ORIGINAL ARTICLE

How to make up your mind

Joost Ziff [ORCID]

University of California, Irvine,
California, USA

Correspondence
Joost Ziff, University of California, Irvine,
California, USA.
Email: joost@joostziff.com

Abstract

This paper develops an account of committed beliefs: beliefs we commit to through reflection and conscious reasoning. To help make sense of committed beliefs, I present a new view of conscious reasoning, one of putting yourself in a position to become phenomenally consciously aware of evidence. By doing this for different pieces of evidence, you begin to make your up mind, making conscious reasoning, as such, a voluntary activity with an involuntary conclusion. The paper then explains how we use conscious reasoning in reflection not just to form and change committed beliefs, but to become aware of existing ones. The paper concludes with an explanation of how the limitations of conscious reasoning require us to maintain committed beliefs in a system. It is our maintenance of this system that allows us to knit together individual episodes of conscious reasoning into one enduring performance as a committed systematic believer.

1 | INTRODUCTION

Much contemporary philosophy has been written about the ways our beliefs fall short of our best-reasoned intellectual endorsements. As a result, there are many conceptual tools for making sense of these shortcomings. Yet, as human beings, we reflect on our beliefs: we ask and answer questions as to what we do or should believe, making intellectual endorsements as we go. We know that doing so is important for eventually changing our beliefs, though it is less clear how. Surely reflection accomplishes something. But what exactly?

What we need, I propose, is a new theory of committed beliefs:[1] a type of belief we commit to through reflection and conscious reasoning. Committed beliefs are a type of belief which we can both form and come to know by consciously reasoning—pursuing and considering our phenomenally conscious evidence—thereby determining what we think is true here and now. We do not decide what we believe or endorse (on pain of endorsing doxastic voluntarism.) But we do make the choice to consciously reason, to put in the effort to find and consciously consider the evidence. As we consciously reason and reflect over time, we form committed beliefs—enduring intellectual endorsements—by which we can begin to take control of our thoughts and our lives.

The task of this paper is twofold:

1. To show conscious reasoning to be phenomenally conscious and directly controllable;
2. To show how conscious reasoning plays a central role in how we exercise control over our committed beliefs.

The first task I will do by providing a new account of conscious reasoning, one that better comports with our phenomenology of reasoning and can deal with extant worries concerning the existence of cognitive phenomenology. This view casts conscious reasoning as an act of putting oneself in a position to become phenomenally consciously aware of evidence. By doing this for different pieces of evidence you begin to make up your mind towards a conscious endorsement. Whether your mind is made up is fundamentally involuntary; it depends on the sense of endorsement you have when you express your conclusion. What you do have control over is how—the processes by which—you make up your mind: conscious reasoning. As such, conscious reasoning is a voluntary activity with an involuntary conclusion.

For the second task, I use a conceptual tool from the self-knowledge literature called the 'transparency method.'[2] According to this method, we discover our beliefs by forming them; I answer the question 'do I believe there will be a third world war' by consciously reasoning as to whether it is true that there will be a third world war. I extend this idea to any case where I answer questions about my beliefs in reflection. Reflecting in this way—even when answering questions about non-committed beliefs or beliefs I have already formed—always involves consciously reasoning to make up my mind. As such, I do not retrieve a committed belief without some sort of assessment; I am always (re)evaluating and (re)committing to a belief when I reflect. A committed belief is not an entity I store in my memory; it is an enduring performance to consciously endorse a certain proposition over time.

In § 2, I discuss reflection, introducing the notion of a committed belief to begin to make sense of the results of reflection. In § 3, I introduce my view of conscious reasoning and argue that it does a better job of making sense of the phenomenology of conscious reasoning than other competing views. Lastly, in § 4, I introduce my strong transparency argument for the performative nature of committed beliefs and discuss our duties to reflect with respect to our system of committed beliefs.

---

[1] My notion of committed belief is altered from Akeel Bilgrami's (2006) and Annalisa Coliva's (2016). I discuss the differences between my notion and theirs in § 4.2. In addition, Keith Frankish (2004, 2018) and Benjamin Winokur (2023) both propose similar more conscious reflective beliefs, calling them 'superbeliefs/celiefs' and 'commissive beliefs' respectively. However, Frankish and Winokur do not explicitly discuss the effortful activities involved in forming, revising, and maintaining these special beliefs, which is what this paper is focused on.

[2] First introduced by Gareth Evans (1982) and further developed by Bilgrami (2006), Boyle (2011a, 2011b, 2019), Byrne (2011), Coliva (2016), Fernandez (2013), and Moran (2001, 2012.) My application of transparency is closest to Moran's, since he too focuses on making up our minds for self-knowledge.

2 | REFLECTION AND COMMITTED BELIEFS

Most of our lives we passively form beliefs, too busy to take care to consider them. For instance, if my partner asks me 'where is the corkscrew?' I will easily report 'It is in the lower left drawer' without much thought. I did not come to believe the corkscrew was in that drawer because I investigated the matter, nor did I report on the location of the corkscrew after much consideration. I simply formed the belief while I was putting away the dishes.

Other times, however, we do try to take more deliberate control of our beliefs. We do this through what I will call reflection:

> REFLECTION: Consciously and carefully asking and answering what we do or should believe.

Descartes did this in the *Meditations*, where he engaged in the "project" of the "general demolition of [his] opinions." (2017, pp. pp. 17–18) His project involved asking and attempting to answer questions about his beliefs: he took his core beliefs and then re-evaluated them.[3] However, there are some important differences between what I mean by reflection and Descartes' reflection. Descartes engaged in this project because he aspired to ground his beliefs in absolute certainty. My understanding of reflection requires no such thing. We can end up deciding that something is almost certainly true, 50% true, true enough for practical purposes, and so on. Nor do we need to confine ourselves to Descartes' *a priori* methods from the armchair. Nothing in my definition of reflection stops someone from using empirical evidence for a belief. What is more important is to think carefully about one's beliefs.

Now, a full account of the value of reflection will have to wait until § 4.2. But as it stands, it seems intuitive that there is some requirement to reflect on our beliefs throughout our lives. This would not have to be every day, perhaps not even every week. But to never reflect and leave one's beliefs to passivity would seem wrong. Moreover, if we are required to reflect then it is plausible that reflection might accomplish some change in our beliefs. That was, after all, the whole point of Descartes' investigations: to change beliefs for the better. The same holds true for us: we reflect because we think that this activity could help change what we believe at least in some sense. Even if one would argue that there is no duty to reflect, it would be tendentious to deny that reflection does not confer any epistemic benefit.[4] (After all, why be a philosopher?)

That said, our beliefs do not always align with what we intellectually endorse. Take this instructive case from Eric Schwitzgebel:

> HYPOCRITICAL PHD STUDENT: Daniel "sincerely says that low-wage workers deserve as much respect as people who are paid handsomely; maybe even more respect." (Schwitzgebel, 2021, p. 354) Daniel has come to this conclusion about

---

[3] Sosa (2015, pp. 233–54) argues that Descartes' re-evaluation involves not throwing out these beliefs *de re*, but rather re-evaluating them *de dicto*. This idea will become more relevant once we get to § 4.2.

[4] Hilary Kornblith (2012) has argued that the value of reflection is severely overblown. However, Kornblith understands reflection as concerning "thinking about one's own first-order mental states in a first-personal way." (p. 28) Kornblith seems to mean by this that reflection solely involves thinking from the armchair. Kornblith is likely right to decry this activity. Trying to determine whether the beliefs one holds are for good reasons while recumbent would likely lead to "self-congratulation." (p. 3) But that is not what I have in mind when I talk about reflection. My definition does not exclude empirical investigations.

low-wage workers after reflection. He has well-articulated reasons for this conclusion, votes the right way, and will speak on the plight of low-wage workers earnestly. However, he is prone to condescending statements to low-wage workers when he interacts with them which he says with an "air of inauthenticity" ("Oh, bussing tables at Denny's Diner is just as valuable as writing philosophy!") And he is apt to have reflexive thoughts favoring the well-dressed rich over the less-expensively dressed poor when he sees them in the real world.

The question Schwitzgebel poses is fair: "Does Daniel believe that the working poor deserve at least as much respect as those of higher social status?" It certainly seems not. Although Daniel intellectually endorses the proposition, he has not fully integrated it into his own life. Motivated by this, we could do as Schwitzgebel has suggested and develop a notion of belief that excludes Daniel because of how his actions do not accord with his self-professed ideals.[5] But we could take a different track: develop a separate notion of belief which makes sense of Daniel's failings. It may still be true, as Schwitzgebel suggests, that "overall patterns of action and reaction" are more important than our "patterns of intellectual endorsement." (p. 359) But the latter still deserve study. Daniel got something right through reflection and does have the belief in some sense. But in what sense?

I propose therefore that Daniel has what I will call a committed belief, which I will define as follows:

> COMMITTED BELIEFS: One has a committed belief that p if and only if (in this world and in nearby possible worlds, for a non-fleeting period of time[6]):
> One treats p as true in conscious reasoning.
> When prompted in reflection, one expresses a sincere endorsement of p being true.
> When prompted in reflection, one can consciously provide reasons (however thin[7]) for this endorsement.

Daniel would qualify as having the committed belief that low wages workers deserve respect. He can presumably consciously reason on the basis that low-wage workers deserve respect. If Daniel is prompted in reflection as to whether he believes that low-wage workers deserve respect, he will say they do, and he can provide some reasons as to why he thinks that.

As defined above, we will need to understand both reflection and conscious reasoning to understand committed beliefs. The definition of conscious reasoning will have to wait until § 3. However, it is worth clarifying reflection now that we have introduced committed beliefs:

---

[5] Brie Gertler (2021) takes a similar approach to Schwitzgebel, having a case involving a contradiction between self-professed beliefs and actions. However, Gertler agrees that we must at some point exercise direct control over our beliefs to change them, "on pain of regress." (p. 344)

[6] This is to deal with the possibility that one meets these conditions accidentally or for fleeting moments. (This objection was pointed out to me by Annalisa Coliva.)

[7] For instance, these could be mere meta-reasons based on our sense that we did a good job when we previously considered the truth of p. See § 4.1 for more.

> REFLECTION (CLARIFIED): Occasionally in our lives we try to consciously exert direct control over our committed beliefs. We do this through *reflection*: consciously and carefully asking and answering what we do or should believe.

Note that although reflection begins by asking questions about committed and non-committed beliefs, it chiefly concerns controlling committed beliefs. We still use reflection to determine what non-committed beliefs we do and ought to have. But we do not necessarily change them through reflection. Daniel, for instance, could reflectively discover that he non-committedly believes the poor do not deserve respect, but he may not be able to change that belief through reflection. What he can, however, directly change is his committed beliefs.

Reflection is neither easy nor effortless. When we reflect, the point is to not take the most immediate answer. When I ask myself whether I believe in quantum mechanics, I should not be content to say 'obviously I do' just because that was the first answer that popped into my mind. That answer is not my committed belief, which requires careful consideration. As I will make clear in § 4, careful consideration does not always mean exhaustive consideration. There are ways of reflecting on our beliefs that allow us to carry forward a belief from a previous reflection. But reflection does enable us to consider our beliefs more exhaustively if necessary.

I have so far defined committed beliefs and reflection in familiar Cartesian terms, which (I hope) many philosophers will find intuitive. Nonetheless, much still needs to be explained. We still do not know how we answer questions about our beliefs in reflection, nor know how we form, maintain, and revise committed beliefs. As long as we do not understand how we go about reflecting and cultivating committed beliefs, we will not understand what the value of these activities is in the first place. To solve these problems, we will need more clarity about a notion that is central to both reflection and committed beliefs: conscious reasoning.

If committed beliefs are the results of conscious and careful efforts to get our beliefs right then we need to understand what those efforts involve. Similarly, if having a committed belief changes how we consciously reason, then we will not understand what committed beliefs are if we do not know what conscious reasoning is. Once we do explain conscious reasoning, however, we will be able to better understand how a. reflection can result in committed beliefs and b. why we cultivate committed beliefs in the first place.

3 | CONSCIOUS REASONING

Conscious reasoning is the activity we go through when we consciously and deliberately try to make up our minds as to what is true. Say I want to figure out who committed a crime. To gain confidence in a conclusion (say, 'Gus killed Daniel') I examine various pieces of evidence and try to become aware of how they speak to the potential conclusion. Looking at the bloody knife, I might notice that it would be unwieldy for, say, Janet to wield, but is perfectly sized for Gus. By examining a piece of fabric I found hanging on Daniel's front door, I might notice that it has a similar pattern to a shirt that I saw Gus wearing a few days ago. As I examine evidence, I verbalize these little conclusions. 'Gus wore that same shirt.' 'That knife is too large for Janet.' As I start to put things together, I find myself unable to escape one conclusion: 'Gus killed Daniel'. Conscious reasoning as such consists of three things:

1. Conscious consideration of evidence—becoming *phenomenally conscious* of pieces of evidence and characterizing that evidence and how it speaks to *p's* truth in a phenomenally conscious

manner. (Having a visual experience of the fabric and realizing that Gus wears a similar shirt.)

2. Investigative processes—anything we do to help ourselves consciously consider pieces of evidence (Looking, smelling, tasting, getting up, squinting one's eyes, asking, searching one's memory, deducing, inferring, etc.)
3. Verbalizing candidate conclusions.

Or, to put my view in a more typical definitional fashion:

> CONSCIOUS REASONING: One consciously reasons—one answers questions in reflection as to whether p is true—by going through investigative processes to consciously consider salient pieces of evidence to p's truth. Conscious reasoning can end when one expresses a proposition in some way with the appropriate involuntary sense of endorsement.[8]

With the following supplemental definitions:

> INVESTIGATIVE PROCESS: Any effortful thing one does to help one consciously consider a salient piece of evidence.
> CONSCIOUS CONSIDERATION OF EVIDENCE: To become phenomenally consciously aware of a piece of evidence and to (effortfully) characterize how it speaks to p's truth in a phenomenally conscious manner.[9]

In §§ 3.1–3.3, I will elaborate on the various features of my view of conscious reasoning in detail. I will put a special focus on how these features enable us to have direct control over the important parts of the conscious reasoning process.

3.1 | Why Conscious Reasoning Needs Phenomenal Consciousness

One important but (potentially) philosophically controversial feature of my view is its inclusion of phenomenal consciousness—a 'what is it like'[10]—in conscious consideration. When it comes to rational processes, philosophers of mind have often focused on "access-consciousness," (Block, 1995, p. 231) thereby avoiding true phenomenal consciousness. Access-conscious awareness makes something available for "reasoning, … rational control for action and … rational control of speech." Prima facie, this might seem good enough. We are, after all, trying to describe how conscious consideration makes evidence accessible. But I contend that access-consciousness

---

[8] Uriah Kriegel (2015) also argues for what he calls a "credal feeling" (p. 32), but he thinks of it being towards a bare proposition, rather than a feeling with which we say our conclusion. He also argues for a "feeling of involuntariness." (p. 35) However, Kriegel says both these features belong to our self-standing cognitive phenomenology, rather than the phenomenology of conscious reasoning in inner and outer speech, as I will contend.

[9] For the purposes of this paper, this definition implies that something is evidence in virtue of it being consciously considered as such. Whether it turns out to be relevant or salient will be determined by conscious consideration and investigative processes.

[10] See Nagel (1974).

misses the role that phenomenal consciousness has in giving us a more direct control over the conscious reasoning process.

To see this, it will be helpful to better understand the role that thoughts and thinking play in conscious consideration. There is a huge debate as to what extent our thoughts truly are phenomenally conscious.[11] At minimum, many philosophers think that our phenomenology is not distinct enough to allow us to distinguish between two different thoughts solely on the basis of that character. To put this more specifically:

> NON-DISTINCT PHENOMENOLOGY: For any thought *A*, there is a thought *B* such that *A* is a significantly different thought, yet the state of thinking *A* and the state of thinking *B* have the same phenomenology—the same 'what it is like' to experiencing them.

This seems to create a problem for my definition of conscious consideration. Characterizing, for instance, how the bloody knife speaks to Gus's potential guilt presumably involves thoughts. But if the thought characterizing the knife as bloody has the same phenomenology as the thought characterizing the knife as not bloody, it would be impossible to characterize how the bloody knife speaks to Gus's guilt in a phenomenally conscious manner. This problem would be even worse if we had to characterize a thought, like a non-visual memory that the knife was bloody. For now, not only are the thoughts with which we characterize our evidence phenomenally non-distinct, but the evidence itself is too.

But reflection and conscious reasoning must involve more than inchoate non-verbalized thoughts. This is in part because non-verbalized thoughts frequently are only fully recognizable when they are expressed. Robert Pippin (2010) made a similar point in interpreting Nietzsche on the primacy of deed over doer:

> If I start out to write a poem, I might find that it does not go as I expected, and think that this is because the material resists my execution, my inner poem, and so what I get is a "poorly expressed poem." This is a very misleading picture on this account (…) The poem is a perfect expression of what your intention turned out to be. It (the expression of what has turned out to be the intended poem) just turned out to be a bad poem; not a bad expression of a good poem. As Nietzsche keeps insisting, our egos are wedded to the latter account; but the former correctly expresses what happened. (p. 59)

To put it another way: though my unverbalized thought of the poem might initially feel very complete and interesting, it is only when that thought is expressed in poem form that we know its worth. The same point could be made about similar 'sudden realization' cases popular in the cognitive phenomenology literature.[12] A supposed sudden realization of a solution to a math problem only proves itself to be a solution once expressed. What this suggests is that being a good conscious reasoner is much more about the skillful expression of one's thoughts than the initial thoughts themselves. Initial thoughts provide us with ideas, doubts, or suspicions, which help steer how

---

[11] Pro-cognitive phenomenology works include Horgan & Tienson (2002), Horgan (2011) Kriegel (2015), Lennon (2022, 2023), and Pitt (2004, 2011). Anti-cognitive phenomenology works include Bayne & McClelland (2016), Carruthers & Veillet (2011), Lormand (1996), Pautz (2013), Prinz (2011), Robinson (2005, 2011), and Tye and Wright (2011).

[12] See for instance Prinz (2011).

we consciously reason or write a poem. But these thoughts must all be expressed to judge their quality. In the case of conscious reasoning, expressions can take many forms: drawing, writing, making lists, even *sub-vocalizing*.[13] (i.e. speaking in one's head[14]) All these expressions can serve to make evidence and its salience perspicuous. Rather than viewing our thoughts as almost impossible to put in language, it is more helpful to see language as providing the conceptual material for our thoughts.[15]

Now, expressing our thoughts does not in itself make them distinctly phenomenally conscious. There are ways of reasoning where one's evidence is access-conscious but not distinctly phenomenally conscious. These ways of reasoning are quicker and more automatic. Take, for instance, Paul Boghossian's[16] 'reasoning 1.25–1.75,' which are versions of personal-level reasoning that are faster and less precise than conscious reasoning. When we reason in such a way, we might still verbalize certain steps, but we make quick leaps between different thoughts to come to conclusions. The evidence we use when reasoning in this way is access-conscious, in the sense that it was accessible to the reasoning and speech processes involved. But because this reasoning is so quick, we might find ourselves saying things and drawing conclusions without being easily able to tell exactly what steps we took and what rules or ideas we relied on; in other words, that the elements of those reasoning processes are non-distinctly phenomenally conscious.

There is nothing wrong with reasoning in this way. Reasoning quickly clearly has its place when time is of the essence. But to truly take control over how we make up our minds, we must take better care. Because of that care, these characterizations must be distinctly phenomenally conscious.[17] We want to be sure of what we say, what our evidence means, and what we think it shows. This does not mean that we always succeed in making fully phenomenally conscious all the evidence that guides our conscious reasoning. We sometimes make inferential leaps in how

---

[13] Munroe (2022) discusses how we use inner speech "as a representational resource to keep track of information, like lemmas and intermediate conclusions." (p. 858)

[14] One might argue that aphantastists provide a counterexample for sub-vocalized phenomenology. However, Lennon (2023) argues that aphantasists likely do have a phenomenology to their thoughts. Aphantasist Blake Ross (2016) describes the phenomenology of their sub-vocalizations as being like a "milk voice"—a flat, inner monologue with no texture or sound, which we use to tell ourselves things like "Remember to pick up milk." This suggests Ross does have a phenomenology, albeit one stripped of color and sound in such a way that only its form remains. This form is likely enough, however, to imbue one's sub-vocalizations with linguistic content.

Pro-cognitive phenomenology philosophers might argue here that this 'form' of one's sub-vocalizations is precisely the non-sensory cognitive content they are talking about. I have no stake in classifying whether this is sensory or non-sensory content. What is important to me is that the content of sentences we express is incredibly sophisticated. And, as I mention in ft. 15, there is good reason to think that the external tools we use for reasoning enable us to reason more successfully internally.

[15] Macbeth (2014) argues that external systems for expressing logic and mathematics affect our internal reasoning. She explains how systems like Arabic numerals enable children to do complex calculations in their heads that previously were difficult for pre-modern humans. (p. 64) She claims there were similar advances in mathematical reasoning with the development of Frege's *Begriffsschrift* and other similar logical systems, which enable one to logically reason in a language. (p. 286) If Macbeth is correct, then the cognitive phenomenology literature grossly underestimates the impact that displayed reasoning has on its internal analogues. Our cultural inheritance of symbolic manipulation is what drives our capacity to think complex thoughts, not the other way around.

[16] See Boghossian (2014, 2018).

[17] Annalisa Coliva (personal communication, 2023) pointed out that since we might use previous beliefs in our evidence that this might present some circularity. However, while conscious reasoners in practice should try to suss out their previous beliefs, there is nothing inherent in conscious reasoning itself that demands that they use previous beliefs. I discuss why this is the case in § 4.2. Thus, when we do try to carefully use previous beliefs in conscious reasoning, we start a separate conscious reasoning episode to determine what that belief was. (As described in § 4.1.)

we process evidence that we do not fully investigate. (Nor arguably should.) But the regulative ideal of this type of reasoning is that we try to make our evidence manifest to ourselves. Without that manifest presence, we do not make up our minds ourselves; our minds make themselves up. The only way to truly bring that manifest presence about is to make phenomenally conscious as much of the evidence as is reasonable to make up our minds.

3.2 | Voluntary Non-voluntary Beliefs

Two other important features of my view are the inclusion of effort and an involuntary sense of endorsement. These features make sense of two contrasting intuitions:

1. (VOLUNTARY) DOXASTIC RESPONSIBILITY: that we are responsible for our beliefs—and that responsibility seems to require voluntary activities.
2. DOXASTIC INVOLUNTARISM: that rational people do not voluntarily believe or intellectually endorse things.

Pamela Hieronymi (2006, 2008, 2009) has written about these intuitions, which seemingly pull in opposite directions. On the one hand, we want to understand our responsibility over our beliefs—and it is natural to use voluntary control to make sense of that responsibility. On the other hand, it can seem like the good believer has no choice over what they believe, at pains of denying their own rationality. Hieronymi's solution is to develop a notion of responsibility over our beliefs qua how "[we] are answerable for reasons that [we] take to show [our beliefs] true." (2008, p. 365) Though not very voluntary, Hieronymi claims that this notion preserves our responsibility over our beliefs. Even though we cannot choose what we believe, we are still accorded praise and blame in, for instance, how we produce reasons when challenged.

When it comes to regular beliefs, Hieronymi's points are well-taken. However, to my mind, Hieronymi does not affirm enough the intuition that many of the processes surrounding belief formation and revision are voluntary. The responsibilities Hieronymi mentions are ones that we can shirk if we do not spend the requisite effort. And spending effort is a voluntary affair. When one is spending effort, one is aware that a. one is making something happen (a *de se* realization[18]), b. if one were to stop spending effort that the world would be different, and, most importantly c. that spending the effort is optional, given its unpleasantness.[19] To be clear, Hieronymi recognizes the voluntariness of effortful actions too, noting that getting rid of a mess in your kitchen requires voluntary spent effort. (p. 372) Moreover, Hieronymi is likely right to "refuse to extend" this idea of direct voluntary control to non-committed beliefs. But spending effort is key to conscious reasoning and (as we shall see in § 4) to committed beliefs. Without effort, we do not do all the investigative processes—looking, smelling, getting up, searching one's memory, inferring, etc—nor any of the characterizing processes involved in conscious consideration. (See § 3.1.) Doing any one of these activities on their own might be fairly effortless. Inferring can happen almost

---

[18] Perry (1979) contains the classic mystery shopper example that illustrates the idea of a *de se* realization nicely.

[19] This is not to say that one's experience of effort rules out any skeptical scenarios. If, for instance, I am trying to lift a couch, my awareness of that effort does not prove that I am actually contributing to the couch being lifted. Perhaps I am lifting at a wrong angle and my friend is picking up most of the slack without me noticing. However, for this paper, I do not need to establish beyond a doubt that my efforts contribute to conscious reasoning. I only need to make it plausible that they do.

spontaneously, as can putting a thought into words. But doing multiple activities over time, forcing one's mind to reinfer, or turning over the knife again to look at its handle—in other words, sustaining one's inquiry—will require sustained mental and physical effort.[20] It is in our persistence to keep investigating and do our best to consider evidence, that we are responsible for our conscious reasoning and our committed beliefs.

Nonetheless, Hieronymi is still right about the involuntariness of our resultant (committed) beliefs and conclusions. So, my view makes sense of intuition 2 by having conscious reasoning end in an involuntary sense of endorsement. We do not end conscious reasoning through an expenditure of effort. Rather, we end it through the involuntary sense of endorsement that we have as we express a candidate conclusion. I can expend effort to express a candidate conclusion whenever I want. But it is far harder to convince myself of a candidate conclusion's truth just by expressing it. In fact, if I do convince myself of a conclusion by expressing it—for example, by repeating 'Dylan is the murder' to myself over and over—I would be consciously reasoning badly[21] (or perhaps not reasoning at all.)

3.3 | Why Inferring Is Not Concluding

Lastly, my view has important but potentially controversial features relating to the role of inference. I hold that a. deducing and inferring are investigative processes, and b. we conclude conscious reasoning through an involuntary sense of endorsement.

The first feature may seem unintuitive because of the dominant "rule-following model of reasoning." (Munroe, 2021, p. 8327)[22] According to this model, good reasoning is done and concluded through *inference*. Take the examples of reasoning from McHugh & Way (2018):

> If Jane had a beer, then there are none left. Jane had a beer. So, there are none left.

---

[20] Sustained effort helps differentiate my view from Galen Strawson's (2003) more skeptical one. Strawson thinks that reasoning, recalling, and judging at the most only need voluntary "prefatory" and "catalytic" action (231) to initiate these processes. Once they are initiated, one just waits for the mind to involuntarily supply the conclusions. I think Strawson is right that for many mental processes the mind spontaneously supplies the results. Indeed, we generally do not recall memories or infer conclusions in simple inferences directly; we initiate these processes and then the results just 'pop' in our heads. But conscious reasoning does not merely involve simple mechanical inferring or memory recall. Rather, it is an effortful agential process that one sustains over time, involving potentially multiple investigative lines and exhaustive (re)characterizing of what evidence shows. (See § 3.1.) Moreover, the process does not terminate when an answer pops in one's head; if one cannot say the candidate conclusion with the right sense of endorsement, then conscious reasoning cannot end. (See § 3.3.) As such, conscious reasoning involves voluntary sustained effort, not just initial "ballistic" (245) effort.

[21] Rowan Holloway (personal communication, 2023) pointed out to me there might be counter-examples involving mental illness. Someone with depression, for instance, might benefit from repeating positive claims about themselves to make them seem more plausible. It is unclear, however, whether these methods should be considered part of conscious reasoning or as coping mechanisms to deal with deficiencies in conscious reasoning.

[22] This article also mentions Boghossian (2014, 2018), Broome (2013, 2014) as supporters of a rule-following model of reasoning. Valaris (2017) also discusses the dominance of a rule-following model of reasoning. Valaris (2017, 2019) and Munroe (2021) argue for non-rule following views of reasoning. Space constraints prevent a proper discussion of these views. However, since they are not discussing conscious reasoning as I define it, there is no pressing need to debunk them.

I shall get beer. In order to get beer, I must go to the shop. So, I shall go to the shop. (p. 167)

In these examples, 'So, there are none left' and 'So, I shall go to the shop' are conclusions licensed by *modus ponens*. These conclusions were not made by an involuntary sense of endorsement, but by a voluntary act of inference. These seem like unproblematic examples where people form beliefs—and thus endorsements—through inference. It is reasonable to ask why conscious reasoning would work so differently.

Clearly, deduction and inference have their place in understanding reasoning, even of a conscious variety. This is why my view includes them as investigative processes.[23] But conscious reasoning has its own goal: gaining confidence in a potential conclusion. To gain that confidence, we must quell potential doubts. While inference can help quell doubt, it cannot do so determinatively. Take this inference that I might make in the Gus murder investigation:

1. 'If I find a picture of Gus wearing that shirt, then he has to be the murderer.'
2. 'Ah, there is a picture of Gus wearing that shirt.'
3. 'So, he has to be the murderer.'

Deductive reasoning might help convince me that Gus is the murderer, but I could just as easily, when I say 3, not feel convinced:

1. 'So, he has to be the murderer…'
2. 'Or does he? That still does not seem right …'

Importantly, nothing need have been wrong at any step of my inference. It is possible that when I verbalized 1 and 2, I fully endorsed them, and that when I said 3, I did so with an understanding that it followed from 1 and 2. Normally, I would unproblematically form the belief. Yet, when I expressed 3—the end of the inference—I ended up not endorsing it. Why I ended up not having a sense of endorsement is underdetermined. Clearly, I doubted something. But my lack of endorsement and feelings of doubt can have a non-distinct phenomenology.[24] We frequently experience feelings of indeterminate doubt; we sometimes start to say things with a confidence that fades as we finish speaking. With quicker forms of reasoning, these feelings do not typically arise. But these feelings can play a role when we are consciously reasoning.

---

[23] An anonymous reviewer pointed out that this might seem to imply self-conscious reasoning in the sense of "being aware of one's beliefs" or "aware of oneself as in the process of determining" them. They worried this was the case "since one must make use of metacognitive clues about one's attitude toward various inferential steps and propositions." This would clash with the transparency method I discuss in § 4.

However, I do not think my view implies self-conscious reasoning. To draw inferential conclusions, we need not conceive of ourselves as taking on propositional stances, taking on beliefs, or applying *modus ponens*. We could just as well say the various steps and feel out whether the conclusion sounds right. (Indeed, this is likely what competent reasoners who have not learned philosophical logic do.)

That said, good conscious reasoning can involve a self-conscious attitude towards one's propositions in inference. Say one consciously reasons from P and $P \rightarrow Q$, to Q, but considers whether one does actually believe P and Q before drawing one's conclusion. However, in such a case, we would be actually dealing with a composite process; two reflections about P and $P \rightarrow Q$ which launch two respective conscious reasoning processes, and an overall conscious reasoning process for the Q conclusion. These 'P' and '$P \rightarrow Q$' conscious reasonings might involve more investigative processes, which might involve more self-conscious reflection, but we ultimately end up at non self-conscious reasoning.

[24] See § 4.1.

As a result, conscious reasoning cannot conclude through inference alone. Instead of viewing these feelings as interfering with the proper work of inference, it makes more sense to think of resolving these feelings as the focus of conscious reasoning itself. We would verbalize these feelings just as we would characterize our evidence; by expressing them, thereby making them specific doubts—expressed contrary evidence—that we can then quell or accept.

This might seem wrong. Surely, feelings should not play a role in conscious reasoning. For instance, one's feeling of disgust towards surgery should not dissuade one from necessary medical treatment. But not all feelings are irrational or should be precluded from guiding our reasoning. If you have a sense that something is wrong with your tentative conclusion, you should examine that feeling and see if it is rationally substantiated. Not every feeling is a good reason to halt inquiry. We teach children, for instance, not to judge people hastily based on their initial feelings of dislike. But we are not actually teaching children to judge dispassionately, without any feeling. Rather, we are teaching children to not let initial reactive feelings dominate their thinking, and instead let their rational feelings—like the sense of endorsement—guide them.

To understand why a sense of endorsement could be a rational feeling, we should understand it not just as a feeling that attaches to expressed propositions,[25] but also as a *sui generis* intellectual seeing. Take for instance the following example:

> CHEATING SIGNIFICANT OTHER: Judy suspects that their partner Lizbeth is cheating on them but does not want to believe it. They consciously reason about the matter; look through their partner's phone, recollect certain instances when their partner said they were working late, cross-reference these memories with text messages their partner sent, etc. As Judy goes through this evidence, it reluctantly starts to dawn on them that it is highly likely that Lizbeth is cheating on them. Judy considers this thought, goes over all the evidence again quickly in their mind, and sees it points only one way: 'Lizbeth is very likely cheating on me'. They assert this to themself with a sense of endorsement and end their conscious reasoning.

What I aim to show with this example is that a sense of endorsement can be a. acquired rationally, b. constitutively linked with how one comes to understand the evidence one cognizes, and yet still be c. a feeling with its own phenomenology. Judy does not want to believe that Lizbeth could be cheating on her. Yet, Judy recognizes in themself the misleading phenomenology associated with wishful thinking,[26] and takes the time to appraise the evidence fairly. As they do so, they experience a feeling of inevitability towards the conclusion. Part of this feeling is brought on by sadness and despair; yet part of it is brought on by their honest desire for the truth. Moreover, their rational feeling—their sense of endorsement—comes about because of how the process of examining evidence holistically instilled that feeling in them.

Thus, the relationship between evidence and conclusion in conscious reasoning is not just about 'operating well' or 'following rules.' Instead, this relationship is holistic, coming to see the

---

[25] Such a feeling would be more like an intuition. Intuitions are plausible pre-theoretically, unlike senses of endorsement, which become plausible through conscious reasoning.

[26] Coliva (2008, 2012) argues that the phenomenology of wishful thinking and judging can be so similar that we cannot determine what state we are in on phenomenology alone. In response, Conor McHugh (2012) argues defeaters will always be present to differentiate wishful thinking from judging. I agree that these defeaters are present. However, I want to tie their presence to the story of how we make up our minds. We might at one moment not be able to distinguish an episode of wishful thinking from judging, but good conscious reasoners suss out their phenomenology as it unfolds over time. (See ft. 27)

evidence as implying certain things being true.[27] Our evidence still justifies our conclusion; we are still expected to refer to some of it when explaining our committed beliefs to others. Nonetheless, our conviction in our conclusion is our own—we obtain it not by blindly applying rules, but by taking up our own unique intellectual point of view.

4 | REFLECTION AND COMMITTED BELIEFS (again)

Now that we understand conscious reasoning, we can finally better understand how we reflect, and how we form, maintain, and revise committed beliefs. Here is the plan for doing so:

First, I will explain how we must use conscious reasoning in reflection to answer questions about our beliefs, and how using conscious reasoning in this way forms committed beliefs.

Next, I will explain not just how we form committed beliefs but how we maintain them. I do this by looking at committed beliefs not individually, but how they hang together in a system. By looking at committed beliefs in this way, we see how they allow us to preserve and correct the results of conscious reasoning. In turn, the value of reflection comes from how we can use it to do much of this preserving and correcting.

Before I go about these tasks, let me briefly explain the idea that will motivate much of what is to come: what is known in the self-knowledge literature as *transparency*. This is the idea that we can come to know our beliefs by determining what we think is true. Take, for instance, Gareth Evans' (1982) famous example:

> If someone asks me 'Do you think there is going to be a third world war?', I must attend, in answering him, to precisely the same outward phenomena as I would attend to if I were answering the question 'Will there be a third world war?' (p. 225)

This idea has an intuitive appeal. By figuring out what I think is true, I make true that I believe it. This works because when I intellectually endorse the proposition 'there will be a third world war' I—in normal conditions—instill a committed belief that there will be a third world war. Recall from § 2 that having a committed belief involves treating a proposition as true, being able to endorse the proposition, and giving reasons when prompted in reflection. Given that we generally remember our endorsements and why we took our evidence to show it to be true, it seems plausible that conscious reasoning can instill such a committed belief.

The basic idea is to expand transparency to almost all operations of reflection:

CENTRALITY OF CONSCIOUS REASONING:

1. One must use conscious reasoning to become aware of a committed belief.
2. One almost solely uses conscious reasoning to form and maintain a committed belief and the reasons for holding it.
3. One can use conscious reasoning to become aware of non-committed beliefs.

---

[27] One might worry that evidence considered in moment *A* cannot figure in a conclusion made in moment *B*, making it impossible for our conclusion to be based on that evidence. This misunderstands how time operates in our phenomenology. We do not operate solely in instant moments when consciously reasoning; we are also in a specious present that extends beyond moments that give us a sense of contiguity. Our evidence as such can still be in our mind to some degree even though it came to conscious awareness in previous moments. MacDonald (2014) and Husserl (2008) make similar points about the specious present.

This thesis would mean (perhaps controversially) that becoming aware of a committed belief[28]—even a settled one—involves making up one's mind. This will be crucial to understanding in § 4.2 that having committed beliefs is an active affair. Combine that with a duty to reflect, and one can begin to see how being a committed believer involves sustained effort over time.

4.1 | Using Conscious Reasoning to Discover and Form Beliefs

Do we always make up our minds when we become aware of a belief? The idea seemed plausible enough in Evans' case above. Part of what made that case work was that the question posed really was a question about what I believe now. It is perfectly legitimate to answer the question 'do I believe now that p' by answering the transparent counterpart question 'is p true?' As such, my intellectually endorsing a proposition through conscious reasoning is pretty much equivalent to committedly believing it.

But this gets harder when we are answering questions about what we believed before. Specifically,

1. asking in reflection what I believe about my past committed beliefs,[29]
2. asking in reflection what I believe now about something when I have previously formed beliefs about that thing before,[30] and
3. asking in reflection what I non-committedly believe about something.

All these cases can be understood as transparently answering questions as to what we believe now.

Take case 1. Suppose we alter the previous Evans case:

> Gareth asks himself in reflection whether he previously believed that there will be a third world war. He carefully considers the question and gathers evidence of whether he had held this belief. He tries to think of memories of him saying things about a potential third world war, but he can find none. He concludes that he never believed that there will be a third world war.

This is a legitimate response to the question. However, Gareth seems to apply not transparency, but what is known in self-knowledge literature as *self-interpretation*.[31] With this method, one treats the task of determining one's beliefs the same as determining the beliefs of another. Take a similar case:

---

[28] Note that it is possible to answer questions non-consciously or non-carefully. However, we do not then become aware of our committed beliefs, but our non-committed beliefs. Coliva (2016) discusses the example of her answering a question as to how old her mother is while "while engaged in a different activity." (p. 28) As Coliva points out, she would still be able to answer the question, but she would not be reporting on her committed belief, but rather a "first-order disposition" to so answer, one that is typically automatically inculcated when one forms a commitment. But such automatic first thoughts are not the same as a considered intellectual endorsement. The only way to get that endorsement is through reflection.

[29] Byrne (2011) originally pointed out this case.

[30] Shah & Velleman (2005) originally pointed out this case.

[31] Gertler (2021) canvasses supporters of this approach to self-knowledge for certain mental states, Carruthers (2011), Cassam (2015), and Lawlor (2009). In addition, Kirsch (2018, 2020) discuss ways in which self-interpretation of emotions and memory allow us to exercise narrative control over our lives.

> Jill, a philosopher with a long and storied career, is writing a paper on hinge episte-mology. Because she has written so many papers, she is not sure what position she used to hold on hinge epistemology—what she used to committedly believe. She reads her old papers carefully, examines some of her memories that reading these papers prompted, and concludes what position she used to believe about hinge epistemology.

In both cases, Jill and Gareth are treating the task of determining what they believe as if they were determining the beliefs of another. Though Gareth bases his understanding of what he said on his memories; it might equally have been memories of what someone else said. We could imagine that Gareth's close friend Smareth who is around him a lot could be in just as good a position to determine Gareth's beliefs as he is.

It might seem difficult to conceive of Jill and Gareth as answering questions about what they believe now—thus consciously reasoning. But note that the evidence that Jill and Gareth use is such that they have no privileged access to what their beliefs used to be. What they do stand in a position of authority to say is what they believe now. For they can make true what they believe now simply by consciously reasoning. If Jill and Gareth are then trying to consciously and carefully answer the question as reflection demands, and using evidence to do so, it makes more sense to conceive of them as transparently answering questions about their beliefs now. After all, if we asked Jill 'Do you believe that you used to believe in hinge epistemology,' she would certainly agree.[32] Just because one is answering questions about one's beliefs in one's past does not mean that one is not making up one's mind.

Take now case 2. We do not want to have to redo all our conscious reasoning every time we reflect on a previously considered belief. Thankfully, this is not the case. Take the following example:

> In reflection, Jack asks whether he believes that he locked his car. In answering this question, he remembers that he previously asked this question before and concluded that he had locked his car. So instead of either going back to his car to check that it is locked or scouring his memory, he decides that his previous inquiry was likely thorough enough. Thus, he concludes that he probably did lock his car.

It might seem that in answering the question of what he believes, Jack is just accepting what he believed previously and thus avoiding consciously reasoning about the question. But that is not the case; it is just that Jack's memories of what he previously decided are relevant in answering the question. It is true that his conscious reasoning is not as thorough as it could be. Instead of investigating whether his car is locked by going outside, he decides that his previous investigation was likely enough. But Jack's behavior is not epistemically suspect. Jack's present committed belief is still formed by considering what is true; it is just that his consideration of what is true can include thoughts like 'I'm pretty sure I investigated this already. No need to obsessively recheck.' as evidence. After all, it would be irrational for Jack to constantly go outside to see whether his car is still on the lot. It is similarly acceptable to use evidence like 'I remember the scientific consensus on mRNA vaccines being such-and-such' to conclude that vaccines are effective. However, in using this kind of evidence in conscious reasoning, one is still obligated to properly evaluate its

---

[32] To be clear, we must have certain conceptual resources in place to even transition from answering a question about my belief to answering a question as to what is true (and vice-versa.) See Coliva (2016, pp. 188–197) and Boyle (2019) for two proposals. Nonetheless, once those conceptual resources are in place, the bulk of our time is spent consciously reasoning.

salience. Moreover, one must also match one's confidence[33] in one's conclusion to the strength of the evidence. My confidence in the workings of the mRNA vaccine can still be strong but should be of a different order had I perused the studies myself.[34]

Finally, let us take the case of uncovering non-committed beliefs in reflection, specifically those that we are counted as having because of our behavior. Take this modified Daniel case from § 2:

> Jill notices Daniel talking down to a low-wage worker and criticizes him for this behavior. This prompts Daniel to reflect, asking himself if he really does believe that low-wage workers are worth respect. He thinks about the behavior Jill pointed out, and suddenly remembers other similar incidents as well. As he thinks through these incidents, he notices unresolved feelings in himself of arrogance and conceit. He realizes that he does not fully believe what he committedly believes, and resolves to address these feelings of resentment to better align his behavior to his best intellectual intentions.

With this kind of belief discovery, we must treat ourselves as third personal agents. We do not have control over what we used to believe, committedly or non-committedly. Nonetheless, we do have direct control over our conscious reasoning and thus over our present committed beliefs about our previous behavior. Daniel's committed belief shows him what he should believe in the fullest sense of the word. By forming a committed belief, he can at least begin to change his non-committed beliefs.

4.2 | Systems of Committed Beliefs

We now know how conscious reasoning enables us to become aware of our present committed beliefs. But we still have no understanding of how we maintain committed beliefs over time. As currently defined, it might seem that committed beliefs only exist when we consciously reason about them in reflection. That would be a far cry from an actual commitment: an enduring self-imposed obligation. However, once we widen our scope from single commitments to our overall commitment to being a systematic believer, this problem disappears. In preserving the results of our conscious reasoning in a loose system, we are able to maintain and correct them over time. The value of our committed beliefs—and thus our obligation to maintain them—comes from how they fit within that system.

---

[33] Here, I do not mean to invoke the formalized idea of a credence or degree of belief. I do not think that our sense of endorsement can be reduced to some numerical value. Instead, we modulate our confidence by making our conclusions conditional on our current understanding of the evidence. Ex. 'Given that I trust the scientific community, I believe that the mRNA vaccine works.'

[34] There is an obvious connection here to the reductionism debate around testimonial knowledge—that is, whether trust in others' testimony is basic or whether it is reducible to dispositive reasons for reliability. (See Leonard (2023) for more.) I do not want to take a strong stand on any side of that debate. However, I do want to point out that since reflection means consciously and carefully answering a question, any deployment of evidence—even if it is testimonial—will require some care to be deployed well. Sometimes that will mean checking the reliability of a source; other times it will involve evaluating the suitability of a source for a situation. Thus, the use of testimony in reflection can be both basic (as in the latter) and reduced (as in the former.) So long as one modulates one's confidence appropriately (by, for instance, hedging one's conclusion), there is no reason to view either of these uses with suspicion.

To start, let me make clear why having a particular committed belief imposes no obligations. To do this, I want to compare my notion of committed belief to Akeel Bilgrami's (2006) and Annalisa Coliva's (2016). Both Bilgrami and Coliva claim that having a propositional commitment (under which fall committed beliefs) involves accepting an obligation to be held "rationally responsible" (p. 32) for that commitment. We live up to this commitment by acting in the right way. To use an example from Coliva, if I desire to not get wet and I have the committed belief that opening my umbrella will prevent me from getting wet, then I "ought to see [myself] as bound to opening the umbrella" (p. 262) if it is raining. Prima facie, this seems congruent with my notion of committed beliefs. According to my definition, having a committed belief that *p* implies one treats *p* as true in conscious reasoning. And since conscious reasoning is an act, it might seem then my version of committed beliefs is also about accepting an obligation to act in certain ways.

However, my notion of committed belief leaves this kind of obligation out; committed beliefs on their own have almost no distinguishable value. We can see why if we consider committed beliefs about trivial truths.[35] Take committedly believing that a jar of salt contains 200,403 grains on January 15th. Aside from highly unusual situations (say a 'Count the Salt!' gameshow[36]), this belief imposes almost no obligations. Just because one might affirm through reflection that the jar has 200,403 grains does not mean they should see themselves as bound to remember this for future reasoning. In fact, forming and maintaining the belief that one once counted a salt jar on January 15th would in many cases be considered unhelpful. (Say, if one suffers from OCD or an anxiety disorder.)

Now, of course, some beliefs of ours do have more value than others. A belief in the theory of gravitation presumably has great value; perhaps in comparison with others, perhaps inherently. But without a prior context of other beliefs or practical goals, there is no way to tell whether any belief is a 'salt jar' belief or, say, a 'theory of gravitation' belief. Thus, to understand the value of our committed beliefs, I claim that we need to see how they fit together in an actively maintained system.[37] But to understand the value of such a system we must understand the value of conscious reasoning. For a system of committed beliefs is, as I will argue, a way of managing the ways we are physically and temporally limited in our ability to consciously reason.

Conscious reasoning has great instrumental—if not intrinsic—value. This value can be hard to see because of our inherent limitations as physical and temporal beings. Take the core acts of conscious reasoning: conscious consideration and investigative processes. As covered in § 4, we use conscious consideration to become aware of evidence and its salience, and use investigative processes to put oneself in a position to consciously consider. But we can only consciously consider pieces of evidence one at a time, and our investigative processes do not instantly put us in a position to consciously consider things. This limits our capacity to consciously reason. The moment we stop consciously reasoning, the evidence we gathered and consciously considered starts to fade from our minds. However, these limitations are not inherent to conscious reasoning itself, and can be conceptually separated. Suppose that we were not limited physically and temporally. Suppose we could put ourselves in contact with all the evidence we needed instantly and could consciously consider every matter at the same time. In this outlandish scenario, conscious reasoning would have great instrumental—if not intrinsic—value. Moreover, we would no longer have a need for committed beliefs. We would always be in a state of conscious reasoning, rendering passive belief states moot.

---

[35] See Duncan Pritchard et al. (2022) for more.

[36] Thanks to Annalisa Coliva for the example.

[37] Thanks to an anonymous reviewer for pressing me on belief values.

As it stands, we are physical and temporal beings. So, to maintain the value of conscious reasoning, we are obligated to do our best to manage these constraints. It is in living up to this obligation that we transform ourselves from conscious reasoners to *committed systematic believers*. If we cannot consciously reason about everything all at once, we must consciously reason smartly, being sure not to duplicate our efforts. The best way to avoid this duplication is to see ourselves as taking up stances on propositional matters that we shift and update over time: committed beliefs. Yet, because we cannot actively think about all these stances at the same time, we must ensure that they do not contradict, be it simply or in what they imply when taken together. To be able to choose which committed belief to give up when they contradict, we must see some committed beliefs as more important than others.[38] Altogether, these considerations are good enough to imply that we are committed to our beliefs not individually, but as a system: as a loose set of propositional stances that we maintain and revise. This system we are then obligated to maintain because of our (self-imposed) obligation to being committed systematic believers: to tracking and coming to know important truths about ourselves and the world. More precisely:

> SYSTEM OF COMMITTED BELIEFS: A system of committed beliefs is a set of committed beliefs (as defined in § 1) that a person *P* has such that:
>
> 1. *P* has a rough understanding of their importance, with the more important ones underpinning certain valuable endeavors that *P* is undertaking or being considered more 'core' to the system.
> 2. *P* has an obligation to do maintenance on those beliefs, due to their (self-imposed) obligation to track and come to know important truths about ourselves and the world.
>
>
> Doing maintenance on that system involves, among other things, ensuring that a. the beliefs do not contradict, be it simply or in what they imply when taken together and b. are updated when new evidence comes to light.

This also explains the value of reflection: it is one of the main ways we do maintenance on our system of committed beliefs. Reflection allows us to target a particular part of our system—a committed belief—and improve it by consciously reasoning. To be clear, we do not necessarily see ourselves as committed to a system in the sense of an explicit structure:[39] for instance, we do not diagram out all our beliefs in order of importance. (Nor should we, necessarily.) But when we prioritize certain committed beliefs over others in reflection, we implicitly must treat them as systematic in the way I defined above.

To avoid contradicting what was said in § 4.1, it is important to keep in mind that we do not 'store' committed beliefs in our memory. The definition of committed beliefs in § 2 still holds: we are counted as having a particular committed belief because of how we consciously reason

---

[38] It is not clear that we could produce a number for how many beliefs we have. What might count as one belief for one endeavor might count as two for another. Moreover, as Harman (1986) discusses, we probably do not want multiple beliefs about the same topic, causing needless clutter. Instead, we will have more general and vague beliefs which we can use to derive less general ones when necessary.

[39] Thanks to J. Adam Carter for pressing me on this point.

when we reflect. To make sense of this, I propose that we think of committed beliefs and being a systematic committed believer metaphysically as follows:

METAPHYSICS OF COMMITTED BELIEVING: Being a systematic committed believer and having a specific committed belief are (what I will call[40]) *Aristotelian habits.*

ARISTOTELIAN HABITS:

Aristotelian habits are states that:

1. one has passively, in that one has them even when asleep.
2. yet one must maintain actively, in that one must a. repeatedly perform certain effortful actions characteristic of the habit to be credited as having them[41] and b. (potentially) actively maintain certain sub-habits.

Furthermore, consider the relationship between the 'being a systematic committed believer' habit and the 'having committed belief *p*' habit. This relationship is akin to the one between being a flute player and knowing a specific flute piece: of habit to sub-habit. Knowing and playing various flute pieces is part-and-parcel of being a good flute player. And flute players must maintain their ability to perform certain pieces through practice. Yet, no individual flute piece is necessarily required for one to be a flute player. What is required rather is the maintenance of the flute-playing practice.

We can now see why we do not 'store' beliefs. One does, to a certain degree, 'store' the dispositional abilities to do these tasks. Just as flute players acquire motoric skills for specific pieces, so do committed systematic believers store evidence in their memory for specific beliefs. But these are not the same as the committed beliefs themselves. The maintenance of evidence is just the way we accomplish one of the main goals of systematic belief formation: making truth-aligned assertions on important propositional matters.

Nonetheless, to be a good committed believer it is not enough to just reflect on one's beliefs in the here and now and leave it at that. Our connection to the evidence fades over time, thereby degrading the strength of a committed belief the next time we take it up in reflection. We must do something to ensure that our system of committed beliefs stays intact. Thus, maintaining our systems of committed beliefs also involves managing our connection to our evidence: our memories, documents, physical materials, and so forth. By managing this connection, we can continue to have certain committed beliefs when we reflect on them again. It is important then to refresh our connection to the evidence every so often. The simplest way to do this is to exhaustively consciously reason about a belief again. But to excel as a reflective believer it sometimes is necessary to actively manage and store evidence in some manner. Being a good philosopher, for instance,

---

[40] My notion is probably not quite the same as Aristotle's. I was inspired here by Boyle (2011b) where he conceives of an active belief as an "[actualization] of [a person's] capacity for doxastic self-determination." (p. 21) Boyle, however, does not specifically put this matter in Aristotelian terms. For more on Aristotle's notion of habit/active condition, see Sachs (2002, pp. xi–xvii).

[41] This means it is possible for one to lose one's status as a systematic committed believer if one does not consciously reason or reflect enough. This would likely only happen in very rare cases, if at all. Being completely inconsistent in one's committed beliefs would draw social censure. However, one could imagine that if one were rich and privileged enough, one could avoid this. Space constraints prevent a detailed explanation, as the matter is complex.

involves not just thinking about philosophical issues, but also writing down what one has thought about. (Lest we forget.)

Thus, committed beliefs are not just ad-hoc entities based on how we have consciously reasoned. By asking and answering what we did believe, do believe, and what we should believe; by refreshing our connection to the evidence through conscious reasoning, and by actively managing our evidence to ensure future success in conscious reasoning; we knit together individual episodes of conscious reasoning into one enduring performance of being a committed systematic believer. The direct control we have over our individual episodes of conscious reasoning can thus be extended to our system of committed beliefs.

5 | CONCLUSION

In this paper, I have tried to provide a thorough view of how we make up our minds. As part of that, I have explained both the actual process of making up our minds—conscious reasoning—and how we ensure our minds remain made up—our maintenance of a system of committed beliefs. Understood this way, making up our minds is a highly intellectual activity, likely only reserved for animals with very developed language capacities. But since my view does not concern ordinary beliefs, this is acceptable. Making up one's mind is not innate or easy; it is something that we achieve through effort.

This paper has not addressed the various difficulties that confirmation bias and cognitive shortcuts present. But our job as reflective human beings is not to see ourselves as infallibly rational. Reflection is very much about fallibility; recognizing the ways that our minds and our access to evidence are limited. Therefore, it is perfectly compatible with my view that we as reflective thinkers have an obligation to become consciously aware of our biases. By becoming aware of a bias, we are hopefully able to recognize when it arises in our conscious reasoning, thus letting our conclusions be solely based on our intellectual perception of the evidence as it is, rather than the evidence as we want it to be.

ORCID
Joost Ziff https://orcid.org/0009-0007-3104-6975

REFERENCES
Bayne, T., & McClelland, T. (2016). "Finding the Feel": The Matching Content Challenge to Cognitive Phenomenology. *Phenomenology and Mind*, *10*, Article 10. https://doi.org/10.13128/Phe_Mi-20090

Bilgrami, A. (2006). *Self-Knowledge and Resentment*. Harvard University Press.

Block, N. (1995). On a confusion about a function of consciousness. *Behavioral and Brain Sciences*, *18*(2), 227–247. https://doi.org/10.1017/S0140525X00038188

Boghossian, P. (2014). What is inference? *Philosophical Studies*, *169*(1), 1–18. https://doi.org/10.1007/s11098-012-9903-x

Boghossian, P. (2018). Delimiting the Boundaries of Inference. *Philosophical Issues*, *28*(1), 55–69. https://doi.org/10.1111/phis.12115

Boyle, M. (2011a). II—Matthew Boyle: Transparent Self-Knowledge. *Aristotelian Society Supplementary Volume*, *85*(1), 223–241. https://doi.org/10.1111/j.1467-8349.2011.00204.x

Boyle, M. (2011b). "Making Up Your Mind" and the Activity of Reason. *Philosophers' Imprint*, *11*.

Boyle, M. (2019). Transparency and Reflection. *Canadian Journal of Philosophy*, *49*(7), 1012–1039. https://doi.org/10.1080/00455091.2019.1565621

Broome, J. (2013). *Rationality Through Reasoning*. Wiley-Blackwell.

Broome, J. (2014). Normativity in Reasoning. *Pacific Philosophical Quarterly*, *95*(4), 622–633. https://doi.org/10.1111/papq.12050

Burge, T. (1996). Our Entitlement to Self-Knowledge. *Proceedings of the Aristotelian Society*, *96*(1), 91–116. https://doi.org/10.1093/aristotelian/96.1.117

Byrne, A. (2005). Introspection. *Philosophical Topics*, *33*(1), 79–104. https://doi.org/10.5840/philtopics20053312

Byrne, A. (2011). Transparency, Belief, Intention. *Aristotelian Society Supplementary Volume*, *85*, 201–221. https://doi.org/10.1111/j.1467-8349.2011.00203.x

Byrne, A. (2018). *Transparency and Self-Knowledge*. Oxford University Press.

Carruthers, P. (2011). *The Opacity of Mind: An Integrative Theory of Self-Knowledge*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199596195.001.0001

Carruthers, P., & Veillet, B. (2011). The Case Against Cognitive Phenomenology. In T. Bayne & M. Montague (Eds.), *Cognitive Phenomenology* (pp. 35–56). Oxford University Press.

Cassam, Q. (2015). *Self-Knowledge for Humans*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199657575.001.0001

Coliva, A. (2008). Peacocke's Self-Knowledge. *Ratio*, *21*(1), 13–27. https://doi.org/10.1111/j.1467-9329.2007.00381.x

Coliva, A. (2012). One Variety of Self-Knowledge: Constitutivism as Constructivism. In A. Coliva (Ed.), *The Self and Self-Knowledge* (pp. 212–242). Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199590650.003.0011

Coliva, A. (2016). *The Varieties of Self-Knowledge*. London: Palgrave.

Descartes, R. (2017). *Descartes: Meditations on First Philosophy: With Selections from the Objections and Replies* (J. Cottingham, Ed.). Cambridge University Press. https://doi.org/10.1017/9781107416277

Evans, G. (1982). *The Varieties of Reference*. Oxford: Oxford University Press.

Fernández, J. (2013). *Transparent Minds: A Study of Self-Knowledge*. Oxford University Press.

Frankish, K. (2004). *Mind and Supermind*. Cambridge University Press.

Frankish, K. (2018, February 18). *Alief, b-lief, and c-lief (2018)*. Keith Frankish. https://www.keithfrankish.com/presentations/alief-b-lief-and-c-lief/

Gertler, B. (2021). Rational Agency and the Struggle to Believe What Your Reasons Dictate. In C. Borgoni, D. Kindermann, & A. Onofri (Eds.), *The Fragmented Mind* (pp. 325–349). Oxford University Press. https://doi.org/10.1093/oso/9780198850670.003.0014

Harman, G. (1986). *Change in View: Principles of Reasoning*. MIT Press.

Hieronymi, P. (2006). Controlling Attitudes. *Pacific Philosophical Quarterly*, *87*(1), 45–74. https://doi.org/10.1111/j.1468-0114.2006.00247.x

Hieronymi, P. (2008). Responsibility for Believing. *Synthese*, *161*(3), 357–373. https://doi.org/10.1007/s11229-006-9089-x

Hieronymi, P. (2009). Believing at Will. *Canadian Journal of Philosophy Supplementary Volume*, *35*, 149–187. https://doi.org/10.1080/00455091.2009.10717647

Horgan, T. (2011). From Agentive Phenomenology to Cognitive Phenomenology: A Guide for the Perplexed. In T. Bayne & M. Montague (Eds.), *Cognitive Phenomenology* (pp. 58–78). Oxford University Press.

Horgan, T., & Tienson, J. (2002). The Intentionality of Phenomenology and the Phenomenology of Intentionality. In D. J. Chalmers (Ed.), *Philosophy of Mind: Classical and Contemporary Readings* (pp. 520–533). Oxford University Press.

Husserl, E. (2008). *On the Phenomenology of the Consciousness of Internal Time* (J. B. Brough, Trans.; Softcover reprint of the original 1st ed. 1991 edition). Springer.

Kirsch, J. (2018). Interpreting Things Past. In P. Pedrini & J. Kirsch (Eds.), *Third-Person Self-Knowledge, Self-Interpretation, and Narrative* (pp. 99–113). Springer International Publishing. https://doi.org/10.1007/978-3-319-98646-3_6

Kirsch, J. (2020). Interpreting our emotions. *Ratio*, *33*(1), 68–78. https://doi.org/10.1111/rati.12251

Kornblith, H. (2012). *On Reflection*. Oxford University Press.

Kriegel, U. (2015). The Character of Cognitive Phenomenology. In T. Breyer & C. Gutland (Eds.), *Phenomenology of Thinking: Philosophical Investigations Into the Character of Cognitive Experiences* (pp. 25–43). Routledge.

Lawlor, K. (2009). Knowing What One Wants. *Philosophy and Phenomenological Research*, *79*(1), 47–75. https://doi.org/10.1111/j.1933-1592.2009.00266.x

Lennon, P. (2022). Cognitive phenomenology: In defense of recombination. *Inquiry*, *0*(0), 1–30. https://doi.org/10.1080/0020174X.2022.2107060

Lennon, P. (2023). Aphantasia and Conscious Thought. In U. Kriegel (Ed.), *Oxford Studies in Philosophy of Mind* (Vol. *3*).

Leonard, N. (2023). Epistemological Problems of Testimony. In E. N. Zalta & U. Nodelman (Eds.), *The Stanford Encyclopedia of Philosophy* (Spring 2023). Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/spr2023/entriesestimony-episprob/

Lormand, E. (1996). Nonphenomenal Consciousness. *Noûs*, *30*(2), 242–261. https://doi.org/10.2307/2216295

Macbeth, D. (2014). *Realizing reason: A narrative of truth and knowing*. Oxford University Press.

Macdonald, C. (2014). In my "Mind's Eye": Introspectionism, detectivism, and the basis of authoritative self-knowledge. *Synthese*, *191*(15), 3685–3710.

McHugh, C. (2012). Reasons and Self-Knowledge. In *The Self and Self-Knowledge* (pp. 139–163). Oxford University Press. https://oxford.universitypressscholarship.com/10.1093/acprof:oso/9780199590650.001.0001/acprof-9780199590650-chapter-7

McHugh, C., & Way, J. (2018). What is Reasoning? *Mind*, *127*(505), 167–196. https://doi.org/10.1093/mind/fzw068

Moran, R. (2012). Self-Knowledge,'Transparency', and the Forms of Activity. In D. Smithies & D. Stoljar (Eds.), *Introspection and Consciousness* (pp. 211–236). Oxford University Press.

Moran, R. A. (2001). *Authority and Estrangement: An Essay on Self-Knowledge*. Princeton University Press.

Munroe, W. (2021). Reasoning, rationality, and representation. *Synthese*, *198*(9), 8323–8345. https://doi.org/10.1007/s11229-020-02575-6

Munroe, W. (2022). Why are you talking to yourself? The epistemic role of inner speech in reasoning. *Noûs*, *56*(4), 841–866. https://doi.org/10.1111/nous.12385

Nagel, T. (1974). What is It Like to Be a Bat? *Philosophical Review*, *83*(October), 435–450. https://doi.org/10.2307/2183914

Pautz, A. (2013). Does Phenomenology Ground Mental Content? In U. Kriegel (Ed.), *Phenomenal Intentionality* (pp. 194–234). Oxford University Press.

Perry, J. (1979). The Problem of the Essential Indexical. *Noûs*, *13*(1), 3–21. https://doi.org/10.2307/2214792

Pippin, R. (2010). Lightning and Flash, Agent and Deed. In *Friedrich Nietzsche: Genealogie der Moral* (pp. 47–63). Akademie Verlag. https://doi.org/10.1524/9783050050287.47

Pitt, D. (2004). The Phenomenology of Cognition: Or What is It Like to Think That P? *Philosophy and Phenomenological Research*, *69*(1), 1–36. https://doi.org/10.1111/j.1933-1592.2004.tb00382.x

Pitt, D. (2011). Introspection, Phenomenality, and the Availability of Intentional Content. In T. Bayne & M. Montague (Eds.), *Cognitive Phenomenology* (pp. 141–173). Oxford University Press.

Prinz, J. (2011). The Sensory Basis of Cognitive Phenomenology. In T. Bayne & M. Montague (Eds.), *Cognitive Phenomenology* (pp. 174–196). Oxford University Press.

Pritchard, D., Turri, J., & Carter, J. A. (2022). The Value of Knowledge. In E. N. Zalta & U. Nodelman (Eds.), *The Stanford Encyclopedia of Philosophy* (Fall 2022). Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/fall2022/entries/knowledge-value/

Robinson, W. S. (2005). Thoughts Without Distinctive Non-Imagistic Phenomenology. *Philosophy and Phenomenological Research*, *70*(3), 534–561. https://doi.org/10.1111/j.1933-1592.2005.tb00414.x

Robinson, W. S. (2011). A Frugal View of Cognitive Phenomenology. In T. Bayne & M. Montague (Eds.), *Cognitive Phenomenology* (pp. 197–214). Oxford University Press.

Ross, B. (2016, May 19). *Imagine a dog. Got it? I don't. Here's what it's like to be unable to visualize anything.* Vox. https://www.vox.com/2016/5/19/11683274/aphantasia

Sachs, J. (2002). Introduction. In J. Sachs (Trans.), *Nicomachean Ethics* (pp. xi–xxv). Focus Publishing.

Schwitzgebel, E. (2021). The Pragmatic Metaphysics of Belief. In *The Fragmented Mind* (pp. 350–375).

Shah, N., & Velleman, J. D. (2005). Doxastic Deliberation. *The Philosophical Review*, *114*(4), 497–534. https://doi.org/10.1215/00318108-114-4-497

Sosa, E. (2015). Descartes's Pyrrhonian Virtue Epistemology. In *Judgment and Agency* (pp. 233–254). Oxford University Press. https://doi.org/10.1093/acprof:oso/9780198719694.003.0011

Strawson, G. (2003). Mental Ballistics or the Involuntariness of Spontaneity. *Proceedings of the Aristotelian Society*, *103*, 227–256.

Tye, M., & Wright, B. (2011). Is There a Phenomenology of Thought? In T. Bayne & M. Montague (Eds.), *Cognitive Phenomenology* (pp. 326–344). Oxford University Press.

Valaris, M. (2017). What reasoning might be. *Synthese*, *194*(6), 2007–2024. https://doi.org/10.1007/s11229-016-1034-z

Valaris, M. (2019). Reasoning and Deducing. *Mind*, *128*(511), 861–885. https://doi.org/10.1093/mind/fzy025

Winokur, B. (2023). How to commit to commissive self-knowledge. *European Journal of Philosophy*. https://doi.org/10.1111/ejop.12862