

Imagining Under Constraints

Amy Kind

The claim that imagining does not provide us with any information about the world has often seemed a truism in philosophical discussion. When I imagine that a pumpkin has been transformed into a stagecoach, or that a motley group of escaped zoo animals have taken refuge in my backyard, or that I've won the Publisher's Clearinghouse Sweepstakes, my imaginings give me no reason to believe that such events have occurred.

Three basic (and related) features of imagining account for its apparent inability to provide us with such reasons. First, imagining is typically under our voluntary control. Second, imagining is not world-sensitive—the content of an imagining is determined by the imaginer, not by the world. Third, imagining is uninformative—an act of imagining can provide us with no new information. As Sartre has claimed, “nothing can be learned from an image that is not already known” (Sartre 1948, p. 12). In all three of these respects, imagining seems quite unlike epistemically respectable sources of reasons such as perceiving. Thus, imagination has generally been taken to be, as Brian O'Shaughnessy has said, “out of the cognitive circuit” (O'Shaughnessy 2000, p. 345).

In my view, however, this conclusion is unwarranted and derives at least in part from the mistaken assumption that all cases of imaginings are roughly comparable to the examples above. In fact, there are all sorts of other examples where imaginings play a

role in justifying the beliefs that they prompt. Many of these epistemically significant imaginings are perfectly ordinary ones. When trying to decide whether to become parents, a young couple might call upon their imagination in various ways to help them make their decision—from imagining themselves grappling with exhaustion after a sleepless night with a crying baby to imagining themselves proudly watching a teenager graduate from high school.¹ When these prospective parents head off to buy a new car, they might again call upon their imagination to figure out what to buy—while looking at the different models in the showroom, they might imagine themselves getting their child strapped into the car seat in the back, fitting the stroller and other gear in the trunk, and cleaning milk spills and crumbs (or perhaps worse) off the interior. And likewise, when they are converting a former office into a nursery, they might imagine the crib first against one wall and then against another. All of these imaginings seem to have justificatory force as they make their decisions—as they determine that they do want, and are ready to have, a child; that they’d do best to buy a mini-van; and that the crib needs to go on the wall opposite the windows.

¹ In her work on transformative experience, L. A. Paul (2014) argues that imaginings cannot help when determining whether to become a parent for the first time; on her view, “if you’ve never had a child, *it is impossible to make an informed, rational decision by imagining outcomes based on what it would be like to have your child, assigning subjective values to these outcomes, and then modeling your preferences on this basis*” (Paul 2014, p. 83). Though I disagree, my argument in this paper does not rely on this particular example. One could think instead of parents trying to decide whether to have a second child.

Once we consider such imaginings, it seems clear that the imagination has considerably more epistemic significance than it has been traditionally assigned. In at least some cases, imagining has a role to play in justifying our contingent beliefs about the world.² But recognizing this fact does not itself provide us with any explanation of how imagining can have the sort of epistemic significance that it does. Developing such an explanation will be the aim of this paper. My project here thus turns on providing an answer to the following question: What distinguishes the epistemically significant imaginings from the epistemically insignificant ones? In my view, the answer has to do with the constraints under which we operate when we engage in the former sorts of imaginings, constraints that are absent in the latter sorts of imaginings. The mention of constraint here might seem puzzling since, as David Hume famously said in the *Treatise*, nowhere are we more free than in our imagination. But the freedom we enjoy when imagining does not show that we must always proceed completely unfettered, and in fact it is our ability to constrain our imaginings in light of facts about the world that enables us to learn from them. In this paper, then, I offer a framework for showing when and how an imaginative project can play a justificatory role with respect to beliefs about the world, an account of imagining that I call *imagining under constraints*.

1. Stage Setting

² Elsewhere, I have given arguments in support of the epistemic significance of imagination and I have attempted to diagnose the failures in arguments to the contrary (see Kind, in press).

To start, it will probably be useful if I say something about what I take imagining to be. Importantly, however, my account of imagining under constraints—and, correspondingly, my defense of the epistemic relevance of the imagination—does not depend on a particular theory of imagination. For our purposes here, we need only understand imagination as distinct from related states such as supposition or the entertainment of propositions. When engaged in a *reductio* proof, someone might suppose that Congress passes an assault weapons ban, or that Texas secedes from the United States, or that Elvis Presley is still alive. Such suppositions can be made without any difficulty at all, and without the exertion of much mental energy. Merely bringing the relevant proposition to mind is enough. In order to imagine these things, however, something more is required. Different philosophers give different accounts of how to understand this requirement; on my own view, for example, it is to be cashed out in terms of mental imagery.³ But I do not here presuppose this imagistic account. Rather, I rely only on the assumption that imagining involves a more active effort of mind than does supposition or entertaining a proposition, the kind of effort that Kendall Walton seems to be gesturing at when he suggests that imagining “is *doing* something *with* a proposition one has in mind” (Walton 1990, p. 20).

It would also probably be helpful were I to say something more about the kind of epistemic significance with which I am here concerned. For example, although philosophers typically deny that imagination can play a justificatory role with respect to our contingent beliefs about the world, that’s not to say that they deny it any epistemic import altogether. No one denies, for example, that an act of imagination can jog one’s

³ This imagistic account is developed in Kind (2001).

memory, as when I'm reminded that there's no dog food in the house as a result of imagining my dog sleeping on the couch. Likewise, no one denies that an act of imagination can lead to interesting innovations, discoveries, or new directions of research. To give just one particularly famous example: At the age of 16, Albert Einstein imagined himself chasing a beam of light, an imaginative exercise that he credits as having played a key role in the development of his theory of special relativity. So my opponent's charge that imagination is epistemically insignificant is consistent with the claim that our imaginings might *generate* various beliefs about the world, beliefs that might even turn out to be true. Rather, my opponent means only to deny that imaginings can justify—or even play a role in justifying—such beliefs. I will not here offer any sort of detailed epistemology, and it is my intention to stay clear of debates about (for example) the nature of justification. So perhaps I can best put my point like this: Just as perception teaches us something about the world, so too can imagination.

Interestingly, there is one domain in which imagination has been assigned precisely the sort of justificatory power in which we're interested, namely, the domain of modal truths. Suppose I imagine that there is a heretofore undiscovered planet orbiting the sun between the orbits of Jupiter and Saturn, that the Los Angeles Angels win the next World Series, or that Arnold Schwarzenegger is elected President of the United States. In each case, my imagining the relevant states of affairs gives me reason to believe that this state of affairs—even if it is not astronomically, athletically, or politically possible—is nonetheless metaphysically possible. Perhaps this reason can be overridden by other reasons; perhaps it can't. But it seems that imagining some state of affairs *S* is, at the very

least, relevant for the justification of the belief that S is metaphysically possible.⁴ This claim, which I will call the Imagining-Possibility (I-P) principle, is often loosely captured by saying that imagination serves as an epistemic guide to possibility.⁵

The I-P principle, though widely held, is by no means uncontroversial. Perhaps the most obvious counterexamples arise when we consider imaginings under conditions of ignorance. Someone lacking even a basic understanding of chemistry might be able to imagine that water is not H₂O, but we surely do not want to conclude that it is metaphysically possible that water is not H₂O. Proponents of the I-P principle thus typically impose conditions that an imagining must meet in order to serve as an epistemic guide to possibility. For example, one important condition is that such claims must not be merely *prima facie* imaginable—that is, imaginable on first appearances—but must rather be *ideally* imaginable—that is, imaginable on ideal rational reflection.⁶ Or, to put things in Cartesian terms, our imaginings must be clear and distinct.

Thus, when philosophers claim that the imagination provides us with reasons for our beliefs about metaphysical possibility, it is by no means intended that every imagining can play this role. I take this to be an obvious point about the role of imagination in modal epistemology, but I raise it now to make a parallel point about the central claim in

⁴ Throughout I'll talk of imagining states of affairs, but nothing should hang on this assumption.

Those who think we imagine sentences or propositions should be able to make the appropriate substitutions.

⁵ Many claim that it is conceiving—of which imagining is a subspecies—that serves as an epistemic guide to possibility. See the discussion in [Gendler and Hawthorne \(2002\)](#).

⁶ See e.g. [Chalmers \(2002\)](#).

this paper. In arguing that the imagination provides us with reasons for additional beliefs (i.e. beliefs in addition to those about metaphysical possibility), it is likewise by no means intended that every imagining can play this role, and the fact that there are all sorts of examples of imaginings which lack justificatory power will not count against our central thesis. Rather, just as it seems obvious that we have to impose conditions on when an imagining will be of the right sort to provide us with reason for belief about metaphysical possibility, it should be obvious that we will need to impose conditions on when an imagining will be of the right sort to provide us with reason for other kinds of belief.⁷

2. Ideal Imagination

Unsurprisingly, it turns out that the notion of ideal imagination—or at least a kind of ideal imagination—turns out to be important not only with respect to imagination’s epistemic relevance for beliefs about metaphysical possibility but also with respect to imagination’s epistemic relevance more broadly. But the sort of ideal imagination with

⁷ That said, having invoked an analogy between my project here and the I-P principle, I should note explicitly that I do not mean to put too much weight on this analogy. Importantly, there seems to be a conceptual connection between imagining and possibility that helps to explain why imaginings can be a guide to possibility, a connection that is lacking between imagining and actuality. As Gendler and Hawthorne note, when we imagine, “the things we depict to ourselves frequently present themselves *as possible*, and we have an associated tendency to judge that they are possible” (Gendler and Hawthorne 2002, p. 1). But when we imagine, the things we depict to ourselves are typically *not* presented to us as actual.

which we will be concerned is not one that contrasts primarily with *prima facie* imagination as described above. Consider again a paradigmatic case in which an imagining lacks the kind of epistemic relevance in which we're interested, such as when I imagine that a motley group of escaped zoo animals have taken refuge in my backyard.⁸ Not only is this state of affairs easily imaginable on first appearances, but there's no reason to doubt that it remains imaginable even under ideal rational reflection of the sort mentioned above. Rather, there must be something else that keeps this imagining from having justificatory power—something else that keeps this imagining from being ideal in the sense required for an imagining to have epistemic relevance. But what is this relevant sense? Answering this question will be the task of this section.

Here, as so often in philosophy, we are helped by an example from science fiction. In particular, consider the imaginative capacities of the machines who populate the fictional world described in “The Last Evolution,” a short story written in 1932 by American science fiction author [John W. Campbell](#):

It was 2538 years After the Year of the Son of Man. For six centuries mankind had been developing machines. The Ear-apparatus was

⁸ Though such an imagining may have epistemic relevance for our beliefs about metaphysical possibility—while it may help to justify my belief that it's metaphysically possible that a motley group of escaped zoo animals have taken refuge in my backyard—the imagining lacks any further epistemic relevance. Hereafter, when I talk of an imagining having epistemic relevance, it should be understood in this latter sense, i.e. I mean to be talking about the epistemic relevance of the imagining *over and above* whatever epistemic relevance it has with respect to our beliefs about metaphysical possibility.

discovered as early as seven hundred years before. The Eye came later, the Brain came much later. But by 2500, the machines had been developed to think, and act and work with perfect independence. Man lived on the products of the machine, and the machines lived to themselves very happily, and contentedly. ...

Machines—with their irrefutable logic, their cold preciseness of figures, their tireless, utterly exact observation, their absolute knowledge of mathematics—they could elaborate any idea, however simple its beginning, and reach the conclusion. From any three facts they even then could have built in mind all the Universe. Machines had imagination of the ideal sort. They had the ability to construct a necessary future result from a present fact. But Man had imagination of a different kind, theirs was the illogical, brilliant imagination that sees the future result vaguely, without knowing the why, nor the how, and imagination that outstrips the machine in its preciseness. Man might reach the conclusion more swiftly, but the machine always reached the conclusion eventually, and it was always the correct conclusion. By leaps and bounds man advanced. By steady, irresistible steps the machine marched forward.

Campbell's story depicts a future in which super-intelligent and powerful machines take care of all human needs. But for all their intelligence and power, and despite their artificial eyes and ears, the machines are said to lack the kind of imagination that enables humans to make progress "by leaps and bounds." Instead, the machines have imagination of a different sort—an imagination that Campbell himself describes as *ideal*. The ideal

imagination of the machines allows them to reach correct conclusions—not only conclusions concerning metaphysical possibilities, but also conclusions concerning contingent facts about the world. Indeed, it is through the exercise of their ideal imagination that the machines are able to figure out not only how to counter a devastating attack from an alien force but also how to transcend their material existence (the titular “last evolution” occurs when the machines evolve into beings composed of pure energy).

Obviously, the machines have considerably more “brain” power than the humans—the mental capacities of even the significantly evolved twenty-sixth-century humans pale in comparison with the sheer computational power of the machines, who are described by Campbell as having the capacity to work with countless trillions of facts. And this immense computational power certainly contributes to the machines’ imaginative capacities. But the computational power itself is not what’s primarily important for Campbell’s characterization of their imagination as ideal. Rather, what matters is how that computational power is put to work. Importantly, unlike the imagination of the humans, the machines’ imagination is not *creative*. They do not imaginatively create a new reality completely disconnected from the current reality; rather, their imaginings are closely guided by reality as it is.

Consider, for example, how the machines employ imagination to stave off the impending alien attack that threatens earth and all its occupants. Here they are not concerned simply to figure out whether it is metaphysically possible that the attack will be thwarted if they take certain courses of events; they need to figure out which course of events will *actually* thwart the attack. Thus, it does not help them to imagine an alien force that is smaller, or less intelligent, or less powerful, than the actual army on its way.

It also does not help them to imagine their own existing defenses as more widespread or advanced than they actually are. They need to imagine the alien force as it actually is, and their own existing defenses as they actually are, and likewise for all other factors relevant to the impending attack. Of course, in attempting to determine whether a given course of action will thwart the attack, the machines can't imagine the world exactly as it is *in all respects*, since the implementation of that course of action would result in various changes to the world as it actually is. To figure out whether a given course of action will be successful, then, their imaginings must appropriately take such changes into account. For their imaginative exercises to teach them something about the world—that is, for their imaginings to have epistemic relevance—the machines aim to imagine exactly the changes required by the implementation of their plan *and no others*.⁹ Because the machines are ideal imaginers, they can do this. Insofar as they are forced to depart from imagining reality as it is, their imaginings proceed by carefully controlled imaginative extrapolation.

⁹ This claim as stated is undoubtedly too strong, since it needs to be temporally modified. A change to the world as it is has a cascading effect through time—there will be all sorts of consequences of that change down the line—and it doesn't seem that they all have to be imagined for an imagination to have epistemic significance. The machines don't need imagine the effects their plan will have 50 years hence; they need to imagine the effects such a plan will have now (or at least, now-ish). Thus, the claim that they must imagine all and only the changes required by the implementation of their plan should be interpreted as something like: all and only the changes required by their plan *in an appropriate time frame*.

We can thus tease out two features of the machines' imaginative process that make them ideal imaginers. First, their imaginings capture the world as it is. Second, when their imaginative projects do require them to imagine a change to the world as they believe it to be, they are guided by the logical consequences of that change. Each of these aspects can be thought of as a constraint on imagination, what we might respectively call *the reality constraint* and the *change constraint*. As I want to suggest, the epistemic relevance of the imagination is tied closely to these two constraints.

Interestingly, some reflections by Campbell on the very nature of science fiction help to support this general line of argument. In an essay that attempts to explain what science fiction is and how it is to be distinguished from related genres such as fantasy, Campbell suggested:

The major distinction between fantasy and science fiction is, simply, that science fiction uses one, or a very, very few new postulates, and develops the rigidly consistent logical consequences of these limited postulates. Fantasy makes its rules as it goes along ... The basic nature of fantasy is "The only rule is, make up a new rule any time you need one!" The basic rule of science fiction is "Set up a basic proposition—then develop its consistent, logical consequences." (Campbell 1966)

Our imagination often works analogously to the way that Campbell describes the genre of fantasy as working—completely without constraint. But ideal imagination works analogously to the way that Campbell describes the genre of science fiction as working. In light of this description, it is perhaps unsurprising that it is often thought that we can learn something interesting about our world from reading science fiction that we can't

from reading fantasy. Likewise, we can learn something interesting about our world when we engage in ideal imagination—when our imagination is governed by the reality and change constraints—that we can't learn from our more ordinary imaginings.

In order to better understand these two constraints, it might first be helpful to note why they are differentiated from one another. Let's consider two different ways in which a machine might malfunction. First, an intermittent glitch in its storage mechanism might cause it to add an extra zero when adding numerical information to its factual database; for example, the computer might represent that there are 500 missiles in a particular defense facility when there are really only 50. Second, an intermittent glitch in its processor might result in faulty inferences; for example, having determined that certain upgrades to existing defensive shields will double their strength, the computer might conclude that upgrading a shield that could previously withstand two missile blasts will enable it now to withstand ten missile blasts. While the first machine's imaginings are not governed by the reality constraint, there is no reason to suppose that they violate the change constraint. In contrast, while the second machine's imaginings are not governed by the change constraint, there is no reason to suppose that they violate the reality constraint.

Having differentiated the two constraints in this way, however, it may be tempting to see them as operating sequentially; in particular, it may be tempting to think of the reality constraint as operating *prior* to the change constraint. But to my mind this would be a mistake. Granted, it's quite natural to think of an ideal imagining as involving three distinct steps: (1) We imagine the world (or some part of it) as it is; (2) we make one or some targeted changes as governed by our overall imaginative project; (3) we adjust our

imagining appropriately in light of all and only the consequences of the changes made. But the respect in which these steps are distinct from one another is merely a *conceptual* one and not a *temporal* one. Certainly steps (1) and (2) can occur at once, and it may be that part of doing (2) is doing (3)—making the relevant adjustments, that is, might well be part and parcel of what it is to imagine some change to the world as it is. All three steps might thus be achieved simultaneously as the imagining is developed. Likewise, then, the reality and change constraints should not be seen as operating one before the other. Rather, they work closely in tandem with one another as the imagining unfolds.

Though Campbell's story suggests that these constraints govern the machines' ideal imagination, and thereby enable the machines to learn about the world via imagination, we might wonder whether such constraints are necessary for our imaginings to have epistemic significance.¹⁰ Consider first the reality constraint. Must imaginings represent the world, or some relevant subset of it, *exactly as it is* in order to have the sort of epistemic significance that interests us? Here it seems clear that the answer is no. Though the machines cannot learn whether a proposed defense system will be effective against the alien attack if they imagine fewer aliens than there actually are, they might well be able to learn this if they imagine more aliens than there actually are. A defense system that's effective against two million aliens will also be effective against one million aliens. In fact, the machines might deliberately want to overestimate the strength of the attacking army to give themselves some room for error. Other kinds of examples make the same point. When imagining whether the stroller I just bought will easily fit in the trunk of a car I'm looking at on the showroom floor, it doesn't matter if I mistakenly imagine the

¹⁰ I am grateful to Tyler Doggett for pushing me on this issue.

stroller as green rather than yellow; it also doesn't matter if I mistakenly imagine that the stroller was made in the United States rather than in China.

Next consider the change constraint. In our discussion of the machines' ideal imagination above, I noted that they tested their proposed plan in imagination by imagining all and only the changes to the world that its implementation would cause. But must all imaginings proceed this way in order to have the sort of epistemic significance that interests us? Suppose that once again to provide themselves with some room for error the machines were to deliberately underestimate how much additional protection would be provided by a proposed upgrade to their defensive shields; they imagine a lesser change than would occur were their plan to be implemented. Here too it seems that they could learn from their imagining. Thus, like the reality constraint, the change constraint also does not appear to be necessary for an imagining to have epistemic significance.

Perhaps it would be possible to reformulate the reality and change constraints so that they would serve as necessary conditions for an imagining to have epistemic relevance. Rather than requiring that the world be imagined as it is, for example, we might say that the world be imagined as it is *in all relevant respects*; rather than requiring that an imagining be constrained by all the consequences of the change imagined, we might say that it be constrained by all and only the *relevant* consequences of the change imagined. The plausibility of these reformulations would then hinge on how we cash out the notion of relevance.

But while I am optimistic about the prospect of coming up with an appropriate reformulation along these lines, I will not attempt to do so here. Rather, for the purposes of this discussion I'd suggest that we think of the reality and change constraints as

aspirational in nature. Some imaginings will come closer to meeting them than others. The closer we come to meeting these constraints—the closer we come to being ideal imaginers—the more likely we will be able to learn from our imaginings. In aiming to learn from our imaginings, we should thus aspire to meet them.

Here an analogy to perception might help. Though ideal perception requires that our perceptions represent the world exactly as it is—that they meet a reality constraint for perception—we can still learn about reality from many perceptions that fall short of the ideal. As a result of my aging eyes and a habit of misplacing my reading glasses, the words on the front page of the newspaper often look blurry to me as I'm reading it over breakfast. Since the words are not actually blurry, my perceptions misrepresent the world. But, despite the blurriness, I can still learn from them what today's *Los Angeles Times* headlines are. Likewise, though ideal imagination requires that our imaginings satisfy the reality and change constraints, we can still learn from imaginings that fall short of this ideal.¹¹

3. An Important Worry: We Are Not Machines

Given the two constraints required for ideal imagination, it shouldn't be surprising that machines would be particularly good at it. Machines, after all, are particularly good at following out logical consequences—that is, they are particularly good at abiding by the change constraint. Moreover, because of their immense storage capacity, it seems plausible that machines might well be able to have a sufficiently complete sense of the

¹¹ I will return to an analogy to perception in Section 4.

world as it is, or at least a relevant subset of it, to meet the reality constraint. But these brief reflections suggest something troubling about the notion of ideal imagination. If we define ideal imagination in terms of what machines can do, then it begins to look as if we humans will inevitably fall short—and in fact, that we will fall so far short that we won't be able to learn from our imaginings at all.

In responding to this worry, it's worth first getting clear on exactly what it entails. Most importantly, it fails to lend support to the proposition that the imagination lacks epistemic relevance. Suppose for a moment that it were true that the cognitive limitations of humans prevent us from being ideal imaginers. Even so, my discussion of ideal imagination would still help to show why the imagination does not *in principle* lack epistemic significance—that is, it is not something about the nature of imagination that renders it incapable of giving us knowledge about the world. Compare hoping, for example. There's no way even in principle to constrain hoping so that it would have epistemic relevance. No matter how constrained hoping is, if it is really a case of *hoping*, it is not something from which we can learn anything about the world. The imagination is different. While the nature of hoping is inconsistent with its having epistemic relevance, our discussion of ideal imagination reveals that the nature of imagining is not inconsistent with its having epistemic relevance. And this remains true even if it turns out that the limitations of humans are such that none of our imaginings meets (or comes close to meeting) the ideal—and thus that, in practice even if not in principle, none of our imaginings has any epistemic relevance.

That said, this would certainly be a disappointing result. Fortunately, however, I think we can avoid it by confronting the worry more directly. In short, my response

consists of two key claims. First, I think that we're often better imaginers than the worry above gives us credit for. But second, I want to claim that even when we do fall short of the capabilities of Campbell's machines, we can still come close enough to the ideal for our imaginings to have epistemic significance.

In fact, many people are very, very good at imagining—and, in particular, at imagining under constraints. Consider two such relatively well-known people: the inventor Nikola Tesla and the animal scientist Temple Grandin. Tesla, who is perhaps best known for developing the alternating current technology in widespread use around the world today, was also responsible for the invention of high-voltage electrical coils, long-distance electrical transmissions lines, hydroelectric generators, bladeless turbine engines, X-ray tubes, and various radio-controlled devices. Of interest to us, however, are not his inventions themselves but rather the creative process underlying their development—a creative process that owed to Tesla's remarkable powers of visualization. As Tesla himself described it, he could perfectly picture his inventions in advance of their creation:

Before I put a sketch on paper, the whole idea is worked out mentally. In my mind, I change the construction, make improvements, and even operate the device. Without ever having drawn a sketch, I can give the measurement of all parts to workmen, and when completed these parts will fit, just as certainly as though I had made accurate drawings. (Tesla 1921)

Also according to Tesla, his method was flawless:

The inventions I have conceived in this way, have always worked. In thirty years there has not been a single exception. My first electric motor,

the vacuum tube wireless light, my turbine engine, and many other devices have all been developed in exactly the same way. (Tesla 1921)

Although Grandin's inventions are perhaps not as well-known as Tesla's, over the past several decades, she has significantly improved the welfare of animals throughout the world by revolutionizing the design of livestock-handling facilities. Approximately half of the cattle in North American meat plants are now handled in systems that she designed. Grandin's description of her own creative process sounds remarkably like Tesla's: "Visual thinking has enabled me to build entire systems in my imagination" (Grandin 1995, p. 19).¹² As she also notes, "in my work, before I attempt any construction, I test-run the equipment in my imagination. I visualize my designs being used in every possible situation, with different sizes and breeds of cattle and in different weather conditions. Doing this enables me to correct mistakes prior to construction" (Grandin 1995, pp. 20–1).

One of Grandin's early design successes came in 1978 when she developed an innovative dip vat design for a cattle-handling facility in Arizona. A dip vat, which is filled with pesticide to rid animals of parasites, is a long narrow pool-like structure in which cattle are completely immersed while proceeding through it single file. Prior to Grandin's designs, cows would often panic both when approaching the dip vat and when exiting it. By taking a "cow's eye view" of the situation, Grandin diagnosed the problems with the existing structures and was able to create an alternative in which the cows would calmly enter and exit the equipment voluntarily, without any use of force. Her design process, however, took place entirely in her mind: "I started running three-dimensional

¹² Interestingly, Grandin explicitly compares herself to Tesla. See Grandin (1995, p. 26).

visual simulations in my imagination. I experimented with different entrance designs and made the cattle walk through them in my imagination. These images merged to form the final design.” (Grandin 1995, p. 23)

When considering these first-personal accounts of Tesla’s and Grandin’s use of visualizations—and here I have included only a few key excerpts—it’s hard not to see the imaginings they describe as eerily machine-like. Like the machines of Campbell’s “The Last Evolution,” Tesla and Grandin each have the ability to abide by the reality and change constraints we discussed above. When coupled with an incredible knowledge base—itself also machine-like—about the relevant domain, each of these individuals was able to engage in imaginings from which they learned something about the world. (Grandin, for example, discovered via imagination what dip vat design was most compatible with the temperament of cows.) Contrary to the worry under discussion then, at least some humans seem capable of the ideal imagination enjoyed by the machines.

Clearly Tesla and Grandin have extraordinarily well developed powers of imagination, so much so that their abilities will likely seem far out of reach to those of us whose powers are more mundane. For this reason, the worry that we’ve been grappling with in this section might still seem to have bite. Perhaps there are a few rare cases where humans can achieve machine-like ideal imagination, but if these cases are so few and far between, and if the rest of us can’t even come close, then the imagination won’t have much epistemic significance after all. But why should we think such cases are so few and far between? Why should we think that the rest of us can’t even come close? Here it will be helpful to recall the examples we saw above in which it seemed plausible that an act of imagination could provide an imaginer with reason for beliefs. A prospective parent—

call her Imogen—might engage in various imaginings in an effort to figure out whether she is ready to have children, or what kind of car to buy, or where to put the crib. If she wants these imaginings to be epistemically useful to her, it seems natural that she will deliberately attempt to line them up with her conception of the world. These aren't mere daydreams, but rather are controlled—that is, *constrained*—imaginings. For example, when she imagines the crib against the various walls of the room that will become the nursery, her imagining is constrained by the actual size of the room, the actual size of the crib, the placement of the windows and doors in the room, and so on. Will the door still comfortably open if the crib is placed against the west wall? This can best be determined if both the reality and the change constraints are met—if she imagines the room as it actually is, but suitably adjusted for the insertion of the crib. Likewise, when she imagines a car seat in the back of a two-door sports car, her imagining is constrained by factors such as the size of the car, the size of the infant seat, and the angle of the front seat as it's tilted forward. With the front seat tilted forward and a car seat in back, will there be enough room for a parent to reach in and strap in an infant? Again, this can best be determined if the reality and change constraints are met—if she imagines the car as it actually is, but suitably adjusted for the insertion of a car seat.

Granted, there may be some folks who can't even do what Imogen does, let alone what Tesla and Grandin do. The only way for such a person to determine whether the car seat really fits in the back of the sports car would be to test it out with an actual car seat and an actual car. But many of us can make these determinations without actually trying it out—we can learn something via an imaginative test rather than an empirical test.

How do we do this? We are not machines, and most of us are not even Tesla and Grandin. We're not as good as the machines in setting the relevant constraints, and it doesn't seem that we're as good as the machines in abiding by them once they are set. But in many cases—and which cases these are will vary from person to person—we do seem to be able to be good enough. In part this is because the sorts of imaginative projects in which we're involved will be considerably less fine-grained than the projects undertaken by the machines. Considerably less precision is required for us to successfully imagine many of the kinds of things that we imagine. In part this is also because the sorts of imaginative projects in which we're involved will be considerably less complicated than the projects undertaken by the machines. In many cases, considerably fewer variables are in play. Thus, even though it seems true that we will typically fall short of the machines, that we can't match their capabilities, we can nonetheless approximate their capacity—and in many cases, that will be enough.

4. Another Important Worry: We Don't Know What We're Doing

Though I hope that the reflections of the previous section have laid to rest the worry that the capacity for ideal imagination is so far beyond our reach that we can't learn from our imaginings, there is a related worry in the vicinity. Perhaps some of our imaginings will be ideal (or close enough), but how can we tell? If we can't, then doesn't this limit the epistemic significance of the imagination?

To flesh out this worry, let's return to Imogen. We assumed that she was a pretty good imaginer—not as good as Tesla and Grandin, perhaps, but good enough that she could plausibly count as approximating the ideal for the limited sorts of imaginings she was engaging in. But what if she's not? Blinded by her love of sports cars, her imagining might wildly misrepresent the car's interior size—almost as if it had been magically transformed by one of Hermione Granger's Undetectable Extension Charms. (When an object is transformed by such a charm, it is made bigger on the inside without being made bigger on the outside; what appears to be a small handbag, for example, becomes capable of holding numerous textbooks, clothes, a tent, and even a sword.) An imagining that misrepresents in this way is not appropriately constrained—it is no longer even coming close to operating under the constraint of reality—and it is precisely this lack of constraint that robs the imagining of its epistemic usefulness. An imagining that wildly exaggerates the size of a car's interior cannot teach Imogen anything about whether an infant seat fits in the car's backseat.

What might seem troubling, however, is that Imogen might very well have difficulty determining whether she has violated the constraint of reality. She might be convinced—even reasonably so—that she has represented the car's interior correctly when she has not. Perhaps the machines can run self-diagnostics to ensure that they are operating within standard parameters, but we cannot. In general, we can't be completely confident about whether we have abided by the constraints of reality and change. But if we can't determine this, if (to put it bluntly) we typically don't know what we're doing, this might well seem to undercut the epistemic significance of the imagination.

But let's think again about what has gone wrong with Imogen's imagining in the case we're now considering. In such a case, Imogen is the victim of what we might call *imaginative illusion*. She takes herself to be imagining the car's interior as it is, but her imagining does not represent the car's interior as it is. Here there is an obvious parallel to certain sorts of perceptual illusions—we might take ourselves to be seeing the tower in the distance or the stick in water as they are, but we're not—the tower isn't really round, and the stick isn't really bent. The fact that we can be victims of perceptual illusions does not rob perception of its epistemic significance. Likewise, then, the fact that we can be victims of imaginative illusions should not rob the imagination of its epistemic significance.

And now let's reflect for a moment further about perceptual illusions. In general, we can identify various conditions in which we are more likely to be the victim of perceptual illusions—when the lighting is bad, or when the objects we are viewing are far away, or when we are under the influence of alcohol or hallucinogenic drugs. Relatedly, we can take steps to minimize the possibility of such illusions—we can turn up the lights, or move closer to the objects we're viewing, or we can abstain from alcohol and hallucinogenic drugs. Likewise, I think, we can identify various conditions in which we are more likely to be the victim of imaginative illusions. If Imogen has never seen the inside of the sports car, or she hasn't seen it for a very long time, then she should be wary about her ability to constrain her imagining appropriately. Additionally, as we suggested above, her passion for sports cars might be distorting her ability to imagine the interior correctly. Here, her desires have a distorting influence similar to the distorting influence

of poor lighting or long distance. In cases where she has such desires, then, she should also be wary about her ability to constrain her imagining appropriately.

Ultimately, however, we cannot *always* identify that we have fallen victim to a perceptual illusion, and likewise, we cannot always identify that we have fallen victim to an imaginative illusion. There may be some skeptics who think this fact about perception robs perception of its epistemic significance. But setting aside such skepticism, as long as we think that the possibility of perceptual illusion does not prevent us from learning something about the world from our non-illusory perceptions of it, the possibility of imaginative illusion does not prevent us from learning something about the world from our non-illusory imaginings of it.

5. Concluding Remarks

According to Campbell's depiction of humankind in "The Last Evolution," our imagination may be brilliant, but it is illogical. Unlike the machines, we can advance by leaps and bounds. But also unlike the machines, our conclusions are not always correct. As I have tried to suggest in this paper, however, our imagination need not be as illogical as Campbell makes it out to be. While many of our imaginings are unable to teach us about the world, we have the power to constrain our imagination and, in doing so, to imbue our imaginings with epistemic significance.

Importantly, I do not mean here to privilege constrained imagining over unconstrained imagining. Both of them have their place in human life. Our unconstrained imaginings are important for many of the activities that are important to us, and rightly so—from engaging with art and literature to fantasizing or pretending. But it's our

constrained imaginings that are important for our attempts to learn about the world—and this kind of knowledge cannot come by leaps and bounds. Rather, in modeling our imagination on the ideal imagination of the machines, we are able to make epistemic progress the way they do, by steady, irresistible steps.

Acknowledgements

In thinking about these issues, I learned a lot from Tyler Doggett and Jennifer Church, both of whom served as commentators when an earlier version of this paper was presented at the APA Pacific Division meeting in 2013. I owe them both thanks. I am also grateful to the audience members during that session for their feedback, as well as to audiences at the University of Kentucky, Denison University, and the Workshop on Mental Imagery and Pretense at the University of Antwerp. Thank also to Peter Kung for helpful discussion and feedback.

References

Campbell, John W. (1932). The Last Evolution. *Amazing Stories*.

<<http://www.gutenberg.org/files/27462/27462-h/27462-h.htm>>.

Campbell, John W. (1966). Introduction. *Analog* 6.

Chalmers, David J. (2002). Does conceivability entail possibility? In Tamar Szabó

Gendler and John Hawthorne (eds), *Conceivability and Possibility*. Oxford:

Clarendon Press, pp. 145–200.

Gendler, Tamar Szabó, and John Hawthorne (2002). Introduction: Conceivability and

possibility. In Tamar Szabó Gendler and John Hawthorne (eds), *Conceivability and*

Possibility. Oxford: Clarendon Press, pp. 1–70.

- Grandin, Temple (1995). *Thinking in Pictures: My Life With Autism*. New York: Random House.
- Kind, Amy (2001). Putting the image back in imagination. *Philosophy and Phenomenological Research* 62(1): 85–109.
- Kind, Amy (in press). How imagination leads to knowledge. In Fabian Dorsch and Fiona Macpherson (eds), *Perceptual Memory and Perceptual Imagination*. Oxford: Oxford University Press.
- O'Shaughnessy, Brian (2000). *Consciousness and the World*. Oxford: Clarendon Press.
- Paul, L. A. (2014). *Transformative Experience*. Oxford: Oxford University Press.
- Sartre, Jean-Paul (1948). *The Psychology of Imagination*. New York: Philosophical Library.
- Walton, Kendall L. (1990). *Mimesis as Make-Believe: On the Foundations of the Representational Arts*. Cambridge, MA: Harvard University Press.
- Wisehart, M. K. (1921). Making your imagination work for you: An interview with Nikola Tesla ... *American Magazine* 91 (April 1921): 13ff.

Abstract

As Hume famously claimed, we are nowhere more free than in our imagination. While this feature of imagination suggests that imagination has a crucial role to play in modal epistemology, it also suggests that imagining cannot provide us with any non-modal knowledge about the world in which we live. This chapter rejects this latter suggestion. Instead it offers an account of “imagining under constraints,” providing a framework for showing when and how an imaginative project can play a justificatory role with respect to our beliefs about the world. That we can be free in our imaginings does not show that they must proceed unfettered; as is argued, our ability to constrain our imaginings in light of facts about the world enables us to learn from them. The important upshot is that imagination has considerably more epistemic significance than previously thought.

Keywords

imagination, modal epistemology, justification, epistemic significance, simulation, Temple Grandin, Nikola Tesla