

# Exchanging IRL for VR

---

## Experience Machines and Opportunity Costs

By Richard Volkman

Virtual Reality (VR) experiences have arrived for a wide audience, and some worry that VR experience entices us to abandon experience In Real Life (IRL). While such worries cannot be dismissed out of hand, they are liable to be imprecise and distorted without careful analysis of the opportunity costs of VR experience. Generally, we are not liable to abandon real experience for its virtual analogue. Even a perfect VR experience is generally not a perfect substitute for a like experience IRL. Instead, the success of VR will depend on its ability to produce experiences valued for what they actually are.

Keywords: virtual reality, value theory, experience machine

Categories: • *Software and its engineering~Virtual worlds software*

Corresponding Author: *Richard Volkman*

Email: *volkmanr1@southernct.edu*

## Welcome to VR. Now What?

In the summer of 2016, like thousands of other tech enthusiasts, I invested a significant chunk of time and money into acquiring and setting up all the necessary gear to try out the first wave of serious, consumer-ready virtual reality (VR). I was suitably amazed by my first look around in VR, I was positively giddy the first time I tried to lean against a virtual table and nearly fell over, and I eagerly explored every tech demo, indie game, and VR experience I could find. However, I eventually came to see how all these short experiences and demos were working out, the novelty of it all wore off a bit, and I started to wonder, “Well, that was kinda neat and everything, but what is this VR stuff really going to be good for? What comes next?”

Unlike most of those thousands of other enthusiasts, this made me think of Robert Nozick. As I will argue, what VR can and cannot deliver is revealed in light of his “experience machine” case: Imagine a machine could make it seem to you that you were having any experience you wanted. Would you use it? Nozick thinks not. Suitably unpacked in appropriate detail, this thought experiment is widely considered decisive against hedonism, showing that most of us value something other than the mere experience of pleasure, something that cannot be reduced to experience alone. What humans value besides pleasure is somewhat obscure, but it has something to do with *truth* and *reality*. Since VR excels at delivering immersive experiences untethered from anything we would ordinarily call *truth* or *reality*, perhaps VR cannot deliver what matters most to us.

That is hardly a new or startling suspicion. Indeed, long before VR was widely available to consumers, Heim suggested, “An unrestrained proliferation of worlds cries out for

sanity, for connection with reality, for metaphysical grounding.”<sup>1</sup> Lurking in all this, however, is a curious concern pulling in opposite directions; VR experience is at once impoverished compared to experience In Real Life (IRL) and at the same time compelling enough to distract from affairs IRL. Fittingly, the most striking statement of this concern comes in a VR experience, “Stor Eiglass,” in which one glides through a virtual world all dazzling and vibrant before falling to ruin and decay unnoticed by the world’s inhabitants, who are obliviously elsewhere behind VR headsets. The message is clear enough: We must guard against the temptation to exchange what we care about IRL for whatever pale imitation we might get in VR.

While we cannot deny that VR *might* become the 21<sup>st</sup> Century equivalent of crack cocaine, there is good reason to doubt individuals will generally exchange VR experience for experience IRL. Nozick fully understood this. He did not raise the experience machine to denounce it as an irresistible temptation but from the conviction that, “We learn that something matters to us in addition to experience by imagining an experience machine and then realizing that we would not use it.”<sup>2</sup> This turns the worry about VR on its head, transforming it into something much closer to the actual worry of today’s early adopters of VR and the content creators and developers hoping to serve them: What is VR really good for? To frame this question, as I will argue we must, in terms of what we should generally expect consumers to do is not to ask an ethical question of narrow interest to philosophers and other moral scolds, but to ask a practical question of special concern to developers and consumers of VR experiences. Nonetheless, a careful answer in these terms turns out to have substantive normative implications. I am a philosopher after all, and not an economist, though I hope to avoid also being a moral scold.

I will borrow standard analytic tools from microeconomics, including especially the notions of opportunity cost and substitution, along with specific claims in empirical psychology, to argue that consumers are not liable to abandon real experience for its virtual analogue. More precisely, even a perfect VR experience is generally not a perfect substitute for a like experience IRL. Developers and content creators will struggle with consumers unless they go beyond trying to simulate reality in their virtual worlds. Instead, the success of VR will depend on its ability to produce experiences valued for what they actually are. This means successful VR will be unlike an experience machine in important respects. Following a more detailed unpacking of these key ideas, the paper turns to examples of current VR titles to illustrate the thesis, indicate its significance, and illuminate what is at stake where the analysis is less than straightforward.

## **An Experience Machine is No Place to Find Yourself**

The Internet Encyclopedia of Philosophy article on hedonism observes that Nozick’s experience machine “has proven to be so convincing that nearly every single book on ethics that discusses hedonism rejects it using only this argument or this one and one

---

<sup>1</sup> The Metaphysics of Virtual Reality. Michael Heim. In *Virtual Reality: Theory, Practice, and Promise*. Ed. Sandra Heslel and Judith Roth. Information Today, 27-34, 1991.

<sup>2</sup> *Anarchy, State, and Utopia*. Robert Nozick. Basic Books, 44, 1974.

other.”<sup>3</sup> However, the widespread excerpting of these few paragraphs from *Anarchy, State, and Utopia* tends to obscure their wider context for many readers. This is not necessarily a bad thing, since a measure of excellence for a thought experiment must consider argumentative fecundity—the degree to which it spawns all sorts of interpretations and spin-off arguments—and the experience machine has proven remarkably capable of inspiring all manner of back-and-forth. However, before proceeding with our discussion of what the experience machine says about VR, we need to pin down a particular understanding of the case, and I will argue that the most appropriate interpretation for our purposes is one situated within Nozick’s wider project in political philosophy.

While Nozick means to indicate something troubling about hedonism, it is a plain mistake to suppose he means to reject hedonism in favor of some alternative value theory. Rather, the case is part of a broader project of defending deontological political theory against *any* approach that would assert “priority of the good.” Nozick generally rejects using any particular account of human welfare to specify the boundaries that define distinct individuals because he believes appeals to welfare cannot justify crossing those boundaries. On this view, the right, conceived in terms of “side-constraints” and ultimately spelled out in the familiar individual rights of Lockean liberalism, must be understood to be prior to the good. In that case, one does not respect persons by seeing to their welfare; rather, what counts as good for one depends on one sorting out for one’s self the nature of the good, such that respect for persons requires above all else giving individuals room to become who they are for themselves by their own lights.

This context is helpful for understanding what the experience machine case says and does not say about VR. Whatever other interpretation one might like to give of the core idea, the case in its original context does not promote “axiological brain-in-a-vatism”<sup>4</sup> nor any desire fulfilment theory of well-being<sup>5</sup> nor any other complete or partial theory of value. To the contrary, Nozick explores the place of value theory in a wider moral and political project, and—whatever else it might also do—the experience machine illustrates the case against using any preconceived value theory as the ultimate basis of moral judgment. The case begins a dozen pages earlier, where Nozick cites the familiar “underlying Kantian principle that individuals are ends and not merely means.”<sup>6</sup> On this view, moral requirements are grounded in capacities to formulate ends and not in the ends themselves. Nozick offers the experience machine case shortly after introducing moral side-constraints and returns to the discussion of side-constraints thereafter, concluding that the ultimate basis of moral side-constraints is found in “the meaning of life. A person’s shaping his life in accordance with some overall plan is his way of giving meaning to his life; only a being with the capacity to so shape his life can have or strive for meaningful life.”<sup>7</sup> Seen in this light, Nozick’s point in exploring the experience machine is to

---

<sup>3</sup> Hedonism. Dan Weijers. Internet Encyclopedia of Philosophy – <http://www.iep.utm.edu/hedonism/> - Accessed 13/01/2017.

<sup>4</sup> Against Brain-in-a-Vatism: On the Value of Virtual Reality. Jon Cogburn and Mark Silcox. *Philosophy & Technology*, 27 (4), 561-579, 2014.

<sup>5</sup> The Value of Virtual Worlds. Jonny Søraker. University of Twente, 169, 2010.

<sup>6</sup> *Anarchy, State, and Utopia*. Robert Nozick. Basic Books, 30-31, 1974.

<sup>7</sup> *Ibid.*, 50.

illustrate and defend unfettered individual agency in the formulation and pursuit of a life plan against alternative views according to which what makes life go well or badly does not depend on such agency but precedes it.

Consider the particulars of the case, where Nozick writes, “you would think and feel you were writing a great novel, or making a friend, or reading an interesting book” and goes on to explain our reluctance to enter the experience machine by emphasizing, “we want to *do* certain things, and not just have the experience of doing them... we want to *be* a certain way, to be a certain sort of person”<sup>8</sup> Although talk of what “we want” invites supposing some desire fulfilment theory at work, it is clear from the full context that this is not what Nozick intends, for he professes studied agnosticism about what constitutes the good life, at least with respect to political philosophy. The *Utopia* section of *Anarchy, State, and Utopia* proposes a filter device for empirically discovering the good life that operates like a marketplace of virtue implemented across competing communities instantiating diverse ways of life and experiments in living, and Nozick does not expect any one way of life or set of values to emerge as ultimately best for one and all. “No one should attempt to describe a utopia unless he’s recently reread, for example, the works of Shakespeare, Tolstoy, Jane Austen, Rabelais, and Dostoevsky to remind himself how different people are.”<sup>9</sup> The explicit agnosticism with respect to the good that Nozick advocates for his “framework for utopias” refuses to endorse hedonist, desire fulfilment, and objective list theories alike. Whatever the good life turns out to be for particular individuals will be an empirical matter to be worked out for themselves on their own individual terms in response to a rich and evolving social and physical environment. Nozick supposes they desire what they variously desire *because* it is good (a point explicitly connected with the experience machine in his later work, *The Examined Life*), rejecting the view that what they desire is good *because* they desire it, but there is no reason to suppose in advance that whatever objective values individuals orient themselves towards will have to appear the same for each individual on any single, finite list. The framework does not take sides on these matters, since that is what it means to respect persons as distinct individuals. Ultimately, the point is to discover rather than impose a system of stable associations that approximates an ideal model of stable association for distinct individuals in which, “All admire each other’s individuality, basking in the full development in others of aspects and potentialities of themselves left relatively undeveloped.”<sup>10</sup>

Proponents of welfarist accounts of ethics and politics are liable to misunderstand Nozick’s broader point if they suppose one cannot do substantive normative work without some determinate theory of value. Hedonist critics of the experience machine case, like Feldman<sup>11</sup> or Silverstein<sup>12</sup>, take it to be an argument against hedonism whose

---

<sup>8</sup> Ibid., 42.

<sup>9</sup> Ibid., 311.

<sup>10</sup> Ibid., 306.

<sup>11</sup> What We Learn from the Experience Machine. Fred Feldman. In *The Cambridge Companion to Robert Nozick’s Anarchy, State, and Utopia*. Ed. John Meadowcroft and Ralf Bader. Cambridge University Press, 59-88, 2011.

<sup>12</sup> In Defense of Happiness: A Response to the Experience Machine. Matthew Silverstein. *Social Theory and Practice*, 26 (2), 279-300, 2000.

premises include some reference to what any person would or should do when given the opportunity to enter the machine, but Nozick need not suppose each would answer the same. He can allow there are hedonists in the crowd, and it has been my experience discussing the case that there is one in every crowd. However, if *anyone* agrees with Nozick that entering the experience machine is a bad idea (and Nozick clearly thinks most people will agree with him), then hedonism is presumptively false not because it proves there is necessarily something intrinsically good other than pleasure but because refusal is evidence that pleasure does not exhaust what is intrinsically good for all people. In short, individuals might reasonably disagree about the experience machine. To find out who is right, we implement something like the framework and let them work it out for themselves. While this involves no conceptual refutation of hedonism, there is equally no conceptual requirement that everyone must be a hedonist. The case reveals most people not to be hedonists. At least, they give the appearance of not being hedonists; determining the truth requires further investigation. Most importantly, to find out whether you are or are not a hedonist will require sustained investigation by you, and this effort to discover the point of your own life cannot succeed if you are rapt in an experience machine carefully designed from top to bottom to satisfy some prior assumption about what counts as good for you. Such an assumption would be the epitome of what Kant called “heteronomy of the will.” Agreeing to enter the experience machine amounts to begging the question *against yourself* on the important matter of what constitutes a life worth living.

Nozick proposes to treat value theory as a matter of ongoing empirical discovery in the course of individuals working out for themselves what counts as the best sorts of life for the diverse individuals they are, and this cannot take place in an environment that presupposes what will make life go well. Hedonism proposes to know in advance what Nozick insists cannot be known (and perhaps cannot even *be*) prior to our experiments in living working themselves out. You cannot discover the good life in the experience machine programmed to satisfy your every desire any more than you could discover new physics in a simulated environment programmed to Newtonian specifications. Moreover, without also rejecting moral side-constraints, any retreat from broad, reductionist value hedonism (e.g., Bentham) to a narrower prudence hedonism (e.g., Feldman, Silverstein, Søraker), according to which hedonism is a theory of human welfare while moral and other sorts of value need not be reduced to pleasure, merely concedes Nozick’s main point that morality requires individuals to sort out the value of their lives for themselves and the agnosticism with respect to the good that this entails in the formulation of policy. Once it is admitted that the right is prior to the good, we will have to await the judgments of distinct individuals regarding the nature of the good as this is revealed in the unfolding of their actual lives in a world that exercises their capacities for formulating and judging conceptions of the good, so the nature of the good is not something that can be settled by conceptual fiat in advance of that process. Nozick is betting the sorting of the framework will result in deep pluralism, but the official stance of the framework is neutral agnosticism. This agnosticism with respect to the good clearly does not lead to normative paralysis, at least not with respect to politics, since it implies that what we *ought* to do is to set up some approximation of the framework for utopias and sort out what ways of life are best by trial and error experimentation, just as we sort out any other empirical matter.

## Pleasure Does Not Work in an Experience Machine

Empirical evidence supports Nozick's suspicion that most folks would not enter the experience machine because pleasure comes to us entangled with our cognitive responses to the full concrete nature of the world. Paul Bloom shows that concerns about how things "really" are matter not only for motivation but for the very experience of pleasure itself. *How Pleasure Works* documents that the experience relies on beliefs, values, and attitudes and cannot be disentangled from these. To introduce his point, Bloom recounts the story of Nazi war criminal Hermann Goering, imprisoned and awaiting execution, discovering that his prized painting, which he had thought was a Vermeer, was actually a forgery. When informed of this deception, "Goering looked 'as if for the first time he ha[d] discovered there was evil in the world.'"<sup>13</sup> Bloom argues that "pleasure is deep. What matters most is not the world as it appears to our senses. Rather, the enjoyment we get from something derives from what we think that thing is."<sup>14</sup>

Bloom gives example after example of how our beliefs and expectations impact not only the value we assign to our experiences but the experiences themselves, including cases relating to gustatory, aesthetic, and sexual pleasures. In one study, subjects thought beer with a bit of vinegar added to it tasted better than unaltered beer, unless they were told the vinegar had been added before they tasted it, in which case they report it tastes worse<sup>15</sup>. In another study, different subjects were told that an abstract painting took either 4 hours or 26 hours to paint. Those told it was more arduous to create rated the painting higher in terms of quality, value, and liking<sup>16</sup>. A series of studies especially relevant to our discussion of VR involved a "duplicating machine," which makes perfect copies of real-world objects. As Bloom wryly notes, "Not all duplicates would be worth the same as the originals... You might put a Picasso into the machine, or your wedding ring, or your Tupac signature, but then you would be careful to keep the duplicates separate, since they would be worth much less than the originals. Copying your hamster, dog, or child would have its own special moral and emotional consequences."<sup>17</sup> Of course, actual duplicating machines do not exist, but researchers constructed a suitably deceiving lookalike to fool their test subjects—young children. Kids got a great kick out of "duplicating" toys and rocks and other things, but when asked to assign values to objects, it mattered what sort of thing was duplicated. If an object (say, a spoon) was said to have been once owned by a queen, children would agree it was worth more than an ordinary spoon, but they did not think that value was duplicated by the machine. A queen's spoon was special, but an exact physical copy of a queen's spoon is just an ordinary spoon. "Substances can be duplicated; history cannot."<sup>18</sup>

Bloom accounts for these results by way of "psychological essentialism" (not to be confused for the metaphysical thesis of the same name), "the notion that things have an underlying reality or nature that one cannot observe directly and it is this hidden nature

---

<sup>13</sup> *How Pleasure Works*. Paul Bloom. W.W. Norton & Company, 1, 2010.

<sup>14</sup> *Ibid.*, xii.

<sup>15</sup> *Ibid.*, 47-8.

<sup>16</sup> *Ibid.*, 141.

<sup>17</sup> *Ibid.*, 108.

<sup>18</sup> *Ibid.*, 110.

that really matters.”<sup>19</sup> This human propensity to care about the hidden nature of things makes the concrete history and origin of a person, experience, or thing relevant to determining its value, the amount of pleasure one may justly take in it, and the amount of pleasure one actually will take in it. A stuffed animal made from the duplicating machine may have the same aesthetic, economic, and use value as the original, but if the original is the attachment object of a specific child, then only the original will do. It may be similar, but it is not the same because it lacks the essence that makes the original the thing it is—or at least this is what our brains reliably project onto certain things. “Even the most mundane things have histories. This is their essence. And for some of these objects...this essence is a source of great pleasure.”<sup>20</sup>

At a minimum, psychological essentialism means that not only do the vast majority of persons care about something other than pleasure, but they have to care about these other things *because* they care about pleasure. More profoundly, it suggests that what folks ultimately care about is sure to be the outcome of some complicated tangle of cognitive and affective states spelled out in the ultimate particular of an individual life, such that any reduction in terms of this or that value theory is sure to miss the mark. Even a robust pluralism is bound to fall short if it appeals to some finite and fixed list of goods, just as one cannot account for all the variety of worldwide cuisines by enumerating what is discerned by our taste buds. This would vindicate Nozick’s main contention that we have no choice but to respect persons by allowing their psychological apparatuses to work out for themselves what is valuable in light of their own rich contexts, contingencies, and histories, since it suggests there is no account of what matters independent of all this that could justify infringing on the working of these apparatuses. In other words, we shall have to respect persons by listening to what they say they value and not by appeal to any antecedent value theory they do not themselves endorse.

Nonetheless, Bloom’s psychology is no knock down refutation of hedonism. There is room for the hedonist to dig in and insist that it our psychology is just mistaken about all this—that our attachment to essences is irrational since all that *really* matters is pleasure qua pleasure. More plausibly, the prudence hedonist can admit that the true history and origin of our pleasures matters to us and even has value (a move that becomes more plausible insofar as one emphasizes the import of “attitudinal” rather than “sensory” pleasures), while insisting that this value is not really part of what we mean by “well-being” or “welfare.” Soraker suggests something like this in his formulation of “Confidence Adjusted Intrinsic Attitudinal Hedonism,”<sup>21</sup> but one has to wonder at some point how much of the original thesis is left when hedonism no longer means, “Pleasure is the only intrinsic good,” and starts to look more like, “There are lots of kinds of different goods, including moral goods and aesthetic goods and epistemic goods, and maybe others, but well-being consists entirely of pleasure as the only intrinsic welfare-enhancing good.” At some point, it looks like we are just defining welfare as pleasure in an empty tautology instead of offering an illuminating account.

---

<sup>19</sup> Ibid., 9.

<sup>20</sup> Ibid., 115.

<sup>21</sup> The Value of Virtual Worlds. Jonny Søraker. University of Twente, 191, 2010.

There is surely much more to be said about this, but we can press forward without settling once and for all how much of hedonism is left over after Bloom. As emphasized above, the main point of the experience machine does not turn on what one says about hedonism. So long as one acknowledges that one must be agnostic about the good in setting up the basic side-constraints that govern social interaction, hedonism can be safely considered as one hypothesis among others competing in a framework for utopias, and even hedonists will have good reason to refrain from entering the experience machine, since they will have to acknowledge they might be wrong *about hedonism*, and it would be unwise to beg the question *against themselves* on this important matter. On this picture of things, further study of the nature of the good life is always warranted, and this seems to be a permanent feature of the human condition. As long as we are living, we are never finished figuring out what makes life worth living.

### **Experience in VR is No Substitute for Experience IRL**

To recap, the experience machine is helpfully conceived less as a deductive argument or an intuition pump and more as an illustration of what is at stake if we fail to respect the distinctness of persons. Significantly, Nozick's discussion proceeds from talk of experience machines to the "transformation machine," which automagically makes one into whatever sort of person one would like to be, and the "results machine," which brings about whatever result one chooses. These machines have nothing to do with hedonism or other experiential value theories and everything to do with re-emphasizing the real point: "What is most disturbing about them is their living of our lives for us."<sup>22</sup> While the case surely pumps certain intuitions about what is valuable, what matters for the argument is that the reader accepts that individuals should work these things out for themselves in a rich world that exposes them to a diversity of challenges and alternative ways of life, something even hedonists can accept if only they eschew violating side-constraints in the name of pleasure. Given our ongoing study of the psychology of pleasure, anyone has good reason to be reluctant to enter an experience machine.

This understanding of the experience machine case has at least two important implications for understanding what VR might be good for: 1) Despite certain obvious similarities, VR as we know it today or expect it in the near future is profoundly unlike the experience machine, and 2) The analytic tools we adopt for figuring out what VR is good for should be agnostic with respect to the ultimate nature of the good.

Remember, what is "disturbing" about the experience machine, the transformation machine, and the results machine is the way these machines subvert agency. Unlike these machines, VR does not typically do this. To illustrate, consider Nozick's first suggestion about what one might like to have happen in the experience machine: "you would think and feel you were writing a great novel." As should be clear by now, what is objectionable about this for Nozick, what makes it a subversion of agency, is that one has not actually written a great novel. That is, Nozick is not imagining a case in which one enters the experience machine and then actually writes an actually great novel on a virtual computer, since that would still be actually writing a great novel. The experience

---

<sup>22</sup> Anarchy, State, and Utopia. Robert Nozick. Basic Books, 44, 1974.

machine is disturbing if it gives you the mere pleasure of *incorrectly* believing you wrote a great novel when you really did not, a pleasure that Bloom teaches us would surely disappear if the truth became known. That would be the plain abnegation of your agency. There is very little danger that VR will simply introduce false evaluations regarding the essential nature of our experiences. Behind the headset, I generally still know what I am actually doing. For all you know, dear reader, I *actually* wrote this essay in VR, with my Vive strapped on and running *Word* in *Virtual Desktop*. How could you tell?

This echoes a point made by Mooradian in his contrast between “grounded virtual karate” and “ungrounded virtual karate.”<sup>23</sup> He argues that a VR karate game that only loosely approximates the physical movements of karate (imagine something like the fighting game *Mortal Kombat* in VR) must not be mistaken for actual karate, since the person who has mastered this ungrounded simulation has not mastered anything like actual karate. However, a more sophisticated VR set up could at least in principle ground certain elements of karate more adequately. Such a system probably could not approximate sparring sufficiently well to make for a grounded experience, since relevant feedback is lacking without a physical target receiving blows, but perhaps the moves that constitute floor exercises could be sufficiently well modeled in VR such that one who learned in the simulation could perform these moves tolerably well IRL. However, this indicates no pressing practical concern, at least with respect to currently available VR, since the typical master of ungrounded virtual karate no more confuses that skill for mastering actual karate than the master of *Mortal Kombat* confuses gaming skills for actually knowing how to behead opponents with his bare hands.

The less obvious and more important implication of the discussion so far is the need to adopt analytic tools that are agnostic with respect to value theory. Fortunately, economists have developed several well-honed analytic tools that help us think about consumer behavior and guide conduct in light of this by sorting out what consumers actually care about instead of proceeding from what anyone thinks they ought to care about. Economists use opportunity costs to determine the rationality of some course of action based on the forgone benefits of alternatives evaluated from the agent’s own perspective. The opportunity cost is the value of the next best good the agent could have realized by choosing to spend her resources other than she did. For example, spending a sum of money on a cup of coffee means forgoing all other uses for that money, so the benefits of having the cup of coffee must equal or exceed whatever benefits could have resulted from other uses if the purchase is rational. Relevant to this accounting for consumer choice, economists take into consideration substitution of some goods for others. One good is a substitute for another insofar as a consumer can choose either good to achieve roughly the same benefits. Substitution admits of degree, since a substitute may not realize the exact same benefits as the good it substitutes for or to the exact same degree, even as it gets the job done well enough to justify any tradeoffs in terms of overall costs and benefits, where costs are spelled out in terms of benefits forgone (i.e., opportunity costs). It follows that as the price of a good increases, demand for its substitutes will rise. For example, consumers might generally prefer Starbucks coffee to the store brand, but for many consumers the store brand will do in a pinch (it may not

---

<sup>23</sup> Virtual Reality, Ontology, And Value. Norman Mooradian. *Metaphilosophy*, 37 (5), 673-690, 2006.

taste as good, but it is a hot caffeinated beverage). As the price of Starbucks coffee increases, more and more consumers will substitute the store brand for Starbucks. If the store brand is the relevant opportunity cost of buying Starbucks at the initial price, then the increase in price will increase opportunity costs, since buying Starbucks will mean forgoing other goods in addition to the store brand, and at some point, consumers will prefer the bundle of goods that includes the store brand plus other goods over the bundle that includes Starbucks and fewer other goods. In this case, the store brand is a *net substitute* for Starbucks, since consumers would prefer Starbucks if the prices were more similar, but they prefer the store brand in light of net costs and benefits. Consumers settle for net substitutes in acknowledgment of tradeoffs imposed by the scarcity of resources; they are not getting everything they would value if they did not have to budget their resources. Demand for net substitutes is sensitive to price, but consumers will generally pay more for their most preferred good than its net substitute. Less common are *perfect substitutes*, which are maximally price sensitive; the opportunity cost of a perfect substitute is zero. For example, there is no reason to demand otherwise similar coffee from this or that pot, so rational consumers will consistently choose whichever costs less.

Putting all this together, we can begin to articulate more precisely what VR is and is not likely to be good for. No matter how perfect the simulation of the real world, a VR experience will never be a perfect substitute for any qualitatively similar experience IRL with respect to which psychological essentialism prevails. Insofar as VR gives us an experience that is a mere duplicate of experience IRL, it will not deliver contact with the purported essential nature of the objects of that experience—the actual concrete history that confers so much of the meaning of things in our actual experience of them—and it will therefore leave something out that matters to us. Furthermore, no VR experience that subverts the exercise of human agency will be a perfect substitute for any qualitatively similar experience IRL. Insofar as VR tries to give us an experience of having done something that we have not actually done, it will leave something out that matters to us. Given the exact same price of each experience, hardly anyone would exchange these experiences IRL for like experiences in VR.

This is a very narrow conclusion, since most substitutes are not perfect substitutes, but it sets the stage for dealing with the more usual case of net substitution by making clear what sorts of experiences are and are not the likely opportunity costs in choosing a VR experience. Given the role of essentialism and the exercise of our agency in what we value, VR experience is only a plausible net substitute if it comes at an extreme price discount, and even then, it will be such an imperfect substitute that the experience IRL is not liable to be a relevant opportunity cost of the VR experience. To illustrate, suppose one would like to see the Mona Lisa. As Bloom documents, no reproduction will be a perfect substitute for seeing the real thing, but most of us cannot afford to travel to Paris just to see the painting, so we settle for images in books or on a screen. This is not exchanging the experience IRL for something else, however, since the experience IRL is not within one's budget. The opportunity cost of looking at the reproduction is not actually traveling to Paris to see the real painting—that is not what one forgoes in order to see the reproduction. Indeed, if her budget suddenly allows, we would not be surprised to see her travel to Paris to see the original, even though she already knows what the

painting looks like. Given typical human psychology, there is no good substitute for an encounter with the original.

Soraker makes a similar point in his analysis of the value of virtual worlds<sup>24</sup>. He correctly notes that experience IRL is very often not the relevant opportunity cost of choosing to engage with a virtual world, but he unfortunately proceeds as if choosing the virtual has no opportunity cost. Economists emphasize that there are *always* opportunity costs for everything we do, since we are always managing scarcity. At a minimum, the time spent finding and enjoying a freely downloaded reproduction of the Mona Lisa comes at the expense of all the alternative uses of that time. The key to understanding what VR is actually good for requires identifying the typical opportunity costs of VR experience, and we have seen why these costs are seldom qualitatively similar experiences IRL, but to say what these costs typically are means discerning what would have been the next best choice after VR. That is the relevant benefit one forgoes in choosing a VR experience. I submit that in most cases the relevant opportunity cost of a VR experience will be some alternative *virtual* experience, in the relevant extended sense in which being immersed in a book or movie or video game puts one in a virtual world. (For further discussion of the various meanings of “virtual,” see Brey<sup>25</sup>.) One does not typically exchange the experience IRL of seeing the Mona Lisa for the experience in VR, but one very plausibly might exchange the experience of seeing a reproduction in a book or on a screen for the VR alternative. That is, the opportunity cost of a VR experience is typically an experience from a book, movie, game, or some other media. This analysis in terms of opportunity cost and substitution not only tells us that VR does not typically substitute for experience IRL but also points the way to understanding what VR is actually good for. Just as successfully getting away from a bear does not require running faster than the bear but only running faster than the others also running from the bear, so successful VR does not require delivering experiences as good or better than IRL; successful VR just requires delivering experiences as good or better, all things considered, compared to other media experiences. Note that this is not a conceptual point but an empirical claim grounded in basic microeconomics. To go beyond the armchair empiricism of all this requires an examination of concrete examples, the beginnings of which will occupy the remainder of this essay.

## Evaluating What Does Not Happen IRL

The following discussion of specific VR experiences is not intended as a critique or review of their overall merits. Rather, these cases illustrate and confirm the main results of the analysis so far by indicating how those results illuminate our independent responses to these experiences. In light of all this, we will be in a better position to finally say what VR is and is not actually good for.

*EverestVR* will never be a perfect substitute for actually climbing mountains. No matter how good the technology gets, the relevant opportunity cost of *EverestVR* will generally

---

<sup>24</sup> The Value of Virtual Worlds. Jonny Søraker. University of Twente, 88-90, 2010.

<sup>25</sup> Virtual Reality and Computer Simulation. Philip Brey. In Handbook of Information and Computer Ethics. Ed. Ken Himma and Herman Tavani. John Wiley & Sons, 361-84, 2008.

be watching a documentary on a flat screen, not actually climbing actual mountains. Actual climbers typically value the accomplishment, not just the experience, and this is not a shortcoming that can be remedied without subverting agency in the manner of the experience machine. Similarly, the opportunity cost of playing VR games like “Longbow” or *Thrill of the Fight* is not archery or boxing IRL. Rather, these and other games are generally in competition with similar games on a flat screen. VR holds up remarkably well in comparison to gaming on the flat screen, since gaming does not generally give rise to essentialist or agency concerns. *Arizona Sunshine* is as good a zombie shooter as I have played.

These games compete favorably as games, and they are not substitutes for any other experience IRL, but they do point to a way VR can substitute for certain activities IRL, since each might count as training of a sort. Getting good at “Longbow” is not the same as getting good at archery IRL, but it may be that getting good at “Longbow” is valued for what it is (a game with its own specific skill set) and not as a substitute for something else (archery). It is not what is *virtual* about “Longbow” but precisely what it *really* is that is decisive. One may get good at *virtual* archery in the game, but one gets *actually* good at “Longbow.” If getting good at “Longbow” is compelling enough, then it may displace other activities IRL as a net substitute, but that will be because it succeeds as a game. The relevant opportunity costs will be other games or game substitutes. In cases where VR explicitly serves as training for some activity IRL (something “Longbow” cannot be supposed to do, but which more sophisticated simulations might accomplish), these grounded VR experiences are better understood as complements rather than substitutes for experience IRL. That is, insofar as one trains in VR for the sake of the experience IRL, demand for each moves together. While VR training may substitute for training IRL, training in VR complements the experience IRL. Even if this results in a net reduction of time spent on the activity IRL, the time spent in VR will be a means to activity IRL and will be valued for what it actually is. VR experience qua experience is not substituted for the experience qua experience IRL.

Perhaps the most plausible candidate for VR experience substituting for experience IRL is VR porn, but even here the typical opportunity cost of VR porn is more likely to be flat screen porn rather than actual sex. Replacing porn on the flat screen for a more satisfying VR experience does not seem a bad trade on its face, but the prospect of forgoing opportunities for actual sex in favor of VR experiences may sound disturbingly escapist. This makes straightforward sense on both essentialist and agency grounds; Bloom spends a whole chapter on “bedtricks,” establishing the salience of essentialism with respect to sexuality, and the distinction between sex and masturbation obviously involves notions of agency. If VR sex substitutes for sex IRL, it will not be as a perfect substitute. To what degree do individuals care less about these essentialist and agency matters and more about the raw feelz of sexual experience, and to what degree can VR deliver those raw feelz as well or better than sex IRL? The answers will determine the extent to which VR sex can serve as a net substitute for sex IRL, but individuals are sure to differ in their answers, and there is no way to tell in advance what proportion of individuals might come to prefer VR sex to the real thing. The foregoing analysis cannot settle this matter, but it does illuminate what is at stake. Supposing the opportunity costs of VR porn turn

out to be flat screen porn, porn may be a “killer app” of VR. There is no disputing that many consumers will pay more for what they regard as a higher quality pornographic experience.

A similar analysis predicts that VR social experiences, like *AltSpaceVR* and *RecRoom*, are more apt to substitute for other mediated experiences (e.g., online chat, television, social media) than for social experience IRL. This concisely explains Facebook’s big bet on VR in its purchase of Oculus.

The picture that consistently emerges from all these examples is that VR is not going to be good for replacing experiences IRL but for doing better or at least differently all the things we do with other media. Saying what that means more precisely will have to be a matter of ongoing, trial and error investigation, as the whole history of art, storytelling, journalism, scholarship, and all the other uses of media can attest. At the same time, this analysis explains why efforts like *The VR Museum of Fine Art* or the virtual travel app *Destinations* so far fall flat as anything more than novelty hawking tech demos. Since these experiences fail along essentialist lines—I still have never *really* seen Michelangelo’s David even though I walked around it in VR—they will have to outperform books and pictures and films to be compelling for more than novelty or as a glimpse of what is yet to come. Compared to other media, the resolution, comfort, and overall clunkiness of current VR systems simply does not measure up. This may change. One can already appreciate how VR gives a better sense of scale and dimensionality than flat screen representations of art objects or places, but the experience still has a long way to go. The key point is that there is no reason to suppose VR experience will ever replace seeing these things for one’s self IRL.

On the other hand, as with books, films, and other traditional media, some of the most effective uses of VR involve experiences that cannot be had in any other way. Just as one cannot get the full experience of Shakespeare unless it is performed in the original Klingon, so one cannot get the same visceral reactions of VR horror translated in any other medium. What VR will be especially good at is delivering for us things like this that have no analogue beyond VR and which therefore cannot be readily substituted for any other thing. Despite its schmaltzy and tired underlying narrative of “boy meets girl,” *LoVR* is a great example of how the medium offers unique opportunities for artistic expression and aesthetic experience. *LoVR* modulates affect to fit its narrative by visual and audio stimulus and not by narrative alone. By flashing and spinning of words and sounds in a richly immersive environment, the viewer is swept into the rapture of love at first sight. No description of the experience in text or images can do it justice.

Perhaps more profound than what VR can do for consumers are the implications for artists and creators of content. VR as a tool of human creativity is of concern not merely to a narrow class of developers and content creators but to everyone, especially since it has long been observed that the information age blurs the distinction between consumers and producers of content<sup>26</sup>. VR as a tool for accessible expressive creativity is already

---

<sup>26</sup>Digital Culture: Liberation that Was Not Meant to Be. Richard Volkman. Ethicomp2003, Universidade Lusitana, Lisbon, Portugal, 2003.

available in apps like *Tilt Brush*, which allows one to dynamically “paint” in a 3d virtual space and save these creations for others to enjoy, and consumer level hardware for producing VR videos and snapshots is on the horizon. It remains to be seen whether substantial numbers of people will prefer a VR remembrance of their vacations, weddings, and birthdays, but it is plain that this sort of VR cannot be accused of subverting human agency or replacing the essential value of the experiences hereby preserved for posterity. In the end, perhaps VR will be especially good for exercising our creative and expressive agency.

## Conclusion

The main implications of this analysis are already becoming apparent to content creators, consumers, and developers. In the early days of VR, “Very simple games brought in a huge amount of sales and positivity because they were the first to offer the ability to do ‘X’ in VR...These ‘X’ in VR games will always impress people who have never experienced this medium before, however offering the ability to try something in VR is no longer going to capture the interest of those that have pushed past that initial ‘wow!’ moment.”<sup>27</sup> To realize the full potential of VR, we need to stop thinking of it as a substitute for experience IRL and engage our full agency to create and appreciate experiences valued for what they actually are.

In some ways, these conclusions seem obvious, but they stand in marked contrast to the expectations of many working in VR. Jeremy Bailenson, Director of the Virtual Human Interaction Lab at Stanford University, has suggested that only experiences that are expensive, dangerous, impossible, or rare will be worth creating in VR<sup>28</sup>. The analysis I have offered implies that this should be no more true of VR than it is of any other medium. The apparent success of social VR, which approximates experiences IRL that are inexpensive, relatively safe, easy, and common, proves there is something amiss in Bailenson’s formulation. Indeed, many of the most successful VR efforts have little to do with being expensive, dangerous, impossible, or rare IRL. This comes as no surprise if the relevant opportunity costs of VR experiences are experiences with other media and not experience IRL. Bailenson’s formula makes sense if we suppose the value of VR must be as a substitute for experience IRL. I submit that is not what VR is especially good for.

---

<sup>27</sup>“X” in VR: Going Beyond. Jordan. Rank17 – <http://rank17.com/2017/01/06/x-in-vr/> - Accessed 15/01/2017.

<sup>28</sup>What Virtual Reality is Good For. Will Oremus. Slate – [http://www.slate.com/articles/technology/future\\_tense/2016/08/the\\_only\\_good\\_reasons\\_to\\_use\\_virtual\\_reality\\_and\\_the\\_current\\_vr\\_renaissance.html](http://www.slate.com/articles/technology/future_tense/2016/08/the_only_good_reasons_to_use_virtual_reality_and_the_current_vr_renaissance.html) - Accessed 15/01/2017.