

*Oysters and Experience Machines:
Two Puzzles in Value Theory*

RICHARD KRAUT

THE TANNER LECTURES ON HUMAN VALUES

Delivered at

Stanford University
April 19–20, 2017

RICHARD KRAUT is Charles and Emma Morrison Professor in the Humanities, in the Department of Philosophy, Northwestern University. His interests include contemporary moral and political philosophy and the ethics and political thought of Socrates, Plato, and Aristotle. Among his publications are *Against Absolute Goodness* (Oxford, 2011), *What is Good and Why: The Ethics of Well-Being* (Harvard, 2007), *Aristotle: Political Philosophy* (Oxford, 2002), *Aristotle, Politics Books VII and VIII: Translated with a Commentary* (Clarendon Aristotle Series, 1997), *Aristotle on the Human Good* (Princeton, 1989), and *Socrates and the State* (Princeton, 1984). He is a member of the American Academy of Arts and Sciences.

TWO PUZZLES

In these lectures, I will juxtapose ideas drawn from Plato and Aristotle and two conundrums of twentieth-century moral philosophy, in the hope that they will shed light on each other.

The first puzzle was raised by the British idealist philosopher J. M. E. McTaggart in *The Nature of Existence*.¹ He asks us to compare two lives: the first he calls “oyster-like” because it has “very little consciousness and . . . a very little excess of pleasure over pain”; the second is that of a human being. His striking thesis is that a sufficiently longer oyster-like life is better than any shorter human life, no matter how wonderful the goods in the human life are, or how full of knowledge, virtue, love, pleasure, and whatever else makes for a good human life. McTaggart holds that the oyster-like life would more than compensate for the hour-by-hour poverty of its simple life by the greater quantity of time during which it possessed the simple, faint pleasures it enjoys—so long as its greater longevity is great enough.

McTaggart uses some very large figures to make his point: one million years of a rich human life versus many more for the oyster-like creatures. The important contrast for him is that the human life has an end point far earlier than that of the oyster-like existence—the latter is stipulated to have a life longer by enough to attain eventual superiority. (He adds that at some point it would have a life one million times better than that of the best human life.) I will stay closer to our familiar world by assuming instead that the marvelous human life is about eighty-five or ninety years long, not a million. The other creature in this comparison will be called “McTaggart’s oyster.” I will assume that this creature feels no pain and only the mildest pleasure. The term “McTaggart’s thesis” will refer to the claim that his oyster eventually, after perhaps thousands of years, or however long it takes, will have more well-being than even the best of human lives of normal length.

The second puzzle was raised by Robert Nozick in *Anarchy, State, and Utopia*.² He asks us to imagine an “experience machine”—a device that induces in a detached brain any illusory experience that the subject chooses. The principal claim of Nozick’s discussion of this device is that we would not (and *should not*) choose to live in such a machine, whatever its surface or initial attractions. Inside the machine we would be merely passive consumers of experience, cut off from the real world. It is better

for us to *live* our lives—something no machine can do for us—and to be in contact with a reality beyond our own brains.

It may seem odd to bring together these two thought experiments of twentieth-century value theory. What have they to do with each other? And what have they to do with Plato and Aristotle? My reply is this: If we reject McTaggart's thesis, we must say what assumption leads him astray, or what he fails to recognize. My diagnosis is that he fails to apply to his thought experiment the idea that a sufficient supply of peculiarly human goods is superior to any amount of the lower goods available to the oyster. That goods belong to different orders of value is a thesis that fits comfortably within the thinking of Plato and Aristotle. They hold that the experiential good of being a good human belongs to a higher order than the good of the pleasures of nourishment available to an oyster. They assume that to give a convincing account of where the value of virtue lies, one must depict the inner life of a good person—what it is about such a person's thoughts, feelings, aspirations, pleasures, memories, and so on that makes that individual's life appealing. They could easily accept what might be called "an experientialist conception of well-being" ("experientialism," for short). And that is where Nozick's experience machine comes in: it is generally taken to reveal the minor importance of the experiential components of well-being.

Experientialism holds that (a) well-being is composed of many goods, (b) all of them are experiential, but (c) pleasure is only one element of good experience.

There is universal agreement in modern and contemporary moral philosophy that at least one component of well-being is the felt quality of experience. What is controversial is whether there are other components, possibly of greater value, that are not experiential. Often this question is treated as a debate about hedonism—the thesis that pleasure alone is good. Like many others, I reject that monistic conception of well-being; well-being, I believe, is composed of many goods. Even so, I think that there is something right about hedonism, namely its restriction of well-being to features of our inner life—our introspectable, felt experience.

Aristotle moves in the direction of experientialism when he makes a subtle and compelling point against the thesis that well-being consists in virtue. What he proposes is not an outright rejection of that thesis but a modification of it. The passage I have in mind, from *Nicomachean Ethics*, says that virtue cannot be *eudaimonia*, because a virtuous person might fall asleep

for the entire duration of his life, and no one would think that an enviable state. Aristotle then offers this modification: *eudaimonia* consists not in the *state* of virtue but in virtuous *activity*.

It sounds like a good point, but what is the idea exactly? I suspect that many of us tacitly read into Aristotle's words the obvious point that sleep contains little in the way of conscious experience—only the brief periods during which one dreams. For the most part, there is nothing it is like to be asleep; a virtuous person who slept away his life would have little in the way of conscious experience, and there would be little good in that. This way of understanding Aristotle's argument takes him to be saying, along with Plato and the Hellenistic philosophers, that the value of virtue lies in the way a good person experiences things—in the felt quality of his inner life as he interacts with his social world.

I don't see *any* value in virtue if it is never active and involves no conscious thoughts, feelings, or desires. The same objection can be made to the claim that knowledge, love, and accomplishment are among the components of well-being. It would be better to say that intellectual exploration, feelings of loving and being loved, and a sense of accomplishment make a life go well, for these are experiences; one has at best a small share of them when one is asleep.

Nozick's experience machine might be taken to show not only that there is much more to well-being than pleasure but also, more generally, that there is much more to it than its experiential component. The contrast between his view and mine is stark: for him, the experiential component of well-being is *small*; for me, it is *everything*.

WHERE DOES MCTAGGART GO WRONG?

One might believe that McTaggart's thesis is correct, and so the problem of explaining where he goes wrong does not arise. But even if one had this conviction, there would still be some value in noticing what assumptions he needs to make or reject to arrive at his conclusion. I assume that you, like me, reject his thesis; but that is not what matters most here. The more important task is to take apart his reasoning, and *if* we reject it, to identify what the best alternative is.

One possibility is this: Although we might attribute to an oyster the desire to eat now, it does not have a desire to eat tomorrow. Its mental life does not include a conception of itself as something that endures. McTaggart can reach his conclusion only if he keeps adding one day's pleasures to the next day's, and the next day's, and so on. But that aggregation is not

something the oyster can perform or care about. So, its life does not keep getting better with each passing day. Its well-being does not extend beyond the present moment. It has a good now, and another good later, but not more good over time. With human beings, it is of course different.

A second possibility is this: The oyster's life does keep getting better for a while. But at a certain point, the good in its life has reached a limit. It continues to experience pleasure for however long it lives—thousands of years—but pleasure is no longer good for it after a few days, or months, or years, for by then it has reached the limit of its goodness. But the peculiarly human goods, which are unavailable to an oyster, retain their goodness for the full length of a normal human lifespan.

A third possibility is a variation on the second. It holds that the value of an oyster's pleasure is always such that *some* further increase in well-being is available to it without end. But the amount by which it can be increased keeps diminishing; it approaches but never reaches a limit, and therefore it can never exceed that limit. With human beings, it is different. If the marginal value of human goods diminishes, the rate or onset of diminution is such that a normal human life filled with the elements of well-being is better than the life of McTaggart's oyster.

The fourth possibility is the one that I have said has some basis in ancient ethics. Goods are not all at the same level; in particular, virtue (or virtuous *activity*) belongs to a higher order than others. What this talk of "levels" or "orders" amounts to is this: to say that a good belongs to a higher order entails that a sufficient amount of it is better than any amount of a good of a lower order. They are "incommensurable" in that there is no common coin or measure applicable to the two orders that would tell us how many more goods belonging to the lower order are needed for them, taken together, to be superior to a smaller quantity of goods of the higher order. Analogously, a student who writes many C papers is still inferior as a student to one who writes a smaller number of A papers. It does not matter how many more papers he writes. Similarly, for Aristotle, an active life of virtue for a complete period of time is a better life than that of McTaggart's oyster, no matter how long the oyster feels pleasure.

It might be thought that this fourth alternative is a non-starter. Someone who is skeptical about it might ask: "What could *explain* why the superiority of the oyster's life with respect to the quantity of the pleasure it contains does not make it on balance the better life? However little pleasure is worth, enough of it accumulated over time ought to outweigh the value of all other human goods, given their limited duration." My reply

is that no explanation is needed. Our experience of different goods is such that we correctly judge that a sufficient amount of one is better than the possession of the other for no matter what length of time. We know that judgment to be correct because it is grounded in our experience. That justifies us in believing this to be the case even if we have no explanation of its truth. Our experience has made it self-evident to us.

This fourth approach is the one that I favor for precisely that reason. My experience of some of the great goods of human life is such that I judge it to be superior to the life of McTaggart's oyster. I know what a mildly pleasant gustatory sensation is like. No matter how many days it is felt, and even if it never stops being pleasant and its value remains undiminished, a life that contains only that single kind of bland experience is inferior to a human life filled with the best of what we can experience, even if our lives have many fewer years than the thousands over which McTaggart's oyster endures.

THE RICHNESS OF HUMAN EXPERIENCE

I will temporarily set aside the three other approaches to McTaggart's oyster in order to elaborate on the superiority of experientialism over hedonism. Pleasure is only one part of our conscious internal lives, and is not the only kind of phenomenal quality it is good for someone to experience. *Within the realm of experience*—the aspect of life that is available in the experience machine—pleasure is only one small part of a proper account of what is good for us. We can see that hedonism is mistaken every time we listen to an orchestra or have any other complex sensory experience; it is to a large extent the quality of the *sound*, not the alleged qualitative aspect of the *pleasure*, that makes us count this as a good experience. If we were excited because we were about to hear a new orchestra whose unique sound we have read about, we would be deeply disappointed if someone were somehow able to induce in us the pleasure of the experience minus its auditory quality.

Even if we supposed that pleasure must be an ingredient of any experience that is good for us to have (a supposition open to question), it would not follow that it is the only or most important ingredient that makes it good. That would be like saying that because the main item on one's dinner menu should always contain one or more spices, it is those spices alone that make one's meal good to taste. *All* of the different taste sensations we experience simultaneously join together in making a good-tasting dinner. In the same way, it is the integrated complexity of the experiential life of a

human being—its unification of so many phenomenological components besides pleasure, blending together and modifying each other—that makes it better than an oyster’s simple consciousness.

Adding depth to our experience of life is our sense of the present as just one temporal portion of an experience that extends backward to the past and forward to the future. At each brief moment spent listening to music, for example, we bring to bear on our present experience our memory of what preceded it, and an anticipation of what is to come. The same point applies outside the aesthetic realm: memories and expectations color our encounter with what we experience in the now.

In addition to perceptual awareness (the world of colors, sounds, smells, tastes), there is also a phenomenal quality to our active cognitive life: we are conscious of making judgments and conjectures, entertaining possibilities, searching through our memory, venturing predictions, encountering intellectual surprises, and the like. Conative phenomenology is no less varied, including our awareness of wanting, choosing, planning, wishing, hoping, and so on. A further phenomenological category includes the emotions: being happy, glad, proud, excited; feeling affection, love, admiration, respect, and lust.

Imagine someone standing in front of a painting that he finds visually stunning. All sorts of *thoughts* are going through his mind; he is *intrigued* by what the figures in the painting are doing; many subtle *emotions* are called forth by the narrative significance of the painting’s subject; he *decides* to linger and to look at it from different angles; his eyes take in the contrast of *colors and shapes*; he is *uncertain* whether he fully *understands* what the artist intended, and continues to *explore* the painting’s meaning and the reason for its emotional impact. This is a good experience for anyone to have—not because of something it later brings about, but simply in virtue of what occurs while it is going on. Pleasure may be part of what our viewer experiences—but he is of course also aware of the painting in front of him and his variegated internal responses to it. The goodness of this experience for him supervenes on *all* of the many aspects of his consciousness—not simply on the fact that he is feeling pleasure.

WELL-BEING AND HUMAN DEVELOPMENT

I should guard against a possible misinterpretation of what I am saying. I do not claim that the best life for us would be one that contained nothing in the way of sadness, pain, frustration, anxiety, stress, and other “negative” states of mind. On the contrary, if they are properly integrated into a larger

whole, these states deepen and enhance our experience of life. Many of our best experiences involve overcoming difficulties, defeating obstacles, rising to challenges that we fear we may not be equal to. I can express this point either by saying that these negative experiences are depleted of their potential badness when they are integrated into a larger whole, or that the best lives are good wholes, some of whose parts are bad. For the sake of simplicity, I will continue to say that well-being consists in good experience.

I can now turn back to the question: Which are the good experiences? My response is that I do not pretend to have a complete list—and that we should not assume that completeness is necessary or desirable.

Suppose someone said, “Well-being consists only in reading good novels.” Obviously, that won’t do—it is too narrow. But it is not a further defect that we have not been given a definitive and closed list of which novels are the good ones. (After all, they keep on being written.) What we would need is some rough-and-ready standard for making aesthetic judgments, and sufficient exposure to novels to evaluate their quality. Similarly, experientialism needs to say *something* about good experience. But that task is not neglected here. (One of my aims, for example, is to discuss the good experiences that come with being a good person.) I will not propose an exact number of types of experience that are good. That open-endedness might be a feature that a theory of well-being ought to have.

When we judge that an experience we are having or have had is immeasurably rich and worth having for itself—as when we are absorbed in a great work of art or surrounded by great natural beauty—we have some basis for valuing this experience precisely because it is our experience and we know it from the inside. Similarly, if we judge that our relationship with our children is one of the best things in our lives, we are basing that assessment on what we have lived through. When we “do philosophy,” it is our fascination, wonder, puzzlement, and sense of depth that convinces us that philosophy is worthwhile. (It can hardly be its ability to achieve definitive results!) The *goodness* of these activities is not experienced (it is not an observational property), but our judgment that they are good is based on experience.

It would nonetheless be misleading or worse to say that the test of a theory of well-being is *simply* experience. That would fail to acknowledge the role of theory construction and philosophical argument in finding and testing a theory of well-being. The power of a theory of well-being, to the extent that it has any, derives from its ability to explain what is good in a

wide range of human lives by appealing to a small number of good-making features. Hedonism would be just such a theory, if pleasure could satisfactorily explain, all by itself, what is valuable in art, child-rearing, philosophy, and so on. But it cannot.

My conception of the experiences that are best for us also has a developmental component. Like Aristotle, I hold that the best things in life are most fully available to us not in childhood but when our powers have fully matured. We should grow up for many reasons, and one of them is that better things are in store for us as grown-ups. That being so, it must be the case that we adults know something about what makes for a good childhood—what is good for an individual *when* he is a child (and not good only because it is preparation for later life). It must also be the case that we know what makes adulthood a potentially better period. We compare that early time of life with what comes after it, and we see that opportunities for well-being increase. A theory of well-being would be ignoring valuable data if it did not build on this common knowledge. When life is good for us, that is because the potentialities that gradually are realized through education and training—our cognitive, emotional, and social powers—are regularly exercised over the many years of adulthood, as we engage in ethical activity (and other sorts of activity as well).

THE EXPERIENCE MACHINE AND A CHILD'S PLEASURES

I have claimed that the right way—the only persuasive way—to recognize the prudential value of ethical virtue is to assess it in experientialist terms. The inner richness of moral life, all by itself, will constitute a good life, if it is not weighed down by long periods of severe and irredeemable suffering. The question to be addressed now is whether we should accept the famous critique of experientialism offered by Nozick's thought experiment.

Nozick claims that we would choose not to plug in, and rightly so. First, he says, "We want to *do* certain things, and not just have the experience of doing them." Second: "We want to *be* a certain way, to be a certain sort of person. Someone floating in a tank is an indeterminate blob." Third: "Plugging into an experience machine limits us to a man-made reality, to a world no deeper or more important than that which people can construct. There is no actual contact with any deeper reality."³

Nozick would have to admit that some people have lives so filled with suffering and with so few prospects for improvement that, for them, life in the machine would be an improvement. Better to be a brain in a vat feeling nothing but bliss than be subject to ongoing torture. Conversely, we should

admit that at times certain kinds of experiences can be had only by venturing out into the real world. If, for example, you want to know what marine life looks like several miles below the surface of the ocean, and that part of the world has not yet been observed by any human being, the only thing that will give you this experience is first-hand exploration. So the question to be asked is not whether one should plug into the machine. There is no general answer to that question—it depends on one's circumstances.

What, then, is the issue that merits discussion? It is Nozick's claim that however good one's experiences inside the machine are, one would have a far superior life if those experiences were produced in the normal manner by contact with reality, even though the machine-induced and the reality-induced experiences are indistinguishable. Life lived in the real world is so much better than the experientially identical life lived in the machine, that one should accept all the risks that attend life outside the machine rather than choose the safety guaranteed inside it.

Let us begin with the obvious point that someone might choose to plug in to the machine in order to have experiences that are largely or purely sensual—and as a result her whole life might be somewhat like that of an oyster. Imagine a three-year-old child who chooses to plug in because she wants an ongoing and uninterrupted experience of eating chocolate cake. In some respects, she is making an excellent choice. In real life, eating too much cake will make you sick. I don't think anything valuable would be missing from a neurally induced experience of eating chocolate cake simply because one was not actually eating anything—not moving one's mouth, not breaking up bits of food with one's teeth, not swallowing them, not digesting them. What we want from our encounter with real pieces of chocolate cake is just a sensory experience, not an action in which we move our mouths and fill our bellies.

Nozick stipulates that “after two years have passed, you will have ten minutes or ten hours out of the tank, to select the experiences of your *next* two years.”⁴ So, we are to imagine our three-year-old emerging, at the age of five, from the tank in which she was immersed. She now has to decide what experiences she will prescribe for the next two years. Suppose she chooses two years of the same experience, and then another two years, and so on. She has never experienced anything better than this, and so she sees no reason to vary or broaden her experience.

But notice that the illusoriness of the child's chocolate cake experience plays no role in explaining the defectiveness of her choice. It is the

narrowness of her experience that does all the work, not its being illusory. Someone whose life has only this one good feature—tasting, day in and day out, and at all hours, *real* chocolate cake (without ill effect)—would have precisely the same deficit as a child who has nothing but these sensations through the mediation of an experience machine.

What goes on inside your head—your inner mental life—is enormously important, but that does not mean that if you are fully satisfied with what you are experiencing (as our three-year-old is) you are living as well as you could be. Introspection by itself will not tell you what sort of inner life it is best for you to have. The individual is not “sovereign” in these matters—the sole and authoritative judge of what is good for him. The question “What are the constituents of well-being?” is not to be answered by introspection alone but by philosophical theorizing—which may appeal at certain points to introspection. We can use McTaggart’s error and our diagnosis of it as a test that must be met by an adequate conception of what is good for us. What explains the fact that human life at its best is superior to the life of his oyster, no matter how much pleasure the oyster has?

MUSIC IN THE EXPERIENCE MACHINE

Next, suppose someone chooses to have the experience machine deliver the illusion that he is sitting in a small, acoustically perfect auditorium in which a musician is playing Bach’s sonatas for unaccompanied violin. Each note and phrase sounds as though it were traveling through space and produced by a musician and her instrument; each is utterly clear, vivid, expressive, and artful. In fact, no sounds are traveling across a room—it merely seems that way. Yet, it would be implausible to suppose that the enjoyment of this illusory experience is a less valuable component of well-being than the phenomenally indistinguishable experience that is produced by a violin and musician. The music lover just described is no worse off because of the error he makes about the source of his musical experience. The same point applies to illusory experiences, produced in a vat by human manipulation, in which one seems to be seeing beautiful landscapes. Those things that seem to be trees, sky, clouds are just as beautiful as the real things, and the enjoyment of them is no less a good thing.

So, Nozick’s attitude towards the experience machine, with its insistence that a life connected to the “real world” is superior to a life devoted to the enjoyment of experiences artificially induced in us, seems to overlook or degrade the value of aesthetic experience. I will now expand on that

point, and extend it to our enjoyment of literary fictions. A life dominated by the love of literature, I will argue, has many of the same features as a life lived within an experience machine. If Nozick's arguments lead to the conclusion that a life absorbed with fictional characters and stories is to be avoided, there must be something wrong with it.

“READING AN INTERESTING BOOK”:
TWO INTERPRETATIONS OF NOZICK

To make the experience machine initially appear attractive, Nozick notes that it could “stimulate your brain so that you would think and feel you were writing a great novel, or making a friend, or reading an interesting book.”⁵ Let's take his last example—reading an interesting book—and ask how we are to understand what it would be to have an illusory experience of this sort.

It might involve having images of ourselves, as we sometimes do in dreams, opening a book, turning its pages, thinking to ourselves “how interesting,” and telling a friend who approaches us that he would enjoy it too. We would not take in to our minds any words in the illusory book; we would just have a picture of ourselves reading it, as an external observer might. Now, which would be better for us: (a) really to take in the content of an interesting book, as we ordinarily do outside the experience machine? (Here we are aware of words that convey meaning, and our encounter with those meanings is accompanied by a variety of emotional and intellectual reactions), or (b) to have the vague impression that we are reading such a book?

Obviously, in this case, reality is superior to illusion. One doesn't get anything of value from a book if one merely has a dreamlike impression of reading it.

But perhaps this does not give the experience machine as much power to simulate the experiences of real life as such a device can have. Let's assume instead that it can make us take in all the words of a book of our choice, in their proper order, and that we would react to it with interest because our experience of it replicates our mental encounter with actual books. We would still be living within a grand illusion. The machine would make it seem as though we were sitting in a comfortable chair, with the light on, holding and reading a book, but all the while we would just be floating in a tank, seeing nothing, and having no book in our hands. We would mistakenly take ourselves to be embodied, in a library or study or café, but we would have only a brain, and be located in some laboratory. Even

so, one's experience would replicate the one we have when we really read a book, taking in words on a page (or reading device) as light reflected from them impinges on our eyes. We could ponder what the author is saying, wonder whether we should accept his point of view, be surprised by his unexpected ideas, entertain mental images of characters described, react with displeasure or delight by the way the book is written or organized, and so on.

In conceiving of the experience machine in this way, are we remaining true to Nozick's intentions? It is hard to know. Interpreted in one way, what one experiences in the machine is confined to what is sensory in nature, along the lines of an oyster's pleasures, or tastes and smells as if of chocolate cake, or sounds as if made by a cello. One has pleasant experiences of one's choosing, but no sophisticated and active mental processes, no higher cognitions, are possible. One could feel as though one were reading great literature, but that would be a dreamlike sensation. There would be cello-like sounds, but no grouping of them into meaningful form.

There is a different reading, according to which Nozick meant to use the experience machine as a way of revealing the greater value of a life in which we regularly and actively make choices that have consequences in a world that exists independently of us. In his thought experiment, the only real choices take place *outside* the machine: they are choices about which experiences to have *within* it. Once you are inside the machine, you are not actually exerting any causal force on the world, although you think you are. That is why Nozick says that "someone floating in a tank is an indeterminate blob. There is no answer to the question of what a person is like who has long been in the tank."⁶ Moral character, he is assuming, requires effective choice between real-world alternatives; in the experience machine, one passively enjoys pre-selected experiences but one *does* nothing. Hence one is a mere blob.

Although this might be the better of the two interpretations, it is still difficult to believe that it is what Nozick had in mind. On this reading, one can use the experience machine to absorb, with perfect understanding, the content of any books of one's choosing. Admittedly, that content would be conveyed by the manipulation of one's brain by a neuroscientist, and not by the act of reading words on a page. But this hardly matters. So understood, Nozick allows that one can learn, with the help of the experience machine, all there is to know about physics, the life sciences, history, and so on. Therefore, contrary to his intentions, one is not locked up in one's own mind—rather, it is the real world, the world that exists

independently of oneself and the neuroscientist, that one has full access to. That could not be what Nozick had in mind when he said that “plugging into an experience machine limits us to a man-made reality, to a world no deeper or more important than that which people can construct.”

AESTHETIC APPRECIATION

Whether or not this second interpretation is true to Nozick’s intent, the experience machine, so understood, is full of philosophical interest. What Nozick would be criticizing, if we read him this way, is the contemplative life—that is, a life in which one has thoughts, emotions, and sensations of whatever sophistication, variety, and agreeableness one likes, but in which one does not interact with any other human beings, and more generally lacks any causal efficacy upon the real things of the world.

If the passive life of isolated contemplation is Nozick’s target, he faces the objection that he fails to recognize the value of a life devoted to aesthetic appreciation—a life fully immersed in the fictive worlds constructed by the authors of great literary works. When you read, you are making minimal use of any part of your body beyond your brain; you might as well be detached from it. You are not trying to change the real world; you are not *doing* (that is, affecting) anything. When you read works of fiction, as opposed to works of science, history, biography, and so on, you become emotionally engaged with and seek to understand characters and situations that are the products of an author’s imagination and would not exist were it not for that author’s creativity. There is nothing you can do to alter an already finished work of fiction—it has an effect on you, but not vice versa. While you are engaged with that work, you can be in total isolation. This is a kind of life that some impressive people aspire to, and that many literature departments prepare people to live.

The issue before us is how to think about the value of an engagement with fictional worlds. We love stories, even as children, and putting events into a narrative order may be crucial to our ability to navigate our social world. But for this purpose we do not need *fictional* narratives of the sort produced by literary artists. Good biographers and historians also give us the pleasures of a good story well told—and they give us a deep understanding of things that really happened and people who really exist. Why bother, then, with the fantastical creations of writers of fiction? There is a sense in which they do not give us—to use Nozick’s phrase—“contact with reality.”⁷ Even if it could be said, truthfully, that Madame Bovary exists—that is, that the fictional character of this name exists—she is still a

fictive entity rather than a “real human being” who exists in space and time. What good does it do to spend one’s time thinking about her, when one could be learning about real people—people who exist in space and time?

In reflecting on this question, we must steer clear of the philistinism that Dickens clumsily satirizes in *Hard Times* through his depiction of Thomas Gradgrind, the educator who values nothing but “facts.” A conception of well-being that is forced to conclude that the arts or playful exercises of a child’s imagination have no place, or at best a merely instrumental role, in human life would have no plausibility. Without composers, musicians, novelists, poets, playwrights, and painters, there would be fewer opportunities for us to enrich our experience to an extent that makes our existence immeasurably superior to that of McTaggart’s oyster.

So, there must be *something* right in the “formalist” approach to aesthetics set forth in the works of Oscar Wilde and other figures of the nineteenth and twentieth centuries who defended “art for art’s sake.” What should we learn from this movement? As a first step, we should acknowledge that great beauty and other important aesthetic virtues can be found in what came to be called “absolute music”—music that is not “programmatic,” not about anything but its own content.

Some of the most treasured compositions of the Western canon—Bach’s suites for solo instruments, Beethoven’s quartets, Shostakovich’s preludes and fugues—fall into this category. No truths about our everyday lives can be learned by listening to them, but that does not demote them to an inferior order of art. In saying this, I am rejecting the thesis—prominent in Schopenhauer and dating to the Pythagoreans of antiquity—that music is valuable only because it reflects and reveals the true nature of the world (whether that be futile striving or mathematical proportion). As the experience of music lovers attests, what is valuable in absolute music is already there in the way it sounds to a trained ear—it need not be sought in something non-musical that lies behind it. If you ask what use such music has, you are looking in the wrong place for its value.

We can infer from this a more general conclusion. The insight to be found in the slogan “art for art’s sake” is that what is true of absolute music is applicable to the other “fine arts” as well. Works of fiction, so understood, need not achieve excellence by teaching us general lessons useful for the conduct of our lives in the real world. When they lack didactic value, that does not make them defective as works of art, nor should they be avoided because they do not help us find our way among flesh-and-blood human beings.

That does not show that music ought not to be programmatic, or that we should avoid novels that will shape our lives when we are not reading them. If a work of fiction teaches us, for example, what life was like in the trenches of World War I, that is a good reason to read it. We will learn something about the real world, and this is not a motive to be despised. A novel might also be more directly useful, for example by describing an unfamiliar part of the world that we are about to visit. Authors have no reason to eschew works that offer readers these sorts of insights about reality. Similarly, it does not count against Beethoven's Ninth Symphony that it might inspire listeners to feel a sense of unity with all humankind. If it succeeds, all the greater is its accomplishment.

The conclusion we should draw is that works of art can be valued solely for their internal aesthetic features, or for the further good they do, or for both reasons. Those that, like absolute music, are to be valued simply because of the rich experience we have when we attend to them, can be, on balance, as worthy of our admiration as those that have some further value as guides to reality. That is the kernel of truth in the slogan "art for art's sake."

With this in mind, we can return to the experience machine. Absolute music can be as fully appreciated inside the machine as outside. Other forms of artistic excellence can be appreciated in the way absolute music is, their value lying entirely in the works themselves and not in what they reveal about a further reality beyond them. Someone inside the machine whose experience is filled with the love of such works as these does not change the world or acquire knowledge of it. But it would be a form of philistinism to hold that the life of such an individual would have little or no prudential value. If Nozick's low assessment of the value of the experience machine entails that we should read only those books or listen only to that music that instructs us, or prepares us to change the world, he is uncomfortably close to Mr. Gradgrind.

UNFINISHED BUSINESS

My defense of an experientialist conception of well-being is far from complete. More needs to be said about McTaggart's oyster and the experience machine. I have said nothing about the interior life of the virtuous person. I have left aside the question of posthumous goods and the phenomenon of false friends who betray us behind our backs. But even if we suppose, for the sake of argument, that these phenomena or the experience machine

show there are non-experiential components of well-being, there is a “fall-back” position I can retreat to: experientialism might be very close to the truth, even if it is not the whole truth. If there are non-experiential components of well-being, but any amount of them, however long-lasting, is inferior to a sufficient amount of experiential goods, then the falsity of experientialism would make little or no practical difference.⁸

NOTES

1. J. M. E. McTaggart, *The Nature of Existence*, (Cambridge, England: The University Press, 1927), volume 2, book 7, chapter 7, sections 868–70, pp. 452–53.
2. Robert Nozick, *Anarchy, State, and Utopia* (New York: Basic Books, 1974), pp. 42–45.
3. *Ibid.*, 43.
4. *Ibid.*, 42–43.
5. *Ibid.*, 42.
6. *Ibid.*, 43.
7. *Ibid.*, 45.
8. I am grateful for the excellent comments and criticism I received when I presented many of these ideas at Stanford University. My commentators were Rachel Barney, Stephen Darwall, Rebecca Newberger Goldstein, and Thomas Hurka.