



# The Experience Machine Objection to Desire Satisfactionism

**ABSTRACT:** *It is widely held that the experience machine is the basis of a serious objection to hedonistic theories of welfare. It is also widely held that desire satisfactionist theories of welfare can readily avoid problems stemming from the experience machine. But in this paper, we argue that if the experience machine poses a serious problem for hedonism, it also poses a serious problem for desire satisfactionism. We raise two objections to desire satisfactionism, each of which relies on the experience machine. The first is very much like the well-known experience machine objection to hedonism. The second asks whether someone who accepts desire satisfactionism should want to form a desire to plug into the experience machine.*

**KEYWORDS:** welfare, hedonism, desire satisfactionism, the experience machine

The experience machine looms large in discussions of welfare. It seems clear to many that, at least for an ordinary person, a life unknowingly lived inside a sprawling virtual reality designed to please the person is not as good a life for the person living it as a life lived in the real world. For this reason, the experience machine is the basis of an objection to hedonistic theories of welfare—an objection that many regard as a serious or ‘near-fatal blow’ (Crisp 2006: 620), if not a truly fatal one.<sup>1</sup>

However, it is widely accepted that desire satisfactionist theories of welfare can readily avoid the objection (Crisp 2014: §4.2; Heathwood 2010: 650; Timmons 2013: 163; Griffin 1986: 10; Brink 1989: 224; Darwall 1997: 162). Thus, it is widely held that the experience machine objection provides both a motivation and a significant advantage for desire satisfactionism. This is perhaps one reason why, as L. W. Sumner notes, ‘Desire theories have come to dominate the welfare landscape in this century the way that hedonism dominated in its own time’ (1996: 113).

In this paper, however, we argue against philosophical orthodoxy and attempt to show that, if the experience machine provides the basis for a serious objection to

We would like to thank Chris Heathwood, Robert Pasnau, Jonathan Spelman, Alexander Zambrano, and two anonymous referees for *JAPA* for their helpful comments on earlier versions of the paper. Joseph would also like to thank the Alexander von Humboldt Foundation for its generous financial support of his current research endeavors.

<sup>1</sup> Those who consider it a definitive objection include Finnis (1980: 95–96), Brink (1989: 223), and Griffin (1986: 9). Of course, some argue that the experience machine does not pose a conclusive problem for hedonism, see, e.g., Crisp (2006: 635–42) and Feldman (2011: 67–72). For a summary of recent responses to the experience machine objection to hedonism, including responses that identify and empirically support the claim that aspects of the experience machine scenario bias readers’ judgments about the case, see Weijers and Schouten (2013).



hedonism, it also provides the basis for a serious objection to desire satisfactionism. We argue for this claim by raising two new objections to desire satisfactionism, each of which relies on the experience machine. The first is very much like the well-known objection to hedonism. The second asks whether someone who accepts desire satisfactionism should want to form a desire to plug into the experience machine. To be clear, we are not taking a position on how serious a problem it is for a theory of welfare to recommend plugging into the experience machine. However, we do show that, contrary to the conventional philosophical wisdom, *if* it is a problem to begin with, then it is a problem for hedonism *and* for desire satisfactionism.

Before turning to those objections, we consider briefly how the experience machine is used to object to hedonism and how desire satisfactionism is thought to deal with a parallel objection.

## 1. The Standard Experience Machine Objection

Since the publication of Robert Nozick's *Anarchy, State, and Utopia*, the experience machine has generally been taken as the basis of an objection to hedonistic theories of welfare (1974: 42–43). In its most general form, hedonism holds that how well a life goes for the one living it is determined by that life's balance of pleasure over pain.<sup>2</sup> If hedonism is correct, then it would presumably recommend plugging into the experience machine (although there are versions of hedonism that might avoid this result, such as what Feldman [2004: ch.5] calls 'Truth Adjusted Intrinsic Attitudinal Hedonism'). And not only that, but it seems that such a life would actually turn out to be the best possible life for the one living it, as the vast majority of hedonism's varieties would have it, at least. This is a problem for hedonism because many people have the intuition that not only would such a life not be the best possible life, but also that they themselves would be worse off inside the machine, even if their lives would have a higher balance of pleasure over pain. Of course, it is worth pointing out that thinking that ordinary people would not be better off in an experience machine is still compatible with thinking that some people might be better off inside an experience machine than they would be if they were to continue to live in the real world.

In contrast to hedonism, desire satisfactionism holds, roughly, that how well a life goes for the one living it is determined by that life's balance of desire satisfactions over desire frustrations.<sup>3</sup> Desire satisfactionism is thought to avoid the problem of recommending plugging into the experience machine because there are a number of

<sup>2</sup> Shelly Kagan (1994) has argued that there is a distinction between how well off an individual is and how well that individual's life is going. If one is convinced by Kagan's arguments, we intend to be discussing how well off an individual is, and we intend to capture that when using locutions like, 'how well a life is going for the one living it'.

<sup>3</sup> There are, of course, different versions of Desire Satisfactionism—e.g., actual versus idealized versions. We will consider the importantly different versions of Desire Satisfactionism proposed by Chris Heathwood, James Griffin, and John Harsanyi. Accordingly, when we talk of Desire Satisfactionism, we intend the most common varieties of Desire Satisfactionism. We do not make the further and far stronger claim that our arguments apply to *any conceivable* version of Desire Satisfactionism.

desires we cannot satisfy in an experience machine. As Mark Lukas suggests, ‘Life off the machine is better because we don’t merely want the pleasures associated with, say, climbing a mountain. We want to climb a mountain, and life on the experience machine frustrates this desire and others like it’ (2013: 3). In other words, ordinary people do not only desire to obtain, do, or experience things *virtually*. They desire to obtain, do, or experience things *actually*, in the real world.

The ability to respond to the experience machine objection has been a major motivation for desire satisfactionism. As Lukas puts it, ‘desire theories have an easy time here [with the experience machine]’ (2013: 3). Chris Heathwood agrees that ‘Preferentism . . . nicely handles the experience machine problem’ (2005: 650). Roger Crisp notes that the ability to avoid the experience machine objection is ‘one motivation for the adoption of a desire theory’ (2014: §4.2). Indeed, Nozick’s own discussion raises the possibility that the problem with life in the experience machine is that it does not satisfy our deepest desires: ‘Perhaps what we desire is to live (an active verb) ourselves, in contact with reality’ (1974: 45).<sup>4</sup> So there is widespread agreement that, although hedonism may struggle to cope with the experience machine, desire satisfactionism has an easy go of it.

In the next section, we question whether the desires of ordinary people are really such that life outside of the experience machine is better from the standpoint of desire satisfactionism. According to that objection, desire satisfactionism does recommend plugging into the experience machine, and thus does not give the result frequently attributed to it. In section 3, we suggest that even if ordinary people’s desires did not weigh in favor of plugging into the experience machine, desire satisfactionism would nevertheless recommend adopting a different desire set that *would* be satisfied by plugging in. According to this objection, desire satisfactionism implies that, even if people do not currently want to plug in, they *should* want a desire set of the sort that would be satisfied by plugging in.

It is important to get clear on the precise conclusion we argue for here. We do not argue that there is *no difference whatever* in how hedonism and desire satisfactionism evaluate a life lived in the experience machine. As we noted, life in the experience machine has the greatest possible balance of pleasure over pain. But even if desire satisfactionism ultimately recommends plugging in, this falls short of the much stronger claim that life in the experience machine has the greatest possible balance of desire satisfactions over desire frustrations. Thus, while hedonism implies that life in the experience machine is unfailingly the best possible life for the person living it, desire satisfactionism does not imply the same. Our conclusion is that hedonism and desire satisfactionism are similar vis-à-vis the experience machine in a more limited respect: namely, that both recommend that

<sup>4</sup> We merely say that Nozick’s discussion raises this *possibility* because Nozick’s own view was decidedly not that this is the *actual* problem with life in the experience machine. Nozick subsequently clarified that he did not take the experience machine to be a reason to endorse desire satisfactionism: ‘I am not saying simply that since we desire connection to actuality the experience machine is defective because it does not give us whatever we desire—though the example is useful to show that we *do* desire some things in addition to experiences—for that would make ‘getting whatever you desire’ the primary standard. Rather, I am saying that the connection to actuality is important whether or not we desire it—that is *why* we desire it—and the experience machine is inadequate because it doesn’t give us *that*’ (Nozick 1989: 106–107).

an ordinary person plug into the experience machine. This is the sense in which we claim that if the experience machine poses a problem for hedonism, it also poses one for desire satisfactionism.

## 2. The Standard Objection and the Weight of Experience

The first argument in short is this. Desire satisfactionism is thought to avoid the experience machine objection because an ordinary person has many desires that cannot be satisfied in the experience machine. But this is too quick, because to avoid the objection it is not enough that an ordinary person merely have *some* desires that cannot be satisfied in the experience machine. Those desires must be more numerous, more intense, or some combination of the two, as compared to the desires that would be satisfied in an experience machine, but not the real world. We argue that this is not the case. For an ordinary person, far more desires overall would be satisfied and far fewer desires frustrated in the experience machine than in the real world. What is more, even if for an ordinary person there were a difference in intensity between desires regarding the real world and desires regarding experiences, that would not be enough to overcome the sheer number of experiential desires that would be satisfied and the sheer number of experiential desire frustrations that would be avoided in the experience machine. Therefore, not just according to hedonism, but according to desire satisfactionism as well, a life in the experience machine is better for an ordinary person than a life in the real world. We defend this argument piece by piece in what follows.

### 2.1. Ordinary Desires Satisfiable in an Experience Machine

In order to see why, for an ordinary person, far more desires would be satisfied and far fewer desires frustrated in the experience machine than in the real world, it is important to recognize, first, that some desires can be satisfied in an experience machine. For example, desires for sensations (such as for the taste of chocolate cake) can be satisfied in an experience machine, as can desires for the experience of engaging in various sorts of activities, such as reading or playing chess. More generally, any desire that we have about things *seeming* to be the case can, in principle, be satisfied in the real world or in an experience machine. Call desires that can be satisfied by an experience machine ‘experiential desires’. Desire satisfactionism is thought to avoid the experience machine objection because not all of our desires are experiential desires. Indeed, it seems that many of our most intense desires are for things actually to be the case in the real world and not merely to seem to be the case. Call desires of this sort ‘nonexperiential desires’. Put in this new terminology, then, desire satisfactionism is supposed to avoid the standard experience machine objection because it maintains that, at least ordinarily, human beings have a number of reasonably strong nonexperiential desires sufficient to weigh against plugging in.

But is this true? Evaluating this claim is made easier by distinguishing further between classes of desire. Consider, first, among nonexperiential desires the class of

desires for things largely outside of our control. Many ordinary people desire things of this kind, such as a significant decrease in the violent crime rate, world peace, or the development of more effective treatments for cancer.<sup>5</sup> If one is an ordinary person and not, for example, an important world leader or a brilliant scientist, whether these nonexperiential desires are satisfied is almost entirely outside of one's control. As a result, one's plugging into an experience machine rather than staying in the real world will make almost no difference with respect to the frustration or satisfaction of these desires. Of course, whether it will *seem* to you as if these sorts of desires have been satisfied may be changed by plugging in or not—but whether these desires are *in fact* satisfied would not be affected one way or the other by your plugging in.

Second, there are nonexperiential desires for things more in our control in the real world. First and foremost, there is the desire to be connected to reality in the first place—a desire satisfied by remaining in the real world and frustrated by entering an experience machine. We might think here also, for example, of the desire to climb a mountain, to have a successful marriage, or to receive a promotion. As we are all aware, desires of this sort are very often frustrated, despite our best efforts. However, some desires like these are satisfied in the real world and would be frustrated in the experience machine, since as nonexperiential desires they could not be satisfied, for example, by virtual mountains, virtual marriages, or virtual promotions. With respect to these desires, as desire satisfactionists point out, life is better outside an experience machine than inside.

The final class of desires to take into account is the category of plain experiential desires. This class of desires includes all desires about pleasures and pains and sensations or purely psychological states more generally (although, if one formed desires concerning the causal sources of these things, then it would be possible for desires about these sorts of things to be nonexperiential - see [section 2.2.2](#) below). For example, the desire to be rid of back pain, for the taste of a decadent chocolate cake, or for the feeling of a tropical breeze are all experiential desires. Also, for example, if one takes being happy to be a purely psychological state, then a desire for happiness is also an experiential desire. Finally, the desires that we have that certain things *seem* to be the case are experiential desires. For example, we do not only desire that we have a partner who loves us. We also want it to *seem* to us to be the case that we have a partner who loves us; surely, we would be disappointed if it seemed to us that our partners did not love us! So too we do not only desire that we be healthy, or good at our jobs, or well-liked. We also want it to seem to us to be the case that these things are so. Again, all of these desires about things seeming to be the case are also experiential desires. Many of the desires in this class could, in principle, be satisfied either in an experience machine or in the real world.

<sup>5</sup> These and some of our other examples are not instances of self-regarding desires. One might think that only the satisfaction of roughly self-regarding desires contributes to one's well-being; Overvold (1980, 1982), for example, argues that only desires for states of affairs that entail the existence of the desirer (when the relevant states of affairs obtain) are relevant to her well-being. However, Chris Heathwood, for instance, points out that, if this were so, the success of a serious fan's favorite sports team would not contribute to her well-being—which seems implausible (2014: 215). We find a view like Overvold's less plausible. However, if one finds such a view more promising, it will make no difference in the overall argument that follows.

However, it is worth noting that, in the experience machine, one could, at least in principle, also satisfy experiential desires for things that are not possible in the real world, such as the feeling of flying without technological aids.

If we think carefully about these three classes of desire and the number of desires of each kind we are likely to satisfy in the real world as opposed to in the experience machine, it seems clear that we would satisfy far more desires in the experience machine than in the real world. After all, when it comes to nonexperiential desires for things largely outside of our control, whether one is in an experience machine or in the real world would have almost no bearing on the number of desires satisfied or frustrated. When it comes to nonexperiential desires for things somewhat more in our control, one would satisfy more and frustrate fewer desires in the real world than in an experience machine, but given the propensity of the world to withhold things that we want, very many of these desires would also be frustrated in the real world. However, when it comes to experiential desires, one would satisfy far more and frustrate far fewer desires in an experience machine than in the real world. Indeed, *all* of our many and varied desires related to pleasure, pain, sensations, psychological states, and seemings that together make up a huge class of desires would be satisfied in an experience machine (assuming they're consistent).<sup>6</sup> But obviously very many of these desires would be frustrated in the real world. When taken together, then, for an ordinary person, far more desires would be satisfied and far fewer desires frustrated in the experience machine than in the real world.

There are two main strategies the desire satisfactionist might employ to resist this conclusion. First, the desire satisfactionist might try to show that the *number* of desires satisfiable in the experience machine is smaller than we argue it to be. [Section 2.2](#) considers two arguments for significantly shrinking the class of desires that can be satisfied in an experience machine. Second, the desire satisfactionist might try to show that for ordinary people, desires satisfiable in an experience machine are far less *intense* than nonexperiential desires. [Section 2.3](#) considers an argument for this view. We consider and argue against each of these proposals in turn.

## 2.2. Arguments Regarding the Number of Desires

**2.2.1. *Conditional Desires to the Rescue?*** To begin, one might think that in ordinary people desires for seemings are not freestanding or unconditional desires. One might think that they are, instead, conditional in nature. For example, one might think that ordinary people do not desire that it seem to them that their children love them, full stop. Rather, ordinary people desire that, *if* it is true that their children love them, *then* it seem to them that their children actually love them. And, more generally, perhaps experiential desires for seemings are often of the conditional form, 'If it is true that the relevant nonexperiential desire is satisfied, then I want it to seem to

<sup>6</sup>Of course, for people with ordinary psychology, unmitigated desire satisfaction may result in something akin to hedonic adaptation. Hedonic adaptation involves the waning of one's affective responses to circumstances that endure (see, e.g., Frederick and Loewenstein 1999: 302–29). In light of this, a well-programmed experience machine would probably include both a variety of experiences and even an occasional desire frustration, which could be instrumentally useful in maintaining the intensity of other desires.

me to be the case that it is satisfied'. And presumably the correlative judgment also applies: 'If the relevant nonexperiential desire is *not* satisfied, then I do not want it to seem to me to be the case that it is satisfied'. If this were so, then a very large number of desires for seemings could not be satisfied in an experience machine because the condition that it be true that the relevant nonexperiential desires be satisfied could not be met in an experience machine.

However, this is not ultimately plausible. To see why, consider a pair of cases in which parents have children who do not love them. Let us assume that the children's lack of love is a permanent condition, and therefore knowing that the children, in fact, do not love the parents does not in any way help the parents to obtain the love of their obstinate children. Second, let us assume that there is nothing nearly so important to the parents as their children. Now, in the first case, it seems to the parent that her children love her. In the second case, it seems to the parent that her children do not love her. Perhaps ordinary people would prefer to be the parent in the second case because they would prefer to know this hard truth, rather than live a lie. Or perhaps they would prefer to be the parent in the first case since then they would at least feel that their greatest desire is satisfied, even if it can never actually be satisfied. For our purposes, it does not matter. What matters is that ordinary people feel at least some pull toward being the blissfully deceived parent. And, indeed, it seems clear ordinary people feel at least *some* pull in that direction. And, if they do, then there must be something about the life described in the first case that they desire, even if ultimately most ordinary people would prefer to know the truth that their greatest desire has been and will forever be frustrated (that is, even if their desire for truth is more intense than their desires for various things connected to the pleasant deception - a possibility we consider in [section 2.3](#)).

There are two possible candidates for what ordinary people desire in the blissfully deceived life described. The most straightforward is that, as we have argued, ordinary people have freestanding or unconditional desires that even important things seem to be the case. And, if that is right, then this subclass of experiential desires can be satisfied in an experience machine. Another possibility is that the parent to whom it seems that her children do not love her is in considerable distress because of it and does not experience the myriad pleasures concomitant to a life with the (seeming) love of her children; on the other hand, the other parent does experience these pleasures and avoids that distress. If this explanation generalizes, then what might initially appear to be desires for seemings could be construed as desires for pleasure and desires to avoid pain. But if this is true, then these desires could also be satisfied by plugging into the experience machine. Thus, regardless of how this subclass of desires is construed—as desires for seemings or as desires for pleasure and avoidance of pain or a combination—given desire satisfactionism, they can be satisfied by plugging in. Accordingly, when we talk about experiential desires paired with nonexperiential desires, such desires can be thought of as desires for things seeming to be the case, desires for pleasures associated with seeming to get what one wants, or some combination of the two. But the important point is that regardless of how this class of desires is understood, they can be satisfied in an experience machine, and so the number of desires that can be satisfied in an experience machine cannot be curtailed in this way.



2.2.2. *The Desire for Reality to the Rescue?* Still, one might pursue a different avenue in order to argue that the number of desires that can be satisfied in an experience machine is significantly smaller than we have suggested. We have already acknowledged that ordinary people desire reality. Perhaps our desire for reality is so pervasive that it finds its way even into our desires for experiences. For example, it has been assumed that a desire for the taste of chocolate cake could be satisfied in the experience machine or in reality. But perhaps a desire for the taste of chocolate cake fully spelled out is a desire for the taste produced by a *real* chocolate cake. If desires for experiences (when fully spelled out) always or even very frequently append information about the causal sources of the relevant experiences and if the causal source generally appended is *reality*, then it would turn out that the experience machine would also frustrate a large number of desires for experiences. Accordingly, the class of truly experiential desires (i.e., those satisfiable in an experience machine) would be much smaller than would be required for the argument to go through.

Again, we acknowledge that ordinary people desire to be in touch with reality. Most people would not only be shocked, but also disappointed to discover that they had been living in a virtual world. However, this is generally because of nonexperiential desires: ordinary persons who wake up from the machine would be disappointed that, for example, their friends and family were merely virtual. However, we take it that they would not be disappointed, for example, that the taste of the chocolate cake they enjoyed came from a virtual source.

This is further illustrated by everyday attitudes toward these desires. Whether my back pain results from an actual slipped disc or a sophisticated virtual representation as of my having slipped a disk, what I want is for my back pain to stop. So too, whether the taste of chocolate cake results from an actual chocolate cake or a sophisticated virtual representation as of chocolate cake, what I want is for that pleasurable sensation to continue. Indeed, if people even now could eat virtual cake, enjoying the same sensations as those associated with eating real cake, there is little doubt that people would do so—the marketing scheme, ‘All the flavor, no calories!’ writes itself. Indeed, it seems likely that such a product would be exceedingly popular. When it comes to pain and pleasure considered in themselves, we take it that it is precisely the experience of pain that ordinary people generally want to avoid and precisely the experience of pleasure that ordinary people generally want to have. Since desires regarding pleasures and pains are among the most important desires for experiences, this radical proposal that desires for experiences, when fully spelled out, are almost always desires for experiences *causally connected with reality* is thus not plausible.

But perhaps there is another, more plausible way the desire for reality might tell against the judgment that the experience machine represents a better life for ordinary people than life in the real world. The desire for reality may not be built into more or less every other desire, but the desire for reality taken in itself, along with other desires that do include reference to reality (i.e., nonexperiential desires), also need to be weighed seriously. Because even if the desire for reality taken in itself plus nonexperiential desires more generally are not as numerous as the radical proposal would have it, perhaps these desires are so intense that they outweigh the



desires for mere experiences and seemings that would be satisfied in an experience machine. We consider this kind of proposal next.

### 2.3. Intensity of Desire to the Rescue?

One might agree with our conclusions about the *number* of desires satisfied in an experience machine relative to the number satisfied in the real world and still maintain that for an ordinary person desire satisfactionism recommends life in the real world. After all, the number of desire satisfactions is not all that matters, according to desire satisfactionism. Intensity matters, too. If ordinary nonexperiential desires for things partially in our control are far more intense than experiential desires, then, according to desire satisfactionism, an ordinary life lived in the real world is better than one lived inside an experience machine.

The plausibility of this response depends in part on how various versions of desire satisfactionism construe the relevant sort of intensity. Here, we consider two well-known options that differ in important ways. First, one might think of the intensity of desire as a kind of felt intensity (Heathwood 2005). On this proposal, a desire is intense when I really, strongly feel that I want something. Second, one might think that the relevant sort of intensity or strength of desire has to do with the place of that desire in one's desire set (Griffin 1986: 14–15). On this view, the most intense desire is a global desire—that is, the desire to live a certain kind of life. The place of global desires in one's desire set is, as it were, at the top, determining the sorts of more local desires that one should have, and so intensity of desire in this sense tracks the structure of the desire set.

As we noted, the desire for reality taken in itself is one obvious candidate for a nonexperiential desire that may be significantly more intense than many of our experiential desires. But on either interpretation of intensity, the desire for reality taken in itself is not sufficiently intense to outweigh the myriad other desire satisfactions available in the experience machine. We can see this by reflecting on how ordinary people spend their free time—that is, time during which they can do what they most want. Very frequently free time is used to 'unplug' from reality. Indeed, many of the most popular pastimes in the developed world today are precisely about escaping reality into virtual or fictional worlds: watching movies, watching TV, reading or listening to novels, and playing video games. If the desire to be in touch with reality were extraordinarily intense in terms of how it feels, then we would not expect people to use their free time in this way. So too, if the desire to be in touch with reality were extraordinarily intense in the sense that it was at or near the very top of ordinary peoples' desire sets, we would expect them to organize their lives very differently. Indeed, we would expect them to use their free time to do almost anything other than unplug from reality. So it does not seem plausible to think that our desire for reality is intense enough to outweigh all of the desires we would satisfy in the experience machine, but which would go frustrated in the real world. Perhaps these trends are lamentable; perhaps they show that people's desire for reality, while extant, is not as strong as it ought to be. But this complaint, whatever one thinks of its merits, is not one that is available to the desire satisfactionist.

Of course, this may just show that one cannot put the burden of this response on a single desire, even the desire for reality. Perhaps it is nonexperiential desires in general that outweigh their experiential counterparts. Indeed, it seems there are good reasons to think that such nonexperiential desires are far more intense. We can see this by appealing to the well-known case of the deceived businessman (Kagan 1998: 34–36). For the businessman, his most global and most strongly felt nonexperiential desires are the same and so these desires count as intense, whichever view of desire intensity one accepts. As is well-known, in the businessman's case, very many things that he intensely desires *seem* to be so. In our terminology, then, the experiential desires paired with his strongest nonexperiential desires are all satisfied. He seems to have a wife who is devoted to him, children who love him, and coworkers who respect him; they treat him precisely as they would if they in fact loved and respected him. However, the businessman's nonexperiential desires for these things are all frustrated; these people all pretend to care for the businessman in order to achieve their own aims, but his wife despises him, his children hate him, and his coworkers loathe him. It seems that, from his point of view, it would be better if his wife, children, and coworkers were in fact to like him somewhat, rather than merely seeming to like him a great deal. This is readily explicable if the businessman's nonexperiential desires for these things are far more intense (in either sense) than the experiential desires paired with those nonexperiential desires. Thus, perhaps nonexperiential desires are far more intense than their paired experiential counterparts, in which case, as was noted, desire satisfactionism still has a way out of the experience machine objection.

However, a variant of the deceived businessman case can be used to show that it is implausible to think that a person's nonexperiential desires, even for things that we intensely desire (in either sense), are sufficiently strong to outweigh the experiential desire satisfactions available in the experience machine. To see this, imagine an inverted deceived businessman case. Instead of being deceived into thinking that his family and coworkers like and love him, imagine that the businessman is deceived into thinking that his family and coworkers dislike and hate him. In other words, imagine that his family and coworkers, in fact, hold the businessman in very high regard and love him a great deal. But perhaps because he has been scarred by past deceptions, it seems clear to the businessman that his family and coworkers do not like or love him. His experiential desire that his family and coworkers seem to like and love him is frustrated, but his nonexperiential desire that they actually like and love him is satisfied.

If our nonexperiential desires were much more intense than the experiential desires with which they are paired (either in terms of subjective pull or in terms of their centrality in our system of desires), then it should be clear to us (from the outside) that the businessman is much better off in this inverted scenario as compared to the original scenario (notice that the number of desires considered in each case is the same). But if it is to be preferred at all, the businessman does not seem to be *significantly* better off in the second scenario than in the first. That is, it does not seem that nonexperiential desires, even those related to our central life concerns, are much more intense than the experiential desires with which they are paired. And so it seems that, for an ordinary person, the difference in intensity

between experiential and nonexperiential desires is not enough to overcome the sheer number of desires that would be satisfied and the sheer number of frustrations that would be avoided in the experience machine.

Finally, it is worth noting that the desire satisfactionist cannot avoid the conclusion that non-experiential desires are not sufficiently intense to outweigh experiential desires that would be satisfied in an experience machine by appealing to idealized or ‘true’ desires—i.e., the desires that one ‘*would* have if he had all the relevant factual information, always reasoned with the greatest possible care, and were in a state of mind most conducive to rational choice’ (Harsanyi 1982: 55). If idealizing desires were to help, it would have to be the case that our idealized nonexperiential desires are ordinarily significantly stronger than our idealized experiential desires. But the two businessman cases considered above give us good reason to think that our idealized nonexperiential desires are not significantly stronger than our idealized experiential desires. After all, we already do, more or less, stand in the place of an idealized counterpart of the businessman in making our judgments about those cases; we have a great deal of salient information that the businessman and the inverted businessman both lack, and we are able to reflect and reason coolly and with great care about their lives. If our judgment is that the inverted businessman is not significantly better off than the first businessman, then we have every reason to believe that that judgment would be shared by the businessman when suitably idealized. This shows that we have reason to believe that our idealized nonexperiential desires are not significantly stronger than our idealized experiential desires. Thus, considering idealized or ‘true’ desires does not help the Desire Satisfactionist.<sup>7</sup>

And so the experience machine objection used against hedonism seems to apply to desire satisfactionism as well. For those who think that ordinary people would be worse off inside the experience machine than in the real world, both hedonism and desire satisfactionism are put into question as theories of welfare.

### 3. The Experience Machine and What One Should Want Out of Life

The argument of the previous section hinged on what desires ordinary persons actually have. In this section, we argue that even if it were not true that ordinary persons have such desires, according to desire satisfactionism they *should* want a desire set of the sort that would be satisfied by plugging in.<sup>8</sup> This objection is thus different in kind from both the experience machine objection to hedonism and the

<sup>7</sup> It may also be worth noting that a combination of Griffin’s view of intensity of desire and idealization of desires does not allow one to escape the problem. We occupy an idealized position vis-à-vis the businessmen. Even if we confine our evaluation of the case to the most global of the businessmen’s desires, one is not significantly better off than the other. So we still have no reason to think that the nonexperiential desires, however global or idealized or both, are significantly more important to well-being than the experiential desires with which they are paired.

<sup>8</sup> That is, if one lives in a world where the experience machine exists. However, this condition is not at issue in the debate over the experience machine; it is not generally thought that Desire Satisfactionism has an advantage over hedonism with respect to the experience machine, but only until the day when the machine is finally built.

argument just developed with respect to desire satisfactionism. Those arguments rely on the claim that, given the actual psychological makeup of ordinary people, they would be better off in the experience machine given either theory of welfare. By contrast, this objection hinges on the sorts of second-order desires, or desires about desires, that we should have according to desire satisfactionism. As we shall explain, if one finds the recommendation to plug into the experience machine disturbing, then one should also find disturbing the recommendation that one should want to have the kind of desires that would be satisfied by plugging in.

In outline, the argument is this. The experience machine provides an extremely reliable source of good experiences. If someone *were* to desire only good experiences and the best life for herself, then plugging into the experience machine would be rational, since the experience machine provides a more reliable source of good experiences than the real world. When it comes to their own welfare, it happens that most people do not only desire good experiences, but from the point of view of desire satisfactionism, this is merely an incidental psychological fact. If that were not true—if, say, people could choose any set of desires whatsoever—then so long as an experience machine were available desire satisfactionism would actually recommend a desire set that consists entirely in wanting good experiences. For if someone were to have such a desire set, she would not only be deeply satisfied by plugging into the experience machine, she would be satisfied extremely reliably—indeed, she would be satisfied much more reliably than if she were to remain in the real world. Accordingly, desire satisfactionism implies that, in a world in which there is an experience machine and one could choose any desire set whatsoever, a person interested in getting the best life for herself should choose to desire life in the experience machine. And if one finds it troubling that hedonism implies that one should plug in to the experience machine, one should also find it troubling that desire satisfactionism implies that one should want to plug in.

At least that is the gist of the argument. We spell it out more precisely as follows.

### 3.1. The Implications of Desire Satisfactionism for the Welfare Value of Desire Sets

Desire satisfactionism is primarily the view that one's welfare depends upon whether one's desires are satisfied or frustrated. But when conjoined with other facts, the theory also has the conceptual resources to say whether a given desire set has a high or low expected welfare value. Desire sets composed of very intense desires, for instance, will *ceteris paribus* have a higher welfare value than desire sets composed of weak desires. (This is true given either of the senses of intensity we have mentioned. Using Griffin's sense of intensity, the claim would be that, *ceteris paribus*, global desires have a higher expected welfare value than local desires.) So too desire sets composed of desires that are more easily satisfied will *ceteris paribus* have a higher welfare value than desire sets composed of desires that are difficult to satisfy.<sup>9</sup>

<sup>9</sup> This captures the relevant features of desire from the point of view of desire satisfactionism. Many theorists agree that the strength or intensity of a desire is relevant to one's welfare—e.g., Heathwood (2005: 490), Griffin

Because desire satisfactionism implies a way to measure the expected welfare value of desire sets, it provides guidance not only in evaluating how well one's life is going, but in what desires one should form, if one wants one's life to go well. In other words, desire satisfactionism can inform us about what we should want to want. Consider an example. Suppose you have been invited to an unusually fancy cocktail party and you are thinking about trying the high-end caviar the hosts are serving. Being a person of modest means, you have never before tried caviar, so right now you do not have a desire for it. (You have a more general desire to eat tasty things, but you are not sure if caviar is one of those.) Yet you hear plenty of the guests proclaim what exquisite caviar it is, and you reasonably believe that if you tried it, you would enjoy it and form a concomitant desire for high-end caviar in the future. In this situation, you are confronted by a question not just about what you do want but what you should want: should you want to eat high-end caviar? On the one hand, you do desire tasty things and it is likely you will find it very tasty. On the other hand, you are unlikely to get more of it in the future; right now you do not know what you are missing, and so you do not have occasional unrequited cravings for the taste of caviar. You reason that, on balance, forming a desire to eat high-end caviar will make you worse off even when you factor in the satisfaction that comes from eating the caviar this time. Desire satisfactionism can explain this judgment in terms of the expected welfare value of the desire for eating expensive caviar on the grounds that while the desire would be intense, it would not be easily satisfiable.

Thus, if one wants to maximize one's welfare, desire satisfactionism, when taken together with facts about the world, implies judgments about which desires one should form. More specifically, it recommends desire sets that are (a) the most intense and (b) the most easily satisfiable.

### 3.2. The Welfare Value of the Desire to Live in the Experience Machine

We can now ask what the expected welfare value is for the desire set consisting entirely in desires for experiences, which can be satisfied in an experience machine. The first axis of evaluation, intensity, is subjective—that is, what determines the level of intensity of a desire is the individual's psychological makeup. This is true whether one takes intensity to be a matter of how a desire feels or where a desire fits in one's desire set. We have argued that from the point of view of desire satisfactionism, the desire sets of ordinary people are such that living in the experience machine would be better than the alternative. But even if one denies this, this is an entirely contingent fact. If people could choose any psychology they wanted, the desire for life in the experience machine could be as intense as any other desire set. And from the point of view of desire satisfactionism, this desire set could then have as high an expected welfare value as any other desire set when it comes to considerations of

(1986: 14), and Lukas (2013: 2)—though as we have seen, there are different accounts of what intensity amounts to. Moreover, desire frustrations are captured by the criterion of ease of satisfaction since desire frustrations occur more when a desire cannot be satisfied easily. Some discussions also include duration (see, e.g., Heathwood 2005: 490), but whether one includes duration or not does not affect the cogency of the argument.

intensity. Obviously, in the real world one cannot just choose a psychology for one's self, but as we shall see, this implication of desire satisfactionism will nevertheless have important consequences for the theory.

By contrast, unlike intensity, the ease of satisfaction of a desire is not entirely subjective.<sup>10</sup> Rather, it depends partially upon factors in the world, factors that the agent does not have complete control over. To recall our caviar example, the desire for high-end caviar is not easily satisfiable for a person of modest means since caviar is expensive. (And more generally, the world supply of caviar is dependent upon a whole host of factors.) When we consider desires for good experiences constitutive of the desire for life in the experience machine, we can see that the experience machine can provide good experiences very reliably. This is not the case with the real world: your favorite breakfast cereal is discontinued; your chosen career proves untenable; the person you love ends up loving another. Even fairly basic desires, such as the desires for food, shelter, and physical safety, can be frustrated by the vicissitudes of life. Of course, some people get lucky and have lives where their desires are largely satisfied. But no one can say with justified confidence that one's life will in fact go that way. Even if all one wants out of life is good experiences, the world is not a wish-fulfillment machine. But where the real world fails, the experience machine succeeds. After all, if all one wants out of life is good experiences, then the experience machine *is* a wish-fulfillment machine. It is true that the experience machine does not *guarantee* only good experiences. It is possible that the experience machine could malfunction, lose power, or have a suboptimal program design. (Although, occasional suboptimal experiences may provide a needed source of contrast or variety and thus instrumentally make the other experiences better.) Nevertheless, it is still the case that the experience machine produces good experiences significantly *more* reliably than the real world. In a world where the experience machine exists, desiring good experiences would be as easily satisfiable a desire set as any.

### 3.3. Desire Satisfactionism and What One Should Desire

We can now see why desire satisfactionism implies that if one could choose any desire set whatsoever, one should choose to desire life in the experience machine. If one could choose any desire set whatsoever, desiring life in the experience machine could be (a) the most intense desire set. And in a world where it really exists, the desire for good experiences constitutive of desiring life in the experience machine would be (b) as easily satisfiable a desire set as any. And as we have seen, if one wants to maximize one's welfare, desire satisfactionism implies that one should adopt the desire sets that are (a) the most intense and (b) the most easily satisfiable. That is, if you had your pick of desire sets, desire satisfactionism would recommend that you desire life in the experience machine.

<sup>10</sup>It is to some extent subjective, inasmuch as the same desire could be more or less easily satisfiable, even given the same external circumstances, depending upon an agent's psychology. For example, if there are two people who are equally situated and both desire to fall in love, that desire may be more easily satisfiable for one than the other—say, if one of them is more willing to trust.

It may be objected that this argument is just a special case of a more general criticism of desire satisfactionism. The more general criticism is that because desire satisfactionism only evaluates desires on the basis of intensity and ease of satisfiability, it recommends adopting easily satisfiable but otherwise impoverished desires. (For a typology of such criticisms, though none rely solely on ease of satisfiability, see Heathwood 2005). For instance, desire satisfactionism might recommend the desire set consisting solely of an intense desire to count blades of grass. (The example here is taken from Rawls [1971: 432], but, to be clear, this is not the use to which Rawls puts the example.) Alternatively, we could run a similar argument to produce the conclusion that desire satisfactionism recommends intensely desiring things *exactly as they are* at any given time. Adopting such a desire set would produce an immense gain in welfare, without necessitating any changes in one's activities, skills, or life plans.<sup>11</sup>

The objection here is correct that the argument of this section is one version of a larger category of objection to desire satisfactionism. However, this version is of special significance because it reveals that such an objection applies to the experience machine. And as we have noted, the philosophical orthodoxy is that the experience machine poses no problems for desire satisfactionism—indeed, that it helps *motivate* desire satisfactionism.

Moreover, for those who believe that it is a serious problem for any welfare theory to recommend plugging into the experience machine, this conclusion casts serious doubt on desire satisfactionism. Indeed, the consequences of desire satisfactionism may even be more disturbing on this score than given the argument of the previous section. For those who agree that, as we are, it would be better for us to plug in, it would be tempting to try to soften the blow by noting what is undeniable—that we are seriously imperfect creatures. One might say that we often desire things we should not desire and take pleasure in things we should not even when considered strictly from the point of view of our own welfare. One might say that while life in the experience machine may not be an ideal life, this is merely a reflection of the fact that we are not ideal creatures. Yet, if the argument of this section is sound, then this strategy is entirely unavailable to the desire satisfactionist. For desire satisfactionism implies that of all the possible psychologies one could have, from the flawed to the faultless, it would be highly advisable to choose the one in which we exclusively desire life in the experience machine. Desire satisfactionism does not just imply that we want good experiences enough that we should plug in; it implies that it is what we should *want to want*.

<sup>11</sup> It might produce a greater gain in welfare than even the desire set for life in the experience machine (followed by plugging in) because although the experience machine is by stipulation well-programmed, it is not perfectly programmed or perfectly functioning, and there may be some mismatch between one's desires and the outputs of the machine. Thus, a desire for things being exactly as they are will, on this argument, be desire satisfactionism's top recommendation. However, the theory will still strongly recommend the desire for life in the experience machine in the sense that even if it is not the very top recommendation, it is (under the conditions already specified) a desire set with a very high expected welfare value, and it is therefore highly advisable.



## 4. Conclusion

It is widely believed that not only does desire satisfactionism avoid any problems with the experience machine, but also that it can explain why many have the intuition they do about a life lived in it. Accordingly, avoiding the experience machine objection is a prominent reason given in favor of adopting desire satisfactionism. However, contrary to philosophical orthodoxy, we have argued that at least for those who think ordinary people would not be better off inside the experience machine, the experience machine provides no reason to prefer desire satisfactionism to hedonism. If the experience machine causes problems for hedonism when it comes to ordinary lives, it also causes problems for desire satisfactionism. And, of course, if it is true that at least for an ordinary person a life in the experience machine is not as good a life for the person living it as a life lived in the real world, then our arguments cast serious doubt on the adequacy of desire satisfactionism itself.

DAN LOWE

UNIVERSITY OF COLORADO AT BOULDER

[daniel.p.lowe@gmail.com](mailto:daniel.p.lowe@gmail.com)

JOSEPH STENBERG

HUMBOLDT-UNIVERSITÄT ZU BERLIN

[joseph.stenberg@hu-berlin.de](mailto:joseph.stenberg@hu-berlin.de)<sup>12</sup>

## References

- Brink, David Owen. (1989) *Moral Realism and the Foundations of Ethics*. Cambridge: Cambridge University Press.
- Crisp, Roger. (2006) "Hedonism Reconsidered." *Philosophy and Phenomenological Research* 73, no. 3: 619–45.
- Crisp, Roger. (2014) 'Well-Being'. In Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. <http://plato.stanford.edu/archives/win2014/entries/well-being/>.
- Darwall, Stephen. (1997) 'Self-Interest and Self-Concern'. *Social Philosophy and Policy*, 14, 158–78.
- Feldman, Fred. (2004) *Pleasure and the Good Life: Concerning the Nature, Varieties, and Plausibility of Hedonism*. Oxford: Oxford University Press.
- Feldman, Fred. (2011) 'What We Learn from the Experience Machine'. In Ralf M. Bader and John Meadowcroft (eds.), *The Cambridge Companion to Nozick's Anarchy, State, and Utopia* (Cambridge: Cambridge University Press), 59–86.
- Finnis, John. (1980) *Natural Law and Natural Rights*. Oxford: Oxford University Press.
- Frederick, Shane, and George Loewenstein. (1999) 'Hedonic Adaptation'. In Daniel Kahneman, Ed Diener, and Norbert Schwarz (eds.), *Well-Being: The Foundations of Hedonic Psychology* (New York: Russell Sage Foundation), 302–29.
- Griffin, James. (1986) *Well-Being: Its Meaning, Measurement, and Moral Importance*. Oxford: Clarendon Press.

<sup>12</sup> We are equal coauthors; the order of authors has been determined alphabetically.

- Harsanyi, John C. (1982) 'Morality and the Theory of Rational Behavior'. In Amartya Sen and Bernard Williams (eds.), *Utilitarianism and Beyond* (Cambridge: Cambridge University Press), 39–62.
- Heathwood, Chris. (2005) 'The Problem of Defective Desires'. *Australasian Journal of Philosophy*, 83, 487–504.
- Heathwood, Chris. (2010) 'Welfare'. In J. Skorupski (ed.), *The Routledge Companion to Ethics* (London: Routledge), 645–55.
- Heathwood, Chris. (2014) 'Subjective Theories of Well-Being'. In Ben Eggleston and Dale E. Miller (eds.), *The Cambridge Companion to Utilitarianism* (Cambridge: Cambridge University Press), 199–219.
- Kagan, Shelly. (1994) 'Me and My Life'. *Proceedings of the Aristotelian Society*, 94, 309–24.
- Kagan, Shelly. (1998) *Normative Ethics*. Boulder, CO: Westview Press.
- Lukas, Mark. (2013) 'Desire Theories of the Good'. In Hugh LaFollette (ed.), *The International Encyclopedia of Ethics*. Malden, MA: Wiley-Blackwell.
- Nozick, Robert. (1974) *Anarchy, State and Utopia*. New York: Basic Books.
- Nozick, Robert. (1989) *The Examined Life: Philosophical Meditations*. New York: Simon & Schuster.
- Overvold, Mark Carl. (1980) 'Self-Interest and the Concept of Self-Sacrifice'. *Canadian Journal of Philosophy*, 10, 105–18.
- Overvold, Mark Carl. (1982) 'Self-Interest and Getting What You Want'. In Harlan B. Miller and William H. Williams (eds.), *The Limits of Utilitarianism* (Minneapolis: University of Minnesota Press), 186–94.
- Parfit, Derek. (1984) *Reasons and Persons*. Oxford: Oxford University Press.
- Rawls, John. (1971) *A Theory of Justice*. Cambridge, MA: Harvard University Press.
- Sumner, L. W. (1996) *Welfare, Happiness, and Ethics*. Oxford: Oxford University Press.
- Timmons, Mark. (2013) *Moral Theory: An Introduction*. 2d ed. Plymouth, UK: Rowman & Littlefield.
- Weijers, Dan, and Vanessa Schouten. (2013) 'An Assessment of Recent Responses to the Experience Machine Objection to Hedonism'. *Journal of Value Inquiry*, 47, 461–82.