrace X-theology than P-theology.

Pascal explored the problem of deciding whether to be a theist, on the assumption that P-theology is true. We have just described the problem of deciding whether to accept P-theology, on the assumption that atheism is true. These two separate problems can be combined into a single problem in which the agent

decides which of four pairs of beliefs to adopt. This fuller problem can be represented as follows:

	G&P	–G&P	G&X	–G&X
B(G)&B(P)	heaven and feel good	rot and feel good	hell and feel good	hell and feel good
B(–G)&B(P)	hell and feel bad	rot and feel bad	heaven and feel bad	heaven and feel bad
B(G)&B(X)	heaven and feel bad	rot and feel bad	hell and feel bad	hell and feel bad
B(–G)&B(X)	hell and feel good	rot and feel good	heaven and feel good	heaven and feel good

The first line in each cell describes the payoffs of the relevant action in the after-life; the second line describes its this-worldly consequences.

The wager that Pascal described involves comparing the first and second acts where the world is assumed to fall in the first or second column. The wager described at the beginning of this section involves comparing the second and fourth acts, assuming that the world falls in the second or fourth column.

What can be said of the fuller 4×4 problem just defined? Everything depends on the assignment of probabilities (based solely on the evidence) to the four possible states of the world. As in Pascal's wager, we will restrict our attention to an agent who starts off thinking that (–G&P) is highly probable.

The first point is that if all the probability is concentrated on (–G&P), there is a tie between B(G&P) and B(–G&X). That is, if you believe, based on the evidence, that (–G&P) is certain, this belief is *prudentially unstable*. An agent who begins with this belief must move away from it, if belief revision is driven by prudential considerations of the sort we have been describing. Once the agent abandons (–G&P), there are *two* belief states to which the agent may migrate. One of these equilibria was described by Pascal; we have described the other.

The second point is that if most, but not all, of the probability is concentrated on (–G&P), what the best act is will depend on how the remaining probability is distributed. If all of it is assigned to (G&P), then B(G&P) is prudentially best. If more is assigned to X-theology than to (G&P), then B(–G&X) is best.¹⁵ Here we recapitulate the point made in Section 3; assigning a high probability to P-theology (or to the conjunction –G&P) is not enough for Pascal to reach his desired conclusion. However, the present point is that Pascal's prudentialism obliges him to cast his net more widely than he supposed; if belief in God should be decided on nonevidential grounds, the same will be true of other beliefs. This means that Pascal's very *local* prudentialism—in which only theism is subjected to prudential scrutiny—is arbitrarily narrow.¹⁶

5. Is Prudential Belief Revision Incoherent?

Although we believe that Pascal's wager is a flawed prudential argument, we do not reject prudential belief revision as such. It is perfectly possible for agents to face well-conceived decision problems in which there is a good prudential reason to believe a proposition that the evidence says is very improbable, as the following example shows.

Professor Smith is in her office, musing on two propositions that she believes. Based on the evidence she possesses, she believes that the President of the United States is, at this very moment, not in Wrigley Field; call this proposition –W. She also believes, again based on considerable evidence, that her student, Fred, is a highly reliable and sincere individual—he says what he means and means what he says; call this proposition S. While Smith is reflecting on this pair of beliefs, Fred bursts into her office and makes the following announcement.

Fred says that Professor Smith must decide whether to change her belief in the whereabouts of the President. After Smith makes this decision, Fred will independently check the President's whereabouts. If the President turns out to be in Wrigley Field, Professor Smith's decision will have been a matter of life and death. If she has changed her belief to W, she will be allowed to live. But if she has not, Fred promises to shoot her with the gun he produces from his backpack. In contrast with these momentous consequences, nothing much is at stake if Fred discovers that the President is not in Wrigley Field. In this case, Fred will allow Smith to live regardless of what she happens to believe.

Intuitively, there is a clear-cut prudential argument that Smith should try her best to believe proposition W—that the President is now in Wrigley Field. Although this proposition is improbable, given the evidence at hand, the circumstances make it prudent to believe what the evidence says is improbable.

We now will describe a Pascalian formulation of Smith's decision problem in which only one of Smith's beliefs is subject to prudential scrutiny—namely, her belief about the whereabouts of the President. We then will describe a quite different problem in which one of Smith's background assumptions is evaluated. Then, finally, we will characterize a fuller decision problem in which both propositions are evaluated together. The pattern will recapitulate what we have said about decisions concerning belief in God and belief in P-theology, except that the net result turns out to be interestingly different.

In the first local argument, Smith continues to believe on evidential grounds that Fred is sincere (S), but subjects her belief that the President is not now in Wrigley Field (–W) to prudential scrutiny. These are the pay-offs she needs to consider:

	W	–W
B(W)	live + e	live
B(–W)	die	live + e

The payoff entry “e” represents the modest benefit of having a true belief; those inclined to think that this “benefit” is no benefit at all are invited to let $e = 0$. It is clear in this case that as long as W has nonzero probability and the difference between life and death swamps the magnitude of e , Smith is better off believing W .¹⁷

The second local argument we want to describe may not be one that Smith would think to consider. In this decision problem, Smith continues to believe on evidential grounds that the President is not now in Wrigley Field ($-W$), but subjects her belief that Fred is sincere (S) to prudential scrutiny. These are the pay-offs that Smith needs to consider:

	S	$-S$
$B(S)$	live + e	live
$B(-S)$	live + p	live + e + p

In this decision problem, Smith continues to believe that the President is *not* in Wrigley field, so she figures she will live whether or not Fred is sincere. As before, “e” represents the benefit of having a true belief. The new item is “p,” which represents the psychic benefit of being free from anxiety. If Smith embraces $-S$, she won’t worry that Fred will shoot her. Of course, whether $p > e$ depends on the values of the person in question; W. K. Clifford might deny this inequality. But Professor Smith, if she is anything like the two of us, will quickly conclude that $p \geq e$. As a result, she will decide to shift from believing that Fred is sincere to believing that he is not.

So far, we have considered two local deliberations, in which the second seems to undercut an assumption made in the first. This is just as unsatisfactory in the case of Professor Smith’s problem as it was in the case of Pascal’s wager. Let us now move to a more global treatment of the problem, in which Smith’s initial belief in $-W$ and her initial belief in S are both subjected to prudential scrutiny. Here are the payoffs:

	$W \& S$	$-W \& S$	$W \& -S$	$-W \& -S$
$B(W) \& B(S)$	live + e	live	live	live
$B(-W) \& B(S)$	die	live + e	live	live
$B(W) \& B(-S)$	live + p	live + p	live + e + p	live + p
$B(-W) \& B(-S)$	die + p	live + p	live + p	live + e + p

If Smith begins, as we said, by thinking that $(-W \& S)$ is highly probable, but allows that the other options have at least nonzero probability, then she should change. She should believe either $(W \& S)$ or $(W \& -S)$, depending on how e and p are related. *Notice that each of these end states includes believing W .*

Here we have the fundamental difference between the decision problem that Pascal describes and the one that Professor Smith confronts. Professor Smith's decision to abandon her belief that the President is not now in Wrigley Field is *robust*. Although her rough-and-ready deliberation may begin with the assumption that Fred is sincere, she will reach the same decision about proposition W even if she allows prudential considerations to be brought to bear on her background assumption. For even if she subjects both her belief in S and her belief in $\neg W$ to prudential scrutiny, she still should try to reform her belief in $\neg W$. No similar robustness attaches to the Pascalian problem. If the agent begins with the assumption that P-theology is certain and holds that belief immune from revision, matters proceed as Pascal described. But if the agent's initial belief in $\neg G$ and his initial belief in P-theology are both subjected to prudential scrutiny, it is far from straightforward that belief in God remains the best action.¹⁸

There are three elements in both these decision problems: a set of *background assumptions* allows a *prudential reason* to be developed for changing one's degree of belief in some *focal proposition*. Indeed, this threefold structure is not unique to prudential reasoning; it is a well-known feature of evidential reasoning that an *observation* typically provides evidence concerning the truth or falsity of a *hypothesis* only in the light of *background assumptions*.¹⁹

The difference between Pascal's wager and the predicament of Professor Smith allows a principle to be stated that the background assumptions in a given problem must satisfy. If Pascal wishes to assume that P-theology is true in developing his prudential argument for theism, then the conclusion that theism is the best act should not be undermined by subjecting P-theology to prudential scrutiny. If Professor Smith wishes to assume that Fred is sincere in developing her prudential argument for changing her belief concerning the whereabouts of the President, then her conclusion that she is best off believing that the President is in Wrigley Field must not be undermined by a prudential evaluation of the assumption that Fred means what he says. And if a scientist uses background assumptions to argue that the observations make some hypothesis plausible, it must be true that the hypothesis would remain plausible in the light of the observations, even if the observations were to provide evidence that the background assumptions are false.²⁰

If a background assumption B allows consideration C (evidential or prudential) to be used to justify X,²¹ there are two ways that this justification of X can be robust. The first is that C occasions no change at all in B; the second is that C does motivate a change in B, but the revised formulation of B *still* allows C to justify X. It is perhaps here that the character of background assumptions in evidential arguments often differs from the role of background assumptions in the prudential arguments we have surveyed so far. Smith's background belief that Fred is sincere *does* get changed when it is subjected to prudential scrutiny, but this change does not undermine the conclusion that Smith *also* has a prudential reason to revise her focal belief concerning the whereabouts of the President. In

contrast, the background assumptions exploited in evidential inference often remain totally unaffected by the evidence at hand. For example, when a scientist infers that the temperature is 98.6° because the thermometer says that the temperature is 98.6°, the background assumption is being made that the thermometer is reliable. Typically, this background assumption is neither supported nor undermined by the fact that the thermometer reads 98.6°.

6. James' "Will to Believe"

The example of Professor Smith strongly suggests that there can be no global and unconditional requirement of objectivity. At least sometimes, it is perfectly sane to allow prudential considerations to overturn the testimony of the evidence. We believe that such a case is encompassed in the spirit, if not the letter, of William James' pragmatic argument for religious belief in "The Will to Believe."

James (1897) focuses on this-worldly consequences, not the possibilities of heaven and hell. He considers someone who would be despondent without a faith in something "higher," but who would be able to live a happy and productive life if such a faith were embraced.²²

The first objection we considered in connection with Pascal does not apply to James' reasoning. James postulates no theology as part of his argument, so it would be irrelevant for us to consider the consequences of alternative theologies. James ignores such issues in the same way that Professor Smith quite reasonably does. The focus is on tangible consequences in *this* world.

Our second objection to Pascal has some bearing on James' argument, but he has a reply. James envisions a person who will be depressed if an atheist and happy if a theist. But this two-part disposition is as much open to reform as the agent's atheism. Instead of becoming a theist, perhaps the agent should consult a psychiatrist who can help him be a happy atheist. If becoming a theist can benefit the agent, other changes may do the same.

James cannot dismiss this suggestion out of hand. However, its evaluation will depend on the details. If it is relatively easy and costless to embrace theism, but rather difficult and expensive to go into therapy, where the end result is in any case less certain, utility may be on the side of James' recommendation. James' fallibilism can lead him simply to agree that theism *may* not be the best solution. But then again, it may be, at least for some people some of the time.

7. Concluding Remarks

The only good traditional objection to Pascal's wager is that Pascal does not establish that his theological assumptions are highly probable. This is true, but it does not go far enough. For even if the theology *were* highly probable, Pascal's argument is vulnerable on two fronts.

Our first criticism of Pascal's wager was that theism is not a robust solution to the problem that Pascal poses; although theism must be the right act if P-theology

is certain, theism can fail to exhibit this property if P-theology is less than certain. Robustness, we maintain, is an important consideration for decision making under uncertainty.

Our second criticism also focused on Pascal's theology, but it has a different character. We explored some of the implications of allowing prudence, rather than evidence, to determine what one believes. Pascal's prudential approach to the problem of theism obliges him to take the same approach in matters of theology. But Pascal's theology is not prudentially stable, at least for an atheist who is bound to remain one. Here we are advancing, so to speak, an argument *ad hominem*; we do not ourselves endorse the relevance of prudence, but merely point out what Pascal's framework obliges him to consider.

Pascal's wager implicitly involves two simplifying assumptions—that $\text{Pr}(\text{P-theology}) = 1$ and that theism is the only belief that the agent will scrutinize prudentially. We do not criticize the use of simplifying assumptions; we have used them ourselves, as any analysis must. However, decisions reached via simplifying assumptions should be *robust*; they should remain in place even when the simplifying assumptions are replaced with more realistic ones.

We have argued that Pascal's conclusion fails this test for each of the simplifying assumptions just mentioned. If we abandon the assumption that $\text{Pr}(\text{P-theology}) = 1$, it is far from clear that theism is more prudent than atheism. And if we subject P-theology as well as atheism to prudential scrutiny, the same doubts arise.

The challenge posed by our consideration of Pascal's wager is to define the scope and limits of objectivity. The defects in Pascal's argument notwithstanding, it isn't *always* true that our beliefs should be regulated by the evidence and the evidence alone. However, this negative result cries out for a positive thesis: Can general conditions be described under which it *is* prudent to be objective? Ideally, a general thesis of this type would also explain contrary cases. We would be able to recognize the generality, if not the limitless validity, of the argument for objectivity; and outside its boundaries we would be able to locate the deviant cases of Pascal, James, and Professor Smith. On another occasion, we hope to begin where the present inquiry must, out of prudence, end.

Notes

*We wish to thank John Earman, Ellery Eells, Berent Enç, Malcolm Forster, William Lycan, Alan Sidelle, Dennis Stampe, and an anonymous referee of this journal for useful comments on earlier versions of this paper.

¹We are not concerned in this paper to decide the historical question of what Pascal actually intended. Rather, we are considering what has generally come to be viewed as "Pascal's wager." Hacking (1972) argues that Pascal actually formulated three quite different arguments. We are discussing the one that we think is the most intellectually challenging.

²"Theism" and "atheism" usually name propositions. We use the terms here to name the "acts" of believing or disbelieving that God exists.

³In the vocabulary used in Sober (1994), (a – c) represents the *importance* of the proposition that God exists and (d – b) the import of the proposition that he does not.

⁴Notice that Pascal's wager, as just represented, treats the assignment of a probability to the proposition *God exists* and the believing of that proposition as separate elements. For modern probabilists who are skeptical of the idea of *acceptance*, what could this distinction mean? If belief is understood in terms of *degree of belief*, then assigning a probability and having a belief are not two acts, but one.

The interpretation we favor is this: First, one assigns a probability to the proposition, based solely on the evidence. Then prudential considerations are adduced that allegedly justify a *change* in probability assignment. Construed in this way, Pascal argues that even if the evidence gives *p* a low value, prudential considerations should lead one to "update" that probability and make *p* quite high.

⁵Or more precisely, Clifford gives a prudential reason (described below) for never allowing anything but the evidence to influence what one believes. This position can be rescued from self-contradiction by introducing a hierarchy of levels. Let level-1 propositions be about the world outside the mind. Let level-2 propositions be about level-1 propositions, and so on. Then a consistent formulation of Clifford's position is that there exists a compelling level-2 prudential argument, whose conclusion is that all level-1 propositions should be accepted or rejected solely on the basis of the evidence.

⁶For example, consider the Dutch book arguments that Bayesians construct to argue that beliefs should be probabilistically coherent. These are consequentialist defenses of a particular conception of rationality.

⁷We think this objection boils down to nothing more than the point that believing is not a basic action. Given the lives that most of us lead, we can't simply decide to believe a proposition, anymore than we can decide to be President of the United States or to run a four minute mile. There is nothing special about belief or mental activity involved here. What is true in all such cases is that we can decide to *try*. Decisions about whether it makes sense to try, of course, should take account of the probability of success.

⁸In their article on Pascal's wager, Lycan and Schlesinger (1989) defend the Anselmian idea that God is an absolutely perfect being by arguing that this position is favored by Ockham's razor. We have two questions about their discussion. First, we are inclined to doubt that it is legitimate to use Ockham's razor in the way that Lycan and Schlesinger propose; see Sober (1988) for discussion. Second, even if God is a perfect being, we fail to see how this helps shore up Pascal's wager. Why expect P-theology to accurately describe the behavior of a perfect God? As it happens, we find far more attractive the theology that says that a perfect God would be inclined to forgive the theological mistakes of his creatures; he would care more about having people lead morally good lives than about their having true beliefs concerning his existence. In any case, we will argue in the next section that even if Pascal's theology *is* defensible (in the sense that it can be shown to be highly probable), this is not enough to license the conclusion that Pascal draws.

⁹The reader may be inclined to object that a 0.1 chance of an infinite payoff is better than a 0.001 chance. We will not try to argue otherwise, but merely refer the reader to the argument that begins in the next paragraph.

¹⁰The idea that the payoffs are finite has an independent motivation. The agent whom Pascal is addressing presumably is allowed to care about the afterlives his loved ones (*x* and *y*) attain as well as about the afterlife he himself (*s*) achieves. In particular, this agent may well prefer that *s* and *x* (but not *y*) go to heaven over having *s* (but neither *x* nor *y*) do so. This is difficult to model if payoffs are allowed to be infinite.

¹¹We don't claim that $\Pr(G)$ is less than $\Pr(X)$, only that this would suffice to undermine Pascal's argument. Note that $\Pr(G \& P)$ can be less than $\Pr(X)$ even when each of $\Pr(G)$ and $\Pr(P)$ is greater than $\Pr(X)$.

¹²For example, consider theology *X'*, according to which you go to heaven if and only if God exists and you are an atheist. Objections to Pascal's wager that resemble the one we have formulated in terms of theology *X* can be developed by focusing on *X'*.

¹³The importance of robustness as a property of mathematical models in biology is discussed in Orzack and Sober (1993).

¹⁴In the problem we are describing here, the agent should abandon belief in P-theology for prudential reasons, even if the evidence says that $\Pr(\text{P-theology}) = 1.0$. This differs from the wager that Pascal discussed; in Pascal's problem, it isn't true that the agent should abandon atheism even when the evidence says that $\Pr(\neg G) = 1.0$.

¹⁵This point requires that we think of heaven and hell as entailing finite payoffs, an assumption that we discussed in footnote 10.

¹⁶Although P-theology is prudentially unstable, the question remains whether some similar theology is stable.

¹⁷More precisely, Smith should believe proposition W if and only if $w(\text{Life} - \text{Death}) > e - 2w$, where w is the probability that proposition W is true.

¹⁸It may be objected that our analysis of Smith's problem is open to the same objection that we leveled at Pascal's argument. We assumed that if Fred is not sincere, *he will do nothing*. But the possibility that Fred is not sincere admits of many conceivable subcases. Perhaps if we specified these more finely, a different decision would result. Our answer is that Professor Smith should exploit her well-confirmed beliefs concerning how people behave. It is the content of these background assumptions that renders Smith's decision about proposition W robust. If this isn't so, then the simple analysis that Smith uses is not satisfactory.

¹⁹For discussion of this well-known epistemological point, see Sober (1988).

²⁰The need for background assumptions in any decision problem raises the question of how *global* the prudentialism is that Pascal can consistently advocate. He focuses on one belief (theism) and we have focused on a second (P-theology). The more general question is whether *all* beliefs could be evaluated on purely prudential grounds. Presumably, it at least *makes sense* to be a 100% evidentialist, though whether that turns out to be the *best* policy is another matter. And even if it turns out that 100% prudentialism is incoherent, more *local* forms of prudentialism can nonetheless be coherently formulated.

²¹The term "justify" can mean justifying a *proposition* or justifying an *action*. We purposely use the ambiguous term here because we want the idea of robustness to encompass both prudential and evidential reasoning.

²²Although James restricts his prudential argument to the case in which the evidence does not rule out the attractive proposition, this hedge appears arbitrary from a pragmatic point of view. If theism is prudentially permissible when the evidence is indecisive, why shouldn't it be permissible, as Pascal insists, even when the evidence goes against God's existence?

James also says that the prudent deliberator need not consider alternatives that are not "live" (his example is "belief in the Mahdi"). If "live" means that the option has a subjective probability greater than zero, then James' restriction, in fact, does not allow him to ignore alternatives that he finds far-fetched, or which he simply does not feel like considering. On the other hand, if "live" means something different, then it is easy to show that attending only to "live" options can lead to conclusions that contradict the dictates of prudence.

References

- Clifford, W. (1879) "The Ethics of Belief." In *Lectures and Essays*, volume 2. London: Macmillan, pp. 177–211.
- Hacking, I. (1972) "The Logic of Pascal's Wager." *American Philosophical Quarterly* 9: 186–192.
- James, W. (1897) "The Will to Believe." In *The Will to Believe and Other Essays in Popular Philosophy*. New York: Longmans, pp. 1–31.
- Lycan, W. and Schlesinger, G. (1989) "You Bet Your Life: Pascal's Wager Defended." In J. Feinberg (ed.), *Reason and Responsibility*, 7th edition, Belmont, California: Wadsworth Publishing Co., pp. 82–90.
- Orzack, S. and Sober, E. (1993) "A Critical Assessment of Richard Levins' 'The Strategy of Model Building in Population Biology.'" *Quarterly Review of Biology*, 68: 533–546.
- Pascal, B. (1662) *Pensées*. translated by W. Trotter, New York: J. M. Dent Co., 1958, fragments 233–241.
- Sober, E. (1988) *Reconstructing the Past: Parsimony, Evolution, and Inference*. Cambridge: MIT Press.
- Sober, E. (1994) "The Adaptive Advantage of Learning and A Priori Prejudice." In *From a Biological Point of View*, Cambridge University Press.