

A CASE FOR EPISTEMIC AGENCY

Dustin OLSON

ABSTRACT: This paper attempts to answer two questions: What is *epistemic agency*? And what are the motivations for having this concept? In response to the first question, it is argued that epistemic agency is the agency one has over one's belief-forming practices, or doxastic dispositions, which can directly affect the way one forms a belief and indirectly affect the beliefs one forms. In response to the second question, it is suggested that the above conception of epistemic agency is either implicitly endorsed by those theorists sympathetic to epistemic normativity or, at minimum, this conception can make sense of the legitimacy of the normative notions applicable to how and what one should believe. It is further contended that belief formation in some respects is a skill that can be intentionally developed and refined. Accepting this contention and the existence of certain epistemic norms provide inconclusive yet good reasons to endorse this concept. Recent challenges to this concept by Hillary Kornblith and Kristoffer Ahlstrom-Vij are also considered.

KEYWORDS: belief formation, belief control, epistemic agency, epistemic normativity

Introduction¹

In this paper I respond to two questions: What is *epistemic agency*? And what are the motivations for having this concept? I argue that epistemic agency (EA) is the agency one has over one's belief-forming practices, which will directly affect the way in which one forms belief and indirectly affects the beliefs one forms. There are a number of reasons to adopt this concept. First, it is implied in extant theories of epistemic normativity. For example, if one accepts that there are epistemic virtues or epistemic duties, then one is implicitly accepting some notion of EA. A second reason to adopt EA is that belief formation is in some sense a skill; we can improve as believers. Doxastic self-improvement and skilled belief formation are marks of EA. I suggest, then, that there are inconclusive yet good reasons to endorse this concept.

Recently some have registered skepticism about the viability of EA. Hilary Kornblith, e.g., regards EA as mythological because we do not control the actual

¹ Thank you to audience members at the University of Rochester and at Cornell University's Ethics of Belief Workshop for helpful comments and questions. And special thanks to the two anonymous referees, Patrick Bondy, Richard Feldman, Kate Nolfi, Sharon Ryan, and John Turri for helpful comments and discussions.

reason-responsive mechanistic processes that bring about a belief.² If we cannot control that aspect of belief formation, he argues, we are not agents with regards to belief formation. A second challenge to EA is offered by Kristoffer Ahlstrom-Vij.³ He argues that even if we can give an account of EA that avoids Kornblith's challenge by appealing to higher-order processes like reflective reasoning and deliberation, there are reasons to reject this concept. Ahlstrom-Vij contends that a key motivation for having a concept like EA is that we can practice epistemic self-improvement; there are important ways that we cannot self-improve epistemically, however. Because of this limitation, we cannot then be said to have EA. I agree with Kornblith that we do not have direct control over the lower-level mechanistic components of belief formation. I also agree with Ahlstrom-Vij that epistemic self-improvement is one of the key motivations for developing and accepting a concept like EA. I disagree, however, with both of their conclusions. This paper proceeds, then, as follows. First I present and develop my conception of EA. I next make a case for why we should accept this conception. And finally, I address and reject the aforementioned arguments against EA.

1. Epistemic Agency

I propose that epistemic agency is distinct from other notions of belief control. It does not require one's ability to choose one's belief or one's ability to decide to believe – e.g. I want to believe that p , so I believe that p . Nor is it a concept that simply requires one's ability to control or influence the formation of specific beliefs – e.g. by A -ing I will come to believe that p .⁴ Rather, I contend that epistemic agency denotes the motivation and ability to refine and alter one's belief-forming methods and subsequent belief-forming practices – these methods and practices can collectively be understood as a *doxastic disposition* or belief-forming abilities, which can also be characterized as one's propensity to form true or false or coarse-grained or fine-grained beliefs within different domains. Epistemic agents have the ability to hone and refine their belief-forming abilities through altering their doxastic dispositions.⁵ In so doing, one can affect one's beliefs, but neither directly nor specifically.

² Hilary Kornblith, *On Reflection* (Oxford: Oxford University Press, 2012).

³ Kristoffer Ahlstrom-Vij, "Why We Cannot Rely on Ourselves for Epistemic Improvement," *Philosophical Issues* 23, 1 (2013): 276-296.

⁴ This process is certainly something an agent does. But, as will be clarified below, I distinguish this type of voluntary belief control from EA.

⁵ Recent works on belief control has made appeals to dispositional characteristics as playing, or potentially playing, an active role in the beliefs or judgments we make. Kate Nolfi, "Why is

The above suggests that EA is distinct from other belief-control concepts. The concepts I have in mind fall under the heading of doxastic voluntarism, which comes in two forms: direct and indirect. *Indirect doxastic voluntarism* (IDV) is the thesis that we can form beliefs through indirect voluntary control. For example, I want to form the belief that the lights are on in my office, so I go to my office and turn the lights on. As long as my perceptual apparatus is working as it should, that I turn on the lights and observe their illumination will result in my believing that the lights in my office are on.⁶ *Direct doxastic voluntarism* (DDV) is the thesis that we can form our beliefs through direct voluntary control. For example, if one were able to believe that p at will – decide to believe that p – then DDV would be true.⁷ Closely related to DDV is *doxastic freedom*. For one to have doxastic freedom is for one to have the freedom to exert direct voluntary control over one's doxastic attitudes.⁸

I argue that EA is not reducible to IDV or DDV, and that EA does not require that the agents in question have doxastic freedom. EA neither reduces to DDV nor requires one to have doxastic freedom. I agree with the majority of philosophers who conclude that both of these theses are implausible. One cannot simply decide to believe; one cannot freely choose one's doxastic attitude. That is, doxastic states are not the types of things that we have direct agential control over.

Epistemic Evaluation Prescriptive?" *Inquiry: An Interdisciplinary Journal of Philosophy* 57, 1, (2014): 97-121, e.g., defends a causal account of doxastic control, wherein one's adoption of how one ought to believe causally effects how we actually believe. One such way this might play out, which Nolfi finds promising but falls short of endorsing, is that our judgment about what we ought to believe exerts "causal power in shaping the ways in which we are disposed to regulate our beliefs" (110). She dubs this disposition to regulate belief as one's *cognitive character*. Conor McHugh, "Epistemic Responsibility and Doxastic Agency," *Philosophical Issues* 23, 1 (2013): 132-157, also appeals to dispositions as way to combine reason responsiveness with agency – or, in his terms, doxastic agency – in his "Reasons-Action Principle: Responding to a putative reason involves conscious agency, at least dispositionally" (146). Where one's disposition is in part the ability to respond to reasons through recognizing them as reasons as such; recognition and responsiveness are not passive and thus require agency. While both of these theories of belief control are not inconsistent with the one presently defended, my proposal is broader.

⁶ This example is from Richard Feldman, "The Ethics of Belief," *Philosophy and Phenomenological Research* 60, 3 (2000): 671-672.

⁷ This view is not widely accepted, although it does have defenders. See, e.g., John Turri, David Rose, and Wesley Buckwalter, "Choosing and Refusing: Doxastic Voluntarism and Folk Psychology," unpublished, Matthias Steup, "Doxastic Freedom," *Synthese* 161 (2008): 375-392, and Patrick Bondy, "Epistemic Deontology and Strong Doxastic Voluntarism: A Defense," *Dialogue* (2015), DOI: <http://dx.doi.org/10.1017/S0012217315000487>.

⁸ A defense of doxastic freedom can be found in Steup, "Doxastic Freedom."

The question then becomes: where is the agential control that I am suggesting we have in the formation of a belief?

The control needed for EA is more closely related to IDV than DDV. I am not proposing that EA is simply a form of IDV, however. In IDV, one acts with the goal to form a specific belief, as in flicking the light switch so as to form the belief that the light is on. EA does not require that a specific belief be formed. One does not exercise one's epistemic agency with the goal of forming a specific belief; when one exercises one's epistemic agency one has the goal of forming beliefs in a specific way - e.g. to maximize truth while avoiding falsity or to form more fine-grained beliefs given a specific body of evidence. EA is similar to IDV, however, in that IDV merely requires normal agency for the desired belief to be formed. Likewise, to adjust one's doxastic disposition, one merely needs agency as we commonly understand it.

Consider, for example, candidate higher-level processes involved in belief formation, processes like deliberation, reflection, attentiveness, and intentions. These processes are the types of things we can have control over. High-level processes coupled with the lower-level mechanistic reason-responsive processes of coming to hold or not hold certain beliefs, then, exemplify EA. That is, we can control certain processes as a means to altering how we believe.⁹ Reflection and deliberation, e.g., are types of practices that can influence belief formation and are things we can control.¹⁰ The role that these processes can have in EA is through one's attempts to achieve one's epistemic goal through reflecting and deliberating on evidence or through what one has accepted as sufficient or insufficient evidence or through what one accepts as an epistemically virtuous principle by which to achieve one's epistemic goals. This reflection does not facilitate one's choosing what to believe; rather, it informs one's belief-forming process for *how* to believe given one's epistemic goals – *viz.* purposefully undertaking steps that will alter our doxastic dispositions. Let us turn to an example.

⁹ Recall that my proposal does not say that we act to form beliefs about specific propositions. In this way I would distance my view from something like Pascal's proposal, as he suggests that we can act so as to believe that God exists and thereby come to believe the specific proposition "God exists." I do not think that we can act to believe in this way. I suspect the only plausible ways in which we can act to believe are uninteresting IDV cases.

¹⁰ There may be one way to slant this so that it seems as though my proposal requires that we do seek out belief in specific propositions. If our world is consistent with classical logic, then if one were to have the goal to form only true beliefs, one's forming a true belief about a specific proposition could be seen as one attempting to believe that exact proposition because it is the only proposition that would satisfy the inquiry in question. That is, one would have intended to form that belief in virtue of its being true.

Consider Campbell's desire to form accurate beliefs about scotch so as to be able to converse with other aficionados.¹¹ Currently, Campbell's beliefs are not fine grained concerning the different qualities he experiences when drinking scotch. His desire to form these more fine-grained beliefs motivates him to undertake certain acts that will refine his palate – e.g. he reads the taster's notes for each scotch he drinks, familiarizes himself with the jargon associated with the different scotch-drinking steps, and is phenomenologically sensitive to the experiences he has while drinking scotch. Eventually Campbell's beliefs concerning the different scotches he samples become fine grained. Thus, rather than simply forming coarse-grained beliefs, like merely having the ability to distinguish between a scotch whisky and a rye whiskey, Campbell can now form more fine-grained beliefs, such as the age, region, and casking processes of the different scotches he samples. Observe, then, that Campbell neither chooses to believe any specific proposition, nor exercises doxastic freedom in order to implement the agency associated with EA. The agency Campbell exercises is the straightforward agency we associate with action planning and goal satisfaction. The present example is just a warm up as to how we might think about intentionally altering our doxastic disposition so as to alter the way we form beliefs. More details emerge below.

2. The Case for Epistemic Agency

In what follows I make the case for EA. In making this case, I first discuss how EA can make sense of epistemically normative notions like epistemic virtues and epistemic duties. Those who reject epistemic normativity will not find this defence convincing; those who accept epistemic normativity will already in some sense endorse my proposal. There are additional reasons to accept EA, however. Some ways in which we form beliefs can be understood as a skill, an ability that can be developed – i.e. we can intentionally attempt doxastic self-improvement and achieve our epistemic goals. As forming beliefs can be a developed ability, it requires some type of control – the type I suggest is found in my formulation of EA. So, even if one remains unconvinced regarding epistemic norms, the fact that we can set and achieve some epistemic goals, i.e., we can undertake steps for doxastic self-improvement and become more skilled epistemically, is further reason to adopt EA.

¹¹ This implies that practical interests can motivate some of our epistemic interests. Some may find this difficult to accept. I don't.

2.1 Epistemic Agency and Epistemic Normativity

One reason to accept the current portrayal of EA is that it can help us make sense of different normative notions in epistemology. It is plausible that we can form goals or have motivations for how we come to form beliefs or the practices we use to form true beliefs. There is also reason to think that many of our candidate theories of justification have a motivational element. Consider two normative epistemic ideas: epistemic virtues and epistemic duties. I propose that if we accept, even if only in part, that there are epistemic virtues or epistemic duties, then some type of EA is required. One reason to accept EA, then, is that it is already implied in a number of theories.

EA and Epistemic Virtues

It is clear that much of virtue epistemology takes for granted the type of agential process that I am here espousing. According to a number of virtue epistemologists, cognitive or epistemic virtues have a motivational component. These virtues are habits developed with the motivation for knowledge, which in virtue terms is true belief that results from a cognitive virtue. Regardless of whether we accept this definition of knowledge, there do seem to be epistemic virtues, and they do seem to be motivated by the desire to acquire true beliefs. Thus, part of the appeal of cognitive virtues is the motivational component to become a more virtuous believer. This motivational component is analogous to the motivational aspect of EA: that one desires - has some motivation - to refine or improve one's belief-forming habits, in general or over a specific domain. Compare this proposal with Linda Zagzebski's description of the motivational component of a cognitive virtue:

The primary motivation underlying the intellectual virtues is the motivation for knowledge. Such a motivation clearly includes the desire to have true beliefs and to avoid false ones, and... such a motivation leads a person to follow rules or procedures of belief formation that are to her epistemic community to be truth conducive.¹²

On this view, then, if it we can acquire epistemic virtues, then we must have the ability to develop these cognitive habits. Or consider John Greco's characterization of a cognitive virtue, which again highlights the importance of the agent's role in acquiring and fostering a virtue:

A virtue, in one important sense, is an ability. An ability, in turn, is a stable disposition to achieve certain results under certain conditions. Further, when we

¹² Linda Zagzebski, "Virtues of the Mind (selections)," in *Epistemology: An Anthology*, ed. Ernest Sosa and Jaegwon Kim (Blackwell Publishing, 2000), 465.

say that a subject *S* has an ability to achieve certain results, we imply that it is no accident that *S* achieves those results. *S*'s disposition to achieve the relevant results is grounded in certain properties of *S*, such that under the appropriate conditions any subject with those properties would tend to achieve those results.¹³

Again, it seems clear that EA captures the idea that we have an ability to achieve certain doxastic goals via developing one's doxastic disposition – or as the virtue theorists term it, to form cognitively virtuous habits.¹⁴

EA and Epistemic Duties

Epistemic deontologism holds that there are things we epistemically *ought to* or *ought not* believe.¹⁵ That is, we have epistemic duties to believe in a certain way. Much like the motivational component of epistemic virtues, epistemic duties also have a motivational component. Here, however, the epistemic goal is to form and maintain appropriate beliefs, according to the duty one has to believe in the appropriate way. That is, one ought to believe that *p* when it is epistemically appropriate to do so. From the evidentialist perspective, for example, if one's evidence supports that *p*, then one ought to believe that *p*; if one's evidence does not support that *p*, then one ought not believe that *p*. Thus, if one is not suitably attuned to one's evidence, then one is falling short of one's epistemic duty.¹⁶ If we have EA, however, then we can develop our abilities such that we do believe in the right way, *viz.* we form beliefs that are sensitive/responsive to the right evidence.

¹³ John Greco, "Virtue and Vices of Virtues Epistemology," in *Epistemology: An Anthology*, ed. Ernest Sosa and Jaegwon Kim (Blackwell Publishers, 2000), 468.

¹⁴ I am not defending Virtue Epistemology. I am merely attempting to show that if there are epistemic virtues – i.e. a motivation and habitual ability to form true beliefs and avoid false ones – then they are best understood in terms of our having EA.

¹⁵ The issues involved in deontological epistemology are complicated. Sorting through these issues would take us well beyond the scope of this paper. For a discussion concerning some of these issues see Feldman, "Epistemic Deontologism." I merely here wish to highlight that if one prefers to think of epistemic justification in deontological terms, then EA is for that person.

¹⁶ The *ought* being proposed here is epistemic; epistemic deontologists are concerned with what one epistemically ought to believe. This *ought* is distinct from other types of *ought*, like moral or prudential oughts, although some have made the case that the two are more closely related than they might seem. The *locus classicus* here is William K. Clifford, "The Ethics of Belief," in *The Theory of Knowledge*, 3rd edition, ed. Louis Pojman (Canada: Wadsworth, 2003), 515 - 518. For a contemporary defense of Clifford's proposal see Sharon Ryan, "In Defense of Moral Evidentialism," *Logos & Episteme. An International Journal of Epistemology* VI, 4 (2015): 405-427.

A common complaint about epistemic deontology is that it seems to require a level of voluntary control over our beliefs that we do not have. The argument runs as follows:

1. If we have a duty to x , then we can be held responsible for whether or not x .
2. We cannot be held responsible for something we do not have voluntary control over.
3. We do not have voluntary control over our beliefs.
4. Therefore we cannot be held responsible for whether or not we believe that p
5. Therefore we do not have epistemic duties.

In response to this argument, I suggest that defenders of epistemic deontology should, or implicitly do, accept my version of EA. EA provides a way to avoid the above criticism without having to deny the quite plausible first premise. Epistemic deontologists can reject premise (2). We can be held responsible for our beliefs because there are things we can do that can affect them - we therefore do have the right kind of control to allow for epistemic duties.

The kind of control that we have over our beliefs is indirect. Furthermore, this control is voluntary only in the sense that we can undertake certain practices of our own volition; it is not the type of voluntary control found in (2). The type of voluntary control being referred to in (2) suggests that voluntary control occurs over the formation of specific beliefs by an act of will – *viz.* DDV. This is not the type of control being defended in EA. Regardless, the type of control being defended in EA is sufficient for responsibility.

Take one who has a skewed view of the evidence for a certain set of beliefs. Suppose Jones, e.g., believes that the university is conspiring against his daughter, which explains why she is doing poorly in her classes. In Jones's mind, his daughter's lack of success is not due to an absence of effort or skill on her part, but rather the result of a conspiracy between faculty and administration to ruin the student's academic career. The evidence that Jones uses to justify these beliefs is based on grade reports and instructor comments. However, he also has competing evidence: he attended orientation; he observed all of the student-success-based programs offered by the university – e.g. a well-stocked library with lots of quiet study space, subject tutors, mental health specialists, and so on; and he also attended college and is thus able to appreciate the rigors of adjusting to freshman year. If Jones were to be more impartial, he would further recognize that it is against the university's interests to have students do poorly. Because the evidence

is such that the father's conclusion is clearly false and because he himself has sufficient evidence supporting the falsity of his belief, we can conclude that he shouldn't believe that the university is conspiring against his daughter. His judgment of the evidence is at once incorrect and something he could change with a more disinterested reflection. This is a case where the believer has all the evidence needed for the correct belief, yet forms the wrong conclusion by unduly giving too much weight to an inconclusive or defeated subset of available evidence. He could and should believe otherwise.

We can also think of cases where a belief is formed too hastily on insufficient evidence or a belief is not formed when it should have been due to one's having sufficient evidence. Suppose Black visits Germany and happens to see a black cow. She infers from this observation that all the cows in Germany are black. This conclusion seems too hasty. Although there is some evidence supporting it, Black should not hold this belief because her evidence is seemingly insufficient. Conversely, suppose Green is an agricultural expert. She has experienced many European cow breeds. She has observed that Germany has almost exclusively Holsteins, which have a spotted coloring.¹⁷ She is quite skeptical however, and does not conclude that most of Germany's cows are spotted, even though she has epistemically sufficient reasons to form this belief. Green should believe that most German cows have a spotted coloring.

Being more objective or disinterested in how one interprets one's evidence upon reflection would prevent Jones forming the false belief about his daughter's university. Adopting a more skeptical attitude would prevent Black from forming beliefs too hastily and becoming less skeptical will prevent Green from unwarranted agnosticism. These dispositional characteristics are the types of things we can alter over time, indirectly affecting the beliefs that we form.¹⁸ We can voluntarily take steps to become more attuned to what the existing evidence is, to be more sensitive to what that evidence supports, and more apt to form the appropriate doxastic attitude given that evidence without having to voluntarily believe a specific proposition. These steps are facilitated by EA. By accepting EA, defenders of epistemic deontology are able to avoid this standard criticism made against their view.

¹⁷ Germany in fact has over 40 breeds of cow, but for the sake of the example we can assume that it has almost exclusively Black Holsteins.

¹⁸ How we can alter these doxastic dispositions and the types of steps we can take in so doing were exemplified in §1, with Campbell's goal to form more refined beliefs about scotch. This type of process is more fully discussed in §2.2 below.

It might be countered that although it is plausible that we can do things to affect how we form beliefs, which may allow for the possibility of epistemic duties, it is strange to praise or blame someone for her belief. As Richard Feldman observes, “praising someone by saying something like, ‘That was a really great bit of believing you did there’ sounds bizarre.”¹⁹ And while it does sound bizarre to praise or blame someone in such a direct manner for what one believes, there are other similar ways in which the praise or blame is implicit. It is not nearly as bizarre sounding to hear something like, “I sure am glad that Judy is on the murder case; she can get to the truth like no other” or “John sure knows how to pick ‘em!” implying that John is a poor judge of character and is bad at forming beliefs about those with whom he associates. We do assess the beliefs of others. That believers have EA allows for such assessments. So although we can recognize that baldly asserting praise for a belief sounds bizarre, there are ways of attributing responsibility and assessing one’s beliefs and the ways in which one forms one’s beliefs that are not so bizarre. I suggest that the appropriateness of this latter possibility is due to EA.

It appears, then, that if we accept some form of epistemic normativity, whether by accepting that there are ways we should believe or ways in which we should form beliefs, we are accepting some form of agency. Accepting that satisfying these motivations does not require the ability to have direct voluntary control over our beliefs, there thus seems to be some plausibility for EA as developed here.

2.2. Doxastic Self Improvement and Skilled Belief

Not all will be convinced by epistemic normativity, however. Fortunately, there is another reason to accept the conception of EA here defended: belief formation is a skills-based endeavor. Some people are better than others at forming beliefs in different areas and in different ways. This observation suggests, then, that some beliefs can be formed skillfully. *A fortiori*, belief formation can be viewed in some respects as a skill. If we accept that some people are more skilled, that is, are better than others in forming true or more detailed beliefs in certain domains, and these skills result in part from the higher-level processes discussed above, then one’s belief-forming abilities can be intentionally refined and developed. If one can intentionally do anything, then agency must be involved somewhere in the process.²⁰ One’s ability to alter one’s doxastic disposition enables the capacity for

¹⁹ Feldman, “Epistemic Deontology,” 353.

²⁰ I assume that intentions imply agency. I therefore leave this claim unsupported. A brief discussion of intentions and doing something intentionally is included in §3.2 below.

one to improve one's skills that affect one's belief-forming abilities. In other words, we can undertake doxastic self-improvement in some ways as a result of EA.

On my formulation, while not sufficient for EA, phenomena like deliberation, reflection, and other dispositional characteristics involved in the belief-forming process are paradigm representations of the types of higher-level processes involved in EA. These processes can influence the lower-level processes, that we do not directly control, which then will affect what one comes to believe through how one forms one's beliefs. Take deliberation, for example. Deliberation is something we do. Focusing our attention on different pieces of evidence, seeking further evidence prior to committing to a belief, and further reflecting on reasons to believe can all be included as elements of the deliberative process. Suppose I am a jury member and must form a belief about the defendant's guilt or innocence. I do not passively accept all of the information that is presented to me during the trial. I deliberate and weigh the evidence presented; I play a more active role in the belief-forming process. Ultimately a belief will emerge from these deliberations, one that at the lower levels I cannot control, one that is reason responsive. Regardless, this type of indirect influence remains something we do and it can affect how and what we believe. We thus have reason to accept the potential influence of our higher-level processes and the mechanistic aspects of belief formation.

That higher-level processes can influence belief formation speaks to another reason for accepting EA, and also involves defending my second claim from above: that belief acquisition can be skillfully undertaken. We can intentionally improve some of the ways in which we come to hold our beliefs, just as we can refine a number of other characteristics. A carpenter's ability to hammer nails into boards with one swing is a skilled movement. By the simple fact that the carpenter's job involves hammering nails often, she improves this skill over time. Suppose, however, that a non-carpenter wishes to strike nails with similar skill and efficiency. What should the non-carpenter do to achieve this goal? The obvious response is that the non-carpenter do what the carpenter has done, namely, hammer a lot of nails. The non-carpenter can intentionally undertake the process of acting like the carpenter to eventually become a skilled nail-striker.

I submit that honing our abilities as belief formers in certain areas is similar to refining one's ability to strike nails. The carpenter's ability to hammer the nail was not something consciously acquired; the ability emerged over time as a result of hammering nails. Most of us form many of our beliefs in ways similar to the carpenter hammering nails. We simply form beliefs as part of our role as belief

formers. In some cases we acquire more skill and refinement in forming beliefs simply because of where our attention is most often focused – e.g. generally a judge will be able to form more accurate and refined beliefs about a person’s character than could a plumber, whereas a plumber will form more accurate and refined beliefs about piping than a judge. And, like the apprentice who desires to improve his skill as a nail-striker, so too can we develop skills as believers by using similar methods of imitation and emulation. If Jay wants to form accurate beliefs in a given field, Jay can emulate the experts within that field, practicing their belief-forming habits, and thus improve her own belief-forming habits. In some cases this improvement may simply be a matter of recognizing and gathering the relevant evidence. In other cases the improvement might concern the deliberative process, where no further evidence is needed; one’s ability to judge one’s extant evidence has simply become more fine-grained.

If improvement in nail striking is found in the repetitive emulation of the journeyman carpenter, how is this type of practice possible for belief formers? Recall the case of Scotch-drinking Campbell. Campbell’s attempt at doxastic self-improvement is to emulate expert scotch drinkers by familiarizing himself with the lingo and by extensive epistemic research – i.e. self-reflectively drinking a bunch of different scotches. We observe that this process is a combination of acquiring more information and also reinterpreting the existing information one has. The new information includes, e.g., the scotch-drinkers’ lingo, the different steps included in sampling scotch (nose, palate, and finish), and the nuances found in different scotch-distilling regions. The extension of EA is not merely acquiring new information, however. Consider the existing information – *viz.* the scotch itself. In addition to the new information presented above, Campbell becomes more phenomenologically aware of the extant experiences. By refining his ability to experience scotch tasting, Campbell refines the judgments and subsequently the beliefs that are formed on the basis of these judgments.

We can extend this type of practice to more abstract belief-forming practices. Suppose I want to form more accurate and refined beliefs about human character. How do I imitate the practices and habits of a proven judge? Suppose Judy, a proven judge, gathers evidence by studying the cues she takes from testimony, body language, and the plausibility of certain cases as described by the person whose character is being assessed. Further, after she has collected this evidence, she then processes it by reflecting on it, comparing it to other cases, and by training herself to do so disinterestedly and with proper proportions. If one were to adopt these tactics, over time one would improve one’s abilities as a judge of character. Here, the development of one’s doxastic disposition does not require

the gathering of more information; rather, the development simply requires one to be able to better understand the information one already has access to, like testimony and body language. Certain traits such as disinterestedness, patience in adopting a belief, and properly proportioning the weight given to certain types of evidence would each be aspects of our doxastic dispositions that can be, to some extent, under our control.²¹

The previous two examples require an extended period of time, or broader extension of EA, to alter one's disposition. This needn't be the case, however. More immediate, narrower, applications of EA can occur also. Consider Sophie, a sophomore philosophy student. She is at a social gathering and overhears a senior student have a slip of the tongue, using *ad hoc* instead of *ad hominem*. From this testimony, Sophie forms the belief that *ad hoc* is the fallacy of attacking the person, not the argument. When misusing this phrase, Sophie is corrected by Carrie, who is Sophie's peer. Sophie argues with Carrie about this point but comes to discover that she has formed a faulty belief. Due to this experience, Sophie decides that she should be more judicious with the amount of credence she gives to different types of evidence. She thereby becomes less likely to form rigid conclusions given limited or weak evidence. This way of improving our doxastic disposition can have a nearly immediate effect. Second-year Sophie's judgments will be much less hasty from this point forward.²²

What can we take away from the above examples? First, we observe that the seasoned judge who forms accurate beliefs about character and the seasoned scotch drinker who forms fine-grained beliefs about whisky are like the carpenter who has developed nailing skills as a result of her job. Likewise, the individual who intentionally approaches character judgment like the seasoned judge or who intentionally refines his palate for the nuances of scotch, so as to form more

²¹ Contrast this example with Jones the suspicious father example from §2.1 above. Jones believes a false proposition on the basis of poor evidence selection. Were he to be more objective towards the situation more generally and less hasty to form his judgement that the university is 'out to get' his daughter, Jones would not have formed this belief. These types of practices are the types that we can influence.

²² John Turri, in conversation, has suggested that this type of decision to not be so open to weak testimonial evidence is sufficient for DDV. Turri suggests that the adoption of this attitude is done directly and has an immediate impact on what one believes. This decision and immediate impact, according to Turri, is analogous to deciding to believe. In other words, my proposal might be construed as a type of direct voluntarism. This is a conclusion Turri would endorse as a defender of DDV. I am not a proponent of DDV and do not wish to take on the burden of defending that view. So if Turri is correct, and my arguments work, then so much the better for defenders of DDV. I do not conceive of EA in this way, however. And I doubt that any of those appealing to this concept would accept DDV either.

accurate or fine-grained beliefs, are akin to the apprentice who intentionally emulates the practices of the carpenter to become better at striking nails. And while many of our skills as belief-forming agents are developed unconsciously, that we can consciously undertake certain processes to become more skilled at forming beliefs - *viz.* intentionally emulating those whose skills are already developed in the desired area - provides us with reason to accept EA.

3. Arguments Against Epistemic Agency

I next consider two challenges to EA. The first is Hillary Kornblith's appeal to the mechanistic lower-level process of belief formation that we cannot directly control. The second is Kristoffer Ahlstrom-Vij's contention that we are unable to epistemically improve in important ways.

3.1 Kornblith and the Mechanism of Belief-Formation

Recently Kornblith has taken issue with the notion of reflection as providing the philosophical fruit that some believe it bears. EA is one of the concepts subject to scrutiny within this grander project. Kornblith is never really clear on what he means by *epistemic agency*, i.e. the view that he is attacking, but seems to conclude that whatever else it amounts to, reflection is one of its essential components. In other words, reflection somehow captures what it means to be an epistemic agent because it enables us to affect our beliefs. In other words, reflection somehow captures what it means to be an epistemic agent on Kornblith's understanding of this concept. He contends, however, that higher-level processes, like reflection, are not marks of EA because none of these processes involve directly committing to or altering a belief - appealing to them does not warrant adopting EA. Indeed, on his account, EA is a mythological concept that should be done away with.

Kornblith proposes that when we consider how it is that we come to acquire beliefs, we do not consider ourselves as actively involved in the formation of the belief itself. However, we can be tempted to think we are epistemic agents when

[w]e consciously entertain alternative views, and we think about which, if any, belief about the situation before us we are justified in holding. In situations like this, we seem to play a more active role. We don't just find ourselves believing things. Rather, we decide what to believe; we make up our minds; we choose to believe one thing rather than another.²³

²³ Kornblith, *On Reflection*, 85.

The suggestion, then, is that because we have the ability to weigh evidence and consider alternative possible conclusions, we are tempted to judge ourselves as epistemic agents.

Kornblith argues to the contrary, claiming that this type of agency cannot occur without some form of voluntarism:

If [defenders of EA] wish to insist that we are agents with respect to our beliefs, that there is, in short, genuine [EA], then how are we to make sense of this idea if it is not by way of some sort of voluntarism about belief?²⁴

Kornblith assumes therefore that EA, whatever it amounts to, must include some form of voluntarism in belief acquisition. We can summarize his argument as follows:

1. Either beliefs acquired from high-level processes, e.g. reflection or deliberation, are under direct agential control, or the appeals to higher-level processes do not legitimize the notion of epistemic agency.
2. All beliefs, whether formed from high-level processes or unreflectively, are no different at the lower, mechanistic, level of belief acquisition.
3. There is no agential influence on the lower level of belief acquisition.
4. Therefore, appeals to the higher-level processes of belief formation do not legitimize the notion of epistemic agency.

Premise 1 includes the different phenomena that one might appeal to in defending the notion of EA.²⁵ Kornblith considers the possibility that if any one of these concepts is discovered to be under our control, then perhaps EA is not illegitimate. However, in support of premises 2 and 3, Kornblith considers the aforementioned higher-level processes, rejecting each in turn as marks of EA. Thus, we are to conclude that EA is an illegitimate idea.

Kornblith's denial of EA is motivated by his conclusion that, regardless of the control one might have over any of these higher-level processes, we have no control over the mechanistic lower-level component of belief formation. Take reflective belief as a mark of agency.²⁶ Even when an agent reflects, there is no agency in the actual formation of the belief. Likewise for deliberation when

²⁴ Kornblith, *On Reflection*, 85.

²⁵ Indeed, I make such an appeal in §1.

²⁶ Kornblith cites Richard Moran, *Authority and Estrangement* (Princeton: Princeton University Press, 2001), Ernest Sosa, *A Virtue Epistemology* (Oxford: Oxford University Press, 2007), Christine Korsgaard, *The Sources of Normativity* (Cambridge UK: Cambridge University Press, 1996), and Stuart Hampshire, *Freedom of the Individual* (Chatto and Windus, 1965) as each suggesting that we have EA and that reflection is an essential part of it.

judging what to believe or intentionally acting in such a way as to form a certain belief.²⁷ According to Kornblith, a belief is produced by a reason-responsive mechanistic process; a belief is not produced through deliberate choice or from one's wants or desires. The mechanistic response to reasons is in no way under agential control.²⁸ The appeal to higher-order processes over which we can have some control, then, does not conceptually legitimize the notion of EA.

Clarifying his proposal, Kornblith considers an analogue, the 'screening wand' - a tool used at security stops in airports. These wands are used by screening agents and react to certain types of external stimuli. The suggestion, then, is that any higher-order processes involved in belief acquisition are like the process of directing the wand. The actual production of the belief is mechanical in the same way that the wand's reacting to external causes is mechanical. Our role as belief formers is like that of the security agent who can manipulate the wand but cannot directly control its responses to external inputs. The cognitive response to evidence is similar to the screening wand, in that the response is exclusively produced by external stimuli and not by anything the agent does. In other words, just as there is no agency in the actual screening mechanism, neither is there agency in the actual production of a belief.

Response to Kornblith

I suggest we can reject the assumption made for premise 1 of Kornblith's argument - that EA requires our beliefs to be under our direct control. The main concern with this argument, most notably in this premise, is that Kornblith reduces EA to direct voluntarism about belief but doesn't provide any positive reasons to accept this reduction. The concern, then, is that Kornblith unwarrantedly limits agency to those things that we have direct control over. I argue that agency needn't be limited in this way, but rather that agency can be captured by describing what it is

²⁷ As an example of the type of phenomena Kornblith has in mind when referring to intentionally acting in order to form a belief, consider Pascal's suggestion that if we act as though God exists, partaking in worship and the sacraments, we will eventually come to believe in God's existence. See Blaise Pascal, "The Wager," in *The Phenomenon of Religious Faith*, ed. Terrence Reynolds (Pearson Prentice Hall, 2005), 141-144.

²⁸ Kornblith acknowledges that we are not always perfectly responsive to reasons: "When we offer reasons for belief, however confident we may be that we are in the right, we do not just assume that our interlocutors will come to share our views" (130). But, like any other mechanism, the reasoning mechanism is subject to interfering factors. However, "whether the reasoning mechanisms are operating well or badly, we need not, and do not, assume that the individual to whom reasons are offered will exert any agency with respect to his or her beliefs" (131).

we do in relation to the ends we are working towards. In this way the agency in EA can be understood in purely epistemic terms although what the agent actually does cannot. I discuss each of these responses in turn.

With regards to EA requiring direct control of our beliefs, it will be helpful to contrast the two conceptions of EA under consideration. Kornblith's characterization of EA requires that one have direct control over lower-order belief formation if we are to demythologize EA. In other words, one must be able to have direct agency over one's beliefs. My suggested understanding of EA is more modest. The agency in EA is not the ability to control the lower-level processes, or mechanisms, of belief acquisition. Rather, the agency in EA takes place between the higher-level processes we can control and lower-level processes where belief simply happens. This conception does not require direct control of over our doxastic states.

Kornblith's airport screening agent example captures the distinction between our two views and highlights where I think Kornblith's conception is misplaced. The metal detecting wand represents the lower-level mechanistic belief-forming process; the ability to manipulate the wand by directing it in various directions represents the higher-level process. Kornblith argues that without the mechanical wand, no screening is possible. Because we do not control the mechanisms of the wand itself, wherever the wand is, *that* is where the actual screening takes place. Analogously, Kornblith argues that the actual formation of belief is at the mechanistic level, where there is no agency. Thus, he concludes that EA is a faulty concept.

This analogy does not work, however. Consider the possibility of there being both good and bad screeners. Poor screening could result from some mechanical issue with the wand itself or from the agent not knowing or adjusting the wand properly to the right kind of external stimuli. EA is not like simply directing the screening wand in any old way; rather, it is like learning to use the wand, or learning to use it better - or perhaps even like making improvements to the wand once one understands how it functions. We can allow that the end result takes place at the level of the wand's mechanisms, but we can also maintain that the screener has exercised agency over the screening process, and is at least partly responsible for successful screenings. Likewise for the agency one exercises over one's beliefs. The believer does not believe at will; the believer can, however, exercise agency over the belief-forming process. The crucial factor here is that what one does at the higher level can influence what occurs at the lower. This influence appropriately captures EA and shows how Kornblith's wand analogy doesn't actually support his conclusion.

EA, as I have developed it, then, is not subject to Kornblith's criticisms. His characterization of EA is a version of direct doxastic voluntarism, an epistemically unpopular view rarely endorsed by sympathizers of epistemic normativity or the possibility of doxastic improvement. So, while we are both addressing EA, the differences between our portrayals of this idea call into question whether we are indeed referring to the same concept. I suggest we are not, and that Kornblith is rejecting a notion that very few actually support.²⁹ To maintain that EA is mythological, then, Kornblith has to either (a) deny that each of the above agent-driven higher-order processes have any effect on our beliefs whatsoever or (b) provide us with reasons to accept his portrayal of EA. Kornblith agrees that if one believes p but acquires evidence for $\sim p$ after deliberation, or reflection, or some intentional action, this added evidence will affect one's doxastic state, even if the effect happens mechanistically. So option (a) is out. What about option (b)? Are there reasons that the agency in EA must be at the lower level? Kornblith does not give any reasons here. He simply asks: If the agency is not at the lower levels, then where? But as suggested above, this assumes that EA is some form of direct voluntarism about belief. So option (b) is out also.

It could be contended that this response does not do justice to the dialectic at hand. Kornblith accepts that there is agency occurring at the higher levels. Despite this type of agency, however, what one is doing when one is supposedly exercising EA can be accurately described without any distinctly epistemic terms. Believing, e.g., is not something that one does; belief is a phenomenon that happens independently of anything that an agent does *qua* agent. So there is nothing *epistemic* that is subject to the agency I suggest EA is based on. What we have, then, is simply a quibble over the term, or worse, an outright avoidance of the challenge that Kornblith presents to EA.³⁰

I do not deny Kornblith's challenge. I do think he is wrong to think that we are not agents with regards to what we believe and how we come to believe because there is a distinction between what one does as agent *qua* agent and how one's beliefs simply come about mechanistically. Kornblith's issue is that we

²⁹ There are numerous examples that could support this conclusion. Feldman, "The Ethics of Belief" and "Epistemic Deontologism," clearly rejects the suggestion that we have the ability to directly control our beliefs, while also accepting a normative component to what we believe. Nolfi also holds that there are prescriptive 'epistemic oughts' over which we have doxastic control, but not of the direct type espoused in DDV (Nolfi, "Why is Epistemic Evaluation Prescriptive"). Ryan likewise holds that we can be held epistemically and morally responsible for our beliefs while not having direct voluntary control over our beliefs (Ryan, "Moral Evidentialism").

³⁰ Thank you to the anonymous referees who highlighted these points.

needn't employ any distinctly epistemic terms when describing the agency involved in EA as I define it. This issue, however, leaves open another way in which we can think of agency, which does cite explicitly epistemic content: the goals and motivations we have when exercising EA.³¹

One of the central characteristics of EA is that one can exercise it when attempting to achieve a predominantly epistemic goal. Thus, it is not that we need to have the agency at the belief level; rather, it is that we have the agency with regards to achieving some epistemic ends. This proposal is obviously somewhat modest, but it does allow for other features that are evident in our own circumstances – like the ability to willfully improve in some epistemic ways. The practical versus epistemic agency distinction is certainly blurred here; indeed, as admitted above, EA is simply the exercise of our normal agency and thus there needn't be a distinction between the two. This conclusion does not matter. The agency exercised in EA reveals that we can be active with regards to what we believe. So while the activity is not the direct willing of a belief, the activity does affect how it is that we come to form and maintain some beliefs. And, while in some sense the dispute is merely terminological, it is nonetheless important as it reflects on deeper issues at work in the background. Taking Kornblith's challenge seriously, we can admit that there is no agency within the strictly epistemic domain, but this admission is not sufficient to deny EA. The agential focus is on achieving one's epistemic ends, like being able to form more accurate, fine-grained, sophisticated beliefs in general or in specific areas. If we can in some ways at some times achieve these goals – or even strive to achieve these goals – we are agents with regards to the results. If those goals are predominantly epistemic, then we are epistemic agents, even if what we do as agents is not directly controlling the epistemic.

Building from the preceding response, we can shift our focus to the issue of agential control and mechanistic bodily processes. Central to Kornblith's rejection of EA is that we have no control over these types of bodily processes. If we cannot directly control f , then we are not agents with regard to f . This assumption is not obviously warranted, however. There are marked differences between functions that we have some control over and functions that we have no control over. Contrast belief-formation with something like hair growth. Hair growth is mechanistic. We neither control the rate at which our hair grows nor the volume

³¹ Ahlstrom-Vij, whose rejection of EA we discuss below, recognizes this point as well: "...there is one type of account of [EA] that remains largely unaffected by Kornblith's critique, namely a type of account that takes [EA] to encompass the full range of things that we do in pursuit of epistemic goals..." (Ahlstrom-Vij, "Why We Cannot Rely," 277).

of hair we naturally have.³² There is a difference between the type of mechanistic processes over which we have no control, like hair growth, and those processes over which we have some control, like belief formation. Hair growth is not something we have any control over; how we form beliefs is something we have some control over.

Consider an analog to the type of indirect control we have over a mechanistic process like belief formation: heart rate. Whether or not one's heart beats and the rate at which it beats is not directly under one's control. That is, one cannot merely will one's heart to beat at, e.g., 72bpm. One can, however, undertake certain actions that we recognize will alter one's heart rate. If one sprints for one minute, one's heart rate will be significantly higher than if one were at rest. Furthermore, and more closely related to the type of control we have over our beliefs, one can affect one's overall heart rate indirectly via actions one can directly control. If I jog with the intention to improve my resting heart rate, and successfully do so over time, then I indirectly control the rate at which my heart will beat. I do not directly will my heart to beat at a certain rate; rather, I affect it more generally.

The simple appeal to mechanistic processes does not undermine the potential for intentionally affecting, albeit indirectly, that process in a goal-oriented way. We do not have the type of indirect control over hair growth that we do in bodily functions like our heart rates – and, as I am suggesting, in belief formation. There is a distinct difference between those processes that we have no control over whatsoever and those that we can indirectly control. Kornblith does not account for this difference; his emphasis is simply on the mechanistic aspects of the processes. There is a difference, however. And it provides further evidence that Kornblith's dismissal of EA is too hasty. That there is a mechanistic component to belief formation that we cannot directly control is not sufficient for the conclusion that we do not have some control over the beliefs we form. Neither is such lack of direct influence sufficient to show that we are not agents with regards to the ends we attempt to reach.

3.2 Ahlstrom-Vij's Rejection of Epistemic Improvement

Ahlstrom-Vij³³ provides a more direct challenge to EA. In contrast to Kornblith, Ahlstrom-Vij accepts that EA is more than simply controlling the lower-order processes of belief formation, but rather can be understood as the ways in which

³² There are, of course, surgical or other external ways to get more hair. But there are not steps that *I* can take in hair growth that are of my volition.

³³ Ahlstrom-Vij, "Why We Cannot Rely."

we attempt to achieve epistemic goals.³⁴ Given this characterization, which captures much of the presently defended conception of EA, Ahlstrom-Vij proposes that we should still reject this concept. His concern is that a central – perhaps the central – motivation in support of EA is that we want to and can partake in epistemic self-improvement. If indeed we are motivated to epistemically improve, “our main focus should be on the ways in which we fail *systematically*, rather than accidentally.”³⁵ That is, if we have systematic epistemic shortcomings, it is regarding those shortcomings that we should attempt to improve. Given this goal, however, we are faced with the problem that there are epistemic shortcomings that we are incapable of improving on, even if we want to. We are to conclude therefore that “we simply cannot rely on ourselves for epistemic improvement.”³⁶ If we cannot rely on ourselves for epistemic improvement, then we do not have EA.

The candidate example Ahlstrom-Vij employs to highlight our inability to epistemically improve involves cognitive biases, which are “systematic, and now well-established tendencies to form inaccurate beliefs.”³⁷ Consider, for example, the well-established cognitive biases involving tendencies to conflate personal traits when making comparative self-assessments. One may believe oneself to be more objective than 80% of a sample group; 60% of this group may assess themselves with the same ranking. At least a significant percentage of these assessments cannot be accurate as the ratio of people to percentiles does not jibe.³⁸ As this phenomenon represents a systematic epistemic failure on our part, we should be motivated to improve in this area. According to Ahlstrom-Vij, however, we cannot self-improve in these areas. We can run Ahlstrom-Vij’s argument thus:

1. If we have EA, then we can epistemically self-improve in areas where we systematically fail epistemically.
2. Cognitive biases are an area where we systematically fail epistemically.

³⁴ Ahlstrom-Vij assumes that “There is one and only one epistemic goal, and that is the dual goal of attaining true belief and avoiding false beliefs” (277). This is a general, all-encompassing, epistemic goal. The examples provided above concerning one’s goals to form more fine- or coarse-grained beliefs may be reducible to one horn of this goal. In any case, whether our respective portrayals of epistemic goals are consistent is inconsequential for the current discussion.

³⁵ Ahlstrom-Vij, “Why We Cannot Rely,” 278.

³⁶ Ahlstrom-Vij, “Why We Cannot Rely,” 293.

³⁷ Ahlstrom-Vij, “Why We Cannot Rely,” 278.

³⁸ Ahlstrom-Vij presents a number of psychological case studies where this type of phenomenon occurs.

3. We cannot epistemically self-improve our cognitive biases.

4. Therefore, we don't have EA.

Premise (1) establishes that we should be able to employ our EA to improve in areas of systematic epistemic failures. Premises (2) and (3), however, reveal an area of systematic failure where we cannot self improve epistemically. Thus, we are to conclude that we do not have EA.

Ahlstrom-Vij proposes three ways that we may attempt epistemic improvements in response to cognitive bias: self-correcting, self-binding, and external constraints. I focus on the first of these three proposals, as that is the focus of my response to the above argument. The self-correcting approach to cognitive bias suggests that we as epistemic agents can take steps to correct those systematic faults in reasoning that lead to inaccurate beliefs. According to Ahlstrom-Vij, there are two problems for the self-correction approach. First, there is a motivational problem resulting from our inability to recognize that we need to improve in this area in the first place – symptomatic of such biases is that we cannot accurately recognize when we are being biased in such a way.

Suppose, however, that we discover our cognitive biases. Could we not then take steps to correct them? Ahlstrom-Vij suggests that this approach won't work either, due to the *proper correction* problem. Even if we did somehow become privy to a tendency to conflate personal traits, self correcting via EA requires that we can do so successfully. Successful self-correction in turn requires that we satisfactorily deal with a number of challenges. The first challenge is that we recognize all and only the times that we are being biased. Another set of challenges are that the corrections made are all properly made. One has properly corrected if one neither over nor under corrects and one has made all and only the necessary corrections. Ahlstrom-Vij refers to numerous cases where at least one of these challenges is not met. While he accepts that the evidence he has presented against the self-correcting approach does not reject the possibility that one may be able to meet all of the above challenges, Ahlstrom-Vij highlights that *possible* does not mean *probable*. He concludes that

even if we assume that the relevant agents are at all motivated to engage in bias correction... there are substantial challenges they need to meet when it comes to doing so correctly.³⁹

What the above tells us is that there are good reasons to think that we cannot meet these challenges. So whether or not EA is marked by higher-order reasoning in conjunction with reasons-responsive lower-level bodily processes, we have

³⁹ Ahlstrom-Vij, "Why We Cannot Rely," 283.

reason to reject this concept as doing any substantial work. Its most appealing feature is undermined. That is, if we cannot epistemically self-improve in important ways – like in areas where failure is systematic rather than accidental – then there is reason to think we do not have EA.

Response to Ahlstrom-Vij

Ahlstrom-Vij's case against EA is compelling, but I think it falls short. The main problem arises in the first premise of his argument above: If we have EA, then we can undertake epistemic self-improvement in areas where we make systematic epistemic failures. If the failures are systematic, then it seems at the outset that we are doomed to fall short of fixing these failures. There are responses we might make to this premise, however. First, there might be areas where we cannot epistemically improve, but this inability in one area does not mean we cannot epistemically improve in other ways. Secondly, it's not actually clear that we cannot improve in the areas where cognitive bias manifests itself.

In support of the first response, in §2.3 above I gave a number of examples where epistemic improvement is possible and how we might go about improving in these ways. Admitting that we cannot improve when it comes to certain systematic biases we might have, does not mean admitting that we cannot epistemically improve *tout court*. It simply shows that there are limits to areas where we can improve. It is not controversial to accept that there are limitations on our abilities as epistemic agents. Indeed, there are limitations on our abilities as agents *qua* agents. Take the piano student who has learned all of the correct notes to Mozart's *Rondo Alla Turca*, but because of a systematic inability to keep time during high tempo songs cannot play it at a consistent speed. When it is pointed out that her timing is incorrect, she simply cannot improve. At slower tempos, however, she can keep time perfectly and can thus play Chopin's *Nocturnes*, e.g., as they were written. I suggest that it is an overgeneralization to propose that she is not an agent with respect to her piano-playing abilities because she systematically fails in one area that she is unable to improve on. Likewise for EA. There may be areas in which we are all systematically incapable of making epistemic self-improvements. All this tells us, however, is that we are not agents with respect to that part of our belief-forming abilities, not that we're not epistemic agents at all. So at most Ahlstrom-Vij highlights some limitation on EA.

A second response to Ahlstrom-Vij's first premise is to simply deny it by denying that the conditions needed for one to qualify as *improving* are too stringent. Recall that when taking steps to improve bias, Ahlstrom-Vij argues that one with EA would only correct for bias when one is actually being biased, would

not over or under correct, and would make all and only the necessary corrections. These conditions again seem reasonable in that they are the types of adjustments that would achieve the goal of forming the correct belief with the appropriate credence. We might ask, however, why do we need to have such fine-grained abilities? The standard seems too high; we do not place such stringent conditions on other less controversial exercises of agency. Consider the archer who misses the target 3' to the left. She adjusts her approach but overcorrects and misses the next shot 2' to the right. Although she over corrected, making more than the necessary corrections, most of us would admit that being a foot closer to the target is an improvement. Likewise if one discovers a personal epistemic shortcoming and overcorrects, but forms a belief that is closer to the truth than the previous one. This correction seems like an improvement. Or suppose one takes some but not all of the necessary steps to making a correct self-assessment in a case of cognitive bias. These steps also seem like improvements. Ahlstrom-Vij appears to have only provided us with idealizations for epistemic self-improvement, which might again simply highlight the limitations we have on becoming ideal epistemic agents.

One might counter that I am not being sensitive to the nuances of Ahlstrom-Vij's argument, most notably with what it takes to be an agent. If the archer, for example, takes aim at the target but has no notion of what the necessary steps are to hitting it aside from 'draw, aim, let go,' then the improvement between shots is merely a lucky accident. The necessary steps to hitting the target were not involved in her action plan and thus she was not an agent in respect to improving her shot – the improvement was mere luck. Similarly, the improvement in belief was lucky in the example concerning under or overcorrection or missing some of the necessary steps to eliminating cognitive bias. The apparent improvement made to achieve the goal of eliminating cognitive bias is also just luck.

Providing a full account of how lucky outcomes affect agency will take us too far astray from the point here. We can acknowledge, however, that one's agency is not undermined simply in virtue of a quasi-lucky outcome. As long as one is motivated to achieve one's goal, and proceeds with the intent to achieve that goal, one can still be an agent with regards to that undertaking, even if luck has some role in achieving the desired outcome. For example, one can intend to try to achieve one's goals despite there being uncertainty as to how to do so – or uncertainty as to whether or not one even can achieve the goal.⁴⁰ The archer can

⁴⁰ Alfred Mele, *Springs of Action* (Oxford: Oxford University Press, 1992), chapter 8, argues, e.g., that intentions to try to *A* can stand in for intending to *A*. One can be an agent with regard to achieving some ends, and thus achieve those ends intentionally, even if one's intentions were

intend to try to hit the target, even if she does not know all of the steps necessary for the successful completion of her goal to hit the target. Likewise, one can intentionally account for and improve cognitive biases in one's reasoning with the intention to try to form true and avoid false beliefs without knowing all of the necessary steps needed to do so, or without even being confident that one can.⁴¹ We have reason to accept, then, that even if an outcome is in some ways lucky, one can still intentionally improve in so far as they are motivated to improve and have taken steps – tried – to achieve that goal. I find it sufficiently plausible that intentionally doing anything denotes agency and so leave that undefended.

We have thus addressed reasons for rejecting Ahlstrom-Vij's rejection of EA. In the first case, if we grant that Ahlstrom-Vij is correct that we cannot improve in some cases of systematic epistemic failings, all we have been given is a limitation on where we can make epistemic self-improvements, not that we cannot make such improvements. In a second response, I propose that Ahlstrom-Vij's proper correction conditions are too stringent. If we held similar conditions for skill-based actions, we would undermine garden-variety agency. Failure to achieve a goal does not mean failure at making some improvements or that we are not agents with regard to those improvements. And even if we are limited in our ability to recognize the necessary steps to eliminate cognitive biases, or to what degree we need to adjust our self-assessments, that we can try to improve in these areas is sufficient to show that we have some agency here.

only to try to *A*. Michael E. Bratman, *Intentions, Plans, and Practical Reasons* (Cambridge Massachusetts: Harvard University Press, 1987), chapter 8, proposes a Single Phenomenon view of intentions and what one does intentionally: the intentional actions and acts done intentionally are both related to intentions, but that does not require that one intends what one does intentionally. In this case, one may intentionally improve by intending to try to improve although one might have constraints on their intending to improve. Hugh J. McCann, "Settled Objectives and Rational Constraints," *American Philosophical Quarterly* 26 (1991): 25-36, rejects Mele's and Bratman's distinction between intending to *A* and intending to try to *A*, arguing that intending to try to *A* is simply intending to *A*. Regardless of which view one accepts – whether or not one is intending to try to improve, or is intending to improve – if one does in fact improve, then one does so intentionally.

⁴¹ Donald Davidson, *Essays on Actions and Events* (Oxford: Oxford University Press, 1980), 92, provides a well-known example of intentionally doing something while also doubting whether or not one is or can be successful in what one is trying to achieve. In his example, Davidson is trying to make ten carbon copies by pushing his pen with some force while writing. He intentionally makes the copies, just in case he is successful in so doing; he is skeptical that he actually is making the copies while writing, however.

Conclusion

I have proposed that epistemic agency is best understood as the control we have over developing and refining our doxastic dispositions or belief-forming abilities, which include the propensity one has to form true or false or coarse-grained or fine-grained beliefs within different domains. In having this control we exercise epistemic agency in ways similar to indirect doxastic voluntarism – i.e. voluntarily undertaking a process so as to form a certain belief. Unlike indirect doxastic voluntarism, however, I have suggested that EA is not exercised so as to form a specific belief; EA merely affects the ways in which we can form beliefs more generally. This concept is motivated by the plausibility of epistemic normativity, skilled belief, and our ability to set and take steps towards achieving epistemic goals for doxastic self-improvement. There are good reasons, then, to accept EA as an actual phenomenon and a philosophically fruitful concept.