Technology and Knowledge

My aim in this paper is to give a philosophical analysis of how, precisely, technology can be a condition for gaining scientific knowledge. My concern is with what scientists can know in practice, given their particular contingent conditions, including available technology, rather than what can be known “in principle” by a hypothetical entity like Laplace’s Demon. I begin with the observation that what we know depends on what we can do. For example, in science, gaining certain knowledge depends of having certain evidence. This makes the ability to gather that evidence a necessary condition for gaining the knowledge. I’ll argue that a scientist is (under certain conditions) expected to seek evidence before making a judgment, meaning that the “epistemic possibility” of attaining scientific knowledge sometimes depends on the possibility of undertaking certain activities. In turn, the possibility of undertaking certain activities depends in part on factors like ethical constraints, economical realities, and available technology. I’ll focus on technology, and in particular on scientific instruments, and introduce a new way to analyze the set of actions made possible by changes in technology. Specifically, I’ll argue that changes in technology make certain activities “technologically possible,” and these activities can under certain circumstances extend our knowledge, for example by making new evidence available to scientists. That is, the epistemic possibility of gaining access to scientific knowledge depends (in some cases) on the technological possibility for the construction and operation of scientific instruments.

Peter Galison argues for the necessity of the digital computer to certain advances in scientific knowledge during the Manhattan Project. “Some kind of numerical modelling was necessary, and here nothing could replace the prototype computer just coming into operation in late 1945: the ENIAC” (Galison 1996, 122). Galison’s claim is that a technological change is what made it possible for scientific knowledge to develop. John Agar argues against the necessity of the technological change for advancing knowledge: other approaches could have sufficed. “Computerization was usually first proposed when the existing practices and technologies were still capable of the computational task at hand” (Agar 2006, 873). Implicit in these statements are crucial assumptions about the relationship between technology and action and between action and knowledge. Galison and Agar agree that certain scientific knowledge became accessible to Manhattan Project scientists with the capacity to run Monte Carlo simulations. They disagree about what changes in the situation of the scientists enabled them to run Monte Carlo simulations. Galison thinks access to digital computers made the difference. Agar thinks that the Monte Carlo method could have been implemented using existing computational approaches.

I propose that the apparent conflict between Galison's and Agar's claims about the digital computer is an illusion arising because these authors invoke different implicit notions of possibility. The commonsense notion of possibility is just *what might happen*, *what might exist*, or *what might be true*. What complicates matters is that we freely embellish the notion of possibility in particular contexts, adding new differentia into the mix. For example, "driving the wrong way down a one-way street," is physically possible, because the laws of physics do not forbid it, but because a municipal law does forbid it, it is regulatively impossible. We commonly indicate which modal concept we mean by affixing a label to the claim[1] or by allowing context to communicate our categories. It is not problematic that a proposition (or corresponding action or state of affairs) is possible according to some modal concepts and impossible according to others. Ordinary language labels are usually sufficient to identify the relevant distinction. For example "traveling to the nearby star Alpha Centauri by 2020," might be theoretically and physically possible, but at the same time technologically and economically impossible.

To return to the case at hand, Agar is saying, roughly, that computers were not necessary for completing the needed work; Galison, that they were. What is at stake in this disagreement is an understanding of the role that a particular piece of technology made in the practice of science at a particular time. According to Galison, the digital computer is the piece of technology that brought certain propositions into the realm of scientific knowledge – that is, they changed what was technologically possible (allowing the Monte Carlo method to be put into practice), and that changed what was epistemically possible for the scientists. Agar, on the other hand, thinks this way of putting things gives too much credit to the material means of accomplishing a task that should be understood as having been made possible by a conceptual change (the Monte Carlo method). That is, the Monte Carlo method was technologically possible without computers.

[1] An informal survey by the author of the philosophical literature since 1950 dealing specifically with possibility counted at least 19 distinct modal concepts, with 47 different names for them.

In order to understand these claims about the difference technology makes to scientific knowledge, we must understand the relationship between technology and action and we must understand the relationship between action and knowledge. The rest of this paper is devoted to these two tasks. The paper is divided into two parts. In the first part, I consider the relationship between action and knowledge. I'll argue that being able to complete certain actions can be a necessary condition for gaining knowledge. For example, in scientific practice, gaining certain knowledge depends on having certain evidence. Being able to gather that evidence is therefore a condition for gaining the knowledge. I'll argue that a scientist is (under certain conditions) *expected* to seek evidence before making a knowledge claim within his or her domain of expertise, and I'll construct an account of "epistemic possibility," the possibility of knowing, that captures the dependency of knowledge on action.

In the second part of the paper, I turn to the relationship between technology and action. A number of other factors affect our ability to act, including economics and ethics, but I will focus just on technology, touching on the others only incidentally. I'll introduce "technological possibility," which depends on the availability of material and conceptual means to bring about a desired state of affairs. The possibility of spanning a river with an iron bridge turns on what the world is like: that iron is available, has certain properties that allow it to be formed into trusses, and so on. But the possibility of spanning a river with an iron bridge also turns on how our concepts fit together. If we did not know that iron has the properties it does, it would not occur to us to attempt the project. What this means is that physical possibility and conceptual possibility are necessary but insufficient conditions for technological possibility.

Part 1. Doing and Knowing

In this section, I will develop an account of epistemic possibility that takes into consideration practical conditions for the possibility of knowing. In addition, because I am particularly interested in scientific knowledge, I will develop my account in such a way that it can accommodate epistemic duties, such as the evidence-seeking duties scientists adopt when they aim to produce scientific knowledge. As I will argue, this requires a definition of epistemic possibility that meets a practicability criterion and a responsibility criterion. My approach will be to begin with a standard definition of epistemic possibility and then elaborate it so that it brings clarity to cases like the Monte Carlo case described above.

Epistemic possibilities are relative to an agent's epistemic position. Thus, for an agent S to claim that the proposition Φ is epistemically possible is for S to say that Φ is possible relative to S's epistemic position. Taking "S's epistemic position" to mean "what S knows" leads straightforwardly to the standard definition of epistemic possibility (see, e.g., Gibbs 1970; Gendler and Hawthorne 2002):

(a) Φ is epistemically possible for S if S doesn't know –Φ.

The idea is that if S doesn't know that Φ is not the case, then S must consider Φ to be possible. For example, if S knows Φ, then S cannot know –Φ, and Φ is epistemically possible. On the other hand, if S knows –Φ, then Φ is trivially epistemically *im*possible. Finally, if S knows neither Φ nor –Φ, then Φ is epistemically possible for S. More concretely, if I have just checked my pockets for my lost keys (and failed to locate them there), then for me, it is not epistemically possible for my keys to be in my pocket. But if I have not yet looked on the table, then it is epistemically possible for my keys to be there (just as it is epistemically possible for them not to be on the table).

Although epistemic possibility is meant to reflect the peculiarities of S's epistemic position, it is possible for S to be wrong about epistemic possibilities. S might deem Φ possible, forgetting to take into account that –Φ, for example. In such a case, we simply say that S is wrong: S should know better than to think Φ is possible. That is, Φ is in fact *not* epistemically possible for S, even if S thinks that Φ is possible. Epistemic possibility is about what S *can* know based on S's epistemic position rather than what S does in fact conclude.

As it turns out, this leads to a strange consequence. Definition (a) presumes that the only factor relevant to S's epistemic position is what S knows at the time, a condition that fails to hold for any proposition S hasn't considered before. Suppose Φ is the proposition that "4+3=9," something S would reject upon even a moment's consideration. Nevertheless, if S has never considered whether "4+3=9," then "4+3=9" is epistemically possible for S, because S has no beliefs about it. One this view, if S blurts that "'4+3=9' is possible" without pausing to consider it, we have no cause to say S is wrong, for "4+3=9" really is epistemically possible for S. Yet if S should later consider whether "4+3=9," S would immediately judge it to be impossible, and could then be blamed for saying that "4+3=9" is possible.

If the goal of epistemic possibility is to reflect S's current epistemic position, it seems like a strange consequence that we can blame S for having forgotten Φ, but not for never having thought about it. Richard Foley suggests that "our everyday evaluations tend to be concerned with whether one has been responsible in arriving at one's beliefs" (Foley 2003, 9). Both remembering and considering seem like relevant responsibilities for what S is in a position to know, but definition (a) treats the two differently because it takes "S's epistemic position" to be identical with "what S knows." If epistemic possibility is to be of use in understanding real situations, we must elaborate our understanding of what "S's epistemic position" is.

I propose to expand the notion of epistemic position to include what S is in a position to know in addition to what S actually knows. Let me give an example of what I mean. In the context of scientific knowledge claims, we routinely expect knowledge claims to carry special weight in light of experimental evidence or theoretical justification. That is, we impose special responsibilities, or *epistemic duties*, on scientists (see, e.g., Kornblith 1983). Epistemic duties are controversial in epistemology, but the idea is straightforward: sometimes, we require that claims be backed up. We see epistemic duties in action all the time. When I call my office to ask a colleague to check my mailslot for a letter I am expecting, I won't be satisfied with the claim that it is "possible" that the letter has arrived – I want to know! I expect my colleague to check before making claims. But expectations have limits: I won't blame my colleague for not noticing that the letter has been delivered to the wrong mailslot. The point is that epistemic possibility should not only reflect what justified true beliefs S has, but should also take into consideration S's responsibilities to gather additional evidence. On the epistemic duties view, S should not always settle for the evidence she has in hand, but must in at least some cases conduct an inquiry or seek new evidence before making a knowledge claim.

I should mention that an alternative to the epistemic duties view states that although we do sometimes have the duty of seeking new evidence, that duty should be understood as being moral, not epistemic (see, e.g., Conee and Feldman, 2004). My account is compatible with either view, but I shall use the term "epistemic duty" to indicate any duty that is a condition for making a knowledge claim (whether or not the duty promotes genuine knowledge). I will also side-step the larger question of whether we always have epistemic duties, and instead distinguish between *weak epistemic possibility*, which does not include duties, and *strong epistemic possibility*, which does include duties. For the remainder of the paper, when I refer to "epistemic possibility" I mean strong epistemic possibility.

Let me digress briefly to clarify a subtlety in epistemic duties. There is a distinction to be maintained between the question of when Φ is justified and when S is justified in claiming to know Φ (see, e.g., Williams 2001). Epistemic duties, as I use the term here, have to do with claiming. S's making a claim about Φ invokes a responsibility, but not an ultimate responsibility. What determines epistemic responsibility is not the actual epistemic risk of a particular claim, but its perceived risk relative to a particular context.

It may seem worrying to sweep aside the epistemic risks associated with a given context, but the idea is not to ignore those context-level risks, merely to focus on the risks associated with a particular claim within that context. Focusing on the contextual claim has the effect of normalizes its risks against a chosen reference background. Background risks don't disappear; they are simply shifted away from center stage.

Background risks can be accommodated in a number of ways. We can demand that they be traced (or be traceable) to basic beliefs (as in foundationalist accounts of knowledge). We can reduce them to mere stipulations (as in some relativist accounts). Or we can recognize that what counts as background is negotiable. As Helen Longino puts it, "as long as background beliefs can be articulated and subjected to criticism from the scientific community, they can be defended, modified, or abandoned in response to such criticism" (Longino 1990, 73-74). The effect of putting background on the bargaining table is to create a sort of "division of labor" for epistemic risks. Even if some risks are unaddressed or unknown at a given time, they can be articulated and worried over at a later date (recalibrating dependent risk budgets accordingly). This negotiation model seems to fit the way that science works, at least some of the time. For example, in order to conduct detailed research, a scientist interested in molecular physics takes on board risks associated with commitments to causality, mass-energy conservation laws, statistical laws, and so on. But setting those risks aside doesn't mean accepting them unquestioningly; scientists decide which risks they need to address before making a claim, and other scientists decide whether the appropriate risks have been addressed before accepting the claim.

Part of what is being negotiated is who is responsible for addressing particular epistemic risks. Responsibilities may be stronger, weaker, or even unrelated to the justification standards for knowledge. Ideally, epistemic duties for S being able to claim Φ would exactly match what justifies Φ. But suppose that there is a mismatch, and responsibilities are insufficient to justify Φ (this being typical of cases that worry philosophers). Even if fulfilling duties is not sufficient for genuine knowledge, fulfilling epistemic duties is still necessary for having knowledge of Φ, because any claim that fails to fulfill responsibilities is a non-starter in the context in which it is made. It would be nice if we knew which responsibilities are relevant to justification, but we simply cannot be certain. Inquiry in science, as in other fields, works on the basis of its internal standards, which sometimes produce genuine knowledge and sometimes not. But in order for a claim to be eligible for consideration, the claimant has to fulfill the relevant epistemic duties. I will refer to this as the 'responsibility criterion' for epistemic possibility.

As we have seen, definition (a) fails to provide an adequate 'responsibility' criterion for epistemic possibility. But perhaps such a criterion can simply be added to the original definition. For example, we might say that:

> (b) Φ is epistemically possible for S if S does not know that –Φ,
> *nor would careful reflection establish that –Φ.*

Here, "careful reflection" is an epistemic duty for S, one that clearly eliminates unconsidered alternatives like "4+3=9," while still being limited to the knowledge S has (plus inferences from that knowledge). As it turns out, however, "careful reflection" is too strong a requirement to be practicable, which means that definition (b) fails to appropriately reflect what S can know in practice. For example, Goldbach's conjecture states that every even integer greater than two can be written as the sum of two primes. It hasn't been proved or disproved, but the axioms of mathematics are such that Goldbach's conjecture, if true, is true necessarily, and if false, is false necessarily. Mathematicians don't yet know the truth-value, and many hours of careful reflection have not solved the problem. Nevertheless, some amount of additional reflection might solve it, as has been the case for many other mathematical conjectures. Suppose that Goldbach's conjecture is true, and that a million hours of careful reflection could establish that fact. Then by definition (b), S would be wrong to say that it is possible Goldbach's conjecture is false – for S has the duty of carefully reflecting on the problem.

Epistemic possibility is supposed to reflect S's particular epistemic position. Defining epistemic possibility in terms of the conclusions S could reach rather than the conclusions S does reach leads directly to the kind of "in principle" considerations I want to avoid. The worry is that "in principle" is in reality "impossible." It isn't clear from definition (b) what standard we should use to judge what is epistemically possible for S.

One plausible solution is to relativize epistemic possibility not to S (or an ideal S), but to S's actual epistemic community:

> (c) Φ is epistemically possible for S if S does not know that –Φ,
> *nor does any member of C*, where C is S's epistemic community.

The advantage to this definition is that it smoothes out some of the peculiarities of S's particular thought processes, while remaining true to human limitations. Even if S hasn't considered whether Φ, perhaps someone else in C (however we wish to define the community) has ruled it out. The aim is not to require that S know everything known to everyone else in S's community, but rather to hold S responsible for judgments that clash with what is known to someone else in the community. In a scientific community, this definition works rather well, because knowledge is (ideally) made available to the entire community by mechanisms such as publication. On definition (c), S can be deemed wrong on the basis of failing to take into account results published by other scientists. Unfortunately, this elaborated version of epistemic possibility has difficulties of its own, as Ian Hacking shows.

> Imagine a salvage crew searching for a ship that sank a long time ago. The mate of the salvage ship works from an old log, makes some mistakes in his calculations, and concludes that the wreck may be in a certain bay. It is possible, he says, that the hulk is in these waters. No one knows anything to the contrary. But in fact, as it turns out later, it simply was not possible for the vessel to be in that bay; more careful examination of the log shows that the boat must have gone down at least thirty miles further south. (1967, 148)

No doubt it seemed possible that the vessel was in the bay until the ship's mate rechecked his calculations. But was it really epistemically possible? To Hacking, it seems not, for the evidence the mate used to justify his belief that it is possible that the vessel is in the bay does not, in fact, support that claim. It supports the contrary claim that it is impossible that the vessel is in the bay. Hacking concludes that "the mate said something false when he said, 'It is possible that we shall find the treasure here,’ but the falsehood did not arise from what anyone actually knew at the time” (1967, 148).

Hacking is pointing out that in many cases there is an expectation that S has checked – and has done a good job – before making a claim about Φ. For Hacking,

> (d) Φ is epistemically possible for S if S doesn’t know –Φ *nor would any practicable investigations by S establish that –Φ.*

Here, the idea is that we expected the mate to successfully complete certain reasonable actions before coming to his conclusion. Since he didn’t complete them successfully, we have grounds to blame him. Hacking’s definition allows us to resolve the Goldbach case satisfactorily: S is now responsible to complete investigations that fall within practicable limits. Exactly where we draw that line is still vague, but at least we now have a principle for drawing one. I will refer to this as the ‘practicability criterion’ for epistemic possibility.

As it turns out, ‘practicability’ alone doesn’t always line up with the sort of epistemic duties we impose on S. Paul Teller gives this rebuttal to Hacking’s practicability criterion: Teller’s wife is pregnant, but he doesn’t yet know the sex of his child. For Teller, it is epistemically possible that his child will be a boy, and at the same time epistemically possible that his child will be a girl, and this is despite the fact that there is a “practicable, in fact quite easy” test to establish the sex of Teller’s child. (Incidentally, according to Teller’s account, the sex test was newly available in 1972. A few years earlier, it would have been technologically impossible.) Teller is claiming that we can’t demand that he have this test performed before he answers whether it is possible that his child will be a boy. Remember, Hacking introduced practicability to indicate what S should be expected to know, given S’s situation. He wants us to conclude that the ship’s mate has said something false in claiming the wreck may be in this very bay because the evidence he has examined should have told him otherwise. The mate made a mistake in his calculation, and it is easy to seize upon this and say that the mate should have known better. But the relevant contrast isn’t between what the mate knew after he examined the log and what he should have known. It’s between what he should have known before checking and afterward. What do we want to demand of the ship’s mate before he has looked inside the logbook for the first time? At that time, the mate’s position is similar to that of our expectant father before a sex test has been performed on the fetus.

The mate need only make calculations from the log and the father need only order the test. In each case, should S successfully complete some activity, knowledge of Φ can be had. The difference is not in the practicability of the task: the father's task is easier, if anything. The difference, Teller surmises, is in the expectations of the community C of which S is a member.

Teller therefore proposes the following emendation of the "community C" version of epistemic possibility:

> (e) Φ is epistemically possible for S if it is *not* the case that:
> *(1) Φ is known to be false by any member of community C,*
> *(2) nor is there a member, T, of community C, such that if T were to know all the propositions known to community C, then he/she could, on the strength of his/her knowledge of these propositions as basis, data, or evidence, come to know that Φ is false.* (Teller 1972, 310-311)

For my purposes, Teller's formulation has a significant problem: it doesn't get at responsibilities to look beyond existing knowledge. Like the original 'community' variation (definition (c)), it addresses only what is already known by the community. On the face of it, Teller's formulation doesn't even solve Hacking's salvage ship problem. Teller fills this hole by counting as "known to community C" facts written down in books available to the community (1972, 312). Recall: no one knows what is recorded in the logbook, for no one but the mate has examined it, and the mate has got it wrong. Hacking's point was that we need to establish some reasonable grounds for saying the mate is wrong. His answer was practicability, Teller's is, essentially, a slightly more detailed version of the responsibility account we saw earlier in definition (c). But responsibility to extant knowledge isn't quite what we're after for understanding responsible knowledge in scientific contexts – we usually want scientists to go out into the world and check.

My diagnosis is that both Teller and Hacking have part of the story right. The difficulty in defining strong epistemic possibility is in correctly balancing the practicability and responsibility criteria. This is difficult to do outside of specific contexts, and the solution is to avoid trying. That is, rather than try to define practicability or responsibility separately in some objective manner, the solution is to observe that communities negotiate and define practicable responsibilities for themselves based on their interests, including assessments of epistemic risk. In the case of scientific communities, practicable responsibilities are (partially) explicit: scientists must meet specific standards of evidence and justification or else withhold judgment or use qualified language. (Let me reiterate that outsiders need not accept a community's standards as being either necessary or sufficient for *knowledge*; rather, the point is that meeting such standards is what legitimates the claim within the relevant community.) Within a given community, C, if an individual, S, makes a knowledge claim and meets C's epistemic standards, E, then community C will accept it (whether it is knowledge or not). That is,

> (f) Φ is epistemically possible for S if S does not know that ‒Φ, *nor do the epistemic standards E of community C demand that S carry out any investigation that would establish that ‒Φ.*

Let's see how my account handles the examples we've just seen. On my account, an expectant father can rightly say that it is possible his child will be a boy even if a definitive test is available, because his community (his family and friends) does not demand that he order the test. By contrast, the mate on the salvage crew *is* expected to eliminate the present bay from the list of possible locations for the wreck, since his community (his shipmates) demands that he glean this information from the log. In each case, the relevant community (the community to which our subject S is presenting a claim) decides which practicable investigations S is obliged to undertake. It is the epistemic standard of the salvage crew that lets Hacking deem the mate wrong when he claims the wreck may be in this harbor. And it is the epistemic standard of family and friends that let Teller deem himself correct when he claims that his child may be a boy (even if his child were a girl).

My practicable responsibilities account of epistemic possibility can also illuminate practical discussions of possibility, such as the one with which I began the paper. Digital electronic computers became available at a time physicists at Los Alamos were butting heads with a difficult and dangerous subject matter: nuclear bombs. In order for their claims to be accepted, Manhattan Project scientists had to fulfill certain responsibilities, or epistemic duties. Specifically, they had to meet precise standards of evidence and justification in order for their work to move forward. Before the advent of the digital computer running Monte Carlo simulations, scientists did not fulfill those responsibilities, and so were stuck – they could not move forward on their bomb work, because needed knowledge was unavailable to them. And that knowledge was unavailable because certain actions hadn't yet been performed. Agar and Galison agree about all of this. But they disagree about whether the requisite actions could have been performed. That is, they disagree about what was practicable. What made the Monte Carlo method practicable, according to Galison, was the digital computer. Agar thinks the method could have been implemented using older techniques, that is, the Monte Carlo simulations were practicable before they were actually put into practice.

This leads to an important question: what makes an activity practicable? Many factors can affect practicability, including economics, ethics, and technology. In the Manhattan Project case, at least in the present incident, economics and ethics were pretty well stabilized. What changed was technology. The technological possibility early computers offered was the performance of Monte Carlo simulation in place of intractable analysis and dangerous experiment. Physicists took the option, and it proved fruitful: Monte Carlo simulations provided scientists with results they deemed useful—more useful in certain respects than analytic or experimental data.

A proposed investigation must be technologically possible for it to be carried out, and in at least some cases, the carrying out of the investigation changes what is epistemically possible. The relevant limitation to epistemic possibility is what we can find out given our technological constraints. In the 1950s, the Manhattan Project team wanted to learn about fusion. There are at least 100 billion samples of fusion strewn about the universe; every star has a fusion reaction at its centre. But none of these was available for experimentation. Such experimentation was (in the 1950s) technologically impossible, and could not therefore factor into epistemic possibility. Once technology changed and made direct experimentation feasible, technology then became relevant to epistemic possibility because epistemic possibility is constrained in part by the actions that can actually be carried out.

In the second part of the paper, I turn to technological possibility and the relationship between technology and action.

Part 2. Technological Possibility

Technological possibility describes what can be accomplished with available technology. Technology, in turn, has material and conceptual components. Thus:

> (g) An activity is *technologically possible* for an agent S if S has access to both the material and conceptual means to accomplish the task.

The key to technological possibility is the phrase "possible for" (see Gibbs 1970 for an analysis of "possible for"). This makes it agent-relative.

Technological possibility depends on what capabilities the materials available to S actually afford. The possibility of using Monte Carlo simulations to do bomb work at Los Alamos turns on what the world is like: material means like digital computers had to be available and had to be capable of certain things – simulations had to be *physically possible*.[2]

> (h) A state of affairs is *physically possible* if it is not precluded by the laws of nature.

Physical possibility is about the world, not our ideas of it; that is, physical possibility is not agent relative. Physical possibility is about what is possible *given the laws of nature*, which means that fusion experiments have been physically possible for billions of years (light from stars billions of light years away substantiates that). By contrast, technological possibility is about what is possible for a particular person (or group of persons) in a particular context: it takes a technological advance for scientists to perform fusion experiments. More precisely, physical possibility is necessary but insufficient for technological possibility.

But there is more to technology than just the material arrangement of parts. Technology has a "dual nature"[3] comprising material and concept. Technological possibility depends on concepts as well as material. As I mentioned in the introduction, the possibility of spanning a river with an iron bridge turns on what the world is like (i.e., that iron is available and has certain properties) as well as how our concepts fit together (i.e., that we think iron has certain properties that we can put to use in making trusses). In addition, instruments can change conceptually without changing materially. The cloud chamber was invented to study cloud formation in a scale model, but was quickly found to be useful in detecting charged particles. And instruments can change materially without changing conceptually. A more highly powered microscope has different lenses, but our understanding of its operations is unchanged.

---

[2] I choose physical possibility rather than, say, material possibility as my level of analysis because physical possibility is widely discussed, while material possibility is not. It seems prudent to explicate a new concept in relation to an established one.

[3] "Dual nature" is a term already common in the philosophy of technology, and now making headway into history and philosophy of science (see, e.g., Kroes & Meijers 2006).

This means that technological possibility depends not only on physical possibility but also *conceptual possibility*.

> (i) A state of affairs is *conceptually possible* for an agent S if concepts invoked by a description of the proposed state of affairs do not produce an inconsistency.

Conceptual possibility is agent relative. Like physical possibility, conceptual possibility is a necessary but insufficient condition for technological possibility.

I should note that conceptual possibility is distinct from logical possibility (a proposition is logically possible if it is not self-contradictory). For example, a “largest prime number” is logically impossible whether I know it or not. But a largest prime number is conceptually possible for me if I am unaware of the proof that there can be no such number. Whether an entity or state of affairs is conceptually possible, then, depends on how much conceptual ramifying an agent actually does. Like physical possibility, conceptual possibility is a necessary but insufficient condition for technological possibility.

The dual nature of technology, and of technological possibility, leads to ambiguities: when I state that it is impossible to travel faster than the speed of light, it is equivocal whether I mean that the laws of nature are such that it isn’t possible to travel faster than the speed of light, or that my concepts of ‘travel faster than’ and ‘the speed of light’ are such that their composition produces an inconsistency, or whether I just mean that I haven’t built the ship yet.

One way to avoid such ambiguities is to deny that there is a distinction between conceptual and physical possibility. Charles Hartshorne argues that because "all possibility is in a sense [conceptual], and all is in a sense [physical]" (1963, 593), no principled distinction can be sustained. I think this is mistaken, but let us see how such an account would work. According to Hartshorne, any *prima facie* distinction between conceptual and physical possibility will collapse under interrogation, revealing itself to be an artifact of the shorthand conceptual description we use to describe the situation in question. When the details are filled in, Hartshorne thinks, the distinction between the physical and conceptual dimensions will disappear. Of necessity, describing a situation involves leaving out many details. Some of these details will be incidental, others not. In discussing whether I can leap tall buildings, we will normally omit apparently irrelevant facts, such as the exact position of Alpha Centauri, but we may also leave unstated some background information that we would deem relevant, such as the laws of physics. The laws of physics make my proposed abilities physically impossible, and it is only by leaving them out of our description that such powers seem at first to be conceptually possible. But in fact we *are* committed to the laws of physics, and our having left them out of the shorthand description doesn't mean they can be left out of our conceptual analysis. "It is only because of lack of clarity or definiteness [about these commitments] that [physically] impossible descriptions appear to us as [conceptually] possible" (Hartshorne 1963, 594).

The evaluation of conceptual possibility involves determining the compossibility of a proposed state of affairs with a particular background context. In most practical cases, the context is vague or incompletely specified. Hacking says that we bring logic into the fold when the terms of individuation produce a contradiction (1975, 333). That is, whether a conceptual incompatibility exists can change depending on the level of detail we give to the terms we use to pick out a situation. In considering whether I could leap tall buildings, we might at first neglect to take into account some relevant details, such as the laws of physics, and wrongly judge the deed possible. But if we carry on filling in details, and Hartshorne argues we must, we will wind up in perfect agreement with physical possibility – in the end, the two are indistinguishable (1963, 595). The initial disagreement is merely an artifact of the level of detail we include in the initial analysis.

We can draw on George Seddon's illustration of how possibility analysis should work. An iron bar that floats on water has been supposed to be conceptually possible,[4] but physically impossible. ("Bar," clarifies Seddon, is meant to rule out needles, which float on surface tension, and the Queen Mary, which floats on "Zurich capital" [1972, 483].) Since it is physically impossible for an iron bar to float on water, filling in our concepts with more information about what it is to be water and what it is to be iron and what it is to float will lead to just the sort of self-contradiction that is characteristic of conceptual impossibility. In other words, the two kinds of possibility eventually become equivalent.

I suggest that Hartshorne is mistaken. Physical and conceptual possibility are not equivalent, and not two sides of the same coin. Conceptual impossibility turns on concepts, not the world. The confusion arises because we usually "fill in" our concepts with our knowledge of the world. For

---

[4] Seddon actually discusses logical possibility, so I have made some minor adjustments.

example, I might update my beliefs about the conceptual possibility of iron floating on water should I learn certain facts about iron, water, and floating. But there is no requirement that all of my beliefs be true. Conceptual analysis of empirically informed concepts is still conceptual analysis; looking for an inconsistency in what I know about iron (empirical research informs me that it has a specific gravity between 7.3 and 7.8[5]), water (that it has a specific gravity of 1.0), and floating (that an object can only float if its specific gravity is lower than that of the fluid in which it is placed), is precisely the kind of analysis required to evaluate conceptual possibility.

Seddon ultimately relativizes his version of conceptual possibility to the best present scientific theories. I think this approach is mistaken, as it will fails to describe many practical situations in which people are not conversant in the best present theories. It may or may not be my aim to make conceptual and physical possibilities match up, but even if that is my aim, there is no guarantee that I will succeed. Instead of relativizing to our best scientific theories, I suggest that a situation's conceptual possibility is better determined by the context in which the discussion arises. Context determines the conceptual framework and level of knowledge expected of the discussants, and this fixes the content of concepts. Everything not ruled out will count as conceptually possible—even when the reality proves otherwise.

We can meaningfully speak about the conceptual possibility of situations we know cannot really happen. For example, we expect fiction to carry with it an internal logic, and we are upset when writers break their own rules. Conceptual possibility turns on ideas, not reality. Concepts can be filled in by empirical knowledge, but in general they need not be. This is the reason physical possibility has to be distinct from conceptual possibility. The question whether it is conceptually possible for an iron bar to float on water requires paring away the physical and epistemic aspects of the situation and focusing on making conceptual relationships clear and distinct, for it is on these relationships that conceptual possibility turns. Someone ignorant of our latest scientific theories (or someone without practical experience with the relevant materials) may grant that it is conceptually possible for an iron bar to float on water because there is nothing inconsistent in the surface concepts. On the other hand, relative to our best scientific theories it is conceptually impossible for iron to float on water. If iron really turns out to float on water, then we will want

---

[5] The objection that iron could exist in some as-yet undiscovered form with specific gravity lower than that of water is a bait-and-switch: certainly the meaning of a word can change, but this isn't the possibility we set out to evaluate.

to revise our scientific theories, but until that happens, it will remain conceptually impossible for iron to float on water relative to our best scientific theories. It's the theory, not the world, that directly matters for evaluating the claim.

Technology has a dual nature: loosely speaking, it is "made of" both material and concepts. It fulfills its functions in virtue of its materiality, and it is intentionally produced to realize certain goals. To fulfill its function, there must be agreement between the material and conceptual aspects of the device. A hammer must have the heft to drive a nail and its wielder must know to swing it. But what sort of agreement is required? Roughly, I take it that "technology" refers to both a material arrangement of parts (and the capacities that arrangement affords) and the knowledge of the function it is supposed to perform. A microwave oven is a boxy metal object which when plugged into a working electrical socket and turned on, I know will heat up the food I have placed inside. A telescope is a tube which when focused properly, I know will present a magnified image of some distant object. Often, we know something about how the technology functions (about electromagnetic radiation and standing waves in the case of a microwave oven and about the laws of optics and lenses in the case of the telescope), but this theoretical knowledge isn't strictly necessary for understanding the technological possibilities it unlocks.

There is much to be said about the relationship between matter and method. For now, let me simply assert that evaluating technological possibility historically means evaluating what practitioners knew how to do with the materials they had. Manhattan Project scientists imagined performing fusion experiments, but they didn't know how to do so with the materials they had. Fusion experiments were not technologically possible. Monte Carlo simulations of fusion reactions, on the other hand, were technologically possible because of the digital computer.

I began this paper by considering conflicting claims about what difference technology makes in the practice of science. My diagnosis is that such conflicts can be understood by analyzing differences in implicit assumptions about possibility. Peter Galison observes that "some kind of numerical modelling was necessary [for completing bomb work], and here nothing could replace the prototype computer just coming into operation in late 1945: the ENIAC" (Galison 1996, 122). Agar thinks Galison is claiming that nothing but the ENIAC could have achieved what the ENIAC did. And if this is what Galison means, then of course he is wrong. Agar is absolutely right that "computerization was usually first proposed when the existing practices and technologies were still capable of the computational task at hand" (Agar 2006, 873). This case is no exception. A human calculator could follow any of the instructions ENIAC performs, though a human would do the job more slowly. Early computer literature is full of direct comparisons. A typical example is that ENIAC could perform a particular calculation in 60msec (or 30 seconds if the result is to be printed out). It would take a human computer 7 hours to solve the same problem (von Neumann, 9). ENIAC did not make the individual calculations physically possible – they always were. Nor did ENIAC make the calculations conceptually possible – scientists had quite specific calculation methods in mind well before ENIAC came along (indeed, a large part of Agar's argument is in substantiating this claim: the Monte Carlo method was known to mathematicians decades before the first computer was constructed). What ENIAC provided was faster calculation – by orders of magnitude.

Technological possibility describes just one constraining affordance for human action. There are others. Some appear to be fundamental: physical and biological laws, for example. Others are changeable: ethics, aesthetics, economics, laws, and technology. One promising avenue for future research would take all of these factors as collective constraints on scientific practice. In this dissertation, however, I touch only peripherally on these issues. What I will say here is that Manhattan Project scientists couldn't perform fission experiments because they were too expensive and dangerous, not because they weren't technologically possible. Fusion experiments, on the other hand, were technologically impossible in 1940.

We should understand the conflict between Galison and Agar as stemming from imprecise possibility statements. Galison is claiming that the computer made the Monte Carlo method technologically possible given the particular time constraints, economic pressures, and allowable risks associated with the Manhattan Project. Human calculators were too expensive, bombs too dangerous, and both of them too time-consuming. There are now two ways to interpret Agar's claim. First, that the Monte Carlo method could have been implemented on some existing mechanism, for example a mechanical computer, even while holding other constraints constant. That is, Monte Carlo was *already* technologically possible even without the digital electronic computer. Second, that scientists might have relaxed some other constraint in order to allow an existing method to succeed, for example by hiring more human computers. That is, *technological possibility* wasn't the only constraint in play in the Monte Carlo case, so focusing on technology obscures important aspects of the story. And indeed, it is worth asking whether any constraints other than technology did change. It turns out that at least one did: the standards of evidence accepted by scientists changed. Agar sees this as the key to the Monte Carlo story: standards of evidence and methodology changed to allow Monte Carlo results in. It so happens that the computer came along at the same time, but this isn't the interesting part of the story for Agar. For Galison, what is of interest is the interplay between the two. In this paper, I have left open the question of how a community selects and enforces its epistemic practices. What is particularly interesting about the case of Monte Carlo simulation is that scientists adjusted their epistemic standards to accommodate a new kind of evidence and a new kind of argument. Simulation became an accepted methodology within their community, but before that was possible, scientists had to renegotiate standards so that they could accommodate simulation results.

Technological possibility depends on the availability of material and conceptual means to bring about a desired state of affairs. This means that physical possibility and conceptual possibility are necessary conditions for technological possibility. But other factors, including ethics and economics affect the availability of needed materials, meaning that physical and conceptual possibility are not jointly sufficient for technological possibility. Technological possibility relates to individuals or groups at particular places and times. This is what makes it useful for analyzing the effect of technological advance on the practice of science. Technology gives us the choice to interrogate some aspect of the world. What we can do affects what we can know.

This brings us back to epistemic possibility. Epistemic possibilities relate to an agent's epistemic position; that is, on what an agent is in a position to know. This in turn depends on what an agent is responsible to know and what investigations are actually practicable for the agent. There is an important aspect of this discussion I have left open: Although a practice must be technologically possible before it can be adopted as a community standard, we (as observers) need not accept those standards as being sufficient to justify an epistemic claim. The point was that the *community* must accept any claim that meets its standard. But is there a standard by which we can evaluate the standards? This question is particularly urgent in the context of practices of knowledge production, but its answer is beyond the scope of the present paper.

I began this paper by considering conflicting claims about what difference technological change makes to the practice of science. My diagnosis was that the conflict is due to differences in implicit assumptions about possibility. In the first part, I argued for the inclusion of practicable responsibilities in the analysis of strong epistemic possibility. In the second part, I introduced technological possibility and discussed its relationship to physical, conceptual, and epistemic possibility, concluding that physical and conceptual possibility are necessary but insufficient for technological possibility, which is itself a necessary but insufficient condition for (strong) epistemic possibility.

Bibliography

Agar, John. (2006), "What Difference Did Computers Make?" *Social Studies of Science* 36 (6): 869-907.

Conee, Earl & Richard Feldman (2004), *Evidentialism: Essays in Epistemology*. Oxford: Clarendon Press.

Foley, Richard (2003), "Justified Belief is Responsible Belief," in M. Steup & Ernest Sosa (eds), *Contemporary Debates in Epistemology*. Available online: philosophy.fas.nyu.edu/docs/IO/1161/justifiedasresponsible.pdf

Galison, Peter. (1996), "Computer Simulations and the Trading Zone," in Peter Galison & David J. Stump (eds), *The Disunity of Science: Boundaries, Contexts, and Power* (Stanford, CA: Stanford University Press): 118–57.

Gendler, T. S., and J. Hawthorne, eds. (2002), *Conceivability and Possibility*. Oxford: Clarendon Press.

Gibbs, B. (1970), "Real Possibility," *American Philosophical Quarterly* 7 (4): 340-48.

Hacking, Ian. (1967), "Possibility," *The Philosophical Review* 76 (2): 143-68.

Hacking, Ian. (1975), "All Kinds of Possibility," *The Philosophical Review* 84 (3): 321-37.

Hartshorne, Charles. (1963), "Real Possibility," *The Journal of Philosophy* 60 (21): 593-605.

Kornblith, Hilary (1983), "Justified Belief and Epistemically Responsible Action," *The Philosophical Review* 92(1): 33-48.

Kroes, P., and A. Meijers, Eds. (2006), "The dual nature of technical artefacts," special issue of *Studies in the History and Philosophy of Science* 37 (1).

Longino, Helen. 1990. *Science as Social Knowledge*. Princeton, NJ: Princeton University Press.

Seddon, George. (1972), "Logical Possibility," *Mind* 81 (324): 481-94.

Teller, Paul. (1972), "Epistemic Possibility," *Philosophia* 2 (4): 303-20.

Von Neumann, John. (1946, unpublished), "On the Principles of Large Scale Computing Machines" in Taub, A. H. *John von Neumann: Collected Works Volume V Design of Computers, Theory of Automata and Numerical Analysis*. Pergamon Press: Oxford, 1961. 1-32.

Williams, Michael. (2001), *Problems of Knowledge: a Critical Introduction to Epistemology*. Oxford: Oxford University Press.