

Less Good but not Bad: In Defense of Epicureanism about the Badness of Death

Aaron Smuts

Abstract

In this paper I defend *innocuousism*—a weak form of Epicureanism about the badness of death. I argue that if we assume both *mental statism* about well-being and that death is an experiential blank, it follows that death is not bad for the one who dies. I defend *innocuousism* against the deprivation account of the badness of death. I argue that recent defenses of the deprivation account, such as those offered by Fred Feldman and Ben Bradley, rest on a suspect notion of extrinsic badness—a notion that erroneously confuses states of affairs that merely could have been better with those that are bad. In reply, I defend an alternate account according to which something is extrinsically bad if and only if it leads to states that are intrinsically bad. On my view, sometimes dying may be less good than living, but it is never bad to die.

Prologue

'Man is snapped off like a reed in a canebrake!
The comely young man, the pretty young woman -
all [too soon in] their [prime] Death abducts them!

'No one at all sees Death,
no one at all sees the face [of Death]
no one at all [hears] the voice of Death,
Death so savage, who hacks men down.

...

'Ever the river has risen and brought us the flood,
the mayfly floating on the water.
On the face of the sun its countenance gazes,
Then all of a sudden nothing is there!

'The abducted and the dead, how alike is their lot!
But never was drawn the likeness of Death,
Never in the land did the dead greet a man.'

The Epic of Gilgamesh (tablet X, trans. Andrew George)

1 Introduction

In a cryptic passage in his "Letter to Menoecus," Epicurus presents what may at first seem to be an absurd argument for the claim that it is irrational to fear death:

Foolish, therefore, is the man who says that he fears death, not because it will pain when it comes, but because it pains in the prospect. Whatever causes no annoyance when it is present, causes only a groundless pain in the expectation. Death, therefore, the most awful of evils, is nothing to us, seeing that, when we are, death is not come, and, when death is come, we are not. It is nothing, then, either to the living or to the dead, for with the living it is not and the dead exist no longer.¹

Epicurus's principal conclusion is that death is nothing to us. Death is nothing to us, because it does not lead to any bad experiences—it is the end of experience. And only experiences are good or bad for a person. Hence, death is not bad for the one who dies. Further, he assumes that it is irrational to fear what does no harm. He concludes that if it is irrational to fear what does no harm, then it is irrational to fear death.

Epicurus's letter raises two distinct questions: (1) Is it rational to fear death? And, (2) Is death bad for the one who dies? In this paper, I will say fairly little about the first question. My principal goal is to support his answer to the second question. I defend both Epicurus's position on the badness of death and his reasons, or at least an argument in much the same spirit. Many call the position that death is not bad for the one who dies *Epicureanism* about the badness of death. Since this label is easily confused with Epicurus's related, but distinct, position on whether it is rational to fear death, I adopt the label *innocuousism* for the claim that death does not harm the one who dies. This position holds that death is innocuous because it does no injury to the departing.

I defend innocuousism in the face of a widely accepted style of refutation—the deprivation account of the badness of death.² The deprivation account holds that death is bad for the one who dies when it deprives her of good experiences that she would have had otherwise. That is, death is bad because it deprives us of the goods of life. This account of the badness of death gives rise to a number

¹ Epicurus, "Letter to Menoecus", trans. by Robert Drew Hicks, *The Internet Classics Archive*, <http://classics.mit.edu/Epicurus/menoec.html> (accessed March 1, 2009).

² The deprivation account is defended by nearly every critic of Epicurus. For example, see: Thomas Nagel, "Death," in *Mortal Questions* (Cambridge, 1979); Bernard Williams, "The Makropulos Case: Reflections on the Tedium of Immortality," in *Problems of the Self* (Cambridge UP, 1993); Robert Nozick, "Happiness," in *The Examined Life* (Simon and Schuster, 1989); Fred Feldman, *Confrontations with the Reaper* (Oxford UP, 1992); and Steven Luper, "Posthumous Harm," *American Philosophical Quarterly* 41:1 (2004): 63-72.

Section 1. Introduction

of well-known puzzles that I will not explore.³ Instead, I attack the theory of extrinsic badness at the core of the most convincing formulations of the deprivation account.

Recent defenders of the deprivation account, such as Fred Feldman and Ben Bradley, hold that although death is not intrinsically bad, it is extrinsically bad. They argue that death is sometimes extrinsically bad, not because it leads to intrinsically bad states of affairs, but because it leads to states that are less intrinsically good. I argue that this account of extrinsic badness conflates things that are merely less good with those that are bad.⁴ I intend to show that if we respect the distinction between states of affairs that are bad and those that are merely less good, the deprivation account fails as an objection to innocuousism.

My argument proceeds in a few steps. I begin by developing a contemporary version of Epicurus' argument that I call the *Dead End Argument for Innocuousism*. I then explain the deprivation account of the badness of death. In response, I raise several objections to the theory of extrinsic badness at the heart of the deprivation account. In support of innocuousism, I defend a competing account of extrinsic badness that avoids these problems. Finally, I respond to several objections to the prima facie absurd suggestion that death is not bad for the one who dies.

³ The timing puzzle is the most difficult. It asks: If death is bad for the one who dies, when is it bad for them? Seemingly, it cannot be before they die, unless we allow for backward causation. And it cannot be after, since the person no longer exists. Feldman (*Reaper*, ch.9) provides a good overview of the problem. In addition, Steven Luper's entry "Death," in *The Stanford Encyclopedia of Death*, provides a clear explanation of the puzzle.

⁴ Kai Draper argues that some comparatively bad things are not reasonably feared. They are simply things that we cannot avoid and should not be considered misfortunes. See: Draper, "Disappointment, Sadness, and Death," *The Philosophical Review* 108:3 (1999): 387-414. I am not concerned with the distinction between the merely bad and misfortunes. Rather, I argue that comparatively bad is not necessarily bad.

2 The Dead End Argument for Innocuousism

The *Dead End Argument for Innocuousism* holds that death is not bad for the one who dies. This is because death leads to nothing—death is an experiential dead end. Since death is the end of experience, it is not intrinsically bad for the one who dies. Neither is it extrinsically bad. To be extrinsically bad, something must lead to intrinsically bad states of affairs. But there are no intrinsically bad states of affairs after death. Only experiences are intrinsically bad or good for a person. Hence, death is neither intrinsically nor extrinsically bad for the one who dies. Death is prudentially innocuous.

Here is a formalization of the argument:

The Dead End Argument for Innocuousism

1. The sole bearers of intrinsic prudential value are mental states. (*mental statism*)
2. Death is an experiential blank.
3. Hence, the state of being dead is not intrinsically prudentially bad.
4. An event is extrinsically bad if and only if it leads to intrinsically bad states of affairs. (*causal hypothesis*⁵)
5. Hence, death is not extrinsically prudentially bad.
6. Therefore, death is not prudentially bad for the one who dies. (*innocuousism*)

The *Dead End Argument* rests on three controversial premises: *mental statism* about welfare, the claim that death is the end of experience, and the *causal hypothesis*. To further narrow the scope, in this paper I only intend to defend the causal hypothesis. But it might be helpful to say a few words in support of the other claims, lest the argument be dismissed as a conglomeration of absurdities.

Mental Statism

The first premise of the *Dead End Argument* is extremely controversial. Mental statism is a general theory of prudential value. Theories of prudential value (also known as "well-being," "welfare," and "self-interest") tell us what makes a life good for the one who lives it. Prudential value is distinct from other ways in which a life may be good. For instance, a life of self-sacrifice spent working for the poor might be both morally and instrumentally good, but it is not necessarily good for the one who lives it. We do not intend to benefit ourselves through self-sacrifice. This is precisely what makes it self-sacrifice. Although it is almost always morally good to sacrifice one's own good for the good of others, it is almost never, or at least not necessarily, intrinsically prudentially good.

Mental statism holds that the sole bearers of intrinsic prudential value are mental states. Hedonism, for instance, is a specific type of mental statism that

⁵ I adopt Feldman's name for the position. Feldman, *Reaper*, p.135.

Section 2. The Dead End Argument for Innuocuousism

counts only the mental states of pleasure and pain.⁶ By mental states, the general theory intends phenomenal experiences. Alterations of unconscious mental operations that do not affect any conscious phenomenal experiences do not have any impact on prudential value. Zombies, on this theory, have no prudential value. Further, if there are no qualia in the universe, then, according to mental statism, there is no prudential value in the universe.⁷

Most controversially, mental statism implies what is known as the *experience requirement*—the claim that what you do not experience cannot harm you.⁸ There is a fairly wide consensus, although not without dissenting opinion, that experience machine-style examples show that mental statism is false.⁹ Although I think that mental statism is correct, I cannot provide an adequate defense here.¹⁰ But it might be worth saying a few words to make the view look a bit less absurd.

The most compelling counter-example to mental statism is Nozick's experience machine.¹¹ Nozick asks us to imagine a machine that would be able to simulate a wide array of fantastic experiences. The experiment gives us, what is by now, a familiar sales pitch: Perhaps you want to write the great American novel. Well, in the experience machine you can have the experience of writing

⁶ Parfit and Wolf reverse this distinction. They seem to hold that all forms of *mental statism* are forms of hedonism. See: Derek Parfit, "What Makes Someone's Life Go Best," in *Reasons and Persons* (Clarendon Press, 1984); Susan Wolf, "Happiness and Meaning: Two Aspects of the Good Life," *Social Philosophy and Policy*, 14 (1997): 207-25.

⁷ Mental statism could be formulated in a way compatible with physicalism. But without an experiential component the theory is far less *prima facie* compelling.

⁸ Since the experience requirement is not always presented in the same way, I have chosen to focus on mental statism. Ivan Soll, for instance, defends a theory called *experientialism* that has both motivational and broad axiological implications. Mental Statism, as I have formulated it, implies only a limited axiological claim about well-being: something can affect someone's well-being only if it makes an experiential difference for that person.

Ivan Soll, "On the Purported Insignificance of Death," in *Death and Philosophy* eds. Jeff Malpas and Robert Solomon (Routledge, 1998): 22-39.

⁹ Nozick's "Experience Machine", Nagel's "Deceived Businessman", Nagel's "Contented Infant", Mill's "Pig", and Nozick's "Mongolian Pornographer" are the most pressing thought experiments presented in opposition to *mental statism*.

Robert Nozick, *Anarchy, State, and Utopia* (Basic Books, 1974), pp.42-5; Thomas Nagel, "Death"; John Stuart Mill, *Utilitarianism* (Hackett, 2002); and Robert Nozick, "On the Randian Argument," in *Socratic Puzzles* (Harvard UP, 1997).

¹⁰ Shelly Kagan confesses his temptation toward mental statism in two excellent articles: "Me and My Life" and "The Limits of Well-being". He argues for the claim that to have an effect on a person's well-being, something must have an effect on the intrinsic properties of the person.

Shelly Kagan, "Me and My Life," *Proceedings of the Aristotelian Society*, 94 (1994): 309-324; and "The Limits of Well-being," in *The Good Life and the Human Good*, ed. E.F. Paul, F.D. Miller and Jeffrey Paul (Cambridge UP, 1992).

¹¹ Robert Nozick, *Anarchy*, pp.42-5.

the most celebrated novel in history. Your work will be praised far and wide. Champion athlete, war hero, legendary lover, you name it—in the machine, you will experience any life that you desire. Your life in the machine will seem as real as any experience that you have ever had. You will never know the difference.

Nozick asks us whether or not we would step into the machine. Intuitions diverge—almost no one under twenty-five says no, and nearly no one over twenty-five would step into the machine. Most of us opt out of the machine because we do not merely want to think that we have written the great American novel; we want to write it. We do not merely want to think that we have accomplished something or have had genuine relationships; we want to accomplish things and form genuine bonds with others. We like to win, but we do not want every game to be fixed in our favor.

Although most non-frivolous, non-terminally ill people would opt out of a life in the experience machine, this does not constitute an objection to mental statism. Yes, the thought experiment clearly shows that we want more than mere experiences, but it does not show that things without experiential impact can affect our well-being. The thought experiment merely confirms what we already knew: We desire many things other than our own well-being.¹² Strict psychological egoism is highly implausible. We often non-selfishly desire things for the good of others. People frequently sacrifice themselves for a cause or for the benefit of those they love. And many people have been known to sacrifice their own well-being for other kinds of goods, such as meaning and significance.

Most of us think that a life in the experience machine would be meaningless. Insofar as we desire significance, we will opt out of the machine. This does not show that we think that we would be better off—that we would have a higher state of welfare—outside of the machine. Life in the machine simply cannot give us everything that we want. Since we want more than what merely increases our well-being, the case against mental statism is inconclusive.

I do not pretend that this brief consideration of the experience machine objection is a decisive reply. Rather, I merely want to sketch a plausible line of defense for mental statism. Further, positive defense of the theory is out of scope.

The Rest of the Dead End Argument

As noted above, the *Dead End Argument for Innocuousism* assumes mental statism and two other controversial premises: the thesis that death is an experiential blank and the causal hypothesis. To say that death is an experiential blank is just to say that there are no mental states after death. I will assume

¹² D. W. Haslett also appeals to this distinction, as does Jeffrey Goldsworthy. See: D. W. Haslett, "What is Utility?," *Economics and Philosophy*, 6 (1990): 65-94; and Jeffrey Goldsworthy, "Well-Being and Value," *Utilitas*, vol. 4, no. 1 (1992): 1-26.

Section 2. The Dead End Argument for Innocuousism

that mental states supervene on some set of subvening brain states. When there are no longer any of the relevant brain states, there are no longer any mental states. Brain death is the end of the subvening brain states. Hence, brain death is an experiential blank. In this article, I simply assume that Utanipishti's description in the *Epic of Gilgamesh* is correct: "Then all of sudden nothing is there!"—that death is the end of experience. Additional positive defense of this claim is far out of scope.¹³

As for the causal hypothesis, the focus of this paper, it is best defended in response to criticism. But it will help make the theory a bit more explicit before we begin. The causal hypothesis holds that what makes something extrinsically bad is that it leads to intrinsic bad. An event is extrinsically bad in proportion to the amount of intrinsic badness it causes. Nothing can be extrinsically bad if it does not lead to intrinsic bad. Since events can lead to both good and bad outcomes, when assessing the overall extrinsic value of an event, we need to consider all of its effects. Nothing can be overall extrinsically bad for someone unless it leads to, or more precisely, is causally responsible for more intrinsically bad states of affairs than intrinsically good states of affairs.¹⁴

¹³ It pays to note that a roughly Epicurean argument against the badness of death does not require, what Feldman calls, the *termination thesis*—the claim that death is the end of our existence. Feldman, *Reaper*, ch. 5.

¹⁴ As stated, the evaluation of an event should include all states of affairs that it causes. We will likely want to distinguish between far distant transient causes and genuine, relevant causes. I am reluctant to say that the copulation responsible for my birth was a cause of my stubbing my toe this morning. The copulation does not seem relevant. The causal hypothesis will likely require a far more restricted notion of "cause," than one that would count my birth. However, it is far out of scope to develop such a theory here.

3 The Deprivation Account

Although death may not lead to any intrinsically prudentially bad states of affairs, many argue that, nonetheless, it is often bad for the one who dies. The reason why it is bad is not that it leads to bad experiences. How could it, since it leads to the absence of experience? Rather, death is bad because it leads to fewer good experiences. The tragedy of dying young is that one misses out on all that life has to offer. A fatal car-crash at twenty is far worse than a fatal fall down the stairs at ninety. The twenty year old misses out on the goods of seventy years of life.¹⁵ By dying at twenty, one would live a life of far less prudential good than had one lived to ninety. Hence, death is bad for the one who dies when it deprives her of future goods. As with all defenders of the deprivation account, Feldman and Bradley reject the causal hypothesis.¹⁶ I will focus on Feldman's account.

Feldman argues that the causal hypothesis is too simple to be an adequate theory of extrinsic badness. Although it accounts for a good number of cases of extrinsic badness, it fails to account for a number of clear examples. Consider Joe College:

Joe College

Joe is admitted to two colleges, College A and College B. Joe chooses to go to College A, which does not have a philosophy department. He earns an accounting degree, becomes a moderately successful accountant, and lives a reasonably good life. However, if Joe had chosen to go to College B, which does have a philosophy department, he would have taken a philosophy course, discovered that he had a passion for philosophy, pursued graduate study, successfully navigated the job market, landed a tenure track job, and lived a very good life. The life Joe would have led had he gone to College B would have been unambiguously better.¹⁷

Feldman argues that Joe's choice go to College A, although it did not lead to a bad life, was indeed bad for Joe. It was bad for Joe to go to College A instead of College B, because if he had gone to College B, his life would have been better. Since Joe's decision led to a life that was less good for him, his decision was bad for him. It was not intrinsically bad: Joe did not suffer. But it was extrinsically bad: Joe's life contained less intrinsic good.

However, according to the causal hypothesis, Joe's decision to go to College A was not bad for him. The causal hypothesis states that nothing can be

¹⁵ I am assuming that Schopenhauerean pessimism is false.

¹⁶ Steven Luper defends a preclusion view of harm that amounts to a rejection of the *causal hypothesis*. Since Luper's argument assumes desire-satisfactionism about welfare, it would take me too far afield to discuss his provocative theory. Luper, "Posthumous Harm."

¹⁷ This is a rough approximation of Feldman's example. Feldman, *Reaper*, p. 137.

Section 3. The Deprivation Account

extrinsically bad for someone unless it is causally responsible for more intrinsically bad states of affairs than intrinsically good states of affairs. And Joe's life was not net intrinsically bad as a result of his choice in schools. But, as the example purports to show, the decision was indeed bad for Joe; it was bad because it led to a life of less intrinsic good. Feldman argues that since the causal hypothesis cannot account for the badness of Joe's choice, it is too exclusive as an account of extrinsic badness.

In response to such difficulties, Feldman offers the following account of extrinsic badness, which he calls "EI":

EI: Something is extrinsically bad for a person if and only if he or she would have been intrinsically better off had it not taken place.¹⁸

Feldman argues that EI does a better job of capturing what it is for something to be extrinsically bad for a person. Unlike the causal hypothesis, EI can account for cases like that of Joe College. It was bad for Joe to go to College A, because he would have been better off had he chosen to go to College B.

As additional support for EI, Feldman presents a structurally similar, but intuitively more compelling, case of a girl raised under fundamentalist Islam, a case that I will call the "Taliban Girl."

Taliban Girl

The Taliban Girl is raised in a repressive, fundamentalist Islamic culture that forbids teaching women how to read. She lives a good life according to most plausible theories of welfare. However, if she had been allowed to learn to read, she would have developed a great love for poetry. Further, she would have become a fine poet. Her life would have been higher in prudential value. Hence, it was bad for the Taliban Girl to have been forbidden to learn to read.¹⁹

Feldman argues that EI, unlike the causal hypothesis, can account for why it was bad for the Taliban Girl to have been forbidden to learn to read. Since her life was not bad overall—the Taliban Girl led a reasonably happy life—her illiteracy did not lead to any intrinsically prudentially bad states. However, it was clearly bad that she was kept illiterate. Hence, the causal hypothesis cannot account for such cases of clear prudential bad. Accordingly, we should reject it as too exclusive.

As presented, the causal hypothesis is merely an account of extrinsic

¹⁸ Feldman, *Reaper*, p.138. In the chapter "More Puzzles about the Evil of Death," Feldman offers a related account of the value of a state of affairs for a person: "The extrinsic value for *S* of *P* = the difference between the intrinsic value for *S* of the life *S* would lead if *P* is true and the intrinsic value for *S* of the life *S* would lead if *P* is false" (p. 150).

¹⁹ This is a rough approximation of Feldman's example. Feldman, *Reaper*, p. 137. Ignore the fact that poets tend to live miserable lives.

Section 3. The Deprivation Account

badness, but it strongly suggests a symmetrical theory of intrinsic goodness: to be extrinsically good something must lead to more intrinsic good. Of course, one could accept the theory of extrinsic badness without the theory of extrinsic goodness, but the resulting position would be suspect. It is hard to imagine that such an asymmetrical theory could be anything but ad hoc. Anyone who holds the causal hypothesis should be prepared to defend both the theory of extrinsic badness and the theory of extrinsic goodness. But the theory of extrinsic goodness is vulnerable to a prima facie fatal counter-example, namely it implies that it is not good to anesthetize patients before they undergo surgery. Nor would it be good to give morphine to those suffering chronic pain.²⁰ In so far as it merely leads to less pain, morphine does not lead to any intrinsically good states of affairs. Neither does anesthesia during surgery. It merely leads to less bad. Hence, according to the causal hypothesis anesthesia is not good for those in surgery. This is clearly absurd. However, the theory of extrinsic goodness that is symmetrical to EI suffers from no such problem. Since anesthesia leads to less intrinsic bad, the analog of EI would hold that anesthesia is in fact good. That seems right. Accordingly, the objection holds, we should reject the causal hypothesis and accept EI.

For our purposes, what is important is that EI supports the deprivation account of death's badness. Just as Joe College's attending a school without a philosophy department deprived him of the intrinsic good he would have enjoyed had he gone to College B, death sometimes deprives one of goods that one would have enjoyed had one continued living. Hence, death can be extrinsically bad for the one who dies.

When exactly is it bad?²¹ Bradley offers a refinement of Feldman's EI, specifying the precise time in which something is bad for a person. His formulation is as follows:

OVT: The overall value of p for s at $\langle w, t \rangle$ = the intrinsic value of t for s at w minus the intrinsic value of t for s at the nearest world to w at which p does not obtain.²²

According to OVT, death is bad for the one who dies at the times in which the person's life would have been higher in intrinsic value. If one's life would have been filled with suffering, dying would not be bad for the one who dies. But if one would have lived a life of positive intrinsic value, death is bad for the one who dies at those times at which their life would have had positive intrinsic prudential value in the nearest possible world. For our purposes, the full set of implications of this subtle refinement of EI is not important. Both EI and OVT

²⁰ Steven Luper raises a similar example involving a temporary coma. See: Steven Luper, *The Philosophy of Death* (Cambridge UP, 2009), p. 104.

²¹ Bradley's account attempts to solve the "timing puzzle". Ben Bradley, "When Is Death Bad for the One Who Dies?," *Nous* 38:1 (2004): 1-28.

²² Bradley, "When", p.9. He presents a similar formulation in "How Bad is Death?," *Canadian Journal of Philosophy* 37.1 (2007): 111-128, p. 115.

hold that death is extrinsically bad for the one who dies, if it leads to less intrinsic value for the person.

If we accept EI or OVT, we should reject the fourth premise of the *Dead End Argument*—the causal hypothesis. Without the causal hypothesis the argument is unsound. Hence, the Epicurean argument for innocuousism stands or falls on the truth of EI and OVT style accounts of extrinsic badness.

In the following section, I intend to show that both EI and OVT have the same problematic implication, namely they confuse states that are merely less good with those that are bad.

4 Defense of the Causal Hypothesis

Although Joe College and the Taliban Girl suggest that the causal hypothesis is inadequate as an account of extrinsic badness, I intend to show that these examples are ultimately unconvincing. In addition, the accounts of extrinsic badness expressed by Feldman's EI and Bradley's OVT suffer from a fatal problem—they confuse states that are merely less good with those that are bad. I offer a few thought experiments in objection to any theory with a similar structure.

Annette Chigurh

Walking down the street, you turn the corner and find a table blocking most of the sidewalk. Behind the table stands Annette Chigurh, the good sister of evil Antoine Chigurh (from *No Country for Old Men*). On the table sit two closed brief cases. Having recently lifted two million dollars from her brother, she plans to give the bulk to a lucky stranger. She flips a coin. "What are the stakes?" you ask. "Heads you get the briefcase on the side you call; tails the other. Both contain wads of cash." You call the right. The coin lands heads and the right briefcase is yours. To your delight you find \$100,000 inside. You scream "Woopee!," pick up your booty, and make your way home. Unbeknownst to you, the other suitcase contained \$1,000,000. Had you called "left," you would have walked away with an additional \$900,000.

The Annette Chigurh thought experiment is structurally identical to that of Joe College. In both cases a choice leads to a state of affairs that is less intrinsically good for a person than the state that would have resulted from the only other available alternative decision. Just as EI and OVT imply that Joe's choice of College A was bad for Joe, they also imply that your encounter with Annette Chigurh was bad for you. More specifically, choosing the box on the right was bad for you, since had you not made this choice you would have chosen left, in that case you would have walked away with \$900,000 more. So, it was bad for you because it led to a state of affairs that was less intrinsically good than the alternative. But this is absurd. Any call in a coin toss where you

Section 4. Defense of the Causal Hypothesis

simply walk away with \$100,000 would be fantastically good for you.²³ How could anyone think otherwise?

If this example is not compelling, consider another.

The Lottery

You walk into the local bodega on your birthday and the owner offers to give you a free lottery ticket. He hits the button to issue a Quick Pick—a randomly generated selection of six numbers. The next day you check the numbers and to your delight you got five of six. You won \$100,000! But with all six numbers, you would have walked away with \$1,000,000.

According to Bradley's notion of extrinsic badness, picking five of six winning lottery numbers is bad for the ticket holder. The set of nearest possible worlds includes a world, equidistant from the rest, where you picked all six numbers. If the ticket selection was by chance, every number was just as likely as any other. Since the world where you win the entire jackpot is in the set of the equally nearest possible worlds, and in that world you would be better off, merely winning a portion of the jackpot is bad for you. But that is absurd. Winning \$100,000 is fantastically good for the winner. Both EI and OVT imply that it was bad to have only won \$100,000 from a free lottery ticket. Hence, we should reject both accounts of extrinsic badness.²⁴

In reply the defender of OVT can say that winning 5 numbers is better than 4, and much better than 0, but it is not better than winning 6. Winning 5 compared to 4 is good; winning 5 compared to 6 is bad. It depends on how you look at the situation. And since there are millions of possible worlds where you win far less, as there are millions of lottery numbers that would have no matches at all, winning 5 is good for you according to OVT. It is not good simpliciter; rather, it is good compared to the majority of the other outcomes, and it is bad compared to a small fraction of the other possible outcomes.

But this is not a satisfying reply. There is just one event and one actual outcome. We want to know if the outcome is good or bad. Every outcome can be good or bad compared to something else that might have happened. But we do not want to say that every event is both good *and* bad. We want to know which. Restricting the comparative class to the nearest possible worlds does not help. The lottery example shows that a set of equidistant worlds contains one that is far better. But this does not make winning \$100,000 bad for you. Hence,

²³ This is to assume, of course, that the money does not bring some unexpected misfortune. Antoine will not come calling.

²⁴ Winning five out of six precludes winning six. The example could be revised slightly if one insists that to be analogous to death that we have a near event that precludes a temporally distant event. I could simply specify that the full winning would be paid out in installments. Winning the five precludes the future goods that would come from winning the six. Although designed as a counter-example to OVT, it also poses a problem for what Steven Luper calls "the preclusion view" of harm.

Section 4. Defense of the Causal Hypothesis

comparative theories of extrinsic badness have absurd implications. The defender of OVT cannot say that the outcome of an event is good or bad simpliciter. This is an unacceptable consequence.

The defender of EI or OVT might reply that that situation is not as bad as it appears. EI and OVT do not imply that a walk home where you gain \$100,000 is good for you, but that choosing a box with \$100,000 rather than one with \$900,000 is bad. That, the objection concludes, seems about right. Certainly, it is bad to have missed the chance to walk away with \$900,000 more. In response, two comments are in order. First, EI and OVT do indeed imply that a walk home where you get \$100,000 is bad for you. According to EI and OVT, it is bad if another walk home would have garnered an even greater booty. According to OVT, the walk is good compared with those that do not encounter Annette Chigurh, and it is bad compared with those where you call left. But, once again, this does not mean that the walk is bad simpliciter. This retraces my previous objection.

As a second line of reply, we should confine our attention to the decision to call right over left, as the objection suggests. The defender of EI or OVT might hold the more refined claim that the decision is bad. But this too is problematic. Even if we grant this restriction, the position is untenable. The choice was not bad; it led to a gain of \$100,000. Why would that be bad? It would have been better to call left, but this does not make calling right bad. The underlying problem is that EI and OVT both imply that for any event only one outcome could be unequivocally good for the person, the maximally good outcome. But this is wrong. Sometimes there are many, mutually exclusive *unequivocally good* outcomes to a single event. Just because one might be better, by either a wide or a narrow margin, does not make all the other possible outcomes bad. If the intrinsic value of the outcomes varies, then OVT will say that some are bad compared to the better outcomes, and good compared to the lesser outcomes.²⁵ But this is wrong: Merely being less good does not make something bad. Why would it? It is not just that EI and OVT cannot say that an event is good simpliciter, they say the wrong thing in a wide range of clear cases. To get at the heart of the problem, consider an event with only two outcomes. EI and OVT imply that only one of the outcomes can be good. But that is clearly wrong. Consider Joe Coffee:

Joe Coffee

Joe Coffee has a choice between two colleges, College A and College B. Joe chooses College A, majors in math, goes to graduate school, and lands a good job at Research University. However, had Joe gone to College B, which has a philosophy department, he would have majored in philosophy. Amazingly, after

²⁵ Feldman says something similar in regards to a happy move to Bolivia. In some cases, the move might be good and bad. Bad, not because it led to intrinsic bad, but because it prevented an even greater good. Fred Feldman, "Some Puzzles about the Evil of Death," *The Philosophical Review* C.2 (1991): 205-227.

Section 4. Defense of the Causal Hypothesis

graduate school he would have landed a good job at the same university. Joe would have found philosophy equally fulfilling as mathematics. His salary would have been the same, his colleagues equally congenial. However, the philosophy department at Research University has a cappuccino machine, whereas the mathematics department only has a drip coffee maker. Had Joe gone to College B, he would have enjoyed a tasty cappuccino every Wednesday. His life would have been a little better.²⁶

According to EI and OVT Joe Coffee's decision to go to College A was bad for him because had he gone to College B, he would have had better coffee. This is a troublesome implication. Yes, EI and OVT imply that it was only a little bad for Joe, but they still consider it a bad decision. However, both outcomes would be good for Joe; neither would be bad, not even a little. It may be less good in comparison, but it is clearly not bad. Although going to College A was not the maximally good option, as he missed out on free cappuccino, it was still very good. It is not bad, merely less good. Any theory that says otherwise should be rejected for having absurd consequences. This is not just a minor troubling implication; it is a significant problem that provides good reason to reject both EI and OVT. This is a bullet too big to bite.

Fear and the Bad

Perhaps a way to make it even clearer that EI and OVT mistakenly imply that good things are bad is to show that both accounts of extrinsic badness clash with our intuitions about the appropriateness of fear. It should make sense to fear something if it is bad. Surely, there are lots of bad things in the past that do not warrant fear, but it would be very odd to find an imminent event that was bad for you or yours that did not warrant fear. If an imminent event is bad for you or yours, it makes sense to fear it. Fear should "fit." It might not be, all things considered, best to fear it, but fear should at least be appropriate.²⁷ We typically evaluate the appropriateness of fear based on the badness of the event. If we find that putatively prudentially bad, imminent events are not the kind of things that we think are fitting of fear, then we have good reason to think that they are, in fact, not bad.

Reconsider the case of Annette Chigurh: She flips the coin, you call right, but you would not feel fear while the coin spins in the air. Fear does not fit a win-win situation. Sure, one may feel tremendous hope, but fear is not appropriate, not even if you knew the exact contents of the two boxes. According to EI and OVT, losing the toss would be very bad for you. In this situation, a putatively prudentially bad outcome is both imminent and likely. However, fear is inappropriate. This should strike us as very odd. The best explanation for the oddity is that, contra EI and OVT, neither alternative would be bad for you. One would merely be less good than the other. And it does not

²⁶ I thank Christy Mag Uidhir for suggesting this line of objection.

²⁷ Explaining "fitting" here would require solving the "wrong reasons" problem.

make sense to fear something because it might simply not be as good as another alternative. There is more to be said about the rationality of fear, but these brief observations give us independent reason to think that less good is not bad.

The unintuitive implications of EI and OVT are due to the fact that both accounts of extrinsic badness conflate states of affairs that are merely less good with those that are bad.

An Asymmetry between Less Bad and Less Good

One might balk at the suggestion that there is a genuine difference between events that are less good and those that are bad. Certainly we should expect a corresponding distinction to hold between states that are less bad and those that are less good. But we seem to think that states which are less bad are sometimes good. This asymmetry suggests that the distinction is strained. Consider the case of a bad fall.

Near Miss

Imagine that while working on the roof, you fall and break both arms. Your head lands just inches away from a rock. Had you fallen just a hair to the left, you would have hit your head on the rock and likely suffered severe brain damage. In the hospital, your friends tell you how lucky you are to have missed the rock. Indeed, it certainly seems like a good thing that you missed the rock. The result of the fall could have been much, much worse.

About such a near miss, we would want to say that you were lucky. It very well might have been much worse. Friends often make note of this kind of thing to cheer us up. We say "It could have been worse," but we seldom, if ever, say "It could have been better." This suggests that there is an asymmetry between how we think of situations that are less bad and those that are less good. If we think that less bad is good, then we should probably think that less good is bad. The objection concludes, this asymmetry suggests that the distinction I draw between less good and bad is artificial.

Yes, it is clear that we seldom say "It could have been worse," but this does not indicate an asymmetry between how we evaluate situations that are less good and how we evaluate those that are less bad. It is not that we think that less bad is good, only that we prefer people who look on the bright side. Few could tolerate the company of someone who after hearing good news always says: "Well sure, but it could have been better." Hence, there is no reason to think that the distinction between less good and bad is artificial.

So, yes, we might say that the fall could have been worse, but we certainly do not think that the fall was good for you. It was better than another fall, perhaps the same fall in the nearest possible world, but this does not make the fall good. You broke both arms! How could it possibly be good for you? It was most definitely extremely bad for you to fall off the roof. However, the

inverse of EI and OVT would entail that the fall was good. Whether we look on the brightside or the darkside, falling is bad and winning five of six numbers in a lottery is good. Sure, nearly anything that happens could have been better, but that does not make everything bad. The fall might be less bad in comparison to an even worse fall, but it is decidedly not good. The mere fact that things could have been worse does not make something bad good.

EI and OVT style accounts of extrinsic badness imply that breaking your arms is good for you and that winning \$100,000 is bad. Hence, EI and OVT have highly unintuitive implications, whereas the causal hypothesis tracks our intuitions. The causal hypothesis respects both the difference between the merely less good and the bad, and the difference between the merely less bad and the good. The causal hypothesis holds that something is good for you only if it leads to intrinsic good, bad only if it leads to intrinsic bad.

Perhaps EI and OVT have absurd implications, but, one might ask, how does this support the causal hypothesis? It supports the causal hypothesis because it shows that only a theory that can track the difference between less good and bad can avoid the absurd implications of EI and OVT. The causal hypothesis is such a theory.

Joe College and the Taliban Girl Reconsidered

But what about the problems raised for the causal hypothesis by the examples of Joe College and the Taliban Girl? I think that these, too, can be put aside. Neither provides a clear counter-example to the causal hypothesis. Consider Joe College: It is not at all clear that we want to say that his decision to go to College A was bad for him. Sure it was less good than the alternative, but why would we think that it was bad? Had Joe been granted a vision of the future, he would have seen that he had good reason to prefer college B, but this does not make his actual choice extrinsically bad. This is not to beg the question, but to merely deny the effectiveness of the thought experiment. Joe College never elicited the right reactions from me. But the case of the Taliban Girl is different. It is a *prima facie* compelling objection to the causal hypothesis.

Remember, the Taliban Girl was raised in a fundamentalist Islamic culture where she was forbidden to learn to read. Although she did not live a bad life, it is clear that enforced illiteracy was bad for her. However, since her illiteracy did not lead to any intrinsically bad states of affairs, the causal hypothesis would have us say that it was not bad for her. This is pretty hard to swallow. It seems that we need to reject any theory with such an absurd implication.

A few things can be said in reply. First, one might question the thought experiment itself. Perhaps we cannot play along. It is exceptionally difficult to imagine that the life of an illiterate woman living in a fundamentalist Islamic culture could be one high in prudential value. One might argue that we simply cannot accept the stipulation that her illiteracy did not lead to any intrinsically bad states of affairs. But I do not intend to take this route. I am willing to accept the thought experiment as it stands. Sometimes the indoctrinated and repressed do adopt the values of their oppressors. This is precisely Amartya

Sen's worry about using self-assessment as a means of measuring well-being.²⁸

So, I intend to take a second line of reply. Even if the Taliban Girl's enforced illiteracy does not lead to any intrinsically bad states of affairs, it would certainly put her in jeopardy of lots of intrinsic bad. It makes her dependent and ill informed. Both of which are clearly instrumentally bad in most situations. So, although her illiteracy did not lead to any intrinsically bad states of affairs, it threatened to. That is, it was dispositionally extrinsically bad. Her enforced illiteracy put her in a state that was prone to lots of intrinsically bad states. If we want to include such dispositional outcomes as relevant to whether something is extrinsically bad, as I think we will probably want to do, the causal hypothesis needs some refinement.²⁹

But, for my purposes, further refinement is not necessary. Death does not dispose us to suffer any intrinsically bad states of affairs. It disposes us to nothing. Any theory that respects the difference between states of affairs that are merely less good and those that are bad will have the same implications for death. At worst, death can be less good for the one who dies. It can never be bad, not even dispositionally extrinsically bad.

Third, it pays to note that the Taliban Girl thought experiment does not cleanly pump intuitions about well-being. Along with complicating concerns about the instrumental value of her life, it also arouses moral disgust. Illiteracy limits one's access to ideas. As far as exposure to a variety of ideas is conducive to self-direction, illiteracy curbs autonomy. And we find brainwashing, indoctrination, and other practices that are designed to limit the development of autonomy abhorrent. I suspect that this is not merely because we think that the autonomous are better off prudentially. No, we seem to value autonomy independently of welfare. Of course, whether we are right to do so is another matter. But for the purposes of assessing the effectiveness of the thought experiment, our prereflective moral commitments concerning the value of autonomy most plausibly account for some of our unease about the life of the Taliban Girl.

The Problem with Anesthesia and Morphine

In the preceding section, I presented three problems for the causal hypothesis—two for the theory of extrinsic badness and one for a symmetrical theory of intrinsic goodness. I argue that anyone planning to defend the theory of extrinsic badness should also be prepared to defend the theory of intrinsic goodness. An asymmetrical theory would have all the trappings of an ad hoc

²⁸ Amartya Sen, *On Ethics and Economics* (Oxford, 1987), pp. 45-6. This is also a serious liability of mental statism. In an intriguing effort to get around this problem, Sumner builds an authenticity clause into his subjective account of well-being. L. W. Sumner, *Welfare, Happiness, and Ethics* (Oxford UP, 1996), chapter 6.

²⁹ As stated, the causal hypothesis suffers from all sorts of problems that also afflict any known theory of causation. Since I do not have a satisfactory theory of causation to offer, I cannot spell out the causal hypothesis much more explicitly.

Section 4. Defense of the Causal Hypothesis

construction. The problem is that the causal hypothesis, as a theory of intrinsic goodness, suffers from a far worse problem than those facing the theory of extrinsic badness; namely, the theory implies that it is not good to give anesthesia to patients about to undergo surgery, or morphine to those in chronic pain. But this would be ridiculous. Certainly it would be very bad to deny someone anesthesia, and it would be good to share one's supply for those without. And, without doubt, it is good to give morphine to those in agony.

Let me first address morphine, since this objection can be countered without too much fuss. According to the causal hypothesis, as a theory of extrinsic goodness, it is not extrinsically good to merely reduce the amount of intrinsic prudential badness; it is merely less bad. Hence, it seems that the theory implies that giving morphine to those in agony is not extrinsically good. But this conclusion does not follow. Morphine does not merely allow those who are suffering to suffer less. In most circumstances it also allows for the opportunity for more prudential good. When one is in agony it is difficult to concentrate on anything but the pain. Although morphine dulls the mind, it can also free those suffering to think about things other than their pain. This can be intrinsically good. Hence, administering morphine also sometimes leads to greater intrinsic good. Therefore, according to the causal hypothesis, giving morphine to the suffering is often extrinsically good. When it comes to morphine, the causal hypothesis is not as counter intuitive as it might first appear.

The situation is more difficult when it comes to anesthesia. My reply may sound a bit awkward at first, but I find that it becomes far less uncomfortable after some consideration. The causal hypothesis correctly implies that it is indeed bad to deny a patient anesthesia before surgery. It is extrinsically prudentially bad because denying anesthesia leads to far more intrinsic prudential badness. Fear certainly fits a forthcoming surgery without anesthesia. But, conversely, the theory holds that it is not good to anesthetize patients; it is just far, far less bad. In fact, giving someone anesthesia may prevent there from being any bad at all. Nevertheless, this does not make it extrinsically good. It is not the kind of thing that leads to any intrinsic prudential good.

Of course, in many cases surgery would not be possible without anesthesia.³⁰ In these cases, anesthesia would indeed be extrinsically prudentially good, since the long term benefits of surgery would not have been possible otherwise. But there are undoubtedly other cases where surgery would be possible without anesthesia, just extremely painful. These are the putative problem cases for the *causal hypothesis*. The theory holds that denying anesthesia is bad, but giving it is not good, just less bad. In fact, I think that this is the correct thing to say. I'm not biting a bullet; I'm merely suggesting that we further ruminates on the situation at hand.

Although our prereflective intuitions may suggest otherwise, when we stop to consider the situation more closely, it is not ridiculous to say that anesthesia is not extrinsically good. It is merely less bad than the alternative. We would

³⁰ I could also reply that anesthesia is a surgical tool. It is better to have better tools. Better tools are good to have because they enable one to do both more good and less bad.

Section 4. Defense of the Causal Hypothesis

prefer to do neither. But if we have to perform surgery, it is better to put the patient under. It is better, but not good. When one comes down to it, the theory seems about right. Anesthesia is not extrinsically good; it is merely less bad. We anesthetize patients to prevent a rush of intrinsic bad, not to promote their good.

I do not mean that it would not be morally good to share a precious supply of anesthesia; I simply mean that taking anesthesia is not extrinsically prudentially good. At this point, I owe at least a cursory account for why it would be morally good to share anesthesia if it is not prudentially good. *Prima facie*, it seems odd to say that it is morally good to make someone less bad off, but not prudentially good for them. If morality smoothly tracks changes in prudential value, then the sharp distinction I draw between extrinsic good and extrinsic bad may be merely verbal. Happily, I think we can avoid this conclusion by taking a closer look at the situation.

We can account for the intuition that it is good to donate precious anesthesia this way: When one has a surplus of anesthesia, denying it from others would be extrinsically bad. It would lead to far worse intrinsically bad states of affairs. In such a case, it might even be morally obligatory to donate one's supply. Unlike death, in the case of surgery we are not faced with an alternative between a prudentially good state (living) and a prudentially neutral state (being dead); rather, we are forced to choose between a prudentially bad state (surgery without anesthesia) and a prudentially neutral state (surgery under anesthesia). Although giving anesthesia to a patient is not extrinsically prudentially good, not giving anesthesia is extrinsically prudentially bad. According to the causal hypothesis, our *prima facie* duty to provide anesthesia is a duty to help prevent imminent prudential bad. It is not a duty to do something extrinsically prudentially good. If we have anesthesia to spare, it is not superogatory, it is *prima facie* obligatory to give it to those who are about to undergo surgery. This seems correct.

It is not clear that we would have an obligation to provide anesthesia if it were merely extrinsically prudentially good. Most plausibly we have far fewer duties to promote the good than we have to diminish the bad.³¹ We have a *prima facie* obligation to diminish the imminent prudential bad of others, but no corresponding duty to increase the prudential good of others. If less bad were indeed good, and less good were indeed bad, we would have trouble accounting for the moral distinction. Hence, rather than suggest a merely verbal distinction, the causal hypothesis helps articulate the intuitive moral difference between reducing the bad and increasing the good.

³¹ Ultimately this claim might not be defensible, but it is at least plausible. Accordingly, it shows that the causal hypothesis does not lead to patently absurd results. Rather, the theory aids in making a plausible distinction. But it does not require the distinction.

5 Accounting for Counter Intuitions

Perhaps the most damning objection to innocuousism is that it is *prima facie* absurd. Most non-clinically depressed people are extremely reluctant to die. And we think that killing people for sport, even if the victims have no friends or family, is awful. However, if death is not bad for the one who dies, it seems that our reluctance to die and our intuitions about murder are irrational. But that would be absurd. Even if we cannot pinpoint the precise problem, any argument that implies that it is irrational to avoid dying is clearly wrong. When you have to pick between an obvious truth and a philosophical thesis, you should probably stick with the obvious truth. No premise could be grounded on intuitions stronger than these: it is perfectly rational to take steps to avoid dying and it is wrong to kill for sport.

I agree. If innocuousism implied either absurdity, we should reject the theory first and find reasons later. But innocuousism implies neither. I will not attempt to fully account for our reluctance to die or kill, but a sketch is required.

Reluctance to Die

The conclusion of the *Dead End Argument*—innocuousism—holds that death is never bad for the one who dies; it is only sometimes less good. But we must ask: If death is not bad for the one who dies, why are we so reluctant to die? Surely, desiring to live is not irrational. In reply, the defender of innocuousism could say that it is perfectly rational to want more good rather than less. When death would give us less, it is perfectly rational to want more.³²

But this does not get to the heart of the objection. We are not merely reluctant to die, putatively, we find death terrifying. If death is not bad for the one who dies, if it is merely less good, then our fear must be irrational. As I noted earlier, fear does not fit an event with an outcome that is merely less good than another. For fear to fit, a likely, imminent outcome must be bad. But if innocuousism is right, death does not lead to a bad outcome. So, according to innocuousism, it must be irrational to fear death. This is absurd, despite Epicurus's agreement.

In order to assess the strength of this objection, we first need to determine both what makes fear rational and whether death meets the criteria. Unlike a mood, as a paradigmatic emotion, fear takes objects.³³ We "fear that" and are "afraid about". One standard for assessing the rationality of an emotion is to

³² Sometimes it seems to be a good thing to put the suffering out of their misery. It is sometimes better for those in tremendous pain to die rather than to go on living; it might be less bad for them. But, just as it is not prudentially bad to die, neither is it good. Euthanasia is not good; at best, it is less bad. This helps account for our reluctance to euthanize our loved ones.

³³ I'm assuming a roughly cognitive theory of the emotions—that emotions are concern based construals.

evaluate the perceived fittingness of the object to the response.³⁴ To say that fear fits an object is to say that fear is appropriate to that object. This does not mean that one should feel fear, only that fear is permitted.³⁵ On this standard, fear is only rational when it is directed at an object thought to be fitting to the response. Conversely, we would call fear directed at an object thought to be unfitting to the response irrational. This picture is inadequate, but for present purposes it will do.

So, what makes fear appropriate? When does fear fit an object? It seems that at least one necessary condition is that the object be capable of harming me or mine, otherwise fear does not fit. We can divide fear into two general types: fear for self and fear for others. Fear for self is only appropriate if the object is capable of harming oneself. If we do not think that an object is capable of harming us, then fear is inappropriate. Most plausibly, phobias are irrational because they lack this feature. If you know that common house spiders are incapable of hurting a person, but fear them nonetheless, then your fear is inappropriate. The phobic's fear takes an object that she does not think is harmful. Hence, the phobic's fear is irrational.

Harmfulness is not sufficient for fear to be appropriate to an object. Harm also needs to be likely. If I developed a fear of komodo dragons, we might find this odd, as I live in a region free of such monsters. But if I was going to vacation on an East Asian island covered in the giant lizards, then we might think that my fear is appropriate. Komodo dragons can run thirty-five miles an hour and their saliva is replete with bacteria that invariably cause fatal infections. And, soon I'm going to be very near some! As this example shows, not only does the object need to be capable of harming me or mine, the harm needs to be likely, if not imminent. We need not concern ourselves with just how likely. These two conditions appear to be necessary and jointly sufficient to make fear fitting.

Since everyone will indeed die, death meets the second condition. What about the first? When we consider the possible bad things that might happen to us, death certainly ranks among the top. Dying appears to be one of the worst things that could happen to anyone. But innocuousism implies that death is not bad for the one who dies. If so, then it is inappropriate to fear death, other than

³⁴ Another standard would be to evaluate the emotion in terms of its conduciveness to promoting our self-interests. Many in the rationality of fiction debate have appealed to this style of evaluation. I find it wanting. For an excellent critical overview, see: Berys Gaut, *Art, Emotion, and Ethics* (Oxford UP, 2007), pp. 216-227.

For more on the rationality of fictional emotions see: Colin Radford, "Fiction, Pity, Fear, and Jealousy," *Journal of Aesthetics and Art Criticism* 53 (1995): 71-75, p.75; Derek Matravers "The Challenge of Irrationalism, and How not to Meet it," in *Contemporary Debates in Aesthetics and the Philosophy of Art*, ed. Matthew Kieran (Blackwell, 2005); and Richard Joyce, "Rational Fear of Monsters," *British Journal of Aesthetics* 40.2 (2000): 209-224.

³⁵ One must be careful not to elide this distinction. "Fit" and "appropriate" are too often used in an equivocal manner, taking them as merely permitted in some contexts and required in others.

for the impact our absence might have on those that care about us. This must be wrong. Indeed, I think it is. But not because I think that innocuousism is wrong; rather, I think that the theory of the rationality of fear is mistaken.

Perhaps some of us fear death because we think that it would be prudentially bad. Innocuousism undermines this basis for fear, but it does not imply that we are wrong to fear death. We fear death because we care about more than just our own welfare. We also care about things such as significance. This is precisely what the experience machine shows—not that we think that it would be better for us to be in touch with reality, but that it would be better. Genuine achievements are among the kinds of things that we desire. People adopt all sorts of projects that they know might make them less happy than some other less ambitious pursuits; indeed, we frequently adopt projects that risk prolonged frustration.³⁶ There is nothing irrational about this. Certainly it is rational to care about more than just one's own self-interest. How could anyone think otherwise? Nietzsche's familiar quip expresses the point nicely: "Man does not strive for pleasure; only the Englishman does."³⁷

We care about our projects and the overall shape of our lives. Although it may not be prudentially bad to die, it can be bad in a variety of ways. Depending on the state of our projects, much of what we want to accomplish would be unrealized. When death would thwart what Bernard Williams calls our categorical desires—those that give us reason to go on—we do not want to die.³⁸ But, if one lacks such desires, one welcomes death; one does not fear what brings relief from boredom.³⁹ Although death is not bad for the one who dies, worrying about death certainly is. This does not mean that it is irrational to worry.

It is rational to care about more than one's own good—to care about things such as the instrumental value of one's life. Sometimes we rationally hold the significance of our life in greater esteem than our own good. Achilles' choice illustrates the point. Achilles knew that if he entered the battle against the Trojans that he would die soon, but that his name would live on for ages; however, if he stayed out of the fight, his name would be forgotten, but his life would be long and happy. He chose a short meaningful life over a long life high in individual welfare. Regardless of whether we think that glory on the battlefield is a "project of worth," given his values, Achilles' decision was

³⁶ Victor Frankl goes so far as to suggest that we have a basic drive for meaning. Victor Frankl, *Man's Search for Meaning* (Beacon Press, 1959).

³⁷ Nietzsche, *Twilight of the Idols* ("Maxims and Arrows", #12)

³⁸ Bernard Williams, "The Makropulos Case".

³⁹ Or, perhaps, as Walter Kaufman argues, "If one loves intensely, the time comes when death seems bliss. [. . .] The life I want is a life I could not endure in eternity. It is a life of love and intensity, suffering and creation, that makes life worthwhile and death welcome." Walter Kaufman, "Existentialism and Death," in *Existentialism, Religion, and Death* (New American Library, 1976), p.214.

Section 5. Accounting for Counter Intuitions

eminently rational.⁴⁰ It is perfectly rational to care about our lives and not just our welfare.⁴¹

Although we do indeed fear death, I suspect that the extent of this fear has been far overstated. Typically, or at least most significantly, the reaction to death is not so much fear as it is sadness. We are not so much terrified of dying, as we are saddened at the prospect.⁴² Epicurus is wrong to say that "Whatever causes no annoyance when it is present, causes only a groundless pain in the expectation". Death pains in the expectation because we do not want to become nothing at all. It is not that we fear being nothing.⁴³ As Lucretius notes, we are not afraid of the period prior to our birth.⁴⁴ No, we are horribly saddened at the prospect of never enjoying the pleasures of being alive: holding our loved ones, feeling the sun on a spring day, or simply eating the "crabs at Sam Woo's."⁴⁵ Nagel aptly notes that:

Perception, desire, activity, and thought, are so general as to be constitutive of human life. They are widely regarded as formidable benefits in themselves, despite the fact that they are conditions of misery as well as of happiness, and that a sufficient quantity of more particular evils can perhaps outweigh them. That is what is meant, I think, by the allegation that it is simply good to be alive, even if one is undergoing terrible experiences.⁴⁶

The prospect of never experiencing anything again saddens us. We do not merely feel sorry for ourselves. Thoughts of leaving home hurt, even though once gone we may never think about those that we left behind. Regardless of whether we genuinely believed that death would be fantastically good, that we were to go on to a heavenly existence, we would still rightly be saddened by the prospect of leaving this life behind.⁴⁷ It is not that death is bad for us, but that

⁴⁰ Susan Wolf argues that meaning in a life is a result of active engagement in projects of worth. Wolf, "Two Aspects."

⁴¹ Kagan makes a distinction between me and my life. I think that this is best expressed in terms of different kinds of value. Our lives can have value for us, and they can also have instrumental, aesthetic, and moral value.

⁴² Kai Draper considers the view that our response to death is typically one of pronounced sadness. Draper, "Disappointment, Sadness, and Death."

⁴³ Unamuno would disagree: "I must confess, painful though the confession may be, that in the days of the simple faith of my childhood, descriptions of the tortures of hell, however terrible, never made me tremble, for I always felt that nothingness was much more terrifying." See: Miguel de Unamuno, *Tragic Sense of Life* (Cosimo, 2005), p.43.

⁴⁴ Lucretius, "On the Nature of Things."

⁴⁵ *Manhattan* (Woody Allen, 1979).

⁴⁶ Thomas Nagel, "Death", p. 2. Nagel expresses similar sentiments in "Birth, Death, and the Meaning of Life" in *The View from Nowhere* (Oxford, 1986). Stuart Rachels offers a compatible criticism of Nagel. He argues that it is not experience itself, but "a background of good feeling [that] supplies the additional weight" (256). See Stuart Rachels, "Six Theses about Pleasure", *Philosophical Perspectives*, 18, *Ethics*, (2004): 247-267.

⁴⁷ For the sake of argument, put aside the possible problems with immortality.

we want to continue living for as long as we enjoy life.

Killing for Sport

Although torturing and taunting the weak for sport might be morally bad, it does not come close to the badness of recreational murder. We do not simply think that killing people is bad because it upsets the victim's friends and family. It does not improve things much morally to target the reclusive homeless. Either way, killing for sport is morally hideous. However, innocuousism seems to imply that recreational murder is morally permissible. If death is not bad for the one who dies, why would it not be permissible to quickly and painlessly kill hermits? If innocuousism suggests that human blood sport is permissible, we should reject it.

Happily, innocuousism implies no such absurdity. Innocuousism alone does not imply that murder is not bad. To reach this conclusion, one must also accept *welfarism*—the view that normative ethics should only be concerned with welfare, that prudential value is the only source of moral obligation. But if we reject welfarism, and we think that normative ethics may have to respect more than prudential value, then this absurd implication can be avoided.⁴⁸

Even if one accepts welfarism, innocuousism would not necessarily imply that killing for sport is morally permissible. The theory holds that death is prudentially innocuous; it does not hold that death is prudentially neutral. Although it is not bad, it is often less good to die. Even though it might not be rational to fear what is merely less good, and we may have fewer moral obligations to prevent things that are merely less good than we do to promote things that are less bad, this does not imply that it is morally permissible to actively promote less good.

⁴⁸ A defense of antiwelfarism is far out of scope. For an excellent overview of the arguments for and against, see Sumner, WHE, ch. 7. See also: Simon Keller, "Welfarism," *Philosophy Compass* 3 (2008).

6 Conclusion

In this article, I provide a defense of innocuousism—the position that death is never prudentially bad for the one who dies. I present an Epicurean argument against the badness of death that I call *The Dead End Argument*. The argument rests on three controversial theses: mental statism about welfare, the claim that death is an experience blank, and the causal hypothesis. I assume the first two and defend the third. The causal hypothesis holds that to be extrinsically bad something must lead to intrinsically bad states of affairs; merely leading to a state that is less intrinsically good is not sufficient. Less good is not bad.

Defenders of the deprivation account of the badness of death reject the causal hypothesis. The deprivation account holds that death is extrinsically bad for the one who dies when it leads to less intrinsic prudential good than living. I defend the causal hypothesis from the criticism of both Feldman and Bradley. I argue that Feldman's EI and Bradley's OVT lead to absurd implications because they confuse states of affairs that could have been better with those that are bad—that is, they confuse less good with bad.

I argue that the causal hypothesis is a better account of extrinsic badness, since it respects our intuitions about several thought experiments: Annette Chigurh, the Lottery, Joe Coffee, and Near Miss. Unlike EI and OVT, the causal hypothesis does not imply that winning \$100,000 is bad for you, nor does it imply that breaking your arms is good. In addition, the causal hypothesis better tracks our intuitions around the fittingness of fear. Fear fits the imminent bad, but not the merely less good. Further, the causal hypothesis tracks the prima facie moral difference between diminishing the bad and promoting the good.

Despite its virtues, innocuousism seems to imply two absurdities: that we are irrational in our reluctance to die, and that it is morally permissible to murder for sport. I argue that innocuousism implies neither. We are reluctant to die because we fear for our projects, and we are saddened by thoughts of becoming nothing. It is rational to fear for more than our own well-being. And it is rational to feel sad at the prospect of departing from the objects of our love. Neither rationally requires that death is bad for the one who dies. This is good, because, at worst, death is merely less good than living.⁴⁹

⁴⁹ Heidi Bollich patiently helped me develop my argument and raised several considered objections. I thank Christy Mag Uidhir for very useful feedback on a previous draft of this paper, and for suggesting a few ways in which I could improve my position. I also thank Steven Luper for extremely prompt feedback, which forced me to raise the morphine objection up from the footnotes. I thank Alan Rubel for pressing me on several difficult counter-examples. Ben Bradley kindly provided helpful feedback on an earlier draft. And Chris Heathwood made several suggestions for improving an earlier draft.