

Meaningless Happiness and Meaningful Suffering

Troy Jollimore
California State University, Chico

1.

The theory I will refer to as the eudaimonistic theory of value consists of three claims: (1) that all happiness is intrinsically valuable, (2) that all suffering is intrinsically disvaluable (i.e., possesses negative value), and (3) that nothing else is either intrinsically valuable or intrinsically disvaluable. By 'happiness' and 'suffering' I mean to refer to certain kinds of subjective states. (Some people, by contrast, mean by 'happiness' something like overall human flourishing.) Thus, the theory could be called the *subjective* eudaimonistic theory of value. However, I will leave the qualifier 'subjective' implicit in what follows.

Combined with an agent-neutral consequentialist ethical framework, the eudaimonistic theory of value results in what I will term eudaimonistic utilitarianism: the theory that moral agents ought always to do whatever they can to maximize the amount of happiness in the world and minimize the amount of suffering.

Robert Nozick's famous thought experiment, the Experience Machine, poses a problem for the eudaimonistic theory of value and, thus, for eudaimonistic utilitarianism.¹ The Experience Machine simulates the sorts of experiences that are had by people who are living actual lives. The simulations are imagined

Troy Jollimore received his Ph.D. from Princeton University in 1999 and is currently Assistant Professor of Philosophy at California State University, Chico. His book, Friendship and Agent-Relative Morality, was selected by Robert Nozick for Garland Publishing's Studies in Ethics series and appeared in 2001. His areas of current research interest include the nature of special obligations, philosophical questions relating to love and loyalty, and the internal-external reasons debate in practical reasoning. His publications include "Friendship without Partiality?" (Ratio 2000) and "Beauty, Evil, and The English Patient (Philosophy and Literature 2004), co-authored with Sharon Barrios.

to be so realistic that the people experiencing them cannot, in fact, tell the difference. Given the option of being hooked up to an Experience Machine that would provide us with a lifelong stream of very pleasant and fulfilling experiences, would we choose such an existence? The fact that most of us answer no seems to indicate that we believe, contra the third claim of the eudaimonistic theory of value, that other things besides happiness possess intrinsic value.

Moreover, the Experience Machine also seems to give evidence against the first claim: that all happiness possesses intrinsic value. The value of some types of happiness seems to be largely or even entirely dependent on certain external circumstances that are connected (or are thought by the subject to be connected) to that happiness. Suppose that Lise believes that something good has just happened in her life. She might believe, for instance, that she has just achieved some valued personal goal, such as running a 4-minute mile or having her novel accepted for publication. Or perhaps she believes that her daughter has just accepted an exciting and rewarding job or that her son has just received a prestigious scholarship. Any of these beliefs, let us suppose, would cause Lise to experience a feeling of happiness. But now suppose that in fact, Lise is hooked up to an Experience Machine and is only experiencing the program it is running for her: thus she has *not*, in fact, run a 4-minute mile, nor has she written a novel; and as for children, she has been hooked up to the machine since her own childhood, and so has none. Not only is it clear that Lise's happiness is *less* valuable in such cases than it would be if her beliefs reflected reality, it seems reasonable to doubt whether such experiences of happiness, being radically disconnected from the states of affairs in which Lise believes them to be grounded, would in fact be valuable *at all*.²

Thus, the Experience Machine seems to provide at least some reason for doubting two of the three tenets of the eudaimonistic theory of value. As of yet, however, we have seen no reason to think that it provides any evidence against the remaining claim: that all suffering is intrinsically disvaluable. One could try to challenge this third claim by arguing that what the Experience Machine shows is that, relative to a person who is not living an actual life, *nothing* is really either valuable or disvaluable. But this claim is implausible. Consider a person hooked up to an Experience Machine designed by a sadist, who is made to believe that she is being physically tortured. Given that the pain of the simulation feels just as real, and just as distressing, as pain brought about by means of actual physical torture, it is not clear that her situation is significantly less bad than that of an "actual" torture victim, and it would be highly implausible to say that her experience is not intrinsically disvaluable *at all*. The negative value of at least some suffering,

then, does seem to be intrinsic to the subjective state and not dependent on the external circumstances in the way that the positive value of (at least some) happiness is.

These considerations suggest the existence of an asymmetry between happiness and suffering, considered as subjective states. Happiness is not always valuable; it is to be valued and promoted only when certain other conditions obtain (when, for example, it is reflected in the objective reality of the individual's life). Thus there can be no blanket duty to promote happiness; rather, any provisional duty to that effect must be accompanied by an account of the circumstances under which the duty will and will not apply. But suffering is different: the mere existence of suffering as a subjective experience is in itself disvaluable, regardless of the relation that experience might bear to external circumstances. Suffering, then, is always disvaluable, and there is always (some) reason to remove it, regardless of its relation to the external life circumstances of the person who experiences it.

2.

It is presumably for reasons such as the foregoing that people who are otherwise not attracted by eudaimonistic accounts of normative ethics may nevertheless accept the existence of a moral duty to prevent or eliminate all suffering. Given consequentialist presuppositions—that moral action is a matter of promoting the intrinsically good and minimizing the intrinsically bad—the existence of such a duty (a *prima facie* duty, at any rate) follows directly from the thesis that all suffering is intrinsically bad. Nor is the appeal of such a position limited to consequentialist frameworks; deontologists, too, may be tempted to include a duty to eliminate suffering in one's list of *prima facie* duties, to add a right to have one's suffering relieved to one's list of fundamental human rights, and so forth. Let us call the thesis that there is a blanket duty to eliminate suffering (i.e., a moral duty to eliminate *all* suffering, regardless of its nature) the Elimination Thesis.

There is an important obstacle standing in the way of accepting the existence of such a duty, but it is one that many moralists have thought could be overcome. The obstacle arises from the fact that in some cases in which suffering could be prevented or eliminated, doing so does not strike us as the morally right thing to do. Indeed, the situation is worse than this, for as we will see in a moment, there are at least some cases where it is precisely the suffering itself that seems to be morally called for.

The mere existence of situations in which morality requires something other than the elimination of suffering, where such an elimination would nonetheless be possible, is not in itself difficult to reconcile with the Elimination Thesis. We need only

point out that the duty to eliminate all suffering is not the *only* duty, and that it can be overridden when something more morally significant or urgent is at stake. A theory of *prima facie* duties such as that suggested by W. D. Ross is one sort of theory that can handle this.³ Rights-based theories seem to do just as well, so long as one holds that rights are not absolute, but can override one another in cases of conflict. Nor does this phenomenon seem to pose an insurmountable difficulty for consequentialist theories. Such theories may well hold that suffering always possesses negative value but that there are other sorts of outcomes that can potentially possess even greater negative value, and which agents may therefore be under an all-things-considered duty to prevent, even when the cost of doing so is that some amount of preventable suffering is allowed to occur. The Elimination Thesis, then, seems to survive this alleged difficulty. Any serious attempt to delineate our moral duties is going to have to provide an account of what happens when these moral duties conflict, and as yet we have seen no reason to expect the completion of this task to give results incompatible with that thesis.

A related objection to the Elimination Thesis arises from the fact that suffering of certain sorts, despite its unpleasantness, often seems to play a positive function for the subject who experiences it. People do not usually like to feel pain, but the inability to feel pain would be even worse, for pain is often an indicator of a potentially serious disorder or injury that, if left unattended, could interfere with one's enjoyment of one's life and, perhaps, even become life-threatening. Thus, we would not choose to eliminate people's capacity for experiencing (this sort of) pain, even if doing so were to become an option; and we certainly would not consider ourselves to be under a *duty* to remove this valuable capacity. Again, however, this does not demonstrate that pain is not bad *in itself* but only that there are outcomes in which pain features that are preferable, all things considered, to available alternatives in which our ability to feel such pain has been erased. (If, on the other hand, one could somehow eliminate a person's ability to feel pain while at the same time providing her with a non-unpleasant "damage indicator" that would fulfill all the useful functions of pain, it is not immediately obvious that one would not have powerful moral reason to do this.⁴)

The most serious challenge to the Elimination Thesis along these lines arises from those cases in which the prevention or elimination of suffering would be morally wrong, not because the value of the pain-free state of affairs is simply overridden by the value of a superior, albeit pain-laden, alternative but because the suffering itself seems to possess positive moral value. If, for example, it is morally appropriate, as some have claimed, that the guilty suffer for their crimes, and if this kind

of suffering is in fact a necessary component of the reform of such individuals, then it would seem to be the case neither that all suffering is disvaluable nor that there is a blanket moral duty to prevent or eliminate suffering. The most promising defense of the Elimination Thesis in the face of this objection attempts to assimilate this sort of case to the type of case previously discussed. The claim, then, is that the suffering *itself* does *not* bear positive value but that it enables or brings about outcomes that are sufficiently valuable to justify the intrinsic disvalue of the suffering itself. Thus Jamie Mayerfeld, who defends a blanket moral duty to eliminate suffering, argues that despite appearances, the value of judicially inflicted suffering can be reduced either to the value of the moral knowledge or insight gained by the person who suffers, or to the value of the redemption that the suffering, when successful, brings about. "Perhaps the overall package is worth having; perhaps the knowledge and virtue are worth their cost in suffering," Mayerfeld writes. "But the suffering *is* a cost. In itself, it seems to me, it cannot be seen as anything other than evil."⁵

Mayerfeld's argument is not conclusive, for it is not clear that all of the value attached to judicially inflicted suffering can be assimilated to its educative and redemptive functions. Many people, after all, would follow Kant in holding it appropriate that a guilty party suffer even where there is no chance of his being educated, reformed, or redeemed. I do not, however, wish to pursue this line of argument here. Instead, I wish to argue that, regardless of the outcome of the question regarding judicially inflicted suffering, the Elimination Thesis must be rejected. For there is a different sort of suffering—one that is typically neglected by defenders of the Elimination Thesis—that is not clearly disvaluable and that does not, as so much other suffering does, directly give rise to reasons for its own elimination.

3.

The defense of the Elimination Thesis we have been considering claims that, when the value of suffering in itself is abstracted away from the value of potentially positive aspects of that suffering, the former value will always be found to be negative. Such a defense presupposes that it *is* in fact always possible, at least in principle, to abstract the positive aspects of suffering away from the suffering itself. This in principle possibility is in turn largely established by insisting that those positive aspects be viewed as *effects* of the suffering rather than as inherent qualities of it. Effects, after all, are always conceptually distinct from their causes, even when, from a pragmatic point of view, they cannot actually be prised apart.

In itself, the claim that every potentially positive aspect of suffering can be interpreted as an effect of the suffering may

strike us as highly plausible. What other sort of positive aspect, we might ask, could suffering possibly possess? It is certainly not the case that people tend to experience suffering as pleasant or desirable. Indeed, it is quite plausible to think that the quality of unpleasantness or undesirability is in fact built into the very definition of suffering. And given that the sufferer's attitude toward her suffering is thus naturally negative and, indeed, antagonistic—the natural and appropriate response to suffering is to want it to go away—it is hard to see how such a mental state, regardless of what positive extrinsic effects it might help bring about, could possess any sort of *intrinsic* positive value.

Nevertheless, this reasoning is too quick. There are after all many examples of mental states that are typically unpleasant to possess but that are not generally looked upon as states that ought to be eliminated. I am thinking here not of sensations but rather of *beliefs*—in particular, such painful, unpleasant, or challenging beliefs as, for instance, the belief that we all someday die. Many of us find it unpleasant to contemplate this fact and suspect that we might well be happier if we were not aware of it or could somehow stop believing it; nevertheless, while we are perhaps willing to take steps to allow us to live somewhat more comfortably with this knowledge, most of us would not be tempted to eliminate the belief altogether if this could be accomplished by means of, say, drugs, therapy, or surgical intervention. The question of the propriety of a given belief is settled not by determining whether it is pleasant but, rather, by determining whether it is true. We might add in this particular case that it is, after all, an important fact about human beings that we are mortal; and if the option of living in denial strikes us as easier and less unpleasant, it also strikes many people as cowardly and inauthentic.

Similarly, it might be difficult to face up to the fact that one's marriage, job performance, or overall satisfaction with life has deteriorated and easier simply to avoid acknowledging such harsh facts; yet even if we suppose that the long-term consequences of such repression would be a net decrease in our unhappiness (this is, of course, more plausible in some cases than others), nevertheless many people would again feel that the truth ought to be faced. Likewise, facing the truth about one's own past behavior or moral character in general can in many cases be highly unpleasant and sometimes promises few if any rewards in terms of happiness or inner peace; again, though, this does not strike many as a convincing case for avoiding the processes that lead to true beliefs on these subjects, still less for eliminating those true beliefs one might already have formed.

This seems to establish at least that there are mental states—in these cases, certain types of belief—the possession of

which is typically unpleasant but that we are justified in possessing (and unjustified in eliminating) nonetheless. This is not yet sufficient, however, to give us an objection to the Elimination Thesis, for it does not seem that these beliefs themselves constitute suffering. The belief that one will eventually die may well *cause* anguish; but if the anguish is a distinct mental state from the belief, then it is open to the defender of the Elimination Thesis to claim both that the value of the former can be distinguished from the value of the latter and that the (epistemological) considerations in favor of forming or keeping the belief may be sufficient to outweigh the (moral) considerations in favor of eliminating the anguish (given that the two cannot, in practice, be separated). What is needed, then, is an example of a mental state that is *itself* a species of suffering but that is also, like belief, a cognitive state and one whose evaluation ought thus to be conducted in terms that are not essentially connected to the issue of the state's pleasantness or unpleasantness.

I propose that grief is such a state. Grief is clearly a species of suffering, as anyone who has experienced it will attest. (Note how misleading it would be to suggest that grief was some *other* sort of state that simply tended to *cause* or *lead to* suffering, and how unnatural it would be to hold that someone could grieve *without* suffering.) Yet grief is also a cognitive response to an objective circumstance in the world. That it is so is partially indicated by the fact that there are standards of propriety for grief: one can be criticized either for not grieving enough or for grieving too much, in response to some specific loss. An episode of grief, like a belief, must be responsive to the facts of the world. This is not, of course, to equate the two: grief is not the same thing as belief, nor is it a type of belief. But the former, like the latter, is a type of cognitive response.

Imagine a person, Melissa, who is incapable of grief and who in particular fails to feel much of anything when her best friend, Bob, dies in an accident. (We must be careful to imagine that her response is genuine and not that she is actually repressing and grieving on a deeper, hidden level.) When we ask her why she is responding this way, she replies, "I know people generally feel terrible when these things happen. But eventually they get over it and realize that things aren't so bad. Life is short: why not start feeling good now? I'll miss Bob, it's true. But I'll make new friends soon enough, and they will be able to provide me with all the goods, emotional and otherwise, that Bob used to provide. In fact, like most people in our society I already have more than enough friends; I can afford to lose one or two."

I would suggest that Melissa, if we take her to be sincere, is displaying a serious cognitive failure. To suggest this is not to say that it is *only* a cognitive failure; we might equally well

cite deficiencies in Melissa's emotional response and, perhaps, in her character in general. (None of this should suggest that separating moral, cognitive, and emotional elements of Melissa's response will be easy, or even that it will be possible. Indeed, cases such as this illustrate very well how closely, conceptually speaking, these elements are tied together.) Nor is it to deny that the strictly *factual* claims cited by Melissa in defense of her response are largely if not entirely true. People who experience grief *do* generally get over it, start making new friends, and eventually return to a more normal emotional state. The failure manifested by Melissa occurs *despite* the fact that she is correct in these claims; it appears, not in her judgments and predictions about how people who lose friends feel and act over the long term but, rather, in her failure to understand *what it is to lose a friend*. But to describe Melissa's failure in these terms is to make it clear that her defect is a cognitive one. A failure to grieve, where grieving is appropriate, indicates a misunderstanding of the nature of the world one lives in; it signals a failure to comprehend the magnitude and significance of the loss that has occurred. Such a failure thus indicates a lack of understanding of the value of friends and, thus, of human beings.

Of course, the claim that the failure to feel appropriate grief constitutes a *cognitive* failure will be highly controversial. Grief is, after all, an emotion, and it is commonly assumed that emotions are to be analyzed in terms that are predominantly if not entirely noncognitive. This assumption, however, is more often asserted than argued for, and consideration of cases such as Melissa's seems to cast a certain amount of doubt on it. At the very least an independent argument is needed. One initially promising strategy is to object that grief cannot be considered a cognitive failure, since two agents who are cognitively identical might not grieve identically. This argument has been made by Stephen Wilkinson:

[Suppose that] two women, *A* and *B*, are faced with exactly similar losses. *A* goes through normal grief; *B* does not. *B*'s failure to grieve is not based on irrational beliefs or desires. She has an accurate picture of her situation (as does *A*). All that differentiates her from *A* is that she fails to have a normal emotional response. In particular, she fails to undergo the kind of mental suffering usually associated with grief. Is *B* less rational than *A*? ... I think that the answer to this is 'No'. Our concerns and criticism would not be about *B*'s (lack of) rationality, but about *B*'s mental health, or about the kinds of undesirable character trait that *B*'s failing to grieve reveals.... [I]t is at least odd to call irrational someone's failure to grieve, in the circumstances described.⁶

This argument, however, is flawed in a number of respects. First, and most importantly, Wilkinson simply assumes without any defense that a person who fails to grieve in a situation that makes grief appropriate is nevertheless capable of possessing “an accurate picture of her situation.” But this is precisely what is in dispute, for if grief is properly regarded as a rational or cognitive response, then the fact that a person who ought to experience grief fails to do so *itself* demonstrates that she is not correctly picturing her situation. Moreover, the claim that “it is at least odd to call irrational someone’s failure to grieve” fails to make its point since, as John McDowell has argued, the fact that a certain failure is a cognitive failure—that is, a failure of reason—does not entitle us automatically to conclude that anyone committing this failure must be guilty of *irrationality*; the charge of irrationality is a much stronger claim.⁷ Finally, Wilkinson’s argument makes use of a false dilemma. Worrying about a person’s mental health, or about character deficiencies indicated by her failure to grieve, does not prevent us from also regarding the person as suffering from a cognitive deficiency. In failing to experience appropriate grief, an individual such as Melissa shows herself to be suffering from a number of inter-related problems, all of which may well concern us.

Perhaps Wilkinson, or some other critic of the view that grief is a cognitive response, would object that a cognitive failure must take the form of a failure to grasp some *fact* and that what Melissa is missing is not a fact. But this objection seems to rely on an excessively narrow and restrictive account of what counts as a cognitive failure.⁸ At the very least, we should be hesitant about accepting the objection so long as we lack a principled basis for distinguishing facts from nonfacts. Moreover, many of the analyses that might be proposed to do this work are themselves too narrow. Suppose, for example, we decide that a fact is something that can be expressed in a proposition—an account that seems to capture the intuition that there is a close conceptual link between facts and propositions. On the basis of this analysis, it might be argued that Melissa’s failure (the failure, as we have said, to appreciate the significance of her loss) does not constitute a cognitive failure, since what she is missing (an understanding of what has been lost) cannot be expressed in a proposition and is therefore not a fact. The problem is that this account seems too restrictive when applied to certain other sorts of facts and knowledge. Frank Jackson and Thomas Nagel have argued that knowing what a certain experience (seeing red, for instance) is like is a genuine species of knowledge; a person who has never seen red does not know what it is like to see red, although she may know a great deal else about, for instance, typical causes of ‘red’ sensations, the physiology of the human ocular system, and so forth.⁹ Such a person is missing some crucial knowledge—

indeed, she seems to be missing a significant fact about the world—but it is not a fact that can be stated in a sentence or expressed in a proposition. The proposed account of facts, knowledge, and cognitive failures, then, is too restrictive.

Of course, the claim that knowing what some particular experience is like constitutes genuine knowledge is controversial. Moreover, it would be perfectly consistent to accept that ‘what it is like to see red’ constitutes knowledge, while denying that the same can be said of ‘what it is to lose a friend’. Adopting such a position need not spell the end for the broad strategy I am pursuing, for even if we decided that the failure to grieve did *not* count as a cognitive failure, there might nevertheless be *other* sorts of reasons for regarding ordinary grief responses as justified and appropriate (under the right circumstances, of course). If this were so then the Elimination Thesis might still fail with respect to grief, although its failure could not then be explained in the way I am suggesting. It seems to me, however, not only that our common attitude toward grief quite clearly *does* regard it as inherently justified and appropriate but that it does so for just the reasons I have been suggesting: because, that is, it regards grief as a cognitive response to events. If this is so, then we cannot reject the idea that knowing what it is to lose a friend is a genuine example of knowledge, without also rejecting along with it a very large part—indeed, quite possibly the core—of our ordinary understanding of grief itself.

4.

Grief is clearly a form of suffering. Yet the duty to eliminate suffering does not apply to it. To say this is not at all to deny the existence of duties to alleviate such suffering in certain ways. When one’s friend is grieving, one ought to be there to give comfort and support, and doing so may help make grief bearable, or at least less bad. In many cases such comfort reduces the amount of suffering that occurs. But to require this sort of amelioration is not to require the *elimination* of the suffering. Indeed, even if this were within our power—if, for instance, I had a pill that I could give my grieving friend that would wipe out her grief—it would, it seems to me, quite likely be wrong for me to offer it. I am no more under a duty, or even permitted, to do this, than I am under a duty, or permitted, to offer a pill that would wipe out my friend’s knowledge of her own mortality.

Two potential objections should be dealt with here. First, the conclusion that morality does not require or permit us to eliminate a person’s suffering is not based on a conflation of the positive effects of grief with the value of the grief itself. There are, of course, positive effects that typically accompany grief. The unpleasant emotions occasioned by great loss must, for

most of us, be dealt with at some point; if we do not experience them soon after a loss, they will in all likelihood return at a later time, at which point they may well be even more debilitating. Alternatively, we might find ourselves experiencing a kind of persistent, low-level suffering that endures well beyond the point when full-blown grief would have subsided, returning us to our normal state. Positive effects, then, often do accompany grief. But it is not because of these effects that we are reluctant to endorse a duty to eliminate grief. It is rather because we think that there is quite literally something wrong with people who are unable to grieve, that feeling grief in response to a loss is not only natural but right and appropriate, and that the justification of grief, like the justification of a true belief, has nothing at all to do with its effects on the happiness of the person who manifests it.

The second objection holds that denying that grief is bad must somehow commit us to the position that grief is good and that this is clearly an unacceptable view.¹⁰ Admittedly, holding grief to be good *in itself* would be very odd (though again, we need not deny that it often has significant instrumental value). But my arguments regarding grief have been concerned, not with the question of whether grief in itself is either “good” or “bad” but, rather, with the question of what reasons we have for eliminating it, or refraining from doing so; and while these two questions surely have something to do with each other, it is not as much as some philosophers think. Moreover the temptation to think that one must regard grief as either good or bad *simpliciter* should be resisted. It is an effect of a regrettably common simplifying tendency in ethics: a tendency to think of ‘good’ and ‘bad’ (or in some contexts, ‘right’ and ‘wrong’) as the fundamental ethical terms, and to see all others (including such so-called thick ethical terms as ‘courageous’, ‘rude’, ‘loyal’, and so on) as minor variants, or else derivatives, of these. Such an approach, as Bernard Williams has argued, ignores the richness and complexity of ethical language and the subtle ways in which the descriptive and the normative are intertwined.¹¹ It should not be assumed that ‘good’, ‘bad’, and ‘neutral’ are jointly exclusive and exhaustive, for there is no reason to think it impossible that a given object of evaluation—an act, a character trait, a practice, an institution—might simply fail to fall neatly into any of these three categories. It seems to me that grief is precisely this sort of phenomenon.

Our discussion up to this point suggests that we should hesitate to assert either ‘grief is good’ or ‘grief is bad’ as a blanket statement. Yet it clearly does not follow from this that grief is evaluatively neutral or irrelevant. Rather, grief, occurring in the right degree and properly directed toward a fitting object, is best regarded as *appropriate*. (I suspect that the word ‘appropriate’ and related terms play a much larger

role in ethical practice than philosophers have realized, but I will not argue that point here.) Given the fact that such grief is appropriate, moreover, there seems to be no difficulty in saying that the ability to feel it, when called for, is itself a good thing; yet at the same time, given the way it feels, there is no mystery what a person means when she says that grief is bad. One almost never *wants* to feel grief, even when it is appropriate that one do so. (Indeed, probably the only people who genuinely *want* to grieve are those who find themselves unable to and who realize that this constitutes a problem.) Grief *is* bad, then, in a certain sense—but not in the sense that is directly relevant to ethics; it thus does not follow from the fact that it is bad that it possesses negative moral value. The questions, ‘Is grief good?’ and ‘Is grief bad?’ are much more complicated than they might at first appear, and to refrain from answering either one of these questions in the affirmative is not in any way to commit oneself to a positive answer in the case of the other.

5.

We noted earlier that the eudaimonistic theory of value was false, since not all happiness is valuable. We now know that the theory is false for a different reason: not all suffering is disvaluable. Moreover there is an illuminating parallel to be drawn between the kind of happiness that presents a counterexample to the eudaimonistic thesis and the kind of suffering that does the same. The worthless happiness we took note of earlier was happiness that was not reflected in the objective circumstances of the subject’s life: a feeling of triumph, for instance, when in fact no triumph had been achieved. The problem with such happiness is that it is *meaningless*: it refers to, and only makes sense in the context of, external circumstances that do not in fact obtain. The ‘triumphs’ and ‘achievements’ of a subject of the Experience Machine are not only empty but indeed nonexistent, and the pleasurable mental states that appear to accompany them are thus (unbeknownst to the subject, of course) entirely inappropriate. It is for this reason that such happiness possesses little or no intrinsic value, and why there is no moral duty or compelling moral reason to promote it.

This sort of meaningless happiness has a converse: meaningful suffering. Meaningful suffering, like all suffering, is unpleasant to experience. But this unpleasantness does not entail that the experience is entirely lacking in value, let alone disvaluable. Nor does it imply that there is a duty, or any moral reason whatsoever, to erase such suffering from existence. As we have observed with respect to the case of grief, meaningful suffering is generally regarded, and should be regarded, as justified and appropriate. (This is true, of course, only so far as it *is* meaningful; inappropriate, excessive grief would not be

regarded as justified, and its elimination through therapy or similar means might well be regarded as desirable.¹²) Meaningful suffering of this sort is, as I have suggested, a cognitive response to objective circumstances: it is a way of understanding what has happened in one's life, and more broadly, of appreciating the nature of the world in which our lives take place. Just as one would not want to be a person who experienced only meaningless happiness, one would also not want to be a person who could not, when it was called for, experience meaningful suffering. Both of these people instantiate defective relationships with the world in which they live.

Three sets of implications can be drawn from this view. First, as we have already noted, the Elimination Thesis has been shown to be false. Some suffering is justifiable and not disvaluable; hence, there exists some suffering that we have no moral reason to eliminate. More generally, any moral theory that is value monistic—any theory, that is, that relies on the idea that moral obligations are determined by the presence or absence of a single sort of value (a value which might, perhaps, occur in varying degrees but that is always of the same basic type)—must also be rejected. As we have shown, the question of whether to eliminate a given mental state, for example, cannot be answered through an exclusive focus on any one sort of value. Indeed, it cannot even be answered by an exclusive focus on any one sort of *reason*, for at least when the mental state in question has a cognitive aspect, a variety of reasons, epistemological as well as ethical, will prove to be unavoidably and irreducibly relevant. Thus, our observations regarding the nature of grief form the basis of an argument for a kind of pluralism that applies both to values and to practical reasons.

Similarly, the proposed view casts doubt on the idea that there is a sharp distinction between practical and theoretical reasoning. Such a distinction is assumed not only in many dominant theories of practical reasoning (instrumentalism, perhaps, being the foremost among them) but also in a number of common approaches to logic and philosophy of mind. But this presumption, no matter how natural, must be questioned if our general line of thought has been on target. Given that a mental state such as grief is at once an ethical entity (considered as a type of suffering) and an epistemological one (considered as a cognitive response), it follows that such states, and actions relating to them, cannot be comprehensively evaluated without taking into account both practical and theoretical reasons. Thus, in such cases there is no way of strictly separating reasoning about *how to act* from reasoning about *what to believe*.

Finally, our observations challenge a common set of views regarding the function of certain mental states in relation to the sources of intrinsic value. Eudaimonistic theories of value, as we have noted, hold that happiness is the source of all

positive value and that suffering is the source of all negative value. Our view, however, suggests that, rather than serving as a *source* for negative value, certain types of suffering, such as grief, are better seen as *indicators* of value. Grief, properly and appropriately experienced, reflects the fact that something very bad has happened, but it would be a serious error to mistake the grief itself for the very bad thing. Moreover the very idea that justifiable grief is to be distinguished from its unjustifiable counterpart on the basis of whether or not it is properly judged to be *meaningful* involves, in a fairly unavoidable way, the thought that events in the world, and responses to them, can be evaluated from a point of view that is both normative and (in some strong sense) objective. (This, of course, is one of the things many eudaimonistic theorists were hoping to avoid, by relativizing value to the pleasure or happiness of each individual.) It is only natural, once we have accepted such a view, to extend it to other sorts of mental states: pleasure, for example. The resulting outlook would imply that pleasure, rather than serving as a source of positive value, is primarily an indicator that allows us to identify independently existing values.¹³ Such views are not, perhaps, without their epistemological and metaphysical puzzles, and I cannot pretend even to have begun to deal with those here. Nevertheless, it seems to me that such an approach is indeed implied by our ordinary understanding of many mental phenomena, including (but, I suspect, not limited to) both grief and pleasure. And if this is correct, then it seems to me a very strong reason for paying more serious and sympathetic attention to such views than many analytic philosophers have tended to in the past.¹⁴

Notes

¹ Robert Nozick, *Anarchy, State, and Utopia* (New York: Basic Books, 1974), 42–5.

² G. E. Moore makes the same point, with regard to pleasure rather than happiness, in *Principia Ethica* (New York: Cambridge, 1903, reprinted 1968), 93–4.

³ W. D. Ross, *The Right and the Good* (New York: Oxford University Press, 1930).

⁴ Although, as Elijah Millgram argues, it is not clear that any state *could* play this functional role which was not, in itself, painful or at least highly unpleasant. See “What’s the Use of Utility?” *Philosophy & Public Affairs* 29, no. 2 (2000): 113–36, at 125–6.

⁵ Jamie Mayerfeld, *Suffering and Moral Responsibility* (New York: Oxford University Press, 1999), 86.

⁶ Stephen Wilkinson, “Is Normal Grief a Mental Disorder?” *The Philosophical Quarterly* 50, no. 200 (2000): 289–304, at 296–7.

⁷ See John McDowell, “Might There Be External Reasons?” in *Mind, Value and Reality* (Harvard University Press, 1998), 95–111, at 102–07.

⁸ Cf. McDowell’s comments regarding knowing “what it means that someone is shy and sensitive” in “Are Moral Requirements

Hypothetical Imperatives?" in *Mind, Value and Reality*, 85–6.

⁹ Frank Jackson, "Epiphenomenal Qualia," *Philosophical Quarterly* 32 (1982): 127–36, and "What Mary Didn't Know," *The Journal of Philosophy* 83, no. 5 (1986): 291–5; Thomas Nagel, "What Is It Like to Be a Bat?" in *Mortal Questions* (New York: Cambridge, 1979), 165–80.

¹⁰ Cf. Loretta M. Kopelman "Normal Grief: Good or Bad? Health or Disease?" and "Rejoinder: If Grief is Not Bad, Is It Good?" *Philosophy, Psychiatry and Psychology* 1 (1994): 209–20 and 225–6; and Thomas N. Wise, "Commentary on 'Normal Grief: Good or Bad? Health or Disease?'" *ibid.*, 223–4.

¹¹ See Bernard Williams, *Ethics and the Limits of Philosophy* (Cambridge, Mass.: Harvard University Press, 1985), chapters 7 and 8.

¹² Moreover, it seems at least somewhat plausible to think that it would be not only ethically permissible but ethically desirable or even mandatory to eliminate the artificially induced grief experiences of the Experience Machine subject, were one in a position to somehow do so. This is, after all, a context where the ordinary (epistemological) justification of grief does not apply: grief, for the Experience Machine subject, is not part of a deep understanding of the nature of the world, for the subject is deeply deceived, and so does not *possess* a deep understanding of the nature of the world. The issue of our ethical duties toward Experience Machine subjects, however, is a vexed one, and one I need not try to settle here.

¹³ Elijah Millgram defends such an account of pleasure in *Practical Induction* (Cambridge, Mass.: Harvard University Press, 1997) and "What's the Use of Utility?" (cited in note 4).

¹⁴ I would like to thank Talbot Brewer and two anonymous reviewers for helpful comments on the paper.