



ARTICLE

The Power of Knowledge

Gabriel Târziu 

Munich Center for Mathematical Philosophy, LMU Munich, Munich, Germany

Email: Gabriel.Tarziu@lmu.de

(Received 26 March 2025; revised 21 June 2025; accepted 26 June 2025)

Abstract

According to a well-known aphorism attributed to Francis Bacon, knowledge is power. But what does it mean for knowledge to *be* power? This paper addresses the question not by offering a new theory of knowledge, but by examining how, and under what conditions, knowledge places an epistemic subject at an advantage over those who lack it. The account developed here does two things: first, it explicates the widely held intuition that possessing knowledge confers an empowering advantage; second, it explains why, in certain contexts, increased knowledge can paradoxically generate a sense of powerlessness. This account diverges from both causal and practicalist views of the power of knowledge: the former takes power to be a causal consequence of possessing knowledge, while the latter understands knowledge as inherently linked with the ability to perform potentially useful actions. In contrast, I argue that the power of knowledge is best understood as a probabilistic advantage, namely, an increase in the likelihood that an epistemic subject will be more successful in interacting with the outside world.

Keywords: The power of knowledge; knowledge is power; knowledge practicalism; knowledge and action

1. Introduction

I know that our planet's climate is changing. I also know why it is changing and what could, in principle, be done to stop it; nonetheless, I feel powerless when it comes to doing something to stop climate change. This feeling of powerlessness appears to contradict the well-known aphorism attributed to Francis Bacon: *knowledge is power*. Variants of this idea have been widely echoed throughout time by scientists, policymakers, and philosophers. But what does it actually mean to say that knowledge is power? And under what conditions does this claim hold?

This paper has two main aims: first, to provide an explication of the concept of the power of knowledge, and second, to explain why, in certain contexts (such as when dealing with the issue of climate change), one can feel powerless despite possessing extensive knowledge. Two topics fall outside the scope of this discussion: trivial or useless knowledge (such as knowing how many grains of sand are on a beach) and political power. What I am concerned with here are situations that fall under the following general principle:

PK: knowing that **p** places an epistemic subject **S** at an advantage compared either to her previous self who did not know that **p**, or to other epistemic subjects who do not know that **p**.

Two clarifications are in order regarding **PK**. First, the advantage it posits is contrastive: it does not assume that knowledge is intrinsically powerful in all circumstances, but rather that it makes a difference relative to the condition of not knowing or knowing less. Second, the notion of ‘advantage’ in **PK** is vague and can be interpreted in several ways: as a political advantage (e.g. exercising power over others due to superior epistemic positioning), an epistemic advantage (e.g. the value of knowledge over other epistemic states), or a practical advantage (e.g. doing more as a result of knowing). While all these readings are philosophically significant, my concern here is only with the practical sense of advantage because what I aim to explicate is the widespread intuition that knowledge improves one’s efficacy in interacting with the world: when one knows that **p**, one can do more than when one does not, or compared to someone who does not know.

So, the advantage referred to in **PK** will be understood here in terms of an epistemic subject’s ability to perform **p**-related actions, that is, actions that either depend on knowing that **p** or are appropriate manifestations of that knowledge. Performing such actions can be construed as using knowledge to enhance one’s practical efficacy in interacting with the world. I will use ‘efficacy’ throughout this paper to refer to an epistemic subject’s context-sensitive capacity to put her knowledge to work in ways that yield practical benefits.

We can treat efficacy as a form of power, if what is understood by power is this:

\mathbb{P} :¹ a measure of the degree to which one manages to impose one’s will on the outside world in a way that is contextually relevant/useful (i.e. in a way that makes sense given the situation in which the epistemic subject finds herself).

This is not intended as a general theory of power. Whether one prefers to call this advantage ‘power’, or reserves that term for other forms of influence (e.g. political), is ultimately a matter of conceptual preference. What matters for present purposes is that knowledge, from the perspective of the intuition we are trying to explicate here, seems to afford a form of context-sensitive practical efficacy which some may consider to be a form of power. While the formal framework developed in Section 4 aims to capture this efficacy via an objective measure, \mathbb{P} , it is important to note that an individual’s subjective sense of power can diverge significantly from this objective metric. In Section 5, I will draw on this fact to explain how increased knowledge can, in certain contexts, lead to a paradoxical feeling of powerlessness.

To summarize: the object of this paper is the idea that knowing that **p** can act as an empowering resource for an epistemic subject (relative to a prior state or another subject) because it may create a practical advantage for **S**; specifically, it may increase the efficacy with which **S** can interact with the world. This potential increase in efficacy can plausibly be understood as a form of power.

Now that I have clarified the object of this discussion, I turn to the main intuition behind the account that I’m trying to provide in this paper. In my view, possessing knowledge is analogous to possessing any other resource. Just as having a particular object grants certain advantages, so too does possessing knowledge. Consider, for instance, a person with a basket of apples. It is obvious that, as a result of having the

¹ \mathbb{P} is introduced here just as a working definition of practical power; in Section 4, it will be given formal precision.

apples, that person can *potentially* do a lot of things she could not have done without them: she could eat apples, she could offer someone else some apples, she could make an apple pie or caramel apples, she could make cider, apple juice, apple cider vinegar, or she could place an apple on someone else's head and use it for target shooting, etc. We can interpret this in modal terms by saying that possession of apples expands one's field of possible actions. Of course, this expansion of one's field of possibilities as a result of having apples does not *necessarily* translate into an actual growth of the set of things one can *actually* do. It may be the case that someone is unable, unwilling, or uninterested in doing any of the things on this list, no matter how simple they are. In such cases, possession only increases what is, in principle, available to the agent.

Most of us will undoubtedly be unable, unwilling, or uninterested in using the apples for target practice, but some people may even be unable or unwilling to eat them because they are, for instance, allergic to apples. So, the extent to which one can take advantage of one's available resources depends heavily on personal and contextual factors. What this should make clear is that the practical advantage, in such cases, cannot be reduced to the space of *x*-related possible actions that the possession of *x* opens up for the agent. Rather, it concerns the probability of practical efficacy, that is, the likelihood that one's interaction with the world, given access to the resources, will be more successful under some context-sensitive conditions. The advantage conferred by possessing a resource is therefore not guaranteed; it is probabilistic because the realization of the possibilities it opens up depends on whether the agent finds herself in the right context and is willing and capable of acting on them.

I believe the same is the case with knowledge: possessing a particular item of knowledge similarly unveils a space of possible actions which, only under the right conditions, can enhance the efficacy with which an epistemic subject interacts with the world. However, the actualization of these possibilities depends on one's abilities, dispositions, and context. Consequently, possessing knowledge does not guarantee a more successful engagement with the world; rather, it merely increases the likelihood of efficacy. The reasons behind this will be substantiated later in the paper.

This paper contends that, in this respect, knowing that something is the case is very much like having any other resources (in our example, apples). Of course, this analogy only goes so far. An important difference is that, while without apples you cannot do any of the things listed above, there are many situations in which you may be able to do, in the absence of a particular item of knowledge, some of the actions within the set of possible manifestations of that knowledge (for instance by pure luck, by having merely true beliefs, or by relying on some other things).

My goal in this paper is to develop and substantiate this intuitive account. The discussion will proceed as follows. The next section will be devoted to setting up the stage for the rest of the discussion. In Section 3, I will discuss a strong view on the relationship between knowledge and power: Stephen Hetherington's practicalist conception of knowledge, according to which knowledge inherently has the potential to be useful because it literally is that potential. This view serves as a useful reference point for evaluating whether and to what extent knowledge can be linked to power; it also provides a central ingredient to the account developed here: the idea that knowledge is normally associated with a range of possible manifestations. To provide a more nuanced account, Section 4 will introduce a modal-logic-inspired framework designed to explicate the concept of the power of knowledge. This framework is needed (a) to formally distinguish between the modal space of possibilities that knowledge opens up for an epistemic subject and the conditions under which those possibilities translate into actual, context-sensitive efficacy; (b) to explain the probabilistic link between knowledge and power; and (c) to formally express the key divergence between the account

developed here and the causal and practicalist accounts. Section 5 will apply this framework to explain how and why, in certain cases, gaining more knowledge may paradoxically lead to feelings of powerlessness rather than empowerment, with particular attention to the issue of climate change, where increasing knowledge does not necessarily translate into increased capacity for successful action.

2. The power of knowledge

What does it mean to say that knowledge is power or that knowledge has power? Over time, the idea of knowledge as a source of power has been interpreted in various ways. For Bacon, for instance, it meant that knowledge enhances human control over nature. In *Novum Organum*, he famously claimed that ‘human knowledge and human power meet in one; for where the cause is not known the effect cannot be produced’ (I, Aph. 3; quoted in Adolf and Stehr 2016). According to some interpretations (e.g. Gaukroger 2001), Bacon associated the empowerment derived from knowledge with political legitimacy: those who possess the relevant knowledge are, in virtue of that knowledge, entitled to wield political authority. So, for him, ‘knowledge’ refers to the study of the causes and effects of natural phenomena, while ‘power’ had to do both with exercising dominion over nature and with legitimate ruling over other people.

More recently, the relationship between knowledge and power became a central theme in the work of prominent poststructuralist thinkers such as Michel Foucault and Pierre Bourdieu. Unlike Bacon, these thinkers are quite suspicious about political power and also see a stronger link between knowledge and power than Bacon did. For instance, Foucault considered that power and knowledge incorporate each other. In his opinion, ‘It is not possible for power to be exercised without knowledge, it is impossible for knowledge not to engender power’ (Foucault and Gordon 1980, p. 52). Foucault understands this relationship to go beyond the obvious fact that acquiring and exercising (political) power depends on securing an epistemological advantage over others. Power not only needs knowledge but also ‘power produces knowledge (and not simply by encouraging it because it serves power or by applying it because it is useful) ... power and knowledge directly imply one another ... there is no power relation without the correlative constitution of a field of knowledge, nor any knowledge that does not presuppose and constitute at the same time power relations’ (Foucault 2007, p. 27).

This discussion about the overlap between knowledge and political power is obviously interesting and important, but will not constitute the topic of the present paper (although a lot can be and has been said, e.g. by social scientists, about the role of knowledge in the political decisions concerning climate change). Instead, our focus will be on a more mundane intuition: that knowledge functions as an empowering resource in the practical sense captured by PK and \mathbb{P} . That is, knowledge gives an epistemic subject a comparative advantage relative to those who lack it. This advantage can be understood as a form of power. To illustrate this intuition, consider the following example:

Sylvie and Irma

Sylvie was born and raised in Paris and so she knows a great deal about the city (e.g. where the best restaurants are, what streets to use as shortcuts between different locations in the city, at what hours the city centre becomes overcrowded, the roads to avoid during rush hour, etc.). Irma, by contrast, is a tourist who arrives in Paris as part of a surprise birthday gift from her friends. Having done no prior

research, her knowledge of the city is minimal. She knows that Paris is home to famous landmarks like the Louvre and the Eiffel Tower, but she has no sense of where they are or how to reach them from her hotel.

It is reasonable to assume that most would agree Sylvie is in a better position than Irma, and that her status has to do with the fact that knowledge can act as a sort of empowering factor in contexts such as this. While Sylvie's knowledge gives her the power to do things like enjoying the city to its fullest and helping others (like Irma) navigate it, Irma's lack of knowledge leaves her completely powerless, and so, (at least apparently) dependent on others for most of the city related activities she would like to do.

This intuition can be expressed a little bit more clearly as follows:

PoK*: an epistemic subject **S** has the power of knowledge if **S** is in the possession of an item of knowledge that somehow opens up the possibility for her to do specific things (things with respect to which that item of knowledge is relevant).

How should we interpret **PoK***, and what does it imply about the relationship between knowledge and power? At first glance, **PoK*** suggests that power is extrinsic to knowing – that it emerges as a *causal consequence of possessing knowledge*. However, this perspective has been challenged recently by Stephen Hetherington. The subsequent section explores Hetherington's practicalist conception of knowledge, not as a target for criticism, though (given that Hetherington's work is concerned with the nature of knowledge), but as a means to deepen our understanding of the relationship between knowledge and power. The aim, therefore, is not to offer a comprehensive or exegetically precise presentation of practicalism, but to draw selectively on its core insights, especially the idea that knowledge is associated with a range of possible manifestations, in order to motivate and structure a framework that shifts the explanatory focus from the nature of knowledge to the conditions under which knowledge affords practical advantage.

3. The practicalist conception of knowledge

'Knowledge is indeed power . . . but not because knowing gives one power. Rather, knowing is power – within itself, part of its nature. To know is to have various abilities or skills. This is not because knowing gives one these abilities or skills, nor is it because knowing is a normative ground of actions arising through exercising such abilities or skills' (Hetherington 2022, p. 255).

Knowledge practicalism is an ability account of knowledge, i.e. an account of knowledge that makes a stronger than usual connection between knowledge and what one can do as a result of having it. In recent literature, two of the best-known advocates of ability accounts of knowledge are Stephen Hetherington (2006, 2011b, 2011a, 2017, 2020, 2022) and John Hyman (1999, 2006, 2015). Both Hyman and Hetherington consider that knowledge is best understood as a kind of ability. But, while for Hyman knowledge is the ability to be guided by facts, according to Hetherington, having propositional knowledge of some state of affairs amounts to having the ability to manifest various accurate representations of that state.

Hetherington goes a step further than the typical ability accounts of knowledge by taking knowledge to be inherently linked with (potential) *usefulness*. In his opinion, 'knowledge has an inherent potential to be useful (to at least some extent, and for at least

some possible person) – because literally, it is the potential to be useful in that way’ (Hetherington 2011b, p. 232). At first glance, this position seems perfectly aligned with the concerns of this paper. Taking knowledge to be the potential to be useful appears to directly support the idea that knowledge confers an advantage of the sort captured by **PK** and **ℙ**.

If Hetherington’s practicalist view on knowledge is correct then knowledge and what one can do as a result of having it (its practical consequences) are the same thing, so it makes no sense to distinguish between knowledge and *its* power (as it is, apparently, done in **PoK**^{*}). What we consider to be the power of knowledge is nothing more than the manifestations of knowledge and, as such, is constitutive of one’s knowledge. From this perspective, the ‘is’ in the ‘knowledge is power’ aphorism stands for a metaphysically constitutive relationship: ‘possessing knowledge is already and always, in itself, one’s being able to perform various actions’ (Hetherington 2020, p. 170).

In Hetherington’s view, between knowledge and action, there is a more intimate link than the accidental and external relationship postulated by the traditional way of thinking. According to him, the possession of cognitive and behavioural skills is metaphysically intrinsic to having knowledge. To know something is just to *know how* to perform various actions. So, knowing does not just causally enable actions but is metaphysically constituted by one’s ability to perform some kind of action (Hetherington 2020). But what specific actions are involved?

Between the things that Hetherington lists as manifestations of knowledge that **p**, we find: accepting **p**, assenting to **p**, having the true belief that **p**, remembering accurately that **p**, answering accurately in **p**-related ways, sensing accurately that **p**, explaining accurately in **p**-related ways, solving a theoretical problem to which **p** is the correct solution, performing actions which would not be appropriate unless **p** were true, etc. Hetherington calls these manifestations ‘**p**’s epistemic diaspora’ (Hetherington 2011b, p. 35) or ‘the commonwealth or bundle’ (Hetherington 2017, p. 5) of skills that constitute the complex ability that knowledge really is. In his opinion, this bundle of skills does not have a privileged centre, a metaphysical core to which having knowledge can ultimately be anchored. Epistemologists traditionally took belief to be such a core, while the rest of things that, according to Hetherington, are part of the complex potentiality of an instance of knowledge, were understood to be present only because belief is. But Hetherington considers this to be an epistemologically arbitrary move because, in his opinion, it can easily be shown that any of the other members of the bundle can act as the core in relation to which the others are mere uses. If, for instance, we take explaining-accurately-in-**p**-related ways as the centre, we can interpret one’s believing, sensing, asserting, etc. as results of the fact that the reasoning is present (Hetherington 2011b, p. 36).

A legitimate concern that one may raise at this point is this: if there is no core, i.e. something that can bring them together, what unites the members of an item of knowledge’s epistemic diaspora? As Hetherington argues, there is a simple answer to this: ‘The members are what would, within standard theories of knowledge, be construed as the various possible ways of expressing or manifesting knowledge that **p**’ (p. 37). This answer also helps if we are concerned with determining the components of the bundle. It can lead to a misunderstanding, though. It can make us think that all the members of a piece of knowledge’s commonwealth of skills have to be present at the same time. But, as Hetherington clarifies, it is not the case that, whenever knowledge that **p** is present, all the members of its epistemic diaspora also have to be present – any of them could be, but no one in particular (p. 40).

So, knowledge that **p** amounts, in Hetherington’s view, to knowing how to perform various actions. If we apply this to the *Sylvie and Irma* scenario, we will say that what we

presented above as the common intuition that Sylvie is in a better position than Irma because her knowledge *gives her the power* to do things that Irma cannot do, is at the same time right and wrong. It is mistaken to the extent that it understands power as something separate from knowledge (something that knowledge only causally engenders). According to practicalism, Sylvie's power has everything to do with and does not go in any way beyond the knowledge that she has about Paris: Sylvie's knowledge *is* her power. What the intuition seems to get right is the fact that power, in these types of contexts, consists of one's ability to perform various actions. But, since knowledge reduces, according to this account, to one's complex ability to manifest accurate representations of facts, there is no way for Sylvie to know something but be unable to act.² For instance, there is nothing more to her knowledge that at certain hours the city centre becomes overcrowded than her ability to express this knowledge by manifesting one, some, or all of the actions associated with this item of knowledge's epistemic diaspora (such as having the true belief that this is the case and performing actions which would not be appropriate unless this fact were true).

So, as it is presented above, **PoK**⁺ can be misleading because it seems to incorporate the assumption that knowledge lies beyond what one can do and only causally interferes with one's abilities. If we replace this assumption with the contention that 'knowledge is the conceptually constitutive – and not (as is usually said) just a causally efficacious – locus for these actions' (Hetherington 2011a, p. 85), then we can get this new view about the way we should think about the power of knowledge:

PoK_{pra}:³ an epistemic subject **S** has the power of knowledge simply in virtue of knowing that **p**, because possessing that knowledge entails the capacity to perform one or more (potentially) *useful actions* that are normally regarded as proper manifestations of it.⁴

The reasoning behind **PoK**_{pra} can be presented as follows: since having knowledge amounts to being able to perform (potentially) useful actions, and performing useful actions can be perceived as a form of power, then possessing knowledge is tantamount to having power. One important clarification is needed at the outset: when viewed through the lens of **PK** and **P**, **PoK**_{pra} concerns only the external manifestations of knowledge, that is, the part of an epistemic diaspora that includes outward, contextually relevant

²The notion of an 'ability to act' is ambiguous here, referring either to internal or external actions or manifestations of knowledge. Hetherington's practicalism accommodates both: an epistemic subject can manifest her knowledge of **p** internally (e.g. by remembering **p** or reasoning from **p**) or externally (e.g. by using **p** as a guide to action in the world). So, when it is claimed above that 'there is no way for Sylvie to know something but be unable to act', the intended point is that knowledge entails the availability of some form of manifestation (internal or external). It is important to note, however, that since the focus of this paper is on whether and how knowledge may generate a practical advantage for an epistemic subject, only external manifestations will be considered below when formulating **PoK**_{pra}, i.e. the practicalist-inspired view of the power of knowledge. In this respect, the paper engages, from this point forward, only a restricted version of practicalism: one concerned solely with outward-facing actions. I am grateful to an anonymous reviewer for prompting this clarification.

³**PoK**_{pra} can be taken to be a stronger version of a different account (**PoK**_{abil}) according to which an epistemic subject **S** has the power of knowledge if **S** is in the possession of a piece of knowledge, which means that **S** is able to perform one, some or all the actions that are regarded as manifestations or expressions of that knowledge. Unlike in the case of **PoK**_{pra}, in **PoK**_{abil} the connection between knowledge and power is made on the basis of what can be taken as an equivalence between having abilities (to do something) and having (the) power (to do it).

⁴See especially Hetherington (2020, 170).

actions. In this respect, $\mathbf{PoK}_{\text{pra}}$ has a narrower focus than Hetherington's practicalism. As discussed above, Hetherington includes among the manifestations of knowledge various inner ones, such as assenting to \mathbf{p} , truly believing that \mathbf{p} , remembering that \mathbf{p} , or accepting \mathbf{p} . By contrast, $\mathbf{PoK}_{\text{pra}}$ brackets these and concentrates on the extent to which knowledge can enable successful interaction with the external world.

Accordingly, and unsurprisingly, Hetherington does not endorse anything like $\mathbf{PoK}_{\text{pra}}$. Its introduction here should be taken only as one possible way of conceptualizing the power of knowledge if one adopts a practicalist view of knowledge, but constrains that notion to the kinds of outward-facing actions that \mathbf{PK} and \mathbb{P} are meant to track.

That said, the question we should now ask is: to what extent is $\mathbf{PoK}_{\text{pra}}$ a plausible account of how knowledge confers power? To answer this question, we need to consider whether there are cases in which an epistemic subject possesses both the relevant knowledge and the ability to perform the associated actions, yet remains unable to benefit from its practical utility. The goal is to determine whether a wedge can be driven between knowledge and what one can do with it or obtain from it, i.e. its power (understood as practical utility, i.e. as a benefit that an epistemic subject may obtain as a result of knowing that \mathbf{p}).

As I will argue in the next sections, there are three types of cases in which the presence of such a wedge becomes evident: (i) cases in which even though one has knowledge and is aware of the proper way to express it (or take advantage of it) in context \mathbf{c} and is also endowed with whatever ability needed to do it, one systematically fails; (ii) cases in which one knows that \mathbf{p} and has the ability to perform whatever action is normally considered a relevant/useful manifestation of that knowledge in context \mathbf{c} , but one fails to manifest it because, for some reason, one cannot link the knowledge that \mathbf{p} with the proper (in \mathbf{c}) \mathbf{p} -related action; and (iii) cases in which, even though the (potential) usefulness of a piece of knowledge that the epistemic subject possesses is quite great, she is unable to take advantage of it because unlocking it is beyond any single person's abilities.

3.1. The failure to take advantage of one's knowledge: intervening factors

Let's start by reconsidering the *Sylvie and Irma* case. Suppose, for example, that Sylvie has a sister, Anna, who possesses all the same knowledge about Paris as Sylvie, but who, unfortunately, is a notoriously lazy person so she is usually unable to plan ahead and follow through on what she sets out to do. Assume further that Anna commutes to work by car and, from experience, knows that Paris is congested between 7 and 9 in the morning and between 17 and 19 in the evening on weekdays, and so that, to avoid traffic jams, she needs to leave home a little bit earlier than she usually does. Even though she makes plans every day to do this, she ends up procrastinating and so she systematically fails to act on her plans and ends up stuck in traffic and late for work.

This scenario should be familiar to many. Most of us know exactly what it is like to know that \mathbf{p} , and know to do Φ (in context \mathbf{c}) because \mathbf{p} (and so as a consequence set Φ -ing as an objective), and be able to do Φ , and, nonetheless, still fail at Φ -ing. Every student knows, for instance, what it is like to know the date at which an essay needs to be submitted, to know also that to meet the deadline they need to start working on it as soon as possible, and nonetheless fail to meet the deadline. What we are dealing with in these situations is a severing of the link between knowledge and what one is able to do (as a result of having it) by some intervening factor. This intervening factor can be thought of as being an *ability-modifier* that is an intrinsic factor that systematically tampers with the proper manifestation of one's knowledge that \mathbf{p} in context \mathbf{c} by affecting how

knowledge translates into action. Ability-modifiers are different from *ability-restrictors*, which are extrinsic factors that temporarily affect an epistemic subject's abilities by suppressing their proper manifestation in a particular context. An example of an ability restrictor would be, let's say, Sylvie's neighbor if he would stop for some reason Sylvie from leaving home at the optimal time and, as a result of this, she would end up stuck in traffic.

3.2. The failure to take advantage of one's knowledge: connection issues

A different kind of situation in which one has knowledge but is unable to take advantage of it (i.e. to benefit from its utility) is one in which the link between knowledge and (the/ an appropriate) action(s) is, for some reason, very hard to establish, to begin with. Consider the following case:

Bob

Bob goes to the supermarket to buy food. The cashier tells him that the purchased items cost \$23.13. So, it is safe to say that Bob knows exactly how much he has to pay. Suppose further that Bob has more than enough money on him to pay for his groceries and that there is nothing physically wrong with him. The standard manifestation of Bob's knowledge in this context is that of giving the cashier the exact amount of money that he owes. Unfortunately, Bob has *acalculia* (i.e. he is unable to process numbers and perform calculations) and so, even though he has the money and nothing is preventing him from reaching his pocket (i.e. he has the physical ability to manifest his knowledge in a relevant way), he is unable to pay for his groceries. To do it, he must ask for help.

Bob's situation exemplifies an extreme case of a more common problem. Like in the previous case, most of us find ourselves in situations in which we have trouble making proper practical use of our knowledge, but this time the problem is not generated by something intervening between knowledge and action but by the lack of something needed to connect the two (something such as creativity, spontaneity, deductive reasoning, understanding, and the ability to extrapolate). So, even though we know that **p**, and performing action **φ** is a proper expression (in context **c**) of knowing that **p**, we fail at **φ**-ing because we lack what it takes to link knowing that **p** with **φ**-ing. The most common situations of this kind are, probably, those in which, even though we have all the knowledge needed for successful action we fail to put it together, to see how it connects, and, as a result, we cannot do much with it. A relatable example is when we encounter a household problem and search online for solutions, only to discover that the fix involves common items we already own and know how to use in other contexts.

3.3. The failure to take advantage of one's knowledge: unlocking knowledge's practical utility

I want to discuss in this section a kind of situation that is different from the two discussed above. Here, the problem we are dealing with does not concern anymore the link between knowledge and its proper manifestation in a particular context. Instead, what characterizes this kind of situation is the fact that the practical utility of a piece of knowledge is not accessible (at least not directly) to an epistemic subject because unlocking it surpasses what that agent is able to do by herself. Consider the following cases:

Stuck in traffic

Imagine being caught in a traffic jam, but, unlike in Anna's case, it is not your fault that you are in that situation. The reason why you are stuck in traffic is not because you failed to avoid driving at those hours at which the roads get crowded usually. Instead, in your case, a big tree happened to fall onto the road. In this situation, it is safe to say that you know what is going on and why. You also know what needs to be done to clear up the traffic, but there is nothing that you can do by yourself to make things better. You just have to wait for the authorities (which, let's suppose, were already noticed by your fellow drivers) to come and clear the road.

Robbery

Consider entering a bank in the midst of a robbery. Again, you know what is going on and you have a pretty good idea why. You also know what needs to be done to make things right. But it is not in your power to do anything in this context (I am assuming, of course, that you, like a great majority of people, do not have the required training to thwart the robbery by yourself). You just have to wait for the police to come and deal with this situation.

A more extreme case⁵ of this kind of situation, that all of us are dealing with nowadays, has to do with climate change. Many people are aware of climate change and of the kinds of actions that are typically recommended to address it (Andre *et al.* 2024). But, although there are a number of actions that we can individually take (like, for instance, reducing one's carbon footprint by avoiding air travel and living car-free) to mitigate climate change, none of these actions has any effect if a great majority of our human fellows do not join us.

What these examples illustrate is that, in some situations, while an epistemic subject may know that *p* and may recognize that performing *φ* (in context *c*) is the expected/required action associated with that item of knowledge, the actualization of *φ* exceeds the subject's individual capabilities. Successful action in these scenarios demands either collaboration among multiple agents or the intervention of specialized organizations, institutions, or governments.

3.4. Power and practicalism

The cases discussed in the previous three sections are meant to illustrate that having knowledge and being able to leverage it to one's benefit in interactions with the outside world are two very different things. This means that, at least in some situations, an individual can have knowledge while remaining entirely powerless (if, of course, power is understood as practical efficacy), as Anna, Bob, and whoever finds themselves in situations such as the ones discussed in the *Stuck in traffic* and *Robbery* cases clearly are. At first glance, this may appear to be at odds with Hetherington's practicalism because, from the perspective of this account of knowledge, 'having knowledge is expressive of being inherently a thing with powers – expressive of person-powers, we might say' (Hetherington 2020, p. 176).

The practicalist is not defenceless against these arguments, though. One natural rejoinder that the practicalist may have to the use of the cases discussed above in an

⁵What makes this case extreme is the level of involvement it requires from others in order for the utility associated with this particular item of knowledge to be unlocked.

argument against her position would undoubtedly be that of pointing out that, e.g. Bob and Anna are actually able to do a lot of things (e.g. to accurately believe that **p**, to accurately answer that **p** is the case, to remember, report, reason accurately, and represent accurately) from the epistemic diaspora of their respective items of knowledge and so they have power.⁶ The strategy here is that of avoiding what Anna and Bob cannot do and concentrating on what they can do, i.e. on the members of their respective knowledge's epistemic diaspora that they can manifest.

A different line of defence could be to point out that what the practicalist may have in mind by 'power', when she says that knowledge is power, is actually disposition. Yet another possibility would be to interpret the relationship between knowledge and power in less extreme terms: knowledge is power, but only partly. There is no need for the practicalist to go to extremes with this relationship. If we take power to depend on or be equivalent to effective agency, then this moderate position seems to be something that Hetherington is actually embracing because he takes knowledge to consist only partly in having effective agency: 'To know is, *in part*, to have effective agency; whereas to be knowing is, *in part*, to be manifesting effective agency' (2011b, 87, my emphasis).

There is little need to engage further in an interpretative exercise about what the practicalist may or may not claim in response to what has been said in these pages. As I made clear from the outset, the aim of the preceding discussion was not to criticize Hetherington's view. I hope it is now clear why. One of the main aims of this paper is to provide an explication of the concept of power of knowledge as exemplified by cases such as *Sylvie and Irma* and as captured by **PK** and **ℙ**. However, the practicalist, being concerned with *the nature of knowledge*, has no reason to commit to something resembling **PK** and **ℙ**.

Within practicalism's broader framework, there is considerable flexibility in accommodating the cases discussed in Sections 3.1 through 3.3. However, from the narrower standpoint adopted in this paper, these cases are not so easily dismissible. We can think of them as generating the following dilemma for **PoK_{pra}**: either it is interpreted as associating power with the full range of possible knowledge manifestations (internal and external), in which case it no longer tracks the kind of practical advantage that interpreting **PK** in terms of **ℙ** is meant to capture; or it is interpreted in the narrower sense suggested in Section 3, where power refers only to the ability to perform useful, outer actions based on one's knowledge. On this latter reading, however, **PoK_{pra}** fails in the sorts of cases discussed in Sections 3.1 through 3.3 – cases where someone possesses knowledge but is unable to derive any meaningful practical advantage from it. It is unclear, for example, what advantage Anna or Bob derive from their knowledge. While there may be things that even they can do, it is difficult to argue that these generate any significant benefit for them relative to agents who lack that knowledge. If these cases hold, then **PoK_{pra}**, so interpreted, proves untenable.

But, to reiterate, nowhere in his work does Hetherington support a thesis similar to **PoK_{pra}** nor do I believe he is forced to do so. The difference between our discussion and the core concerns of practicalism can be seen as one of emphasis and conceptual preference: the practicalist is concerned with what knowledge is, whereas I am concerned with when and how knowledge produces practical utility or advantage. I choose to call this latter aspect 'power', but one might very well reserve the term for manifested knowledge, as Hetherington arguably does. But, of course, the terminological choice is secondary; what matters is the explanatory goal. From my standpoint, **PK** and

⁶It is important to emphasise that practicalists are not compelled in any way by their account of knowledge to interpret the kind of advantage referenced in **PK** in terms of **ℙ**. Accordingly, they retain considerable flexibility in how they understand the notion of 'power'.

\mathbb{P} capture the structure of an intuitive distinction between the epistemic subjects empowered by knowledge and those who are not that the practicalist, focused on the nature of knowledge, has no theoretical need to preserve.

That said, the discussion in the previous sections has been far from fruitless. Exploring the practicalist account of knowledge has played a productive dialectical role in helping to isolate and clarify the conceptual terrain. It has provided important insights into the relationship between knowledge, its possible manifestations, and the practical advantages that may or may not follow from it. One key lesson is that knowledge and power (understood along the lines of \mathbb{P}) do not always come together and so having knowledge is not equivalent to having power. An additional, indirect lesson is that knowledge also does not cause power (again, if power is understood along the lines of \mathbb{P} as exerting dominion over the external world). The knowledge that people such as Anna, Bob, or the people in the *Stuck in traffic* and *Robbery* scenarios have does not cause any significant change in the way in which they interact with the world. It does not make them more successful in their actions, for example, and so it *does not seem* to give them any advantage compared to those who lack such knowledge.

4. Possible worlds, p-actions, and the power of knowledge

The main lessons about the relationship between knowledge and power derived in the previous section notwithstanding the situation described in **PK** does seem to make a lot of intuitive sense. Possessing knowledge does appear to give an epistemic subject an advantage over those who lack it. But how else can we understand this advantage if not in the terms suggested by practicalism? Another option, mentioned at the end of Section 2, is to understand it in causal terms: being at an advantage is a causal consequence of possessing knowledge. However, as was explained in the previous section and as I will argue in this section, this way of conceiving the advantage that knowledge can generate for an epistemic subject is also flawed. This is because, as will be discussed below, there are two important types of contexts in which possessing knowledge may not result in a change in power (with power understood along the lines of \mathbb{P}). In the first context (formally described by formula 9 below), the space of possible actions that an item of knowledge opens up for an epistemic subject, given her cognitive profile and background knowledge, may not contain any actions that are contextually relevant or useful. In the second context (formal condition 10), although one may be aware of some contextually relevant or useful actions related to a particular known proposition, one may lack the necessary abilities to perform that action.

Given all this, instead of understanding the power of knowledge in either practicalist or causal terms, I propose to understand it in terms of the space of possibilities that knowledge opens up for an epistemic subject, and how access to this space improves the chances that the subject's interaction with the external world will be successful.

The rest of this section is dedicated to constructing a modal-logic-inspired formal framework to give precise meaning to this suggestion. Such a framework is necessary to provide semantic precision to key claims of this paper. Specifically, to clarify what it means for knowledge to open up a space of possible actions, what it means for an epistemic subject to 'have access to' those possibilities, and what it means to 'fail to actualize' them.

More precisely, what I hope to achieve with this framework is to provide a clearer understanding of why one's knowledge does not necessarily cause a change in one's power and explain why, despite this, **PK** can still make sense. As we will see in the next section (i.e. Section 5), there are also two additional (less intuitive) results we can derive from this framework: first, the evaluation of one's power is always a relative undertaking

that can sometimes create a false impression of power; and second, there may be situations in which the more knowledge one acquires, the less powerful one may feel.

The reason I believe modal logic and the semantics of possible worlds are ideal for this task is that, as expressed in **PoK**⁷, an intuitive way to understand how possession of knowledge affects an epistemic subject is through the space of possibilities it opens up for her.⁷ Specifically, what possessing knowledge *prima facie* does to an epistemic subject is to grant her access to a range of possible actions, which can be understood as being proper manifestations of that knowledge, that the subject may attempt if she finds herself in the appropriate context and has the right disposition.

In constructing this framework, I will begin not only with the intuition mentioned in the previous paragraph but also with an important lesson about knowledge that can be recycled from the discussion on practicalism, namely that associated with each item of knowledge there is an epistemic diaspora, i.e. a set of possible manifestations. However, unlike Hetherington's practicalist account, the framework developed here does not identify knowledge with various powers, each of which has distinctive manifestations. Although I take knowledge to be modally connected to a range of potential actions, it remains conceptually and metaphysically distinct from them.

A natural way of thinking about the bundle of manifestations is in terms of possible worlds – that is, complete and consistent ways the world could be, as in standard modal logic. We can start from that and take each proposition **p** (and, indeed, combinations of such propositions) that is known by an epistemic subject **S** to give rise to a distinct set of possible worlds W_p , which are characterized by the fact that all possible manifestations of **S**'s knowledge that **p**, or *p-related* actions, are exhibited within them.

A *p-related* action may be defined in this context as any action that is a possible manifestation of an epistemic subject's knowledge that **p** or that directly involves or is based on the knowledge that **p**. So, for example, if **p** stands for 'there are apples in the fridge', then for each *p-related* action (for a list of such possible actions, see Section 1) that **S** can possibly perform because **S** knows that **p**, like, e.g. eat some apples, make apple pie, etc., there is a possible world that belongs to W_p at which **S** performs that action.

We can then define the following accessibility relation **KAAR** (knowledge-action accessibility relation) that ties an epistemic subject's knowledge to the possible worlds they can access, where actions relevant to that item of knowledge have been taken: for a subject **S** and a known proposition **p**, **S** has access to a set of possible worlds $W_p = \{w_1, w_2, \dots, w_n\}$ from the actual world (w_0) at which **S** knows that **p** (i.e. $w_0 \models K(S, p)$) such that each w_i corresponds to a world at which **S** has performed a *p-related* action.

In more formal terms, we get:

1. $\forall S, \forall p, w_0 \models K(S, p) \Rightarrow \exists W_p (W_p = \{w_1, w_2, \dots, w_n\} \wedge \forall w_i \in W_p ((w_0, w_i) \in \text{KAAR} \wedge S \text{ performed a } p\text{-related action in } w_i))$

As it stands, **KAAR** is problematic because it does not integrate any of the lessons we learned from the discussion so far. Most importantly, it does not reflect the fact that, as made clear in Section 3.2, linking knowledge and action (i.e. its proper manifestation in a particular context) is not always easy. One, like Bob, for instance, may know that **p** is the case but be unable to connect the fact that **p** with doing **φ** in context **c**. In Bob's case, the failure is generated by the lack of a particular ability, but we can think of cases in which even one's background knowledge can lead to the same outcome. To see this, let's return

⁷I will refrain from taking a stance here on the best way to understand the access to possibilities provided by knowledge, such as whether it should be considered an ability or something else, since this is not essential to the primary objective of this paper, which is to explicate the concept of the power of knowledge.

to the case discussed in Section 1. If we have apples, there are many things we can potentially do with them, but not all of us will know that some of the things on that list are possible things one can do with apples, though. For example, someone who has no knowledge about cider, how it is made, or what it is made of, will not know that it is possible for her to use apples to make cider.

What this should make clear is that not all possible actions associated with an item of knowledge are automatically known to the epistemic subject possessing that knowledge and so that **KAAR** needs to be relativized to an epistemic subject's cognitive profile (i.e. how creative **S** is, how good is **S**'s deductive reasoning, the level of **S**'s understanding, etc.) and background knowledge. Amended this way, **KAAR** provides a more restrictive access to \mathbf{W}_p . In formal terms, it looks like this (where \mathbf{C}_p is meant to represent the cognitive profile of **S** and **B** her background knowledge):

2. $\forall S, \forall p, w_0 \models K(S, p) \Rightarrow \exists \mathbf{W}_p (W_p = \{w_1, w_2, \dots, w_n\} \wedge \forall w_i \in W_p ((w_0, w_i) \in \text{KAAR}(S, p, \mathbf{C}_p, B) \wedge S \text{ performed a } p\text{-related action in } w_i))$

Based on 2, we can now define \mathbf{W}_p as the collection of all possible worlds that **S** can access from the actual world, according to the knowledge-action accessibility relation that takes into account **S**'s knowledge of **p**, **S**'s cognitive profile, and **S**'s background knowledge. Formally, we get:

3. $\mathbf{W}_p = \{w_i \in W | (w_0, w_i) \in \text{KAAR}(S, p, \mathbf{C}_p, B)\}$ (where **W** stands, of course, for the set of all possible worlds).

We can further refine our framework by incorporating the lessons we learned in Sections 3.1 and 3.3, namely that the fact that an epistemic subject is aware of some particular ways in which her knowledge that **p** may be put to work, does not mean that she can actually perform those actions. What one knows about what could in principle be done (based on **p**, \mathbf{C}_p , and **B**) and what one can actually do are quite different things. What this tells us is that not every member of \mathbf{W}_p which is accessible to **S** is in **S**'s reach, i.e. is something that can be performed by **S** at w_0 . We can use this to enrich our framework with an ability function \mathcal{A} , which maps the worlds in \mathbf{W}_p and the set of epistemic subjects to subsets of \mathbf{W}_p . This is meant to open up a potential pathway for the epistemic subjects to actualize possible worlds by performing at the actual world whatever **p**-related action is characteristic of these possible worlds based on the epistemic subjects' choices and abilities.

If we take Σ to represent the set of epistemic subjects, and $\mathbf{P}(\mathbf{W}_p)$ to denote the power set of \mathbf{W}_p , then we can define \mathcal{A} as

4. $\mathcal{A} : W_p \times \Sigma \rightarrow \mathbf{P}(W_p)$.

We can use 4 to give the following definition of actualizable worlds (\mathbf{W}_a), i.e. the set of possible worlds whose characteristic **p**-related action **S** can actually perform at w_0 , given their abilities:

5. $W_a = \mathcal{A}(w_0, S) = \{w' \in W_p | S \text{ is able to perform } \varphi_{w'} \text{ at } w_0\}$ (where by $\varphi_{w'}$ what is meant is whatever **p**-related action is characteristic of w')

We can use all these to express the relationship between the worlds accessible to **S** based on their knowledge and those that **S** can actually bring into reality through their actions (where Φ' is used as a generic term for a **p**-related action that **S** can actually perform):

$$6. \forall S, \forall p, \forall \varphi', w_0 \models K(S, p) \wedge \varphi' \Rightarrow \exists W_p \exists W_a (W_p \wedge W_a = \mathcal{A}(w_0, S) \subseteq W_p)$$

Put simply, what 6 tells us is that for every **S** that knows that **p** and can perform some **p**-related action, there is a collection of possible worlds W_p associated with that knowledge and an actualizable subset W_a that is contained within W_p .

With all these in place, we can now express \mathbb{P} (i.e. power, as defined in Section 1) as a measure that quantifies an epistemic subject's efficacy in putting her knowledge to work in order to interact successfully with the outside world (i.e. to perform contextually⁸ relevant/useful actions). If we take \mathfrak{d} to stand for the set of context-specific or situational demands, $\mathcal{T}_{\mathfrak{d}}$ to represent the subject-specific⁹ contextually relevant **p**-related actions possible (i.e. $\mathcal{T}_{\mathfrak{d}} = \{w \in W \mid \varphi_w \text{ meets } \mathfrak{d} \text{ for } S \text{ at } w_0\}$), and $[0, 1]$ the normalized scale measuring the degree of advantage taking or benefit derivation, we have:

$$7. \mathbb{P} : \Sigma \times \mathcal{P}(W) \times \mathfrak{d} \rightarrow [0, 1]$$

Of course, 7 provides only the 'signature' of the power function without specifying how to calculate its value. In order to fill this gap, we can use the formula below which is meant to capture the efficacy with which an epistemic subject translates potential into actual, useful actions by computing power as the ratio of the number of actualizable worlds (where the **p**-related action meets the context's demands) to the total number of contextually relevant **p**-related actions possible:

$$8. \mathbb{P}(S, W, W_a, \mathfrak{d}) = \frac{|\{w' \in W_a \mid \varphi_{w'} \text{ meets the requirements of } \mathfrak{d}\}|}{|\mathcal{T}_{\mathfrak{d}}|}$$

What we have done until this point is to try to utilize all the lessons learned in the previous sections about the power of knowledge in constructing a formal framework. With this framework now in place, we are better positioned to extend our understanding further. One important point that should be a lot clearer if we are using this framework is why having knowledge cannot be equated with having power (if power is to be understood along the lines of \mathbb{P}): because possessing knowledge provides the epistemic subject, at the same time, with a lot more (i.e. access to the set of W_p worlds) and a lot less (i.e. nothing really in terms of the special skills required for actualising those worlds whose characteristic **p**-related actions meet the requirements of \mathfrak{d}) than what is required by \mathbb{P} .

Another important point that becomes clearer with this framework is that one's knowledge does not necessarily cause an increase in power (as measured by \mathbb{P}) in every instance. One way in which an item of knowledge may fail to affect \mathbb{P} is if the intersection between the set of all contextually relevant **p**-related actions $\mathcal{T}_{\mathfrak{d}}$ and the set of all possible worlds associated with **p** is empty, i.e. if:

⁸The assumption that stands behind this formulation of \mathbb{P} is that there are no universally useful actions, i.e. the utility of an action is always relative to a particular context. If this is the case, then there have to be factors that are specific to each context and which determine whether or not an action is useful/relevant, and so which are crucial in judging an epistemic subject's efficacy in interacting with the world. Utility/relevancy is, of course, a matter of degree, but for our discussion it makes no sense to further complicate the framework developed here by also taking this aspect into account.

⁹What the 'subject-specificity' is meant to do in this context is to restrict the number of possible contexts only to those that an epistemic subject may find herself in during her life. Crucially, subject-specificity should not be confused with 'subjectivity'. In Section 5, I will introduce the idea of subjectivity, which relates to $\mathcal{T}'_{\mathfrak{d}}$. While $\mathcal{T}_{\mathfrak{d}}$ represents an objective set of all potentially relevant actions (whether **S** knows them or not), $\mathcal{T}'_{\mathfrak{d}}$ captures an epistemic subject's understanding about the total number of contextually relevant actions she can perform.

$$9. \mathcal{T}_d \cap W_p = \emptyset$$

In this kind of context, one's knowledge, given C_p and B , is not enough to give her access to a set of worlds in which there are p -related actions that also meet some contextually specific demands.

A different situation in which although one has knowledge one does not also experience a change in \mathbb{P} is if the following happens to hold:

$$10. \mathcal{T}_d \cap W_p \neq \emptyset \text{ but } \mathcal{T}_d \cap W_a = \emptyset$$

In this situation what happens is that although one may be aware, based on p , C_p , and B , about a set of possible contextually relevant actions that she may perform, it is not in her power to perform those actions because she lacks the necessary abilities to do it.

But if knowledge is not power (in a strong, practicalist sense) and it does not even cause power (in the sense of always¹⁰ leading to an increase in the value of \mathbb{P}), what should we make of the intuition behind **PK** (see its statement in Section 1) that having knowledge about something gives one an advantage over those who lack it? It all depends on what we understand by 'advantage' in this context. If we take 'advantage' to have to do with generating a guaranteed boost in the value of \mathbb{P} , then, provided what we said above about the situations expressed by propositions 9 and 10 is true, then the intuition is erroneous. But I believe there is a different, better, way of thinking about the advantage that having or gaining more knowledge can provide one with (in comparison to those that lack it), namely in probabilistic terms. From this perspective, the advantage that having or acquiring a new item of knowledge provides any epistemic subject with can be taken to amount to an increase in the likelihood that the value of \mathbb{P} will be greater for S with that item of knowledge than it would have been in its absence. Let's substantiate this.

According to the framework developed here, what coming to know that p is the case does to an epistemic subject is, first, to open up for her the space of possibilities by providing her with (limited) access (given her background knowledge and cognitive profile) to the set of W_p worlds specific to that item of knowledge. Given her particular set of abilities, our epistemic subject will be able to perform only a subset of the possible p -related actions characteristic of W_p . We used W_a to refer to this subset of possible worlds. But, since exerting dominion over the outside world (or gaining some practical benefit from one's knowledge) cannot be equated with performing whatever actions one can perform irrespective of the context one finds oneself in, we cannot use the cardinality of W_a as a measure of one's power. What we need, then, for measuring \mathbb{P} for a particular epistemic subject is to determine how much of what she can do also satisfies a particular set of context-specific or situational demands. For example, our subject may be able to cut an apple into extremely thin slices, but she may never find herself in the context in which performing such an activity would make sense for her (by, e.g. generating some sort of benefit for her or anyone else). So, we cannot take this particular activity to improve in any way our subject's efficacy in interacting with the world.¹¹ If we take this point into account, then we realize that it is always possible for one's knowledge

¹⁰Remember that we decided from the beginning that we are not taking into account in this discussion knowledge that is in principle useless, such as, for instance, one's knowledge that there are n grains of sand on a beach.

¹¹Just to be clear, we are not talking here about how useful having a particular ability may prove to be for a subject. What we are concerned with is if a p -related action may ever be useful for a particular epistemic subject given all the contexts that she may find herself in during her life.

that **p** to lead to situations such as the ones expressed at points 9 and 10. So, we can confidently claim that it is not necessary that if **S** knows that **p**, then **S** will experience an increase in her power:

$$11. \neg \Box (K(S, p) \rightarrow \text{increase in the value of } \mathbb{P})$$

However, knowing that **p** does *necessarily* increase the probability that, for a **p**-related action that **S** can perform (and knows about) and a particular context in which **S** may find herself, that action will satisfy the requirements of the context and, as a result, will increase the value of \mathbb{P} for **S**. The reason why this is the case is easy to understand with the help of the framework developed here. It all has to do with the fact that knowledge expands an epistemic subject's perspective on what can be done (i.e. it provides access to some set W_p) and, by doing this, it also increases the likelihood that the subject will identify actions within W_p that she is both capable of performing (i.e. that do not fall outside the scope of what she can do based on her specific set of abilities) and that are contextually relevant and conducive to practical efficacy (i.e. given $K(S, p)$, it is likely that $T_d \cap W_a \neq \emptyset$). Consequently, given our definition of power, an epistemic subject's efficacy in interacting with the world, as measured by \mathbb{P} , is more likely to be greater when **S** knows that **p** compared to when **S** does not know that **p**.

If we use the standard notation for the conditional probability of an event ($P(A|B)$), then we have:

$$12. \forall S, \forall p, \forall d (P(\exists w' (w' \in W_a \wedge \varphi_{w'} \text{ meets the requirements of } d) | K(S, p)) \geq P(\exists w' (w' \in W_a \wedge \varphi_{w'} \text{ meets the requirements of } d) | \neg K(S, p))) \quad (\text{we can take } W_a \text{ to stand for a set of actualizable worlds that } S \text{ can access without knowing that } p)$$

But, since the event of interest in 12 appears as the numerator of 8, we can take 12 to imply the following thesis (let's call it the power of knowledge thesis or **PoK**):

$$13. \text{PoK} \equiv \Box \forall S, \forall p (P(\text{increase in the value of } \mathbb{P} | K(S, p)) > P(\text{increase in the value of } \mathbb{P} | \neg K(S, p)))$$

I take 13 to be the best way of expressing the intuitive idea behind **PK**, i.e. that knowledge always provides an epistemic subject with an advantage (albeit only probabilistic).

Before moving on to the next section, it is important to reiterate that, while this view retains from Hetherington's practicalism the idea that knowledge is associated with a range of potential manifestations, it diverges in a crucial respect: I do not take these manifestations to constitute knowledge itself, nor do I think they must be seen as anything more than possible actions the subject is aware of and might be disposed to attempt under the right conditions.

5. Knowledge and the feeling of power

The previous section was dedicated to achieving one of the main objectives of this paper, i.e. providing an explication for the widespread intuition that the possession of knowledge has an empowering effect on the possessor. This section will be devoted to the second main objective, namely explaining why, in some situations (such as the climate change issue), not only is an increase in knowledge not associated with an increase in power for the epistemic subject, but the more one knows, the more powerless one may end up feeling.

To explain why this may be the case, let us consider for a moment a very important result that we can derive by using the framework developed in Section 4, namely that evaluating one's power is always a relative matter because it depends on the size of the set of \mathbf{W}_p worlds that an epistemic subject can access based on her knowledge and cognitive profile. The reason for this has to do with the fact that **KAAR** restricts \mathbf{S} 's access to a subset of W^{12} ($W_p \subset W$). Given this, any epistemic subject's idea about the total number of contextually relevant actions that she can perform is restricted to the portion of \mathcal{T}_d that lies within \mathbf{W}_p (i.e. to $\mathcal{T}'_d = \mathcal{T}_d \cap W_p$).

Thus, when assessing our own or someone else's power – rather than using formula (8) – we are actually using the following formula:

$$14. \mathbb{P}'(S, W_p, W_a, d) = \frac{|\{w' \in W_a | \varphi_{w'} \text{ meets the requirements of } d\}|}{|\{w \in W_p | \varphi_w \text{ meets the requirements of } d\}|} = \frac{|\mathcal{T}'_d \cap W_a|}{|\mathcal{T}'_d|}$$

Unlike \mathbb{P} , \mathbb{P}' is only a measure of an epistemic subject's feeling of power. But, if we can only use \mathbb{P}' (because of **KAAR**), not \mathbb{P} , when assessing power, then the following two situations may occur. The first one has to do with the false impression of powerfulness. Since our idea about \mathcal{T}_d depends entirely on the \mathbf{W}_p worlds that we can access, and since it makes perfect sense that for an epistemic subject $\mathcal{T}'_d \approx W_a$, then, if 14 is indeed the way in which we are usually assessing power, $\mathbb{P}' \approx 1$ when in fact \mathbb{P} may be closer to 0 than 1 for that subject. Thus, depending on her cognitive profile (\mathbf{C}_p) and background knowledge (\mathbf{B}), an epistemic subject may be under the false impression that, for a particular item of knowledge, she can make the most in terms of using it to impose her will upon the world. So, knowledge can give one a false sense of power. To illustrate this, let us consider once more someone who has a basket of apples but, this time, let us assume that this person knows only two ways to use them: eating them raw or making apple pie – both of which she is fully capable of doing. Given her limited perspective, this person naturally assumes that she is making the most of her resources. However, in reality, she is merely scratching the surface of what can be achieved.

At the opposite pole, knowledge can also make one feel powerless. If one's window into the size of \mathcal{T}_d is dependent on what is revealed by \mathbf{W}_p , then the more one knows (the greater one's \mathbf{B} is), the more likely it is that \mathbf{W}_p covers a larger portion of \mathcal{T}_d , although the same may not be the case with the relationship between W_a and \mathcal{T}_d . Consequently, there can be situations in which, by gaining more knowledge (i.e. by increasing \mathbf{B}), \mathcal{T}'_d may increase while $\mathcal{T}'_d \cap W_a$ may decrease. In these contexts, an epistemic subject's feeling of power, as measured by \mathbb{P}' , will also decrease.

We can use the case of climate change to exemplify this type of situation. If we take \mathbf{p} to stand for the proposition that 'our planet's climate is changing', then \mathbf{W}_p will serve as our window into the range of potential actions that could be taken in response to climate change. As our knowledge about climate change deepens, \mathbf{W}_p expands, revealing an ever-growing array of possible responses. However, given the scale of the problem, most if not all of these actions require coordinated, large-scale efforts. Consequently, although \mathbf{W}_p grows in this case and so does \mathcal{T}'_d , the scale of the problem and the collaborative nature of the solutions cause $\mathcal{T}'_d \cap W_a$ to diminish, and thus, according to \mathbb{P}' , our feeling of power decreases. In other words, as we accumulate more climate knowledge, our awareness of the need for collective action and institutional intervention intensifies, which can diminish our personal sense of power.

¹² \mathbf{W} stands here for the exhaustive set of \mathbf{p} -related possible worlds, i.e., for the set of possible worlds characterized by the fact that all possible \mathbf{p} -related actions, irrespective of what a particular epistemic subject can conceive of, are exhibited within them.

6. Conclusions

This paper had two main aims. The first was to provide an explication of the widespread intuition that the possession of knowledge has (at least sometimes) an empowering effect on the possessor. This objective was achieved in Section 4, where it was argued that the primary advantage we can rely on knowledge to always generate for an epistemic subject is probabilistic: knowing about a particular fact or situation can empower an individual by expanding her perspective on what can be done and, in so doing, by enhancing the likelihood that she will identify and perform actions that are both feasible given her abilities and contextually relevant. This, as argued in the paper, is how the intuition that knowledge confers an advantage for an epistemic subject (that is, that knowledge comes with some sort of power) should be understood. A crucial role in achieving this first aim was played by the modal-logic-inspired formal framework developed in Section 4.

The second objective of this paper was to provide an explanation for why it is the case that in some situations (such as with the climate change issue) gaining more knowledge, instead of providing one with more power, can instead make one feel rather powerless. Section 5 was devoted to using the framework developed in Section 4 to show when and why this may be the case. As argued there, the assessment of one's power is always relative because it is unavoidably based on the accessible space of possibilities. As knowledge increases, so does the awareness of the potential actions one could perform. However, if the set of actions that the epistemic subject is both capable of performing and that is contextually relevant does not grow proportionally, her perception of power may actually decrease.

Acknowledgements. I am grateful to an anonymous reviewer for their insightful comments, which have significantly improved this paper.

Funding. This work has received funding from the European Union's Horizon Europe research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 101067782. The views and opinions expressed are those of the author only and do not necessarily reflect those of the European Union or the European Research Executive Agency (REA). Neither the European Union nor the granting authority can be held responsible for them.

References

- Adolf M. and Stehr N. (2016). *Knowledge: Is Knowledge Power?* London: Routledge.
- Andre P., Boneva T., Chopra F. and Falk A. (2024). 'Globally Representative Evidence on the Actual and Perceived Support for Climate Action.' *Nature Climate Change* 14(3), 253–59.
- Foucault M. (2007). 'Discipline and Punish: The Birth of the Prison.' In B.B. Lawrence and A. Karim (eds), *On Violence: A Reader*, pp. 445–71. New York: Duke University Press.
- Foucault M. and Gordon C. (1980). *Power/Knowledge: Selected Interviews and Other Writings 1972–1977*. New York: Pantheon Books.
- Gaukroger S. (2001). *Francis Bacon and the Transformation of Early-Modern Philosophy*. Cambridge: Cambridge University Press.
- Hetherington S. (2006). 'How to Know (that Knowledge-that is Knowledge-How).' In S. Hetherington (ed), *Epistemology Futures*, pp. 71–94. Oxford: Oxford University Press.
- Hetherington S. (2011a). 'Knowledge and Knowing: Ability and Manifestation.' In S. Tolksdorf (ed), *Conceptions of Knowledge*, pp. 73–100. Berlin, Boston: De Gruyter.
- Hetherington S. (2011b). *How to Know: A Practicalist Conception of Knowledge*. Oxford: Wiley-Blackwell.
- Hetherington S. (2017). 'Knowledge as Potential for Action.' *European Journal of Pragmatism and American Philosophy* IX(2). <https://doi.org/10.4000/ejpp.1070>
- Hetherington S. (2020). 'Knowledge as Skill.' In E. Fridland and C. Pavese (eds), *The Routledge Handbook of Philosophy of Skill and Expertise*, pp. 168–78. London: Routledge.
- Hetherington S. (2022). 'Knowledge-Practicalism.' In S.F. Aikin and R.B. Talisse (eds), *The Routledge Companion to Pragmatism*, pp. 252–63. New York: Routledge.

Hyman J. (1999). 'How Knowledge Works.' *The Philosophical Quarterly* **49**(197), 433–51.

Hyman J. (2006). 'Knowledge and Evidence.' *Mind* **115**(460), 891–916.

Hyman J. (2015). *Action, Knowledge, and Will*. Oxford: Oxford University Press.

Gabriel Târziu is a postdoctoral fellow at the Munich Center for Mathematical Philosophy (LMU Munich). His research explores the epistemological particularities of scientific understanding in climate science, with special attention to the role of model hierarchies. He works broadly in epistemology, the philosophy of science, and the philosophy of mathematics, with wide-ranging interests in scientific explanation and understanding, the applicability of mathematics in science, and ontological debates in mathematical philosophy. Before joining the MCMP in 2022, he held research fellowships in Munich, Brussels, Essen, and Ghent and was a postdoctoral researcher at the Romanian Academy.