

# Is a subpersonal virtue epistemology possible?

Hadeel Naeem

## Abstract

Virtue reliabilists claim that knowledge requires responsibly employing one's reliable belief-forming process. Responsible employment requires that the agent is either aware that her process is reliable or is sensitive to her process's reliability in some way. Recent philosophy of mind literature proposes that in some cases a cognitive mechanism, i.e. precision estimation, can ensure that a belief-forming process is only employed as long as it's reliable. This means that epistemic responsibility can sometimes be explained entirely on the subpersonal level. In this paper, I argue that the mechanism of precision estimation—the alleged new variety of epistemic responsibility—is more ubiquitously present than epistemic responsibility. I show that precision estimation operates at levels that have nothing to do with the epistemic domain. Lastly, I explain how all subpersonal epistemologies are likely to fall prey to worries like *the problem of demarcating the cognitive agency* and *the problem of attributing beliefs*.

## Introduction

For virtue reliabilists, knowledge requires that an agent ought to *responsibly* employ a *reliable* belief-forming process (Greco and Sosa 1999; Greco 2003, 2010; Pritchard 2007, 2008, 2010, 2016, 2018a; Palermos 2011, 2014; J. Adam Carter 2018).<sup>1</sup> A reliable belief-forming process generates far more true beliefs than false and to responsibly employ such a process requires that we connect to its reliability in a particular way (Pritchard 2010, 2018a; Palermos 2011, 2014). Simply put, such a connection (or epistemic responsibility) will look like employing it when it's reliable and stop employing it when it turns unreliable.

One way we can responsibly employ a reliable process and connect with its reliability

---

<sup>1</sup>Responsibly employing a reliable process is more commonly called cognitive ability. And, knowledge requires some sort of cognitive ability. See Sosa (1988); Sosa and Plantinga (1993); and Greco (2007) for more on cognitive ability.

is by becoming aware that it is reliable (Pritchard 2010, 2018a).<sup>2</sup> For example, Lele has a magic ball that always gives correct answers about the whereabouts of the president. She reads the manual that came with the magic ball and finds that it's 100% reliable. Now she can responsibly form true beliefs by looking at her ball. However, becoming aware that her process is reliable every time she employs it is a very strict condition for epistemic responsibility. It also does not depict how we employ our internal neural faculties, such as perception, memory, etc.

A weaker<sup>3</sup> condition for epistemic responsibility can be formulated in terms of *cognitive integration* (Greco 2010; Palermos 2011, 2014). The main idea is that we can responsibly employ all belief-forming processes that integrate into our cognitive systems. What makes this account weaker or more minimal is that we don't need access to the reliability of the target process for it to integrate into our cognitive system. As long as the target process sufficiently interacts and cooperates with other processes in the agent's cognitive system, it can integrate into her cognitive system. When a reliable process is frequently employed over a significant period of time, it forms dense and complex connections with other processes and beliefs. The interconnectedness of the processes in our cognitive system helps monitor the reliability of the target process. If an integrated process becomes unreliable, the belief it produces will most likely not cohere and fit with other beliefs and processes of our cognitive system. We will receive meta-cognitive cues indicating a disturbance in the cognitive system. On this account, the agent will become aware when something has gone wrong with her integrated process's reliability.

Recent philosophy of mind literature allegedly proposes an even more minimal account of epistemic responsibility. Andy Clark (2015) argues that, in some cases,<sup>4</sup> a cognitive mechanism can ensure that a process is only employed as long as it's reliable (Clark 2015). In those cases and when this mechanism is operating properly, we can *responsibly* employ our reliable belief-forming process (and generate knowledge). So, (in some cases) we neither need access to our process's reliability nor do we need to be in a position where we can become aware when our process's reliability fails, and still we can responsibly employ our process. This new proposal is based on the predictive processing model

---

<sup>2</sup>See Greco (2010) for a discussion on our need for a perspective on the reliability of our belief-forming process. He discusses how philosophers have tried to debate the skeptics who argue that without a perspective on how one's forming true belief knowledge is not possible. And, demanding a perspective also leads to infinite regress, circularity or arbitrary assumption.

<sup>3</sup>The word weak only refers to the conditions required for this kind of integration. So, the integration isn't weak in any way, but the conditions for it are not very demanding (hence, weak).

<sup>4</sup>Note that Clark does not take issue with the strong and weak cognitive integration accounts as such. He agrees that stringent epistemic hygiene practices are necessary in some cases. But, for some cases, like employing ones trusted internal cognitive faculties (and extended cognitive faculties), a much weaker notion of epistemic responsibility can be sufficient (Clark 2015).

of the brain, and the mechanism that is (in some cases) purportedly sufficient for epistemic responsibility is called *precision estimation* (Clark 2015).

In this paper, I argue that the new proposal is inadequate and unconvincing as an account of epistemic responsibility. Precision estimation, the purported new variety of epistemic responsibility, is present in cases that have nothing to do with the epistemic domain. I also explain how subpersonal epistemologies like Clark's lead to issues like the *problem of demarcating the cognitive agency* and the *problem of attributing beliefs*.

The structure of this paper is as follows. Section 1 outlines what epistemic responsibility is and what the condition for it are. Here I also describe the most minimal requirements for epistemic responsibility (i.e. weak cognitive integration) according to virtue reliabilism. In section 2, I introduce the new proposal that claims to be an even more minimal requirement for epistemic responsibility. I discuss how it's different from weak cognitive integration. Section 2.2 locates the new proposal—subpersonal epistemic responsibility—in the predictive theory of the brain and unpacks the precision estimation function that captures the new variety of epistemic responsibility. In section 2.3, I argue that precision estimation cannot adequately describe epistemic responsibility because it's present on all levels of the brain's functional hierarchy, even on the ones that have nothing to do with forming beliefs. In section 3, I briefly describe other issues that all subpersonal epistemologies can run into, namely the problem of demarcating the cognitive agency and the problem of attributing beliefs.

Before I begin, there is a preliminary detail that I would like to discuss. The setting of this debate is the extended cognition and more specifically the extended knowledge debate. The extended cognition thesis proposes that we sometimes use external resources in a way that extends our cognitive processes outside our skin and skull boundary (Clark and Chalmers 1998; Clark 2001, 2003, 2008; Rowlands 2010; Rupert 2009; Menary 2010; Menary 2013; Drayson 2010; Wheeler 2010; Colombetti and Roberts 2015; J. Adam Carter et al. 2018). The extended knowledge thesis argues that sometimes these extended cognitive processes are belief-forming and can therefore generate knowledge that should be regarded as our extended knowledge Vaesen (2011). Clark (2015) suggests that the way extended knowledge is outlined presents a dilemma. Cognitive extension requires that there is no agential involvement, while knowledge (according to virtue epistemology) demands that one's cognitive agency is relevantly and significantly manifested. So, Clark (2015) puts forward his revised thesis of extended knowledge—knowledge generated with the most minimal agential involvement. He calls it subpersonal virtue epistemology. While this is a very interesting setting and the birthplace of the debate this paper is about, for this paper I want to evaluate how Clark's proposal fares as an epistemology.

# 1 How we responsibly employ reliable belief-forming processes

## 1.1 Epistemic responsibility

Temp has a thermometer in her room that works strangely but reliably.<sup>5</sup> When her thermometer reads 21 degrees Celsius, she forms the belief that it is 21 degrees. Unbeknownst to Temp, her thermometer is broken. However, an invisible genie adjusts the thermostat in the room to match the thermometer reading. When she reads 21 degrees on the thermometer, the genie ensures that the room is 21 degrees by adjusting the thermostat. When the thermometer display changes, the genie modifies the thermostat such that the temperature reading is true. Temp, oblivious to the genie, has a reliable belief-forming process and forms true temperature beliefs. However, does Temp have knowledge that the temperature of the room is 21 degrees?

Epistemologists, almost unanimously, agree that Temp does not have knowledge. One major reason for this is that Temp does not seem to have any connection to her reliable process (Pritchard 2016). Her process's reliability seems accidental and strange. Tomorrow, if the genie took a break and didn't adjust the thermostat, she would believe the reading on her broken thermometer, and she would be wrong.

Virtue reliabilists maintain that knowledge is a kind of virtue that we achieve when we manifest relevant and significant cognitive agency. When we luckily or accidentally land a true belief, we don't manifest relevant agency for the true belief to be creditable to us.<sup>6</sup> While Temp has a reliable process and is forming true beliefs, she has no connection to the reliability of her process. Her beliefs are true because of sheer luck and not because she manifested any relevant agency to form these true beliefs. Hence, Temp does not have knowledge.

Put another way, knowledge requires that an agent responsibly employ her reliable belief-forming process. While Temp has a reliable process, she does not have knowledge because she is not employing her process with *epistemic responsibility*.<sup>7</sup> To responsibly employ her process, she must avoid being in a place where she continues to employ it when it has turned unreliable. In other words, Temp should have a kind of link or some

---

<sup>5</sup>The Temp case is an adaptation of Lehrer (1990)'s Mr Truetemp.

<sup>6</sup>See Pritchard (2018b) for a discussion on evidential environmental luck that is compatible with knowledge.

<sup>7</sup>This paper is specifically concerned with the virtue reliabilist notion of epistemic responsibility. That is, how an epistemic agent can responsibly employ her reliable belief-forming to generate knowledge. Following are some other ways the concept of epistemic responsibility can be found in epistemology literature. Virtue responsibilists argue that epistemic responsibility is cultivated by virtuous character traits. Some epistemologists investigate what epistemically responsible beliefs are and how we form them (Peels 2017). Others inquire if and how justification is related to forming responsible beliefs (Kornblith 1983).

sensitivity<sup>8</sup> to the reliability of her process such that she employs it when it's reliable, keeps employing it as long as it's reliable, and quickly stops relying on it when it turns unreliable. If and when Temp establishes such a connection with her process, she can manifest relevant and sufficient cognitive agency for her true belief to be creditable to her.

One way for Temp to responsibly employ her process is to become aware that her process is reliable. Suppose the genie reveals himself to Temp and explains how he will adjust the thermostat to the next reading on her thermometer. In finding out the source of her process's reliability, Temp has established a connection to it. Temp's cognitive success can now be creditable to her because she is forming it responsibly.<sup>9</sup>

However, reflective access to the reliability of one's process is not necessary for responsibly employing it.<sup>10</sup> Consider how effortlessly and fluently we rely on our perceptual or memorial faculties without always first wondering or confirming that they are reliable (Greco 2010, 142). This indicates another route to epistemic responsibility.

## 1.2 Weak cognitive integration

Suppose Temp has relied on her genie-mediated process for 10 years and uses it approximately twice a day. Even if Temp doesn't become directly aware that her process is reliable, she should grow sensitive to her process's reliability (if it is reliable, that is). Temp has employed her process for a very long time and has formed several beliefs with it, and those beliefs have cohered with several other beliefs that Temp has. For instance, when she reads 40 degrees on her thermometer, she feels hot and forms more beliefs about how she needs a fan in the room or that opening the windows will feel nice. When the reading on the thermometer is 12 degrees, Temp feels cold and wants to switch on the heater. She has used her process frequently and for a long time so that the said process has become interconnected with other processes in her system. And, the beliefs formed by this new process have cohered well with other beliefs in her system.<sup>11</sup>

---

<sup>8</sup>The word sensitivity here is not meant to indicate the sensitivity principle which claims that S knows p only if S would not believe p if it was false (Nozick 1981). I only mean to indicate a kind of responsiveness to one's reliability.

<sup>9</sup>Pritchard (2012) cautions us to use knowledge in this case. Temp does not necessarily meet conditions for knowledge by simply meeting the ability condition. There is another intuition on knowledge—the anti-luck intuition. See (Pritchard 2014, 2017, and 2018b) for more on anti-luck virtue epistemology.

<sup>10</sup>For arguments in favour of a strong epistemic responsibility criteria, see Chisholm (1977), Bonjour (1985), and Steup (1999). Also, see Greco and Sosa (1999) on why such a strong concept of epistemic responsibility is vulnerable to skepticism.

<sup>11</sup>This variation of the Temp case is borrowed from Pritchard (2010). He asks to imagine 'Tempo,' who has been fitted with a smart device that records ambient temperatures. He is brought up in a culture where relying on such a device is very common. We wouldn't demand that Tempo be aware of the reliability of his process in this case, just like we wouldn't demand that children have meta-beliefs about the reliability of their

The coherence and interconnectedness of Temp's cognitive system allow her to monitor the target process's reliability in the background. She may not be directly aware that her process is reliable but the interconnectedness of her system puts her in a position to become aware if something is off about the reliability of her target process. For instance, if the thermometer were to read 15 degrees and she was feeling extremely warm, she would notice that something isn't right.

The phenomenon described above is called cognitive integration (Greco and Sosa 1999; Greco 2003, 2010; Palermos 2014, 2020; D. Breyer and Greco 2008; D. S. Breyer 2013). It is the 'function of cooperation and interaction with other aspects of the cognitive system' (Greco 2010, 152). For many virtue reliabilists, cognitive integration explains how we seamlessly employ our internal faculties without worrying if they are reliable. They are integrated into our cognitive systems. A new process can also integrate into the agent's cognitive system if it 'forms a track record of beliefs' that both cohere with other beliefs and interact with other processes of the system (Pritchard 2010, 148).<sup>12</sup> An interconnected cognitive system allows the agent to be counterfactually sensitive to a target process's reliability. Strong cognitive integration requires that an agent become aware that her process is reliable (Pritchard 2010). In the absence of such awareness, *weak cognitive integration* can still come about when the target process is employed frequently over a period and forms dense and deep interconnections with the agent's cognitive system.

For a process to integrate, there should be a *significant* degree of interaction and cooperation with the agent's system (Greco 2010). This ensures meeting the epistemic responsibility requirement. How? Strong and deep interconnections allow the agent to monitor (in the background) the reliability of the target process. When we employ the target process frequently over a period of time, processes and beliefs that are already a part of our cognitive system become familiar with the beliefs of the target process, so to speak. If the target process were to become unreliable, i.e. form several false beliefs, it would stop cohering and interconnecting with the beliefs and processes in the agent's cognitive system. For instance, if Temp reads 50 degrees Celsius on her thermometer but she feels chilly, she'll notice that something is amiss. Had she employed her process few and far between, she might not have noticed that 50 degrees Celsius is not supposed to feel cold (she might not have many beliefs or reasons to counter the false belief). Strong interconnections are required for the agent to become counterfactually sensitive to her process's reliability, which

---

memory and perceptual faculties. And yet, we do credit children with knowledge of most of the things they see or the things they remember. This indicates that there is a weaker requirement that may capture epistemic responsibility (Pritchard 2010, 146).

<sup>12</sup>According to Greco (1999, 287), our cognitive characters are not only made of our internal neural faculties that make our central nervous system. Skills of perception, methods of inquiry, advanced technologies, etc. can also become a part of our cognitive characters by integrating properly.

in turn makes sure that the agent will responsibly employ it. In other words, the epistemic agent will employ her process when it's reliable, as long as it's reliable, and stop doing so when it turns unreliable.

A key aspect of weak cognitive integration, relevant to the issue of this paper, is that it puts the agent in a position to become aware of the integrated process when it turns unreliable. One way to put it is that integration allows agents the power to monitor the fluency of their cognitive systems (Proust 2014). When a process is no longer reliable, the agent receives met-cognitive cues against her system's fluency (Proust 2014).<sup>13</sup> When Temp, for instance, reads 50 degrees on the thermometer but feels chilly, she becomes aware that something is off.

## 2 The new variety and its problems

### 2.1 Subpersonal epistemic responsibility

Recently, Andy Clark has argued that weak cognitive integration is not weak enough to explain how we employ our neural processes (Clark 2015). Our cognitive faculties are employed so seamlessly that we don't become aware when they stop working reliably. Say you are in a room, and it suddenly turns very dark. You don't stop to reflect on some process being unreliable; you automatically and quickly start to rely on auditory and tactile cues to get around. Clark presents a new variety of epistemic responsibility, one that is weaker than weak cognitive integration. To evaluate it, let's first unpack it.

Temp2 is in the same situation as Temp. However, one day, her genie begins to take Wednesdays off. So from that Wednesday onwards, Temp2's belief-forming mechanism works fine on all days except Wednesdays. After some time, Temp2's cognitive system begins to associate Wednesdays with false readings on the thermometer, but Temp's not aware of these associations. Some Wednesdays later, Temp2 begins to simply ignore the thermometer readings on Wednesdays. On these days, even when she looks at the thermometer, she fails to register or attend to the readings (Clark 2015, 3768–69).

Clark claims that Temp2 can exhibit epistemic responsibility without ever becoming conscious of her process's failed reliability. She can simply learn to ignore her process on Wednesdays. He writes, 'Importantly, Temp2 can display this kind of epistemic virtue even if she is never consciously aware of (for example) her Wednesday-distrust of the thermometer (Clark 2015, 3769).

---

<sup>13</sup>Palermos (2021) mentions how these meta-cognitive cues are produced when parts of our cognitive system interact in a way that acts a specific kind of defeater. The latter alerts the agent that there might be something wrong with the way have formed their belief. Also, see Palermos (2021) to understand epistemic responsibility in terms of epistemic defeaters and how our cognitive systems react to them.

What makes Temp2 different from Temp is that (1) she doesn't become aware that her process has turned unreliable, and (2) she stops employing the unreliable process after some time. The Temp2 cases assumes that Temp (discussed in the previous section), the weak cognitive integration agent, will always become aware when her process turns unreliable and stop employing it immediately.

Before further unpacking this proposal, I will attempt to argue that (1) and (2) do not distinguish the new account from weak cognitive integration. The following are some scenarios in which Temp, who represents weak cognitive integration, will also exhibit (1) and (2). First, consider that there might be nothing wrong with the reliability of Temp2's process. Recall that reliable belief-forming processes are ones that produce far more true beliefs than false. They don't have to be 100% reliable; they may just be 80% reliable. This means that Temp2's process may still be reliable even if the genie takes a few Wednesdays off. A weak cognitive integration agent, like Temp, will not become aware or immediately abandon her process if her genie takes a few Wednesdays off, as long as her process remains generally reliable.

Secondly, it's possible that Temp's process turned unreliable but not in a very obvious way. For instance, if the actual temperature of the room is 27 degree Celsius, both 25 and 125 degrees are false readings. However, 125 degrees is false in a very obvious way, and it will likely not cohere with the agent's other beliefs (like the belief that water boils at 100 degrees, that the room does not feel extremely hot, that I am still alive at this temperature, etc.). In contrast, 25 degrees is false but not conspicuously so, and there is a good chance that the agent's cognitive system will not have beliefs against (or that defeat) such a trivially different belief. This may be why Temp (an example of weak cognitive integration) doesn't become aware and doesn't immediately stop employing her process.

Thirdly, perhaps when Temp reads the thermometer, she is too distracted to pay proper attention to what it says. If she doesn't focus her attention on the thermometer readings, she will likely not register her system's meta-cognitive cues against its reliability either. Suppose she hasn't had food all day, and she is only thinking about what to cook. And so, when her thermometer reads 125 degrees, she glances at it and forms a belief but not one that she is conscious of.<sup>14</sup>

These three scenarios show that even on the weak cognitive integration account, we shouldn't expect the agent to always (1) become aware and (2) immediately stop employing her process when it first turns unreliable. Even so, these proposals fail to show that the new proposal is the same as the weak cognitive integration account. While these proposals do show that Temp might also not become aware or immediately abandon her process in some cases, they don't explain why Temp would simply start ignoring her process after a while.

---

<sup>14</sup>See Audi (2015) to understand different kinds of believing: occurrent and dispositional.



One of the key aspects of the new proposal, therefore, is that after a few Wednesdays have passed, Temp2 starts to ignore her process which turned unreliable.

Arguably, the Temp2 case meets the commonsensical criteria of epistemic responsibility. She employs her process when it's reliable, keeps recruiting it while it's reliable, and stops employing it (albeit after some time) when it turns unreliable. We just saw with the help of the scenarios above that even Temp (an example of weak cognitive integration) may sometimes not become aware or not immediately stop employing her process after it produces a false belief.

Does this mean that we have a new variety of epistemic responsibility, one that can be satisfied with more minimal criteria than weak cognitive integration?<sup>15</sup> The case is complex and requires further unpacking at this point. After all, isn't it odd that an agent who has relied on her process for a long time would simply start to ignore it one day without even realising she is doing so? Wouldn't she at least stop to wonder why she doesn't want to read the thermometer in her room anymore?

On Clark's account, Temp2 doesn't become conscious of her process turning unreliable because she wasn't employing it consciously in the first place. For example, if the room we are in were to suddenly turn dark, we would simply begin to focus on our auditory and tactile cues to get around the room, instead of first realising, 'oh, there is something wrong with my perceptual faculty.' This is characteristic of how we recruit our internal faculties, and this is the sort of epistemic responsibility Clark's trying to capture.

If Temp2 isn't employing her target process consciously then the best way to study such an employment is to understand it on the subpersonal level. When a phenomenon is explained on the subpersonal level, it's explained in terms of parts of the person (Dennett 1969; Drayson 2012). These parts may be atoms, molecules, organs, or processes in an agent's cognitive system. We can explain respiration subpersonally in terms of inhalation and exhalation through one's nose and mouth, but it can also be explained subpersonally as the chemical reactions that take place in the mitochondria of our cells that produce ATPs (adenosine triphosphate) from food molecules. Similarly, when I explain how our perceptual faculties acquire wavelength signals from receptors in the eyes and process them to generate information about the colour of the object we are looking at, I am giving a subpersonal level description of colour detection. Contrastingly, if I were to describe colour detection on the personal level, I might say that I am wearing red-tinted glasses and so everything appears red to me. A personal level description describes in terms of the whole person (Dennett 1969; Drayson 2012).

---

<sup>15</sup>In J. Adam Carter and Kallestrup (2020), we learn how perhaps epistemic integration is different from metaphysical integration. The latter allows an external vehicle to become a part of our cognitive structure such that it extends our cognition into the environment. This kind of integration can be an entirely subpersonal phenomenon. Epistemic integration, however, requires more agential involvement from the agent.

Parts of Temp2's cognitive system send information signals and receive necessary signals to employ the target process. It's in this subpersonal employment of the target process that Clark's finds the most minimal representation of epistemic responsibility—subpersonal epistemic responsibility. According to Clark's view, an entirely subpersonal cognitive mechanism can in some cases ensure that one's target process is employed responsibly. This, he says, is how we employ our internal neural faculties most of the time. There is hardly any 'agential scrutiny' in such employment because it occurs entirely at the subpersonal level (Clark 2010). Note, however, that he is not rejecting the weak cognitive integration account or the strong integration account. Clark concurs that more agential scrutiny is required for epistemic responsibility in some cases, but in other cases, there is as little as a subpersonal cognitive mechanism ensuring epistemic responsibility.

It is now clear how the new proposal differs from weak epistemic integration. The latter requires that the agent (on the personal level) is in a special (counterfactual) relation with her integrated reliable process, such that if the process turns unreliable the agent becomes aware that there is a problem. While the subpersonal epistemic responsibility does not require the agent to be in a particular relationship to her process's reliability. A subpersonal cognitive mechanism in the agent's cognitive system takes care of responsibly employing the target process.

One might wonder, at this point, what it means for epistemic responsibility to be met entirely on the subpersonal level when it sounds like all events (like digestion, respiration, etc.) can be explained on both the subpersonal and personal level. In other words, the subpersonal isn't an actual physical level of some sort; it's only a level of explanation. So even if epistemic responsibility is being on the subpersonal level, there must be something happening on Temp2's personal level that can also explain epistemic responsibility.

While the personal and subpersonal are only levels of explanation, some phenomena only make sense on one of these levels. For instance, take awareness. Parts of my cognitive system are in a specific state when I am aware of something so a subpersonal explanation of awareness is possible. However, awareness is essentially a personal level phenomenon. *I* am aware of my surroundings; it's not a collection of neurons in my head that is aware of my surroundings. Like this, in Clark's account, epistemic responsibility is an entirely subpersonal phenomenon.

Similarly, knowledge, like awareness, is also a personal level quality. There is a subpersonal picture of the state my cognitive system is in when I know something. But, it is not parts of me that know something; it is me. And perhaps because of knowledge epistemologists have almost always considered epistemic responsibility also a personal-level phenomenon. To exhibit epistemic responsibility, agents ought to have either the right reasons for their beliefs or an epistemically virtuous character or a reliable belief-forming process

they are sensitive to, etc. All these are personal-level accounts of epistemic responsibility. This makes subpersonal epistemic responsibility a rather radical position in epistemology.

We have established why Clark's account is understood as a fully subpersonal level epistemic responsibility. But, then, why does it matter what happens on Temp2's personal level? If a phenomenon is being managed entirely subpersonally, it shouldn't matter what happens on its personal level. On the personal level, Temp2 may become aware that her process has become unreliable or she may remain unaware and simply start to ignore the thermometer after a few Wednesdays. Both of these as well as many other personal-level scenarios are possible. If epistemic responsibility is explained entirely on the subpersonal level, it shouldn't matter what Temp2 does on the personal level. Yet, Clark ties his subpersonal view to Temp2 not becoming aware when her process turns unreliable. He doesn't (at all) show that the personal level is simply not significant for understanding epistemic responsibility in the Temp2 case. The case misleadingly emphasises that Temp2 doesn't become aware when her process turns unreliable.

## 2.2 Precision estimation and predictive processing

I have discussed what distinguishes the new account of epistemic responsibility from weak epistemic integration. This section describes the second-order function that the subpersonal account finds sufficient (in some cases) for epistemic responsibility: *precision estimation* (Clark 2015, 3766). It's a mechanism that mediates the context-sensitive involvement of various processes. Precision estimation units assign reliability weightings to processes according to the environmental cues and previous information present to the cognitive system (Clark 2015). Processes that receive higher weightings are employed, and others are ignored (so to speak).

Clark's Temp2 responsibly employs her process because her precision estimation function assigns sufficient reliability to the genie-mediated process. But, when the genie begins to take Wednesdays off, Temp2's cognitive system begins to pick up on cues that link false readings with Wednesdays. After a few Wednesdays, precision estimation units begin to assign insufficient reliability weightings (because of environmental cues present to them) to the genie process. Hence, Temp2 begins to ignore her thermometer. To unpack this further, we need to first situate precision estimation in the predictive theory of the brain (Clark 2013, 2015, 2016, 2017; Hohwy 2013, 2014).

As the name goes, predictive processing suggests that brains are predictive machines. Our brains attempt to predict future incoming stimuli, compute the error difference between the prediction and the actual signal, and improve future predictions by minimising the error. You might recall the dress debate that broke the internet in 2015. Some people saw a blue and black dress while others saw a gold and white dress. Research now sug-

gests that people who expected to see the dress in daylight saw a white and gold one, and those who expected to see it in artificial light saw a blue and black dress (Wallisch 2017; Witzel, Racey, and O'Regan 2017). So, people's perception of colour was modulated by the way their brains made predictions.

The best way to understand cognitive processing in a predictive brain is to notice the hierarchical levels. On the lowest levels, receptors receive basic sensory data and transmit it to other parts of the brain. For instance, photoreceptors in the retina receive specific wavelength data and send it to other parts of the brain. More complex processes, such as edge, colour, and shape detection, occur on intermediary levels. Wavelength data from the lower levels are processed into more complex information about the shape, colour of the object, etc. On higher levels, information processing from intermediary levels is used to give rise to more abstract representations, for instance, conceptual distinctions. The well-functioning of the overall hierarchy allows us to form accurate beliefs. Information about wavelengths, edges, hues, sizes, shapes, etc. all come together in my belief that I see a red table, for instance.

Predictive processing is essentially an interplay of two chains (Clark 2015). Hierarchical levels at the top make predictions about the information they are going to receive from the levels below (top to bottom), like the higher levels that process more abstract information make predictions about the shape and colour of the object. Intermediary levels predict information they are going to receive from the lowest levels. These predictions are supposed to perfectly match the information received so the agent perceives accurately. But, when the match is not exact, an error difference, *prediction error*, is generated, and its value is communicated by the lower levels to the higher levels to improve future predictions. In essence, the interplay between the top-down (predictions making) and bottom-up (prediction error relaying) chains helps us make sense of what we perceive (Clark 2015). We are concerned with the role precision estimation plays in this predictive machine.

As Clark says, 'it would not do to always see what we expect'(Clark 2015, 3767). Precision estimation guides our predictions by selecting which processes will get involved in predicting the external stimuli. They assign reliability weightings to the processes in light of environmental cues and prior information. On a foggy day, a driver's precision estimation units will assign more reliability to previous knowledge of the road instead of visual cues (like signboards). While on a sunny day, her precision mechanism will go back to assigning more reliability to visual cues (Clark 2015; J. Adam Carter, Clark, and Palermos 2018).

For the argument I want to make in the next section, the most important aspect of the predictive brain is that precision estimation operates on all its hierarchical levels. On the lower and intermediary levels, precision estimation determines which processes and sections of the brain will participate in receiving and sending wavelength data and processing

colour, shape, edge, etc., respectively. Similarly, these units also determine which higher-level processes will engage with the information from the lower levels. Precision estimation is a cognitive mechanism that operates on all of the cognitive hierarchical levels.

### 2.3 Why precision estimation does not capture epistemic responsibility?

Here I will show that precision estimation is unsuitable and inadequate as an epistemic responsibility account, and in the next section, I will discuss why all subpersonal theories of epistemic responsibility might share the same fate.

Precision estimation is ubiquitous in ways that epistemic responsibility is not. We saw in the previous section that this mechanism is present on all levels of the brain's functional hierarchy. On lower and intermediary levels, it determines which cognitive processes will contribute to processing wavelengths, detecting edges, perceiving depth, etc. So, the predictive brain model itself puts precision estimation on all levels of cognitive processing. Epistemic responsibility, on the other hand, is a concept tied only to epistemic processes, which are cognitive processes aimed at forming true beliefs (Pritchard 2009). They are a subcategory of cognitive processes.

Precision estimation brings about subdoxastic states, and these are importantly different from states in the epistemic domain. When parts of our cognitive systems are in subdoxastic states, they process information about wavelengths, edges, shapes, proprioception, hues, etc. Subdoxastic states and the information produced in these states are not available for introspection. According to Stich, these states 'occur in a variety of separate, special purpose cognitive subsystems'(Stich 1978, 508).<sup>16</sup> On the other hand, doxastic states, such as beliefs and desire, arise when these subsystems further integrate. Doxastic states are consciously accessible to us. Epistemic responsibility, as epistemologists have mostly described it, is what makes a doxastic state like true belief an epistemic state like knowledge. Precision estimation—the subpersonal variety of epistemic responsibility—places epistemic responsibility on the subdoxastic level.

The claim that precision estimation is a type of epistemic responsibility therefore raises the following questions. How does a phenomenon sometimes sufficiently explain epistemic responsibility and, at other times, not even contribute to belief formation? Or, how do norms that govern lower-level processing also govern the level that generates knowledge? And further, what role does a type of epistemic responsibility play at levels that are not involved in producing beliefs? In sum, if precision estimation is a kind of epis-

---

<sup>16</sup>In this paper, Stich (1978) is arguing against Gilbert Harman's (1973) claim that retinal simulations provide inference to our perceptual beliefs. Harman also claims that inference is a relation exclusively amongst beliefs, and so he regards subdoxastic states as kinds of beliefs. Stich (1978) distinguishes between beliefs and subdoxastic states and claims that inference is not a relation exclusively amongst beliefs.

temic responsibility then why is it present on levels that are not epistemic?

Some may argue that since lower-level processing is causally involved in forming beliefs, it is an essential part of the epistemic process. For example, for me to responsibly form the belief that I see a red car in front of me, all information processing on the lower levels ought to be carried out responsibly. That is, the likes of wavelength data processing and shape recognition should all take place responsibly. Hence, it makes sense (for those who argue this way) that precision estimation is present on all levels of the cognitive hierarchy.

While it's true that lower-level processing causally influences higher-level processing, not all lower-level processing contributes to beliefs. Some of it influences mental states other than beliefs, like desires or other emotions, and most lower-level processing doesn't culminate in mental states at all. Consider walking towards the fridge, navigating your way around the kitchen table, reaching into the fridge, and grabbing a beer. Several cognitive processes help us perform these cognitive tasks. For instance, our cognitive system is involved in detecting edges, identifying shapes and colours, perceiving depth, etc., but it may not form any beliefs about the table, fridge, or beer. Hence, it is unclear how a norm that governs the lower level, i.e. precision estimation, can also sufficiently explain how agents responsibly form beliefs.

### 3 General worries with subpersonal epistemologies

It's not just that Clark's precision estimation makes a particularly inadequate case for subpersonal virtue epistemology but that all subpersonal epistemologies are somewhat doomed. In what follows, I briefly discuss two worries that I think all subpersonal epistemologies will be vulnerable to.

The first of these worries is *the problem of demarcating the cognitive agency*. Suppose your neural precision estimation units were damaged somehow and were replaced by an artificial implant. This implant assigns reliability weightings that determine which processes will be involved in producing beliefs. According to Clark's proposal, the precision estimation function can sometimes ensure that our belief-forming process is being employed responsibly. Hence, in this case, where the precision estimation function is being performed by an implant, epistemic responsibility is also being met by the said implant. Something non-internal and non-neural is now responsibly employing our belief-forming process. One should wonder, at this point, where the cognitive agency lies. If the cognitive agency is one that responsibly employs a reliable belief-forming process then it lies entirely in the artificial implant. But if the implant is using the information from our neural cognitive system to assign reliability weightings to its targets then there is perhaps a hybrid cognitive

agency. Subpersonal epistemologies, therefore, need to outline a framework that explains hybrid cognitive agencies.<sup>17</sup>

For a more complicated situation, suppose the implant in your brain is actually a Google-powered AI device.<sup>18</sup> It's not only monitored by Google, but it operates on a machine learning algorithm that improves the functionality of the device. Just like the non-AI implant, this one also manages precision estimation weightings, thereby determining which process will be involved in generating beliefs. Being an AI-powered device, a Google implant doesn't have to collaborate with the human cognitive agency to employ a belief-forming process. It doesn't need information from our neural processes to assign reliability weightings to target belief-forming processes. If it's AI-powered, it can simply learn on its own and perform even better than biological precision estimation. In this particular case, it is even harder to make sense of where the cognitive agency lies.

While I have used the precision estimation function as an example of a subpersonal epistemology here, the problem of demarcating the agency is not only an issue for this kind of subpersonal epistemology. Any account that claims that a subpersonal mechanism can fully capture epistemic responsibility will run into similar problems. If epistemic responsibility is not connected to the person, there will be a need to understand where the agency lies and why that is so.

Further, the problem of demarcating where the cognitive agency lies also becomes the *problem of attributing beliefs*. This is also a general worry for all subpersonal epistemologies and especially for a subpersonal virtue epistemology. Recall that virtue epistemology considers knowledge a kind of virtue that one achieves when one manifests relevant and significant cognitive agency. Clark has called his proposal a subpersonal virtue epistemology (Clark 2015, 3768). So, sometimes an agent can achieve knowledge because their precision estimation units manifest sufficient relevant cognitive agency. However, as shown above, subpersonal epistemologies make it difficult to delineate cognitive agency. If it's not obvious where the agency lies, it's not clear who true beliefs should be attributed to or who achieved knowledge.

---

<sup>17</sup>There is another worry that is related to subpersonal epistemology, but I will explore it in more detail in another paper. Clark writes, 'I believe that a piece of new cognitive technology could be so well-designed as to be immediately assimilated into our daily routines, requiring no reflective window before properly being counted as delivering knowledge' (2015, 3764). On Clark's account, a belief-forming process can develop into an epistemic process immediately. Whereas, the cognitive integrationists, Greco and Pritchard, have argued that an agent ought to employ their belief-forming process frequently over a period for it to integrate into their cognitive system.

<sup>18</sup>In conference talk, Hernández-Orallo and Vold (2019) propose that AI extenders, or external AI resources that extend our cognition outside our bodies ought to be treated differently from regular extenders. There is no doubt that with artificial intelligence extending our cognition, the complications will be tenfold. Also, see Hetherington (2012) for a discussion on how knowers or epistemic agents might extend into the environment.

## Conclusion

In this paper, I argue that Clark's new subpersonal virtue epistemology presents a strange understanding of epistemic responsibility. I start with outlining the new proposal and describe how it's a more minimal account than the weak epistemic integration. The subpersonal mechanism that ought to ensure that a reliable belief-forming process is being formed is also present on subdoxastic levels that are not concerned with producing beliefs. In the end, I discuss the problem of demarcating the cognitive agency and the similar problem of attributing beliefs that all subpersonal epistemologies will be vulnerable to.



## References

- Aizawa, Kenneth. 2012. "Distinguishing Virtue Epistemology and Extended Cognition." *Philosophical Explorations* 15 (2): 91–106. <https://doi.org/10.1080/13869795.2012.670718>.
- . 2018. "Extended Cognition, Trust and Glue, and Knowledge." In *Extended Epistemology*, edited by Duncan Pritchard, Andy Clark, J Adam Carter, Orestis Palermos, and Jesper Kallestrup. Oxford, UK; New York, NY: Oxford University Press. <https://doi.org/10.1093/oso/9780198769811.003.0004>.
- Alfano, Mark, and Joshua August Skorburg. 2018. "Extended Knowledge, the Recognition Heuristic, and Epistemic Injustice." In *Extended Epistemology*, edited by Duncan Pritchard, Andy Clark, J Adam Carter, Orestis Palermos, and Jesper Kallestrup. Oxford, UK; New York, NY: Oxford University Press. <https://doi.org/10.1093/oso/9780198769811.003.0004>.
- Audi, Robert. 2015. *Rational Belief: Structure, Grounds, and Intellectual Virtue*. Oxford ; New York, NY: Oxford University Press.
- BonJour, Laurence. 1985. *The Structure of Empirical Knowledge*. Cambridge, Mass.: Harvard Univ. Press.
- Breyer, Daniel S. 2013. "Knowledge, Credit, and Cognitive Agency: Knowledge, Credit and Cognitive Agency." *Pacific Philosophical Quarterly* 94 (4): 503–28. <https://doi.org/10.1111/papq.12011>.
- Breyer, Daniel, and John Greco. 2008. "Cognitive Integration and the Ownership of Belief: Response to Bernecker." *Philosophy and Phenomenological Research* 76 (1): 173–84. <https://www.jstor.org/stable/40041158>.
- Carter, J Adam, and Duncan Pritchard. 2019. "The Epistemology of Cognitive Enhancement." *The Journal of Medicine and Philosophy: A Forum for Bioethics and Philosophy of Medicine* 44 (2): 220–42. <https://doi.org/10.1093/jmp/jhy040>.
- Carter, J. Adam. 2018. "Virtue Epistemology and Extended Cognition." In *The Routledge Handbook of Virtue Epistemology*, edited by Heather Battaly, 1st ed., 420–32. Routledge. <https://doi.org/10.4324/9781315712550-35>.
- Carter, J. Adam, Andy Clark, Jesper Kallestrup, and Duncan Pritchard, eds. 2018. *Extended Epistemology*. First edition. Oxford: Oxford University Press.
- Carter, J. Adam, Andy Clark, and Orestis Palermos. 2018. "New Humans?" In *Extended Epistemology*, edited by J. Adam Carter, Andy Clark, Jesper Kallestrup, and Duncan Pritchard, First edition. Oxford, UK; New York, NY: Oxford University Press.
- Carter, J. Adam, and Bolesław Czarnecki. 2016. "Extended Knowledge-How." *Erkenntnis* 81 (2): 259–73. <https://doi.org/10.1007/s10670-015-9738-x>.

- Carter, J. Adam, and Jesper Kallestrup. 2020. "Varieties of Cognitive Integration." *Noûs* 54 (4): 867–90. <https://doi.org/10.1111/nous.12288>.
- Chisholm, Roderick M. 1977. *Theory of Knowledge*. 2d ed. Prentice-Hall Foundations of Philosophy Series. Englewood Cliffs, N.J: Prentice-Hall.
- Clark, Andy. 2001. "Reasons, Robots and the Extended Mind." *Mind & Language* 16 (2): 121–45. <https://doi.org/10.1111/1468-0017.00162>.
- . 2003. *Natural-Born Cyborgs: Minds, Technologies, and the Future of Human Intelligence*. Oxford ; New York: Oxford University Press.
- . 2008. *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. Philosophy of Mind. Oxford ; New York: Oxford University Press.
- . 2010. "Coupling, Constitution, and the Cognitive Kind: A Reply to Adams and Aizawa." In *The Extended Mind*, edited by Richard Menary, 81–99. The MIT Press. <https://doi.org/10.7551/mitpress/9780262014038.003.0005>.
- . 2013. "Whatever Next? Predictive Brains, Situated Agents, and the Future of Cognitive Science." *Behavioral and Brain Sciences* 36 (3): 181–204. <https://doi.org/10.1017/S0140525X12000477>.
- . 2015. "What 'Extended Me' Knows." *Synthese* 192 (11): 3757–75. <https://doi.org/10.1007/s11229-015-0719-z>.
- . 2016. "Attention Alters Predictive Processing." *Behavioral and Brain Sciences* 39: e234. <https://doi.org/10.1017/S0140525X15002472>.
- . 2017. "How to Knit Your Own Markov Blanket: How to Knit Your Own Markov Blanket:: Resisting the Second Law with Metamorphic Minds: Resisting the Second Law with Metamorphic Minds." *Philosophy and Predictive Processing*. <https://doi.org/10.15502/9783958573031>.
- Clark, Andy, and David Chalmers. 1998. "The Extended Mind." *Analysis* 58 (1): 7–19.
- Colombetti, Giovanna, and Tom Roberts. 2015. "Extending the Extended Mind: The Case for Extended Affectivity." *Philosophical Studies* 172 (5): 1243–63. <https://doi.org/10.1007/s11098-014-0347-3>.
- Dennett, Daniel C. 1969. *Content and Consciousness*. London: Routledge.
- Drayson, Zoe. 2010. "Extended Cognition and the Metaphysics of Mind." *Cognitive Systems Research* 11 (4): 367–77. <https://doi.org/10.1016/j.cogsys.2010.05.002>.
- . 2012. "The Uses and Abuses of the Personal/Subpersonal Distinction." *Philosophical Perspectives* 26 (1): 1–18. <https://doi.org/10.1111/phpe.12014>.
- Greco, John. 2003. "Further Thoughts on Agent Reliabilism: Replies to Cohen, Geivett, Kvanvig, and Schmitt and Lahroodi." *Philosophy and Phenomenological Research* 66 (2): 466–80. <https://doi.org/10.1111/j.1933-1592.2003.tb00275.x>.
- . 2007. "The Nature of Ability and the Purpose of Knowledge." *Philosophical Issues*

- 17 (1): 57–69. <https://doi.org/10.1111/j.1533-6077.2007.00122.x>.
- . 2010. *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity*. Cambridge: Cambridge University Press. <https://doi.org/10.1017/CBO9780511844645>.
- Greco, John, and Ernest Sosa, eds. 1999. *The Blackwell Guide to Epistemology*. Blackwell Philosophy Guides. Malden, Mass: Blackwell.
- Hernández-Orallo, José, and Karina Vold. 2019. “AI Extenders: The Ethical and Societal Implications of Humans Cognitively Extended by AI.” In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, 507–13. Honolulu HI USA: ACM. <https://doi.org/10.1145/3306618.3314238>.
- Hetherington, Stephen. 2012. “The Extended Knower.” *Philosophical Explorations* 15 (2): 207–18. <https://doi.org/10.1080/13869795.2012.670721>.
- Hohwy, Jakob. 2013. *The Predictive Mind*. First edition. Oxford, United Kingdom ; New York, NY, United States of America: Oxford University Press.
- . 2014. “The-Self Evidencing Brain.” *Noûs* 48 (1).
- Hutton, Eric L. 2018. *Extended Knowledge and Confucian Tradition*. Vol. 1. Oxford University Press. <https://doi.org/10.1093/oso/9780198769811.003.0011>.
- Kornblith, Hilary. 1983. “Justified Belief and Epistemically Responsible Action.” *The Philosophical Review* 92 (1): 33–48. <https://doi.org/10.2307/2184520>.
- Lehrer, Keith. 1990. *Theory of Knowledge*. Boulder: Westview Press.
- Lynch, Michael Patrick. 2014. “Neuromedia, Extended Knowledge and Understanding.” *Philosophical Issues* 24: 299–313. <https://doi.org/10.2307/26611108>.
- Menary, Richard, ed. 2010. *The Extended Mind*. Life and Mind. Cambridge, Mass: MIT Press.
- . 2013. “Cognitive Integration, Enculturated Cognition and the Socially Extended Mind.” *Cognitive Systems Research* 25–26 (December): 26–34. <https://doi.org/10.1016/j.cogsys.2013.05.002>.
- Nozick, Robert. 1981. *Philosophical Explanations*. Cambridge, Mass. : Harvard University Press. [http://archive.org/details/isbn\\_9780674664791](http://archive.org/details/isbn_9780674664791).
- Palermos, Spyridon Orestis. 2011. “Belief-Forming Processes, Extended.” *Review of Philosophy and Psychology* 2 (4): 741–65. <https://doi.org/10.1007/s13164-011-0075-y>.
- . 2014. “Knowledge and Cognitive Integration.” *Synthese* 191 (8): 1931–51. <https://doi.org/10.1007/s11229-013-0383-0>.
- . 2015. “Active Externalism, Virtue Reliabilism and Scientific Knowledge.” *Synthese* 192 (9): 2955–86. <https://doi.org/10.1007/s11229-015-0695-3>.
- . 2016. “Spreading the Credit: Virtue Reliabilism and Weak Epistemic Anti-Individualism.” *Erkenntnis* 81 (2): 305–34. <https://doi.org/10.1007/s10670-015-9741-2>.

- . 2020. “Epistemic Collaborations: Distributed Cognition and Virtue Reliabilism.” *Erkenntnis*, April. <https://doi.org/10.1007/s10670-020-00258-9>.
- . 2021. “System Reliabilism and Basic Beliefs: Defeasible, Undefeated and Likely to Be True.” *Synthese* 199 (3-4): 6733–59. <https://doi.org/10.1007/s11229-021-03090-y>.
- Peels, Rik. 2017. *Responsible Belief: A Theory in Ethics and Epistemology*. New York, NY: Oxford University Press.
- Pritchard, Duncan. 2007. “Anti-Luck Epistemology,” 22.
- . 2008. “Virtue Epistemology and Epistemic Luck, Revisited.” *Metaphilosophy* 39 (1): 66–88. <https://www.jstor.org/stable/24439957>.
- . 2009. *Knowledge*. <http://dx.doi.org/10.1057/9780230242241>.
- . 2010. “Cognitive Ability and the Extended Cognition Thesis.” *Synthese* 175: 133–51. <https://www.jstor.org/stable/40801360>.
- . 2012. “Anti-Luck Virtue Epistemology.” *The Journal of Philosophy* 109 (3): 247–79. <https://www.jstor.org/stable/43820700>.
- . 2014. “The Modal Account of Luck: The Modal Account of Luck.” *Metaphilosophy* 45 (4-5): 594–619. <https://doi.org/10.1111/meta.12103>.
- . 2016. *Epistemology*. 2nd edition. Palgrave Philosophy Today. New York: Palgrave Macmillan.
- . 2017. *Knowledge, Luck, and Virtue*. Vol. 1. Oxford University Press. <https://doi.org/10.1093/oso/9780198724551.003.0004>.
- . 2018a. “Neuromedia and the Epistemology of Education: NEUROMEDIA AND THE EPISTEMOLOGY OF EDUCATION.” *Metaphilosophy* 49 (3): 328–49. <https://doi.org/10.1111/meta.12295>.
- . 2018b. “Anti-Luck Virtue Epistemology and Epistemic Defeat.” *Synthese* 195 (7): 3065–77. <https://doi.org/10.1007/s11229-016-1074-4>.
- . 2018c. “Extended Virtue Epistemology.” *Inquiry* 61 (5-6): 632–47. <https://doi.org/10.1080/0020174X.2017.1355842>.
- Proust, Joëlle. 2014. “Epistemic Action, Extended Knowledge, and Metacognition.” *Philosophical Issues* 24 (1): 364–92. <https://doi.org/10.1111/phis.12038>.
- Rowlands, Mark. 2010. *The New Science of the Mind: From Extended Mind to Embodied Phenomenology*. Cambridge, Mass: MIT Press.
- Rupert, Robert D. 2009. *Cognitive Systems and the Extended Mind*. Philosophy of Mind. Oxford ; New York: Oxford University Press.
- Sosa, Ernest. 1988. “Beyond Scepticism, to the Best of Our Knowledge.” *Mind, New Series* 97 (386): 153–88. <https://www.jstor.org/stable/2255164>.
- Sosa, Ernest, and Alvin Plantinga. 1993. “Proper Functionalism and Virtue Epistemology.” *Noûs* 27 (1): 51. <https://doi.org/10.2307/2215895>.

- Steup, Matthias. 1999. "A Defense of Internalism." In *The Theory of Knowledge: Classical and Contemporary Readings, 2nd Edition*, edited by L. Pojman.
- Stich, Stephen. 1978. "Beliefs and Subdoxastic States." *Philosophy of Science* 45 (4): 499–518.
- Vaesen, Krist. 2011. "Knowledge Without Credit, Exhibit 4: Extended Cognition." *Synthese* 181 (3): 515–29. <https://doi.org/10.1007/s11229-010-9744-0>.
- Wallisch, Pascal. 2017. "Illumination Assumptions Account for Individual Differences in the Perceptual Interpretation of a Profoundly Ambiguous Stimulus in the Color Domain: 'The Dress'." *Journal of Vision* 17 (4): 5. <https://doi.org/10.1167/17.4.5>.
- Wheeler, Michael. 2010. "In Defense of Extended Functionalism." In *The Extended Mind*, edited by Richard Menary, 245–70. Cambridge, MA: The MIT Press. <https://doi.org/10.7551/mitpress/9780262014038.003.0010>.
- . 2018. "Knowledge, Credit, and the Extended Mind, or What Calvisius Sabinus Got Right." In *Extended Epistemology*, edited by Duncan Pritchard, Andy Clark, J Adam Carter, Orestis Palermos, and Jesper Kallestrup. Oxford, UK; New York, NY: Oxford University Press. <https://doi.org/10.1093/oso/9780198769811.003.0004>.
- Wikforss, Åsa. 2014. "Extended Belief and Extended Knowledge." *Philosophical Issues* 24: 460–81. <https://doi.org/10.2307/26611116>.
- Witzel, Christoph, Chris Racey, and J. Kevin O'Regan. 2017. "The Most Reasonable Explanation of 'the Dress': Implicit Assumptions about Illumination." *Journal of Vision* 17 (2): 1. <https://doi.org/10.1167/17.2.1>.