

Analytic Hedonism and Observable Moral Facts: A Précis of The Feeling of Value

Sharon Hewitt Rawlette

Guest essays represent only the views of the author(s).

Table of Contents

1. Introduction
2. The Basics of Analytic Hedonism
 - Intrinsic vs. Instrumental Goodness
 - The Intrinsic Goodness of Pleasure and Intrinsic Badness of Pain
 - The Conceptual Connection
3. Objections and Clarifications

Introduction

Many people are skeptical that there are any objective moral truths. They think it much more likely that ethics is a matter of personal or cultural opinion, a set of preferences that we happen to have about the way we would like the world to be, or the way we would like people to act.

After all, people around the world differ in their ethical beliefs. In fact, people in the same *family* often differ in their ethical beliefs. And there doesn't seem to be any objective way to decide who is right. When there's disagreement about scientific questions, we can conduct experiments to test the validity of

rival hypotheses, but it doesn't seem like we can do anything like that in ethics. Goodness or rightness just doesn't seem to be something we can empirically investigate.

Or can we?

In this essay, I'll argue for a view I call "analytic hedonism". According to this view, which I defend at much greater length in my 2016 book [*The Feeling of Value*](#), we can indeed observe basic moral facts—and do so all the time. These basic moral facts are the intrinsic goodness and badness of certain of our own experiential states, like pleasure and pain. From our direct acquaintance with the intrinsic value of these good and bad experiential states, combined with further knowledge about what actions and states of affairs are conducive to producing these states, we can build an entire ethical system that is fully grounded in observable fact.

The Basics of Analytic Hedonism

The core claim of analytic hedonism is that the positive and negative qualities of conscious experience are intrinsically good and bad *by definition*. Since philosophers call facts that are true by definition "analytic" truths and call the view that only pleasure is intrinsically good "[hedonism](#)", I refer to this version of moral realism as "analytic hedonism".

Let's unpack the details of analytic hedonism's core claim and consider what reasons there are for believing it's accurate.

Intrinsic vs. Instrumental Goodness

One very important thing to understand is the difference between intrinsic goodness and other types of goodness.

The intrinsic goodness of something is the value it has based purely on its own, internal properties. That is, its intrinsic goodness does not depend on its ability to help create, cause, or constitute some other thing outside it. For

example, the goodness of a piece of cake depends on how much pleasure it causes in people's taste buds (perhaps weighed against the effects it has on people's health). This means the goodness of the cake is not *intrinsic* to it but rather dependent on some external factor or factors. It has "instrumental" goodness, because it's instrumental in the production of something else that is good: a positive taste experience. Intrinsic goodness, on the other hand, is value that something has as an end in itself.

Analogously, intrinsic badness is the badness something has based purely on its own, internal properties. It is something that is never worth pursuing as an end in itself and in fact is worth avoiding, all else being equal.

The Intrinsic Goodness of Pleasure and Intrinsic Badness of Pain

According to hedonism, the only things in the world that have intrinsic goodness or badness are experiential states. The experiential states that are intrinsically good are those that *feel* good, like pleasure, joy, and happiness. The experiential states that are intrinsically bad are those that *feel* bad, like pain, suffering, and sadness.

This doesn't mean that lots of other things in the world can't be good or bad in an instrumental way. It also doesn't mean that every instance of pleasure will be good overall, since a particular instance of pleasure might lead to lots of pain down the road, which would more than outweigh its intrinsic goodness. And some instances of pain could be good overall, because in the long run they produce enough pleasure to outweigh their intrinsic badness. Hedonism doesn't rule out any of these things. It just says that positive experiential states are the only things that are *intrinsically* good and negative experiential states are the only things that are *intrinsically* bad.

The Conceptual Connection

Analytic hedonism, however, says even more. The analytic hedonist says that these experiential states are intrinsically good and bad because having the

experiential quality of goodness is *what it means* for something to be intrinsically good and having the experiential quality of badness is *what it means* for something to be intrinsically bad. If we never experienced pleasure or pain or any other positive or negative experiential quality, we wouldn't have any idea what it meant for something to have intrinsic value or disvalue. We might still desire or avoid certain things, but we wouldn't know what it would be like for something to intrinsically *justify* our desiring or avoiding it. This is because positive and negative experiential qualities give our concepts of intrinsic goodness and badness their essential content.

Now, many philosophers have claimed that there can't be a conceptual equivalence between experiential qualities and intrinsic goodness and badness, because experiential qualities are *descriptive* properties (properties about how the world is) and intrinsic goodness and badness are *normative* properties (properties about how the world ought to be) and these two types of properties can't overlap or be derived from one another. David Hume, for instance, claimed that you can't derive an "ought" from an "is".¹ And G. E. Moore famously argued against definitional links between descriptive and normative properties by claiming that, no matter what descriptive property one selects, it will always be an "open question" whether that property is in fact normative.²

However, I believe it's a mistake to think that the descriptive and the normative can never overlap. Imagine that you are a scientist taking an inventory of all the various qualities present in conscious human experience. You've written down the qualities of experiencing various colors, sounds, and smells. But there are two distinct experiential qualities that you can't quite figure out how to describe. In the end, you realize that the only way to describe the one is to say that it's "good" or "positive", and that you can only describe the other by saying it's "bad" or "negative". That is, you have to mention the normativity of the experiences in order to describe them accurately. The qualities of these experiences are simultaneously normative *and* descriptive.

Moore would counter this claim with the observation that the question “Is pleasure good?” feels more “open” than the question “Is pleasure pleasant?”, so pleasantness (understood here as a shorthand for positive experiential valence) and goodness cannot be definitionally equivalent. Allan Gibbard has recently made a similar argument by pointing out that, if pleasantness and goodness meant the same thing, the statement “Not all that is pleasant is good” would be as incoherent as saying “Not all that is pleasant is pleasant”.³

Note, however, that the analytic hedonist doesn’t actually say that pleasantness and goodness mean the same thing, but rather that pleasantness and *intrinsic* goodness mean the same thing. As noted above, the word ‘goodness’ has different senses, including not only intrinsic goodness but also instrumental goodness and all-things-considered goodness. The statement “Not all that is pleasant is good” can be used to express the truth that not all that is pleasant is instrumentally or all-things-considered good as well as the truth that certain pleasures are signs of instrumentally bad dispositions (such as the inclination to take delight in the pain of another). So it’s perfectly coherent to say “Not all that is pleasant is good”, even if pleasantness and *intrinsic* goodness are definitionally equivalent.

What wouldn’t be coherent is to say something like “Pleasant feelings are intrinsically bad in every way” or “Unpleasant feelings are intrinsically good in every way”.⁴ If we encountered someone who earnestly made these statements, I believe we would start to wonder if they were referring to the same experiential qualities that we are when we use these terms. If they don’t think there’s any sense in which pleasure is intrinsically good, they must be talking about a feeling other than what we mean when we say “pleasure”! Which means that we do indeed find the kind of incoherence in these statements that we ought to if pleasantness and intrinsic goodness are conceptually equivalent, and that Gibbard’s coherence test (and, I believe, Moore’s related Open Question Argument)⁵ fails as an argument against analytic hedonism.

But if pleasantness is conceptually equivalent to intrinsic goodness and unpleasantness is conceptually equivalent to intrinsic badness, this means that normative properties are part of the empirical world. And it means that, if we want to know which states of affairs are intrinsically better than others, there is a way to find an objective answer to this question: by determining how much pleasant and/or unpleasant experience those states of affairs contain.

Objections and Clarifications

Doesn't the Fact That Different Things Cause Different People Pleasure and Pain Make Hedonism Relativist/Subjectivist?

It's certainly true that the particular circumstances that are conducive to pleasure and pain can vary from person to person (and from culture to culture), but this does not make it any less objective that pleasure is intrinsically good and pain intrinsically bad wherever they are found.

It also doesn't change the fact that it is a fully objective truth that whatever circumstances promote positive experience for a certain person are instrumentally good *for that person*. And whatever circumstances cause a person negative experience are objectively instrumentally bad *for that person*. Since it is an objective fact about the world that different people have different experiences in the same circumstances, it is actually an advantage of analytic hedonism that it recognizes this and naturally accommodates it, instead of arbitrarily declaring the same sorts of behaviors or states of affairs to be good for all people no matter the experiences they cause.

On the other hand, we should be clear that it doesn't matter what a person *believes* will produce positive or negative experience, for themselves or others. What matters, from an ethical point of view, is what actually *will* produce positive or negative experience. People often arrive at their ethical opinions by imagining a state of affairs and seeing whether they get a positive or negative feeling about it. However, the intrinsic goodness or badness of a state of affairs is determined by the quality of the feelings that *constitute* that state of

affairs (the feelings that are experienced when that state of affairs actually pertains) and have no necessary connection to the feelings people have when they merely think *about* that state of affairs. That is what makes moral facts objective—the fact that they don't depend on people's beliefs or thoughts or attitudes or judgments about them but on the way the world actually is.

Surely the Pleasure of a Torturer Isn't Intrinsically Good

While it may seem counterintuitive at first, analytic hedonism does assert that all pleasure—[including the pleasure enjoyed by a torturer](#)—is intrinsically good. But let's take a closer look at what this means.

According to analytic hedonism, pleasure is intrinsically good because its *internal quality* is good, and good purely in virtue of what the quality itself feels like, regardless of how it may have come about and regardless of what further things it may cause. If the pleasure felt by a torturer feels equally as good as the pleasure someone else takes in nursing an injured bird back to health, then these two pleasures—considered purely in terms of their own internal characteristics—are of equal *intrinsic* value.

However, this does not mean that the two pleasures are of equal *all-things-considered* value. And it definitely does not mean that it is good for someone to take pleasure in torture. Having a disposition to take pleasure in others' suffering is of tremendous instrumental disvalue to the rest of the world. In fact, it is one of the worst character traits possible, because of the way it puts one's own interests at complete odds with those of others. But the tremendous instrumental disvalue of this disposition does not prevent the pleasure it produces from having the same intrinsic quality of goodness had by other experiences of pleasure. This intrinsic goodness is just swamped by the extreme badness of the disposition that produces it.

Doesn't the Badness of Pain Depend on Our Attitude Toward It?

Moral antirealists object to the idea that any experience has positive or negative qualities intrinsically, rather than because of the value we ourselves place upon it. Sharon Street, for instance, proposes that “[p]ain is a sensation such that the creature having the sensation unreflectively takes that sensation to count in favor of doing whatever would avoid, lessen, or stop it”,⁶ but, she emphasizes, “the badness of pain does in fact depend on our evaluative attitudes”.⁷

Christine Korsgaard expresses a similar opinion, writing that “[p]ain... is less horrible if you can curb your inclination to fight it. This is why it helps, in dealing with pain, to take a tranquilizer or to lie down. Ask yourself how, if the painfulness of pain rested just in the character of the sensations, it could help to lie down? The sensations do not change. Pain wouldn’t hurt if you could just relax and enjoy it.”⁸

There’s some truth to this. Pain is a complex experience, and it normally has at least two components. There are the sensations of nociception, which vary depending on the kind of harm that’s being done to the body (whether it’s being cut or burned, for example). These sensations can be sharp or dull, pulsing or constant, and felt to be at different bodily locations. But, in addition to the sensations of nociception, the experience of pain also normally includes a feeling of *badness*.

This division of the experience of pain into at least two basic parts is widely accepted by scientists because of the existence of cases of “reactive disassociation”, where one can feel the sensations of nociception without the accompanying feeling of badness. This phenomenon is known to occur not only under the influence of opiate analgesics such as oxycodone and morphine but also as a result of prefrontal lobotomies and leucotomies, and as a result of certain brain lesions.⁹

Moral antirealists insist that what is missing in these cases is simply the disposition to avoid the nociceptive sensations, but there are reasons to think

that the missing element is more than that, that it's in fact an experiential quality, a *feeling* of badness. Consider:

- The badness of pain can't just consist in avoidance behavior, because there are documented cases of reflexive avoidance responses without a conscious feeling of pain.¹⁰
- The badness of pain can't just consist in a conscious plan for avoidance, since even those without the cognitive ability to form plans seem capable of feeling pain, such as newborn babies and people with severe mental handicaps.
- And the badness of pain can't just consist in a certain bodily or hormonal state, because the same state can be associated with both positive and negative emotions, depending on the context. Injecting people with adrenaline, for instance, doesn't change their mood all on its own. A situation with a slightly cheerful or slightly irritating valence is required to determine the emotion elicited by the adrenaline (and the situation alone isn't enough to elicit the emotion either, as shown by the lack of response in those injected with an inert substance).¹¹ What is more, we ought to *expect* our dispositions of attraction and aversion to have a qualitative feel to them, for the simple reason that so much other important information that we use in conscious decision making manifests in this way. Our visual system may let us know that a lion is up ahead, but unless we also have access to our brain's verdict about whether the proximity of a lion is good or bad, this purely visual information is useless to promote our survival. We need to be able to consciously reflect on the best escape strategy, and it would seem consistent with our other knowledge about conscious thought that we require an experiential manifestation of our brain's verdict of "Bad!" with regard to the lion's proximity.

Now, you might think that the brain could use an experiential marker for badness that was not itself bad, but there is evidence that the brain needs a marker that is simultaneously experiential *and* intrinsically motivating. Consider the experience of Paul Brand, who spent years developing artificial sensing systems for victims of leprosy.¹² Since leprosy makes a patient insensitive to pain in their hands and feet, Brand developed gloves and socks containing pressure sensors capable of indicating when some action the patient was taking was having a harmful effect on their body. Unfortunately, these indications that they were performing an action such that it could lead to the destruction of a finger or toe did not stop patients from performing the action. They were not sufficiently motivated by this information about their future well-being, and the project ultimately failed.

Brand reports the comment of a colleague, Professor Tims, who said to him, “Paul, it’s no use. We’ll never be able to protect these limbs unless the signal really hurts. Surely there must be some way to hurt your patients enough to make them pay attention.”¹³ In fact, Brand and his colleagues tried an alternative system, which responded to danger to the limbs by inducing pain in a still sensitive body part. Unfortunately, patients preferred to turn this system off rather than heed its warnings!

It appears that, whatever experiential quality represents instrumental badness in our conscious thought processes, it needs to be simultaneously experiential *and* motivating. And I believe that is precisely what we find in the qualitative experience of pain (as well as in pleasure and other positive and negative experiential states). Our minds represent instrumental badness in the form of something that is *intrinsically* bad: negative experience.

Do All Good and Bad Experiences Share A Common Quality?

It might seem unlikely, with all the many possible types of pleasant and unpleasant experience one can have, that there could be a single experiential

quality that they all share, and according to which they can all be compared. However, scientific research suggests this could very well be the case.

In an essay entitled “Social Pain, Support, and Empathy”, Jaak Panksepp summarizes evidence for the claim that “emotional pain, such as that which accompanies grief and intense loneliness, does share some of the same neural pathways that generate the affective sting of pain.”¹⁴ And the same opiate analgesics that I already mentioned as having the effect of making people feel that their pain was no longer bad turn out also to be effective at removing signs of separation distress in many different animals, including primates.¹⁵ There are also chemicals naturally produced by the brain that reduce both pain and separation distress: namely, oxytocin and endorphins.¹⁶ These chemical and neural correlations between pain and feelings of sadness and emotional distress suggest there’s nothing implausible about a qualitative similarity in the experiences.

There is also evidence for our positive experiences’ sharing a common experiential quality, one which is the exact opposite of that shared by negative experiences. Research conducted by Michel Cabanac and others shows not only that subjects are able to rate various pleasures and displeasures (both physical and mental) on a common scale but that their subsequent behavioral choices tend to produce the highest algebraic sum of pleasure as defined by these ratings.¹⁷

How Much Should We Care About Other People’s Pleasure and Pain?

One might think that, because the goodness of pleasure and badness of pain are experiential qualities that can only be experienced by a particular individual, we can only derive an egoistic form of hedonism from them: a view according to which each of us has reason to promote positive experience *for ourselves* but no reason to care about the experiential states that anyone else will experience. However, there are three good reasons to think that a specifically utilitarian version of hedonism is true: that is, that all agents

ought to promote the greatest balance of good experience over bad, no matter who will be the subject of said experience.

First, there is the fact that our positive and negative experiential states seem to present themselves to us as having agent-neutral value. Consider the experience of pain. It seems to me that, when I feel pain, I don't feel that its negative character is only a reason for *me* to avoid it. I don't feel that my pain is merely something *I* have reason to get rid of but which is no reason for anyone else to help me. Rather, I feel that if anyone else can do anything to help me get rid of the pain, then they ought to do so, all else being equal. And I assume that other people's pain has a similar quality. I don't take the fact that a person's pain can only be *felt* by one individual to mean that it has a normative claim only on the actions of that individual. Pain is bad in a way that makes a claim on anyone who is in a position to help.¹⁸

Second, there is the fact that positive and negative experiential qualities don't contain reference to any particular actions they require or to the agents who are required to take them. The experiential goodness or badness is only a feeling of value or disvalue. The concept of action—especially of goal-directed action—arrives much later in our mental development than our first experience of pleasure or pain. Nevertheless, these experiences do have implications for action, which become apparent when we add to our knowledge of the goodness and badness of experiential states some information about their causal relations to the rest of the world. It is when we start to realize that other objects, states, and events could be conducive to the production of pleasure and pain or could prevent these experiences that we realize that experiential qualities have implications for the actions that produce or prevent them. The connection of the experiential qualities to actions exists because those actions are able to produce or prevent the experiences, and this is just as true of actions taken by people other than the experiencing subject as it is by actions taken by the experiencing subject themselves.

Third, the evidence we currently have about the nature of personal identity gives us little reason to think that personal identity is significant in the way that would be necessary to support the view that experiential states are only reason-giving for the agent who experiences them. Specifically, it doesn't look like there is any metaphysically distinct subject who both performs an action and then later experiences the results. A reductionist account of personal identity, according to which personal identity is just some particular web of continuity among experiences, desires, and beliefs through time, appears most plausible.¹⁹ But it doesn't seem that continuity in the characteristics identified by the reductionist account is the right sort of thing to affect the degree to which we ought to promote particular future pleasures or avoid particular future pains. So, if we believe (1) that positive and negative experiences are reason-giving at least for the person who experiences them and (2) that there is no property that binds the moments of a single person's life in a more (or less) reason-giving way than the moments of different people's lives, then we ought to conclude that positive and negative experiences are equally reason-giving for all agents.

Now, if it's true that we have reason to promote the future pleasure and avoid the future pain of all other experiencing subjects, one might think that this is just too much to demand of a single individual. That is why utilitarians generally acknowledge that the mere fact that [the interests of everyone everywhere count equally](#) doesn't mean that each of us will produce the best consequences by always *consciously thinking* about every individual and every consequence every time we make a decision. As Stefan Schubert and Lucius Caviola note in [Virtues for Real-World Utilitarians](#), our psychological limits are very relevant to the way that we put utilitarianism into practice. The fact that we are not physically or mentally capable of attending to all of the potential happiness and misery in the world makes it *instrumentally* important to concentrate our efforts where they are [likely to do the most good](#). But this limitation in the scope of our efforts is a practical concession and not

based on any deep metaphysical fact about our duties to some individuals over others.

What About Duties and Rights?

The above discussion of practical limitations leads naturally to an answer to [another objection frequently brought against utilitarianism](#) (of both the hedonistic and non-hedonistic varieties): that focusing purely on producing the best consequences goes against our intuitions that we have duties to keep promises, tell the truth, and respect individual rights, even when doing these things leads to worse consequences.

As stated above, the fact that we ought to do what will produce the best possible overall consequences doesn't mean that each of us will produce the best consequences by always *consciously thinking* about every individual and every consequence every time we make a decision. The best decision procedure to use is going to depend on which decision procedure produces the best results given the real-world, practical limitations with which we are confronted.

Given our cognitive limitations, it is clear that the best decision procedure will not be one that requires us to consciously calculate all the expected consequences of all our possible actions. Even if it were possible to collect enough information to do this, the act of collecting it and making calculations based on it would waste time and energy that would produce more good employed in some other activity. So we *have* to use shortcuts, because doing so produces the best consequences.

But what kind of shortcut decision procedures should a utilitarian employ? This is determined by at least three general features of actual situations: uncertainty, need for coordination, and motivational limitations.

Uncertainty

All of our predictions about the consequences of our actions are necessarily based on our knowledge of the consequences of similar actions taken in the past, by ourselves or others. We have to rely on generalizations of the form “Actions of Type *A* in situations of Type *B* on average produce *x* units of value”. But to know which generalization to rely on in a particular case, we have to know how broadly or narrowly to characterize the type of action and type of situation.

Here, we have to balance two considerations. The more similar all the comparison cases are to the present case, the more likely they will be to yield an accurate prediction. But, at the same time, the more similar the comparison cases are to the present case, the smaller the overall sample will be, and the easier it will be for a random factor to distort our prediction.

Interestingly, research suggests that the most accurate way to predict complex phenomena is very often to focus on the one factor that is most closely correlated with the effect in question, as secondary factors are more likely to distort than to improve the accuracy of the prediction.²⁰ This implies that we have the greatest chance of bringing about the best consequences if we make decisions according to *very general rules*.

For example, if the single best predictor of how happy one will be in marriage to a particular person is one’s average happiness over the first year of dating that person, then (presuming that personal happiness is the goal one is seeking) one should make one’s marriage decision based solely on this factor, without regard for any other variables.

Similarly, if evidence shows that, in cases where one is faced with a decision whether or not to take an innocent person’s life, refraining from taking the life is the factor most closely correlated with the outcome one is seeking (for instance, long-term maximization of well-being across society), then one ought to follow the general rule of refraining from taking an innocent life and

disregard other variables, including any potential benefit that taking an innocent life seems like it could procure in an individual situation.

Note, however, that the factor most closely correlated with optimal consequences might not be present in every situation. For instance, when the expected benefit of taking an innocent life exceeds a certain threshold—say, when we can be almost certain that taking one innocent life will save 10,000 other innocent lives—choosing to save so many other lives could turn out to be even *more* closely correlated with optimal consequences than is refraining from taking an innocent life. If that correlation is stronger, then we should follow the general rule “Take an innocent life when you can be almost certain that doing so will save 10,000 or more other innocent lives”. And we should only follow the “next-best” rule of refraining from taking innocent lives when such overwhelmingly positive consequences are absent and thus irrelevant to the decision at hand.

Need for Coordination

Not only is making decisions according to general rules the best way to deal with uncertainty about the consequences of our actions, it can also *reduce* this uncertainty. Much of the uncertainty about our actions’ consequences results from the fact that their consequences depend on the actions of others. And if *their* actions also depend on their complex calculation of the consequences, it could be very difficult to figure out how they’re going to act. A practice of making decisions according to general rules can help resolve this dilemma. For instance, two useful rules for coordinating our actions with others are to keep promises and to tell the truth, unless the overall consequences will obviously be very bad.

Motivational Limitations

Probably the clearest of our motivational limitations as human beings is a bias towards ourselves and those we love. It seems likely that it is not possible for

human beings to be equally concerned about the interests of all experiencing subjects. But that means that the optimal decision procedure is going to have to produce the best consequences *given our ineradicable biases*.

One thing we can do to greatly improve consequences given such limitations is to create a strong causal connection between each person's actions and the positive results that individual is most motivated to bring about. An optimal society will be ordered in such a way that each person's actions have a strong link to their own happiness and that of their loved ones.

Each person's actions might seem to *always* have a strong effect on their own happiness and that of those they care about, but consider that such an effect would not be assured in a society in which zealous utilitarians were always interfering in others' lives in an attempt to maximize total well-being. If we didn't generally recognize a duty to refrain from interfering in others' projects and respect certain property of others' as off-limits to us, then each of us would find our lives a chaotic mess of interference from others. We would have very limited possibilities for promoting our happiness because of the way any investment we made in our future could be taken away from us without warning. In such a world, we would have minimal motivation to undertake any projects other than those producing instant gratification, because they would almost certainly be doomed to failure. On the other hand, by generally being disposed to respect others' "rights"—rights to certain degrees of non-interference with their lives, limbs, and property—we foster the kind of environment in which each person is motivated to work and invest in their future happiness. This is why utilitarians have historically offered emphatic defenses of the importance of rights, with [John Stuart Mill](#)'s 1859 essay "[On Liberty](#)" being the most famous.²¹

In conclusion, given the three general features of actual situations just described—uncertainty, need for coordination, and motivational limitations—utilitarianism [actually requires us](#) to adopt a decision procedure that respects certain duties and rights and only allows us to ignore these duties and rights in

cases where it is overwhelmingly clear that the consequences would be much better. But this is more or less in congruence with our moral intuition.

Should We All Just Plug Into the Experience Machine?

Practical considerations are also very relevant to determining whether hedonistic utilitarians should want everyone in the world to spend their lives plugged into an “experience machine” that gives them maximally pleasurable experiences but cuts them off from having any real relationships, knowledge, or accomplishments.

Robert Nozick published his “experience machine” objection to hedonism in 1974,²² and repugnance for life in the experience machine has led many philosophers to reject hedonism as a plausible theory of the good. But this rejection is unwarranted. Our negative feelings about the experience machine are not at odds with hedonism but are in fact very much in line with strong hedonistic reasons for *rejecting* life in the experience machine, due to the important instrumental benefits of being in contact with the world outside our own minds.²³

The experience machine isolates us from information about our future welfare and renders us incapable of any action that could improve our future balance of pleasant over unpleasant experience. We are completely at the mercy of the machine and anyone who might come into contact with it or us. We must trust the machine to look after our interests as well as we ourselves could if we were living disconnected from the machine and even in the face of highly novel circumstances. And we must trust that, if something goes wrong with the machine or with our bodies, such that we are no longer having pleasant experiences but unpleasant ones, this will be noticed and fixed.

For the sake of argument, let’s imagine that we *could* somehow create such an incredibly intelligent machine and be *absolutely sure* that we would be forever safe abdicating all of our decision-making capacities to the machine’s. If this is the case, what is it exactly that we plan to do if we refuse to plug into the

experience machine? Faced with a choice between living in a world where our intelligence and our ability to help others are superfluous²⁴ and plugging into the experience machine where at least we have the *experience* of doing worthwhile things, the experience machine starts to look a lot less crazy.

And we shouldn't ignore the fact that many of us already spend a fair amount of time in something rather like the experience machine: by watching television, playing video games, reading books, or sleeping. These activities avoid some of the major disadvantages of plugging into the experience machine because (1) even while we are engaged in them, they allow us to remain at least somewhat alert to things in the real world that may require our attention and (2) they absorb only a portion of our lives and thus still allow us to remain agents capable of promoting our own happiness and that of others. However, to the extent that these activities do not meet these criteria, we usually do consider engaging in them worrisome. So, in this area as well, our intuitions about maintaining connection to the external world seem to be generally consistent with hedonistic reasons.

Is Animal Pleasure Just as Valuable as Human Pleasure?

If intrinsic goodness and the experiential quality of pleasantness are the same thing, it seems quite possible that a very happy pig is living a more intrinsically valuable life than an only barely happy human being. This is at odds with many people's intuitions—including those of hedonistic utilitarian [John Stuart Mill](#), who advised that “[i]t is better to be a human being dissatisfied, than a pig satisfied”, and defended a view according to which some types of pleasure (the uniquely human ones, as it happens) were more valuable than others.²⁵

However, there are good reasons to be skeptical of our intuitions about the superiority of human life and experience. As [Peter Singer](#) has pointed out, if we look at humanity's history of defining the boundaries of the moral community, we see that we have a decided tendency to privilege those who are like us, in

race, sex, class, religion, and any number of other ways. And yet, over time, we have [come to realize](#) that our biases toward those most similar to us do not indicate an objective fact about the relative value of races, sexes, and cultures.

²⁶ I believe, with Singer, that the time has come to concede that discounting the positive and negative experience of nonhumans based purely on the fact that they are not of our own species is [similarly indefensible](#).

Now, recognizing the equal intrinsic value of positive experience no matter the species that enjoys it does not mean that we can't also recognize that human lives may have some *instrumental* value that other animal lives lack. Humans have abilities for complex thought and coordinated action that allow us to solve problems that pigs can't, and this certainly makes other humans useful to us in a way that pigs are not. But just how instrumentally useful our species turns out to be will depend on the extent to which we use our abilities in service to the greater good.

About the Author

[Sharon Hewitt Rawlette](#) has a PhD in philosophy from New York University, where she studied under Thomas Nagel and Sharon Street and won the Dean's Outstanding Dissertation Award in the Humanities for the work that became her 2016 book [The Feeling of Value](#). After serving as Florence Levy Kay Fellow in Ethics at Brandeis University for two years, she left academia for a freelance writing career and has spent the past decade publishing on the topics of consciousness, ethics, and spirituality. Currently, with support from the [Future Fund Regranting Program](#), she is at work on a new scholarly book exploring the relationship between analytic hedonism and contemporary debates in philosophy of mind.

How to Cite This Page

Rawlette, S.H. (2023). Analytic Hedonism and Observable Moral Facts. In R.Y. Chappell, D. Meissner, and W. MacAskill (eds.), <<https://www.utilitarianism.net/guest-essays/precis-of-the-feeling-of-value>>, accessed 3/19/2024.

1. David Hume (1740). *A Treatise of Human Nature*, Book 3, Part 1, Section 1. ↩
2. Moore actually formulated his Open Question Argument in terms of natural properties rather than descriptive ones, but there's nothing that would limit its scope to naturalistic analyses of normative concepts rather than descriptive analyses in general, and it makes the central issue clearer to talk in terms of descriptive properties. See G. E. Moore (1903). *Principia Ethica*. Cambridge: Cambridge University Press, pp. 5–21. ↩
3. Allan Gibbard (2003). *Thinking How to Live*. Cambridge, MA: Harvard University Press, Ch. 2, especially pp. 23–29. ↩
4. I add the phrase “in every way” to help remove any apparent coherence that might arise from someone’s using a slightly different sense of the term ‘intrinsically’ than what I have been employing in this essay. ↩
5. Because there is an obvious sense in which the question “Is pleasure intrinsically good?” is settled by the meaning of the terms involved. ↩
6. Sharon Street (2006). A Darwinian Dilemma for Realist Theories of Value. *Philosophical Studies*, 127, 146. ↩
7. Street, 145. ↩

8. Christine Korsgaard. The Sources of Normativity, in Grethe B. Peterson, ed. (1994), *The Tanner Lectures on Human Values*. 15. Salt Lake City: University of Utah Press. Excerpted in *Moral Discourse and Practice: Some Philosophical Approaches*, edited by Stephen Darwall, Allan Gibbard, and Peter Railton (1997). Oxford: Oxford University Press, 402. ↩
9. See D. D. Price, A. Von Der Gruen, J. Miller, A. Rafii, and C. A. Price (1985). Psychophysical Analysis of Morphine Analgesia. *Pain*, 22: 261–69; E. L. Foltz and L. E. White (1962). Pain ‘Relief’ by Frontal Cingulotomy. *Journal of Neurosurgery*, 19: 89–100; J. M. Gybels and W. H. Sweet (1989). *Neurosurgical Treatment of Persistent Pain*. Basel: Karger.; and P. D. Wall (2000). *Pain: The Science of Suffering*. New York: Columbia University Press. ↩
10. Robert Coghill. Pain: Making the Private Experience Public, in Murat Aydede, ed. (2005). *Pain: New Essays on Its Nature and the Methodology of Its Study*. Cambridge, MA: MIT Press, 301; and M. M. Morgan, M. M. Heinricher, and H. L. Fields (1994). Inhibition and Facilitation of Different Nocifensor Reflexes by Spatially Remote Noxious Stimuli. *Journal of Neuropsychology*, 72: 1152–60. ↩
11. Stanley Schachter and Jerome Singer (1962). Cognitive, Social, and Physiological Determinants of Emotional States. *Psychological Review*, 69: 379–99. ↩
12. Paul Brand and Philip Yancey (1993). *Pain: The Gift Nobody Wants*. New York: Harper Collins, especially pp. 194–96. ↩
13. Brand and Yancey, 194. ↩
14. Jaak Panksepp. Social Pain, Support, and Empathy, in Murat Aydede, ed. (2005). *Pain: New Essays on Its Nature and the Methodology of Its Study*. Cambridge, MA: MIT Press, 373. ↩

15. Panksepp, Social Pain, 376; J. Panksepp, B. H. Herman, T. Villberg, P. Bishop, and F. G. DeEsquinazi (1980). Endogenous Opioids and Social Behavior. *Neuroscience and Biobehavioral Reviews* 4: 473–87; and J. Panksepp, S. M. Siviy, and L. A. Normansell. Brain Opioids and Social Emotions, in M. Reite and T. Fields, eds. (1985). *The Psychobiology of Attachment and Separation*. New York: Academic Press, 3–49. ↩
16. Panksepp, Social Pain, 376; Panksepp, Herman, et al.; and Panksepp, Siviy, and Normansell. ↩
17. For a summary of this research and a discussion of its theoretical implications, see Michel Cabanac (1992). Pleasure: the Common Currency, *Journal of Theoretical Biology*, 155(2): 173–200. ↩
18. This view on the agent-neutrality of the badness of pain is also defended by Thomas Nagel (1986). *The View from Nowhere*. Oxford: Oxford University Press, pp. 158–62. ↩
19. For arguments to this effect, see Derek Parfit (1984). *Reasons and Persons*. Oxford: Clarendon Press, Part III. Richard Yetter Chappell summarizes Parfit’s arguments [here](#). ↩
20. Gerd Gigerenzer (2007). *Gut Feelings: The Intelligence of the Unconscious*. New York: Viking. ↩
21. John Stuart Mill, [On Liberty](#), reprinted in Albert William Levi, ed. (1963). *The Six Great Humanistic Essays of John Stuart Mill*. New York: Washington Square Press. ↩
22. Robert Nozick (1974). *Anarchy, State, and Utopia*. New York: Basic Books, 42–45. ↩
23. For a more detailed version of my thoughts on Nozick’s experience machine thought experiment, see Chapter 7 of [The Feeling of Value](#) or Sharon Hewitt (2010). What Do Our Intuitions About the Experience

Machine Really Tell Us About Hedonism?. *Philosophical Studies*, 151(3): 330–49. ↩

24. Thanks to Robert Wiblin for pointing out that the relevant difference here is our superfluity. ↩

25. John Stuart Mill, *Utilitarianism*, reprinted in Albert William Levi, ed. (1963). *The Six Great Humanistic Essays of John Stuart Mill*. New York: Washington Square Press, 252. ↩

26. Peter Singer (1981). *The Expanding Circle: Ethics and Sociobiology*. New York: Farrar, Straus, & Giroux. ↩