

# Technology and Epistemic Possibility

Isaac Record

Published online: 19 November 2013  
© Springer Science+Business Media Dordrecht 2013

**Abstract** My aim in this paper is to give a philosophical analysis of the relationship between contingently available technology and the knowledge that it makes possible. My concern is with what specific subjects can know in practice, given their particular conditions, especially available technology, rather than what can be known “in principle” by a hypothetical entity like Laplace’s Demon. The argument has two parts. In the first, I’ll construct a novel account of epistemic possibility that incorporates two pragmatic conditions: responsibility and practicability. For example, whether subjects can gain knowledge depends in some circumstances on whether they have the capability of gathering relevant evidence. In turn, the possibility of undertaking such investigative activities depends in part on factors like ethical constraints, economical realities, and available technology. In the second part of the paper, I’ll introduce “technological possibility” to analyze the set of actions made possible by available technology. To help motivate the problem and later test my proposal, I’ll focus on a specific historical case, one of the earliest uses of digital electronic computers in a scientific investigation. I conclude that the epistemic possibility of gaining access to scientific knowledge about certain subjects depends (in some cases) on the technological possibility for making responsible investigations.

**Keywords** Epistemic possibility · Epistemic modality · Technological possibility · Responsibility · Practicability

Let me begin with a historiographical debate about how to understand the earliest uses of digital electronic computers in scientific investigations. Peter Galison argues for the necessity of the digital electronic computer to certain advances in scientific knowledge

---

I. Record (✉)

Faculty of Information, University of Toronto, Toronto, ON M5S 3G6, Canada  
e-mail: isaac.record@utoronto.ca

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 made it possible for scientific knowledge to develop. John Agar argues against the necessity of that technological change for the advancement of 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 knowledge. Galison and Agar agree that certain scientific knowledge became accessible to Manhattan Project scientists once they began to run “Monte Carlo” calculations, named after the famous gambling establishment because of the method’s use of random numbers. They disagree about which changes in the situation of the scientists made those investigations possible. Galison thinks access to digital electronic computers made the difference. Agar thinks that the Monte Carlo method could have been implemented using existing computational approaches.

Both analyses turn on what was possible under the circumstances rather than on what actually transpired. The conflict between Galison’s and Agar’s claims about the digital electronic computer arises because they invoke different implicit notions of what was possible. The commonsense notion of possibility is just *what might happen*, *what might exist*, or *what might be true*. But in practice, we freely constrain these generic notions of possibility to reflect narrower concerns. For example, suppose I claim that driving the wrong way down a one-way street is impossible. The truth of that claim turns on the kind of possibility we invoke in evaluating it. Driving the wrong way down a one-way street is physically possible, because the laws of physics do not forbid it (making my claim false on this analysis). Alternatively, because a municipal law does forbid it, it is regulatively impossible (making my claim true). Similarly, my traveling to the nearby star Alpha Centauri by 2020 is (let’s say) theoretically and physically possible, but simultaneously technologically and economically impossible. To return to the case at hand, Agar is saying, roughly, that digital electronic 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 played in the practice of science at a particular time. According to Galison, the digital electronic computer 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 consequently what was epistemically possible for the scientists. Contrariwise, Agar thinks this way of putting things gives too much credit to the material means of accomplishing a task. Instead, a conceptual means (the Monte Carlo method) made those advancements possible—and could have done so without being implemented on a digital electronic computer.

In order to understand and evaluate Galison’s and Agar’s claims about the difference technology makes to knowledge, we need an account that explicitly recognizes that technological changes make it possible for individuals to undertake different actions, and some of these actions make it possible for those individuals to gain different knowledge. The rest of this paper is devoted to this task, and the paper is divided into two parts. In the first part, concerning the relationship between possible actions and possible knowledge, I’ll argue that we need an account of epistemic possibility that captures the dependency of knowledge on being able to take the appropriate action. It is, I think, uncontroversial to say that being able to complete certain actions can be a necessary condition for gaining knowledge. In scientific practice, for example, gaining knowledge depends on having relevant evidence, which makes being able to gather the evidence a condition for gaining

the knowledge. My contribution is to argue that because a scientist is (under certain conditions) *expected* to seek evidence before making a knowledge claim within her domain of expertise, we need to build this expectation into our account of knowledge—and because expectations are not always fulfilled, the appropriate philosophical concept is epistemic possibility.

In the second part of the paper, I turn to the relationship between technology and possible actions. A number of practical factors affect our ability to act, including economics and ethics, but I will focus on technology, touching on the others only incidentally. I'll introduce an analysis of technological possibility, which depends on the availability of material and conceptual means to bring about a desired state of affairs, and argue that the epistemic possibility of gaining access to scientific knowledge depends (in some cases) on the technological possibility for carrying out certain investigations. In such cases, technological possibility can be seen as a necessary condition for epistemic possibility. Finally, I will return to the disagreement between Galison and Agar and show how my analysis of epistemic and technological possibility resolves the conflict.

## 1 Doing and Knowing

My overarching aim in this paper is to give an account of the relationship between contingently available technology and the knowledge that it puts within “epistemic reach,” to use Egan’s vivid phrase (2007, 8). The relevant philosophical concept here is epistemic possibility, which is meant to reflect epistemic reach by distinguishing between what a subject can and cannot know given her epistemic circumstances. The Agar-Galison example illustrates that some knowledge claims require the gathering of evidence, which suggests an understanding of epistemic reach that is responsive to the actions a subject can actually accomplish. Canonical accounts of epistemic possibility tend to be insensitive to this issue, as I will show. I will develop a novel account of epistemic possibility that takes into consideration practical conditions for knowing, taking particular care to develop my account in such a way that it can accommodate epistemic responsibilities 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 includes both a practicability criterion and a responsibility criterion. My approach will be to begin with a canonical definition of epistemic possibility<sup>1</sup> and then elaborate it.

The usual starting point for epistemic possibility is that it should somehow reflect a subject’s epistemic position. Thus, for a subject *S* to claim that the proposition  $\Phi$  is epistemically possible is for *S* to say that  $\Phi$  is possible relative to *S*’s epistemic position. Taking “*S*’s epistemic position” to be, in the simplest case, “what *S* knows,” leads straightforwardly to the canonical definition of epistemic possibility (Hacking 1967; Teller

---

<sup>1</sup> The main thread of philosophical accounts of epistemic possibility traces at least to Moore (1962), and the topic has enjoyed a recent flourishing (see, e.g., the collected volume by Egan and Weatherson 2011). According to the canonical account, when a subject says that something is possible, we should understand this claim as being made relative to a person, a group, or a set of information. But precisely which person, group, or set of information? Egan and Weatherson (2011) identify three clusters of answers: contextualists think the context of utterance somehow specifies an answer (see DeRose 1991; von Fintel and Gillies 2011), relativists look to the context of assessment (see Egan et al. 2005; MacFarlane 2011), and expressivists think modal statements express the speaker’s mental state of uncertainty (see Yalcin 2011). The discussion in this paper centers on the canonical contextualist accounts, though I will make note of some divergent views.

1972; DeRose 1991 and the individual contributors to Gendler and Hawthorne 2002 and Egan and Weatherson 2011 all take this as their point of departure):

(a)  $\Phi$  is epistemically possible for S if S doesn't know  $-\Phi$ .

The idea is that if S doesn't know for certain that  $\Phi$  is not the case, then S must consider  $\Phi$  to be possible. For example, if S knows  $\Phi$ , then S cannot know  $-\Phi$ , and  $\Phi$  is epistemically possible. On the other hand, if S knows  $-\Phi$ , then  $\Phi$  is by definition epistemically impossible. Finally, if S knows neither  $\Phi$  nor  $-\Phi$ , then  $\Phi$  is epistemically possible for S (as is  $-\Phi$ ). More concretely, if I have just checked my key hook for my lost keys (and failed to locate them there), then for me, it is not epistemically possible for my keys to be on the key hook. But if I have not yet looked on the table, then it is epistemically possible for my keys to be there, assuming I have no other reasons for excluding that possibility.

Note that, on most accounts, epistemic possibility is sensitive to what S *could* know given S's epistemic position rather than reflecting what S actually believes.<sup>2</sup> Suppose S deems  $\Phi$  possible, forgetting to take into account that  $-\Phi$ . In such a case, S is wrong; S should know better than to think  $\Phi$  is possible. That is,  $\Phi$  is in fact *not* epistemically possible for S, even if S thinks that  $\Phi$  is possible. In such cases, epistemic possibility provides grounds to blame S for a misjudgment.

As it turns out, taking S's epistemic position to mean "what S knows" leads almost immediately to results that confound the intuitions of some philosophers. (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 (at least for any view of knowledge that includes "belief" as a necessary condition). Suppose  $\Phi$  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, according to (a), "4+3=9" is epistemically possible for S, because S has no beliefs about it whatever. On this view, if S blurts that "Perhaps '4+3=9'" 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 (for further discussion of cases like this, see Huemer (2007) and Yalcin (2011)).

If the goal of epistemic possibility is to reflect S's epistemic position, it seems like a strange consequence that we can blame S for failing to recall  $\Phi$ , but not for failing to consider it. Several accounts of epistemic possibility attempt to close the gap between failing to recall and failing to consider by expanding "epistemic position" to include everything "within epistemic reach," yet as we will see, it has been difficult to specify just what counts as being within epistemic reach.

If (a) fails to capture what is within epistemic reach, perhaps we can simply add such a description to the original definition. For example, we might say that:

(b)  $\Phi$  is epistemically possible for S if S does not know that  $-\Phi$ , *nor would careful reflection establish that  $-\Phi$ .*

Here, S isn't allowed to simply blurt out that "perhaps '4+3=9.'" She must first carefully reflect upon the proposition. This eliminates the problem of unconsidered cases like "4+3=9," while still being limited to the knowledge S has (plus inferences from that

<sup>2</sup> Expressivists would presumably disagree. For expressivists, epistemic possibilities express a subject's uncertainty rather than describing a relationship between the subject's knowledge (or information) and the state of the world (see, e.g. Yalcin 2011). I must set this view aside in this paper.

knowledge). Unfortunately, “careful reflection” is too vague a requirement to capture what is practicably within a subject’s epistemic reach, which means that definition (b) fails to appropriately reflect what S is in a position to 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 its truth-value, and many hours of careful reflection have not resolved the situation. Nevertheless, some amount of additional reflection might solve it, as has transpired for many other mathematical conjectures. The point is that “careful reflection” doesn’t distinguish between 5 minutes, 5 hours, or 5 years of reflection. [Stanley’s (2005) suggestion that S take into account “obvious entailments” of what S knows seems to me to do a little better than (b), but it fails to respond to Hacking’s criticism of (c), below]. But is (b) merely too vague in describing epistemic reach, or are we on the wrong track altogether?

I propose to expand the notion of epistemic reach to include *practicable responsibilities*. Careful reflection remains a plausible starting point, though to be complete we would need to say how much reflection a subject is responsible to perform, and how much reflection is practicable. After considering some other proposals, I will argue that the limits of practicable responsibility are determined by context. But before continuing, let me make the case for responsibility, since this suggestion will strike some as tendentious.

Richard Foley suggests that “our everyday evaluations tend to be concerned with whether one has been responsible in arriving at one’s beliefs” (2003, 9). Let me give two examples of responsibly arriving at one’s beliefs. First, in the context of *scientific* knowledge claims, we routinely expect these claims to carry special weight in light of experimental evidence or theoretical justification. Accordingly, we impose special responsibilities, sometimes called *epistemic duties*, on scientists [see, e.g., Kornblith (1983)]. Lest we conclude that role responsibilities are a special case, I point out that epistemic duties appear in everyday cases too. When I call my office to ask a colleague whether a letter I am expecting has arrived, I won’t be satisfied with the claim that it is “possible” that the letter has arrived—I want to know one way or the other! I want my colleague to check the incoming mail. But reasonable expectations have limits: I won’t blame my colleague for not noticing that the letter has slipped behind a desk or was delivered to the wrong recipient. The point is that epistemic possibility should not only reflect what knowledge S already has, but should also take into consideration S’s responsibilities to gather additional evidence. On the epistemic responsibilities 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. These facts are a part of a subject’s epistemic circumstances, and so should be reflected in our analysis of epistemic possibility.

While the idea behind epistemic responsibilities—that sometimes we must back up our claims—is fairly straightforward, the existence and nature of epistemic responsibilities are controversial in epistemology. I should mention that an alternative to the epistemic responsibility 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 responsibility” to indicate any duty that is a condition for making a knowledge claim (indeed, I shall use the term whether or not the duty promotes genuine knowledge). I will also set aside the larger question of whether we always have epistemic responsibilities, and instead distinguish between *weak epistemic possibility*, which does not include

responsibilities, and *strong epistemic possibility*, which does include responsibilities. For the remainder of the paper, when I refer to “epistemic possibility” I mean strong epistemic possibility.

Because responsibility is a novel contribution to the discussion of epistemic possibility, let me briefly describe some features of what I take to be a plausible account of epistemic responsibility. I won’t defend such a view here; rather, I merely want to show that some account might be made to work with my version of epistemic possibility. First, context-relevant risks may be distinguished from background-level risks. Second, epistemic responsibilities need respond only to context-relevant risks. And third, background risks may be converted into context-relevant risks (and vice-versa) through negotiation.

There is a distinction to be maintained between the question of when  $\Phi$  is justified and when S has warrant to claim  $\Phi$  (see, e.g., Williams 2001). Epistemic responsibilities, as I use the term here, have to do with claiming. S’s making a claim about  $\Phi$  invokes a responsibility, but fulfilling this responsibility does not guarantee that a claim is justified. What determines epistemic responsibility is not the actual epistemic risk of a particular claim, but its perceived (context-relevant) risk relative to a particular set of background commitments that S need not defend. The focus on context-relevant risks allows us to bracket the epistemic risks associated with the background commitments in order to stay focused on foreground issues. To give an extreme example, in deciding whether to accept a stranger’s testimony about the local bus schedule, we tend to disregard the possibility that the external world is an illusion. The idea is not to ignore those background-level risks, but merely to focus on the risks associated with a particular claim within the relevant context. Focusing on the contextual claim has the effect of “normalizing” or rescaling its risks against a chosen 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 (as I prefer) we can recognize that what counts as background is negotiable. As Helen Longino puts it (with respect to propositional scientific knowledge), “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 (and dependent foreground risks can be recalibrated 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 the risk of being wrong about 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, S’s fulfilling the epistemic duties for being able to claim  $\Phi$  would be necessary and sufficient to justify  $\Phi$ . But suppose that there is a mismatch, and fulfilling responsibilities is insufficient to justify  $\Phi$ . Nevertheless, fulfilling responsibilities is still necessary for having knowledge of  $\Phi$ , because any claim that fails to fulfill responsibilities is a non-starter in the context in which it is made. That is, fulfilling

responsibilities is a necessary but insufficient condition for  $\Phi$  being strongly epistemically possible for S. It would be nice if we knew which responsibilities are relevant to justification, but we simply cannot be certain. Inquiry in fields like science works on the basis of their 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 responsibilities.

Whether or not the preceding sketch of how epistemic responsibilities are generated is exactly right in its details, I think it is plausible that subjects often have a responsibility to go beyond their present beliefs, such responsibilities depend on a subject's knowledge context, and these responsibilities are relevant to evaluating what is epistemically possible for them. I will refer to this as the "responsibility criterion" for epistemic possibility. The question is how to incorporate responsibility into the definition of epistemic possibility.

One plausible solution is to include S's epistemic community in the definition, since this group negotiates the boundaries of context-relevant responsibilities. It turns out that, for other reasons, the inclusion of community is a common proposal among contextualists as well, so there is a robust literature to work from. For those accounts, the usual idea is that if our concern is that (a) and (b) don't adequately reflect what is in epistemic reach of the subject S, we must make epistemic possibility sensitive to the knowledge or information available to the entire group to which S makes her claim. To put it another way, the knowledge of everyone in the group is within the epistemic grasp of any member: she need merely ask.

(c)  $\Phi$  is epistemically possible for S if S does not know that  $\neg\Phi$ , *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  $\Phi$ , 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.<sup>3</sup> Variations on (c) abound. Indeed, von Fintel and Gillies (2011, 108) identify as "canon" the view that "epistemic modals quantify over the information available to a contextually relevant group. The context decides the group (and perhaps the standards by which they know)". In a scientific community, this definition works rather well in principle, because knowledge is (ideally) made available to the entire community by mechanisms such as conferences and 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

<sup>3</sup> Exactly how to define the relevant community is a matter of earnest debate. The basic premise of relativist accounts of epistemic possibility is that contextualist accounts cannot readily handle cases of agreement, disagreement, eavesdropping, or temporally dislocated responses to purported possibility claims. See Egan (2007) and MacFarlane (2011). See also von Fintel and Gillies (2011) for a review of some responses.

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 for him? 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 (1967, 148) 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".

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  $\Phi$ . For Hacking,

(d)  $\Phi$  is epistemically possible for S if S doesn't know  $-\Phi$  nor would any practicable investigations by S establish that  $-\Phi$ .

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 adjudicate the Goldbach case satisfactorily: S is now only 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 poses 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 (1972, 307) (Incidentally, according to Teller's account, the sex test was newly available in 1972. A few years earlier, it would not have been practicable). 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. Put another way, practicability may be a necessary condition for a duty to be imposed on S, but it is not a sufficient one.

Recall that 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 should have known after he examined the log and what he actually knew. 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  $\Phi$  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 [definition (c)]:



(e)  $\Phi$  is epistemically possible for S if it is not the case that:

(1)  $\Phi$  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  $\Phi$  is false. (Teller 1972, 310–311)

The idea is to restrict epistemic possibility to what some member, T, of the community would be in a position to know if T had all of the relevant communal knowledge at hand. For my purposes, Teller's formulation has a significant problem: it doesn't accommodate responsibilities that would have a subject look beyond existing knowledge. Like the original "community" variation, it addresses only what is already known by the community (von Fintel and Gillies arrived at a similar point quite independently; see their 2011, 112–113 fn. 9).

Consider a slight variation on Hacking's salvage ship problem. Suppose the mate's mistaken calculation is the result of his having skipped a line in the log. This means that neither the mate, nor any other member of the salvage crew knows the relevant propositions about the location of the treasure. Yet we would still blame the mate for this mistake. Teller acknowledges this gap in his account, and fills it in by counting as "known to community C" facts written down in books available to the community (1972, 312). But responsibility to access 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.

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). My diagnosis is that both Teller and Hacking have part of the story right.<sup>4</sup> The difficulty in defining 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 removing context from the analysis. 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. Within a given community, C, if an individual, S, makes a knowledge claim and meets C's epistemic standards, E, then C will accept it. That is,

(f)  $\Phi$  is epistemically possible for S if S does not know that  $\neg\Phi$ , nor do the epistemic standards E of community C demand that S carry out any practicable investigation that would establish that  $\neg\Phi$ .

Let's see how my proposal 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

<sup>4</sup> Others have attempted to amalgamate parts of Hacking's and Teller's accounts. DeRose's (1991) "relevant way" and Egan's (2007) "epistemic reach" are two of the better examples. Both remain quite vague about how to find out what is a relevant way within epistemic reach.

shipmates) demands that he glean this information from the log. In each case, the relevant community (the community to which 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). In sum, epistemic possibility lies at the intersection of epistemic responsibilities and practicable actions. There may be responsibilities that are not practicable and practicable actions that are not responsibilities.

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 within their epistemic community, Manhattan Project scientists had to fulfill certain epistemic responsibilities; for example, they had to meet precise standards of evidence and justification in order for their work to move forward. Before the advent of the digital electronic computer running Monte Carlo calculations, scientists did not fulfill those responsibilities, and so were stuck—they could not move forward on their bomb work, because they needed knowledge that was unavailable to them. That is, they had epistemic responsibilities they could not discharge without specific knowledge about bombs, 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 before the advent of the digital electronic computer. That is, they disagree about whether fulfilling those responsibilities was practicable given the specific situation Manhattan Project scientists were in.

If Monte Carlo was practicable before the advent of digital electronic computers, then the knowledge the scientists sought was within their epistemic grasp—that is, it was epistemically possible. But if Monte Carlo was impracticable before digital computers, then the knowledge they sought was not within their grasp. What makes an activity practicable? According to Galison, a technology, the digital electronic computer, made practicable for Manhattan Project scientists the gathering of the required evidence to fulfill their epistemic responsibilities. By contrast, Agar thinks that the Monte Carlo method could have been implemented using older equipment; that is, that Monte Carlo calculations were practicable before they were actually put into practice. The limiting factor was a lack of theoretical guidance without which nuclear experiments were too dangerous and expensive to be performed. The scientists' theoretical efforts were stymied by intractable analytic equations. Progress slowed. Then available technology changed and progress resumed. But was the technological change decisive or merely coincidental? That is, did the digital electronic computer offer new technological possibilities that made practicable a method that, prior to the digital electronic computer, had been impracticable?

Let me be clear: the relevant constraint on epistemic possibility in a case like this is whether (and when) the scientists could meet their epistemic responsibilities with practicable investigations. If they could do so both before and after the advent of digital electronic computers, then Agar is right, and computers should not be credited with making the investigations possible. But if the availability of the digital electronic computer is what made particular fission investigations practicable, then Galison is right, and computers can be credited with making the bomb work possible. Either way, it was actually carrying out these investigations that made crucial knowledge epistemically possible. The question is whether a change in technology played the deciding role.

I turn to technological possibility and the relation between technology and action in Part 2.

## 2 Technological Possibility

In the first part, I argued for an account of epistemic possibility that can accommodate the epistemic responsibilities that can require a subject to take action. At the same time, my account is sensitive to practical limits to fulfilling those responsibilities. A given subject cannot undertake just any investigation. She will be competent to perform only some investigations, and her technological, economic, and ethical circumstances will allow still fewer. Impracticable responsibilities put limits on what knowledge is within a subject's epistemic grasp. Changes to situational constraints on practicability can change what is epistemically possible for a subject. In the present part, I focus on *technological possibilities* as a hard constraint on practicability and therefore epistemic possibility. Technological possibility depends on availability of the material and conceptual resources required to complete some action or produce a desired state of affairs. Thus:

(g) A course of action is *technologically possible* for a subject S if S has access to both the material and conceptual means to accomplish it.<sup>5</sup>

The possibility of *my* spanning a river with an iron bridge turns on both what the world is like (i.e., that iron is available to me and has certain properties, and that I have certain capabilities with respect to iron) and how my concepts fit together (i.e., that I think iron has certain properties that I can put to use in making trusses). Without the material means, the bridge would fail. Without the conceptual means, I would never attempt it. Given this definition, the connection between technological possibility and practicability is clear: a subject's technological tools are a determinant of what is practicable. An action is only practicable if it is technologically possible.

There are two ways to rule something technologically impossible for a given subject: either the subject doesn't have access to the material means of accomplishing it, or the subject doesn't have access to the conceptual means of accomplishing it. The burden in assessing whether a course of action is technologically possible is in making a sensible determination as to which conceptual and material means should be considered 'accessible' to a subject given the peculiarities of her situation. A course of action that would exhaust a subject's material resources, tax her creative faculties, and take a long time to construct would be *difficult*, but nevertheless accessible. Note that in the case of complex investigations like scientific experiments, the most challenging aspect of completing an inquiry is often in determining whether the equipment has functioned properly, whether the desired intervention actually occurred, or whether a particular inference is actually warranted by the data. All of these tasks should be included in the calculus of the technological possibility of the inquiry as a whole. In the following, I will consider how to draw the line between accessible and inaccessible material and conceptual means.

Let me begin with material means. I said above that the possibility of my spanning a river with an iron bridge depends on my having access to iron, iron having certain

<sup>5</sup> The phrase "possible for" makes my notion of technological possibility subject-relative (see Gibbs 1970 for an analysis of "possible for"). We sometimes speak of technological possibilities in a more abstract sense. For example, there is a sense in which it is technologically possible (for some unspecified someone) to travel to the moon. I have no principled objection to this usage, but in the present paper I am concerned with cases in which S is specified—and locked to a particular time, space, and circumstance.

properties, and my having certain capabilities. This suggests a way to divide material means into three further considerations. The first, access to the material itself, is simply a logistical consideration that depends on a subject's situation—roughly, what a subject has, or can beg, borrow, or steal. The other two, the properties of that material and the subject's capabilities, ultimately rest on *physical possibility*.

(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 subject relative. Physical possibility is about what is possible *given the actual laws of nature*, not our account of those laws. This means, for example, that fusion experiments have been physically possible for billions of years (light from stars billions of light years away substantiates this). 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 like the construction of the Tokomak fusion reactors in the 1950 s and 60 s for scientists to perform fusion experiments. Physical possibility is a hard constraint on technological possibility because technologies cannot subvert the laws of nature, and neither can users of technology. That is, my ability to build an iron bridge depends on the physical possibility of iron taking the form of a bridge and on the physical possibility of my body (and other available technology) giving it that form. Physical possibility is a necessary condition for technological possibility, but physical possibility is not sufficient for technological possibility. There are many actions that are physically possible and yet beyond the means of a particular subject to bring about, even if the subject has access to the requisite material.

Let me be clear about the relationship between physical possibility and a subject's capabilities. I have just claimed that a subject's capabilities are 'ultimately limited' by physical possibility, where physical possibility is to be understood in terms of the laws of nature. This means that physical possibility is sensitive to physiological limits in addition to mechanical ones. To give a concrete example, suppose that my physiology is such that I cannot leap tall buildings in a single bound (unaided), but I can easily jump a hurdle. But what about edge cases? Suppose that I cannot presently dunk a basketball, but that with enough conditioning, I could. Then, strictly speaking, dunking is an accessible course of action for me. However, if I am *presently* considering whether to dunk this basketball or just lay it in, then I had better lay it in, because my time-sensitive context does not afford the months of training and conditioning necessary for me to dunk *today*. Evaluating technological possibility requires specifying a context more or less precisely, and this situated analysis is often helpful in identifying the relevant limiting factors, which may include time, money, material, equipment, or skills.

The second basic component of technological possibility is a subject's access to conceptual means. Here, the relevant hard limit is epistemic possibility.<sup>6</sup> Anything that is epistemically impossible for subject S is also, necessarily, technologically impossible for her. In other words, considerations of epistemic possibility allow us to deem technologically impossible those courses of action that a subject knows she cannot accomplish. In cases in which a subject's reasons for ruling out a course of action have to do with physical considerations, as in the basketball example above, this seems redundant. But in other cases, a subject may know that a course of action costs too much or is unethical or is simply beyond her cognitive means. In such cases, a subject can rule out the course of

<sup>6</sup> I owe a considerable debt to an anonymous reviewer for urging me to reconsider an earlier, wrongheaded approach to analyzing conceptual means.

action on the basis of its epistemic impossibility. We can also rule out those courses a subject has a responsibility to determine that she cannot accomplish. These courses of action are precisely the ones that are conceptually unavailable.

There are two ways for a subject to be mistaken about availability. Either she thinks a course of action is *not* available when it *is* or she thinks it *is* available when it is *not*. For example, there may well be courses of action that, because of mistaken beliefs, seem conceptually unavailable. While a subject would probably not attempt a course of action that she thinks is unavailable, it is nonetheless the case that such a task is technologically possible for her if what ruled it out was a mistaken belief. Indeed, only those tasks that she *knows* to be unavailable (or is required to find out are unavailable) are genuine technological impossibilities. The converse mistake is for a subject to fail to fulfill a responsibility to find out that a course of action is impossible. In such a case, a subject may believe a course of action is available to her, when in fact she should have (for example) made some investigation that would have eliminated that candidate action. For example, before embarking on a laboratory experiment, a scientist might be expected to perform a back-of-the-envelope calculation to ensure that the experiment could have the expected result. If she fails to do the calculation (or gets the wrong answer), she might attempt the experiment. But this does not change the fact that it is technologically impossible for her, and it does not change the fact that she should have known it was technologically impossible (because it was epistemically impossible).

Epistemic possibility is a necessary condition for technological possibility, but not a sufficient condition for it. That is, just because an action is epistemically possible doesn't make it technologically possible. Physically impossible actions that I do not know are impossible are epistemically possible, but not technologically possible. In other words, physical and epistemic possibility are both necessary conditions for technological possibility.

Let me now address four potential concerns about the role epistemic possibility plays in determining technological possibility. First, I concluded the earlier part of the paper with the claim that technological possibility is an enabling condition for epistemic possibility. Now I have proposed that epistemic possibility is an enabling condition for technological possibility. This raises the spectre of circularity.<sup>7</sup> Second, "available conceptual means" depends in large part on the conceptual analysis a subject performs. It might not seem like epistemic possibility is the concept to capture this. Third, we might be concerned about the role community standards play in specifying responsibilities. For example, if a subject moves between epistemic communities with different expectations, a course of action may switch, seemingly willy-nilly, between being technologically possible and impossible. Finally, it might be objected that epistemic possibility doesn't correctly capture what it means for the conceptual means for a course of action to be "available."

The first potential concern results from a conflation of the epistemic possibilities *regarding* the completion of a task and those *based on* the completion of a task. For example, "it is possible to reveal surface details of the moon through telescopic investigation" is a quite different epistemic possibility than "it is possible that the moon's surface is smooth." The line between epistemically possible and impossible claims about the geology of the moon has the potential to shift with the advent of any number of technological advances, including the telescope and space travel. The claim that "the moon's surface is smooth" was epistemically possible (and indeed widely believed) before Galileo's telescopic investigations showed surface features. Galileo's instrument provided new

<sup>7</sup> My thanks to an anonymous reviewer for bringing this concern to my attention.

evidence that could make a difference in determining what propositions about the moon's surface were epistemically possible for various subjects, and in addition the use of the telescope was available as a candidate responsibility for some subjects who wanted to make claims about the surface of the moon. (Exactly who was responsible for what depended in part upon the expectations of the relevant community.) In short, however, the epistemic possibility that "telescopes can reveal distant features" enables the technological possibility of viewing lunar surface details through a telescope, which in turn makes the claim that "the surface of the moon is smooth" epistemically impossible. To put the point more generally, an epistemic possibility regarding the completion of a task enables a subject to pursue it, while the line between epistemic possibilities and impossibilities can shift based on the completion of that task.

The second worry inquires into the relationship between epistemic possibility and conceptual analysis. For example, in the telescope example just described, I noted that the epistemic possibility of using a telescope to investigate the surface features of the moon has to do with "supposed properties" of glass and brass. Conceptual analysis means determining the compossibility of a proposed state of affairs with a particular background context. This may seem orthogonal to epistemic possibility. But as Hacking says, we bring logic into the fold when the "terms of individuation" produce a contradiction (1975, 333), and this allows us to rule out contradictory situations on the grounds that they are epistemically impossible. Let me be clear about how this works. 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, I might at first neglect to take into account some relevant details, such as what I know of the laws of physics, and on the basis of that incomplete picture judge the deed possible. But if I carried on filling in details, says Hartshorne, I would wind up in perfect agreement with physical possibility—in the end, the two are indistinguishable (see Hartshorne 1963, 595). This contention is mistaken for two reasons. First, it makes an unwarranted demand on epistemic responsibilities, and second, it conflates physical possibility with scientific theories.

To better illustrate Hartshorne's contention, we can draw on George Seddon's example of how the relevant analysis should work. An iron bar that floats on water has been supposed by some philosophers to be conceivable,<sup>8</sup> 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 would allow it to be ruled epistemically impossible on conceptual grounds. But for someone ignorant of the latest scientific theories and without practical experience with the relevant materials, there is no such additional information to fill in the concepts. It may well be conceivable to her for an iron bar to float on water because there is nothing inconsistent in the concepts she has. Assuming she has no epistemic duties requiring her to investigate further, floating iron bars are epistemically possible for such a subject. On the other hand, for anyone with relevant common experience or a passing acquaintance with our best scientific theories, it is conceptually impossible for iron to float on water. Furthermore, iron floating on water would be epistemically impossible for anyone with a countervailing epistemic duty.

The third potential anxiety about the role of epistemic possibility in technological possibility is that differing community standards would seem to make courses of action

<sup>8</sup> Seddon actually discusses logical possibility, so I have made some minor adjustments.

switch haphazardly between being epistemically possible and impossible, and therefore between being technologically possible and impossible. According to my account, different community expectations can result in the same action being technologically possible in by one community's standard and impossible by another's. But the difference is not haphazard. One community may require that a subject take action that would rule a course of action epistemically impossible (and therefore technologically impossible as well), while another community has no such requirement. The concern is that the specified course of action would actually be impossible for the subject to carry out in either case, and my definition of technological possibility doesn't correctly reflect this because it gives different answers for the two communities. But remember, the argument isn't that a subject can actually accomplish every technologically possible course of action. It's that a subject cannot accomplish any course of action that is technologically impossible.

The fourth concern is related to the third, but is more general. It might be objected that epistemic possibility doesn't correctly capture what it means for the conceptual means for a course of action to be "available." At the outset of the discussion of access to conceptual means, I stated that one of the ways to be mistaken about availability is to mistakenly consider a course of action available when it is not. Similarly, there may be cases in which a subject does not know that a course of action is unavailable, nor will she have a responsibility to rule it out. In such a case, the course of action will be epistemically possible for S (even if she doesn't explicitly think the course of action is available). If such a course of action is also physically possible, then it will be technologically possible, even though it could never actually be accomplished. This seems to suggest that epistemic possibility is the wrong measure to determine whether a conceptual means is available. If so, it would appear that we are left with three alternatives (besides starting over). First, we could deny that cases like the one I just constructed actually exist. But I have no sturdy basis for making such an argument.<sup>9</sup> Second, we could try to shore up technological possibility by adding some additional condition, but I have no suggestions as to what that condition should look like. Third (and this is the option I prefer), we can accept that physical possibility and epistemic possibility are not quite jointly sufficient for technological possibility after all. Even so, technological possibility is a useful and principled means of drawing a hard line between practicable and impracticable actions, because it is still a necessary condition on practicability. Put another way, my account admits as technologically possible some courses of action that we might wish it deemed impossible. But it deems no course of action impossible that we should wish it to deem possible.

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. That is, we should understand the conflict between Galison and Agar as stemming from imprecise expressions of what was possible. Peter Galison observes that "some kind of numerical modelling was necessary [for completing fission bomb work], and here nothing could replace the prototype computer just coming into operation in late 1945: the ENIAC" (Galison 1996, 122), while John Agar argues that "computerization was usually first proposed when the existing practices and technologies were still capable of the computational task at hand" (Agar 2006, 873).

---

<sup>9</sup> It might be objected that what makes a conceptual means "too complex" to be conceivable necessarily has to do with brain structures, and therefore could be ruled out on grounds of physical impossibility. But this is merely a plausible retort, not a definitive one.

Let's put these claims into the language of epistemic and technological possibility. According to Galison and Agar, Manhattan Project scientists considered three approaches to the problem: they could perform fission experiments to learn crucial facts, they could solve difficult analytic equations, or they could perform a brute-force numerical attack on the bomb equations. Any of these three approaches could satisfy their epistemic responsibilities.<sup>10</sup> The question was whether they were practicable, and this was a matter of some debate. Fission experiments were considered impracticable given the particular time constraints, economic pressures, and allowable risks associated with the Manhattan Project. Dozens of the world's best mathematical minds were basically stumped by the intractable analytic equations, so that approach also appeared impracticable (whether it was or not). And finally, it was not at all clear that there was time enough to run a numerical brute force attack on the bomb equations. What is now clear is that, before the digital electronic computer, numerical analysis had not been successful, but with the computer, such methods were successful. The question is whether it was the computer that made the difference.

Could Monte Carlo calculations have been performed using existing computational methods? A human computer could follow any of the instructions ENIAC performs, and the Manhattan Project employed many such computers. But a human would do the job much more slowly, so there is some question as to whether the calculations could be completed within the required time constraints. Early computer literature is full of direct comparisons between human and digital electronic computers. A typical example is that ENIAC could perform a particular calculation in 60 milliseconds (or 30 seconds if the result was to be printed out). It would take an individual human computer 7 hours to solve the same problem (von Neumann 1961, 9). ENIAC did not make the individual calculations physically possible—they always were physically possible. Nor did ENIAC make the calculations epistemically possible—scientists had quite specific calculation methods in mind well before ENIAC came along (indeed, the bulk of Agar's account goes to substantiate this claim: the Monte Carlo method was known to mathematicians decades before the first computer was constructed). What ENIAC provided was faster calculation—by many orders of magnitude as compared to individual human computers, and by smaller orders as compared to other existing computational methods.

Before ENIAC, "a single hydrodynamics problem in an implosion simulation required passing a deck of punched cards through a dozen machines," a process requiring a full month, even after Richard Feynman and his computing group cut the process by two thirds by devising a parallel computing method, and even with the machines running 24 hours a day (Seidel 1998, 34). By contrast, von Neumann estimated that "one criticality problem requires following 100 primary neutrons through 100 collisions (of the primary neutron or its descendants) per primary neutron," which, computed using the Monte Carlo method on ENIAC, "should take about 5 hours" (Richtmyer and von Neumann 1947, 752). These aren't identical calculations, but the examples give a sense of the dramatic change in speed offered by the Monte Carlo method on the ENIAC as compared with prior computational approaches.

---

<sup>10</sup> The third approach, numerical attack, particularly when it took the form of Monte Carlo calculation, required some advocacy before it was considered acceptable throughout the scientific community, but it was almost immediately accepted within the Manhattan Project itself, perhaps in part because of enduring frustrations with the other two approaches—an interesting story in its own right, but beyond the scope of the present paper. Galison (1996) is a good starting point for that story.



It is on this significant difference in speed that the debate finally turns. For Galison to be right, we must accept that ENIAC's faster computations moved this particular application of the Monte Carlo method from the "too slow to consider" category into the "can't rule it out" column. For Agar to be right, we must accept that scientists could have implemented this application of Monte Carlo using older methods of calculation. This remains a matter of debate, but it is a much more precise debate than the one we started with. I tend to side with Galison in this particular case, because if Manhattan Project scientists had considered implementing Monte Carlo using traditional methods, their back-of-envelope estimations would not (in my estimation) have suggested any advantage over their current approaches. It is only with the considerable improvement in speed that came with ENIAC that the problem could be solved in practical time. It is worth noting that if the computational problem were even larger—say, fusion bombs rather than fission—the case would be even stronger.

It is also worth noting that once the method was suggested to him by Ulam, von Neumann immediately thought to implement it using ENIAC (see Richtmyer and von Neumann 1947, 751–752). Perhaps this is a case of the "Birmingham screwdriver"—with a hammer in hand, everything looks like a nail. And when faster computers are available, more problems begin to look susceptible to a numerical approach. But this is to suggest a psychological mechanism by which the scientists became cognizant of the fact that Monte Carlo was a promising approach—it is not to say that Monte Carlo was inconceivable (or unconceived of) before the computer. It is clear that scientists had the conceptual resources necessary before the advent of the digital electronic computer—no matter what role the computer might have had in reminding them of this possible solution.

I began this paper by considering conflicting claims about what difference technological change makes to the pursuit of knowledge. 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 epistemic possibility. In the second part, I introduced technological possibility, which depends on access to the material and conceptual means of bringing about a desired state of affairs, as one constraint on practicability, making technological possibility a necessary but insufficient condition for epistemic possibility.

**Acknowledgments** I would like to thank the audience at the 2011 meeting of the Society for the Philosophy of Science in Practice for comments on portions of this paper. I also owe debts to Anjan Chakravartty, Boaz Miller, Eleanor Louson, an anonymous reviewer, and especially Greg Lusk for their comments and suggestions that have substantially improved the paper. This paper was completed while I was a Postdoctoral Research Fellow at the Faculty of Information, University of Toronto.

## References

- Agar, J. (2006). What difference did computers make? *Social Studies of Science*, 36(6), 869–907.
- Conee, E., & Feldman, R. (2004). *Evidentialism: Essays in epistemology*. Oxford: Clarendon Press.
- DeRose, K. (1991). Epistemic possibilities. *The Philosophical Review*, 100(4), 581–605.
- Egan, A. (2007). Epistemic modals relativism and assertion. *Philosophical Studies*, 133(1), 1–22.
- Egan, A., Hawthorne, J., & Weatherston, B. (2005). Epistemic modals in context. In G. Preyer & G. Peter (Eds.), *Contextualism in philosophy* (pp. 131–169). Oxford: Oxford University Press.
- Foley, R. (2003). Justified belief is responsible belief. In M. Steup and E. Sosa (Eds.), *Contemporary debates in epistemology*. Available online: [philosophy.fas.nyu.edu/docs/IO/1161/justifiedasresponsible.pdf](http://philosophy.fas.nyu.edu/docs/IO/1161/justifiedasresponsible.pdf).
- Galison, P. (1996). Computer simulations and the trading zone. In P. Galison & D. J. Stump (Eds.), *The disunity of science: Boundaries, contexts, and power* (pp. 118–157). Stanford: Stanford University Press.

- Gendler, T. S., & Hawthorne, J. (Eds.). (2002). *Conceivability and possibility*. Oxford: Clarendon Press.
- Gibbs, B. (1970). Real possibility. *American Philosophical Quarterly*, 7(4), 340–348.
- Hacking, I. (1967). Possibility. *The Philosophical Review*, 76(2), 143–168.
- Hacking, I. (1975). All kinds of possibility. *The Philosophical Review*, 84(3), 321–337.
- Hartshorne, C. (1963). Real possibility. *The Journal of Philosophy*, 60(21), 593–605.
- Huemer, M. (2007). Epistemic possibility. *Synthese*, 156(1), 119–142.
- Kornblith, H. (1983). Justified belief and epistemically responsible action. *The Philosophical Review*, 92(1), 33–48.
- Longino, H. (1990). *Science as social knowledge*. Princeton: Princeton University Press.
- MacFarlane, J. (2011). Epistemic modals are assessment-sensitive. In A. Egan & B. Weatherson (Eds.), *Epistemic modality* (pp. 144–178). Oxford: Oxford University Press.
- Moore, G. E. (1962). *Commonplace book, 1919–53*. London: Allen and Unwin.
- Richtmyer, R. D., & von Neumann, J. (1947). Statistical methods in neutron diffusion. In A. H. Taub (Ed.), *John von Neumann: Collected works* (pp. 751–764). Oxford: Pergamon Press.
- Seddon, G. (1972). Logical possibility. *Mind*, 81(324), 481–494.
- Seidel, R. W. (1998). ‘Crunching numbers’ computers and physical research in the AEC laboratories. *History and Technology*, 15, 31–68.
- Stanley, J. (2005). Fallibilism and concessive knowledge attributions. *Analysis*, 65(2), 126–131.
- Teller, P. (1972). Epistemic possibility. *Philosophia*, 2(4), 303–320.
- von Fintel, K., & Gillies, A. S. (2011). ‘Might’ made right. In A. Egan & B. Weatherson (Eds.), *Epistemic modality* (pp. 108–130). Oxford: Oxford University Press.
- von Neumann, J. (1961). On the principles of large scale computing machines. In A. H. Taub (Ed.), *John von Neumann: Collected works volume V design of computers, theory of automata and numerical analysis* (pp. 1–32). Oxford: Pergamon Press.
- Weatherson, B., & Egan, A. (2011). Epistemic modals and epistemic modality. In A. Egan & B. Weatherson (Eds.), *Epistemic modality* (pp. 1–18). Oxford: Oxford University Press.
- Williams, M. (2001). *Problems of knowledge: A critical introduction to epistemology*. Oxford: Oxford University Press.
- Yalcin, S. (2011). Nonfactualism about epistemic modality. In A. Egan & B. Weatherson (Eds.), *Epistemic modality* (pp. 295–332). Oxford: Oxford University Press.