

EPISTEMIC CONTROL WITHOUT VOLUNTARISM

Timothy R. Kearl

University of Glasgow

[forthcoming in *Philosophical Issues*]

Intentional actions are reasons-responsive; to act intentionally is (perhaps among other things) to behave for a complex of reasons, ones that favor bringing about some option in a space of possibilities.¹ According to a popular, “bottom-up” model of practical agency, this feature of intentional actions is located in an agent’s *intentions*, reasons-responsive mental states that initiate and sustain behavior so as to render it intentional. To the extent that practical agents—things that *act*—exercise a distinctive kind of control over their behavior, it might be called “intention-level control”, without which there are no actions at all. There may, of course, be other interesting and perhaps more robust conceptions of control within the realm of practical agency, such as the control exhibited by libertarian free agents, but those will involve more complex practical agents, not practical agency *per se*.

Inferences, like intentional actions, are reasons-responsive; to infer is to adopt a belief-like attitude for a complex of reasons that favor accepting some hypothesis in a space of possibilities. It is tempting to think that epistemic agents—things that believe or know—exercise a distinctive kind of control over their mental behavior, without which there are no inferences at all. But it is controversial whether our inferences manifest intention-level control, and this controversy has prompted a great deal of avoidable skepticism about epistemic agency itself. My aim here is to sketch a “bottom-up” model of epistemic agency according to which an agent’s *conditional beliefs* are reasons-responsive mental states that initiate and sustain mental behavior so as to render it inferential; epistemic agents manifest “belief-level control”, without which there are no inferences at all.

The idea that epistemic agents manifest belief-level control in making inferences is controversial in a number of ways. One is that epistemic agency is itself a suspicious notion; some wonder whether we have any more control over our mental behavior than we do over our digestive behavior, which is to say none, and without a distinct kind of control, there isn't a distinct kind of agency. Another is that belief-level control could be suited to be the kind of control that makes epistemic agency distinctive. And yet another is that, of all the richness of our mental lives, *mere* inferences could be the locus of our epistemic agency.

But once the thesis of belief-level control is laid out and defended, I hope to show that it adds a great deal of explanatory value to epistemology. Not only can we see important structural similarities between theories of practical and epistemic agency insofar as both bottom out in a distinctive sort of reasons-responsiveness, but, as I'll sketch in closing, we can also theorize about distinctly epistemic notions of responsibility and authorship that might otherwise wither on the vine, responsibility and authorship often being treated as phenomena squarely within the domain of practical agency.

1. Bottom-up Models of Action

Practical agency is characterized by intention-level control, which is to say that the kind and degree of control that practical agents exercise over their behavior is (at least) a matter of their intentions initiating and sustaining it. This way of thinking about action starts at the bottom; actions in general are understood in terms of their connection to this bottoming out point.

To explain intention-level control, I'll introduce a widely endorsed, "bottom-up" model of action. I'll then show how accounts of more complex phenomena within the sphere of practical agency—*identification* and *authorship*, in particular—are built either as

supplements to the bottom-up model, or as departures from it in response to its perceived inadequacies.

1.1. Intention-level control

In the bottom-up model of action, originating in Davidson (1963), actions are events in which one's outward behavior is caused (in the right way) by one's antecedent mental states, namely one's beliefs and pro-attitudes (desires, urges, and so on). Consider a familiar tale: *Tim, sitting at his desk after a long afternoon of grading papers, realizes that nothing would quite hit the spot like a cold beer. Remembering that he's got beer in the fridge, he gets up and walks to the kitchen, opens the fridge, and grabs a beer.* According to the bottom-up model, Tim's grabbing a beer is an action because it is an event in which Tim's outward behavior (making certain reach-and-grab gestures towards beer) is caused in the normal way by the earlier event of Tim's desiring beer, remembering he's got some in the fridge, and intending to get some. The reasons *for which* Tim acts—the complex mental state embedded in that earlier event causally antecedent to his behavior—are his “motivating reasons.”

There are a few explanatorily powerful aspects of this model of action. One is that we can make assessments about the (practical) rationality of a certain course of action by evaluating an agent's motivating reasons in light of her total reasons. For instance, Tim's grabbing a beer seems practically rational in the fairly limited vignette given above, but if we were to fill in more details, perhaps pointing out that Tim is well past the deadline to submit grades and taking much longer to submit them puts his job at risk, and he can't afford to lose his job, and having a beer (or three) greatly increases the odds he'll take much longer to submit them, *and* Tim knows all of this, then his grabbing a beer is an irrational action. Again, it is an action in virtue of its causal connection to his beliefs, desires, and intentions but it is an irrational one because his motivating reasons are

swamped by other considerations that, as a matter of fact, failed to move him. Speaking loosely (but not falsely), we might say that Tim's intention to get a beer against this backdrop encodes a poor plan for action, poor because it is insensitive to or misconstrues what is at stake in so acting.

Another explanatorily powerful aspect of this model is its ability to explain a wide range of actional behavior that is not intentional in terms of other behavior that is. For instance, when I arrive home late at night, I flip the light switch to illuminate the room. Unbeknownst to me, I thereby alert the burglars of my presence.² Of course, this latter event is not something I bring about intentionally. It is also not something that merely *happens* to me; alerting the burglars is something I *do*, albeit unintentionally, and the explanation of this fact is that the event of my alerting the burglars depends on³ the event of my illuminating the room, which is something I do intentionally.

Other actional behavior falls somewhere in between paradigmatically intentional (my illuminating the room) and paradigmatically unintentional (my thereby alerting the burglars). For instance, Harman (1976) discusses a case in which a sniper, firing his gun in trying to kill a soldier, knowingly alerts the enemy to his presence. While alerting the enemy to his presence is something the sniper *does*, it is not something the sniper does intentionally. Still, because the sniper is fully aware that he will bring about this event in firing his gun, there is pressure to explain the nature of this action in different terms than the case of unwittingly alerting the burglars by illuminating the room. Mele (1992a, 1992b) and Mele and Moser (1994) call these sorts of middling cases "nonintentional actions," where the category of non-intentionality applies, perhaps among other things, to actions when they are the foreseeable consequences of intentional actions.

Taken collectively, these considerations make the bottom-up theory of action attractive for explaining our practical agency and the kind of control that is distinctive of it; anything that counts as actional does so in virtue of its connection to something that is

an intentional action. But there are aspects of our agency that, at least at first blush, seem too complex or robust to fit into this simple model.

1.2. Primitivism or reductionism about agency and authorship

One central commitment of the bottom-up model of action is that event-causation is all the causation that is needed for action explanations. Some authors have reacted to this ontologically parsimonious Davidsonian picture by insisting it leaves out what is perhaps the most important part of the story it was supposed to tell: *the agent herself* being the author of her own behavior, not merely the various mental events that occur within her.⁴

On Chisholm's account, agents are another primitive relatum in the causal nexus; there is familiar event-causation, of course, but when it comes to accounting for expressions of agency, or what I've called "authorship," there must also be causation of an irreducibly different kind, namely agent-causation.⁵ Rather than assess the merits of primitivism and its more expansive ontology, I just want to note one impetus for finding it attractive, which is a sense that the agent in event-causal theories "disappears" from her actions; *she* is not their author.

Others take a more conservative stance on whether and how the bottom-up model can accommodate an agent's contribution to her own behavior. Velleman (1992) claims that the "proper goal for the philosophy of action is to earn the right to make jokes about primitive agent-causation, by explaining how an agent's causal role supervenes on the causal network of events and states" (469). Citing Frankfurt (1988) as his intellectual predecessor in this task, he proceeds to offer a reductive account of agent-causation in terms of a special, agency-constituting desire. This presence and function of this special, agency-constituting desire is a part of the event-causal nexus, but it encodes something that other, mere desires within the agent do not: namely, what the agent identifies with

when she acts, and thereby the agent's "place... in the explanatory order of the world" (455).⁶

Velleman's Frankfurt is one thing, but Frankfurt's Frankfurt is another. Frankfurt's Frankfurt's claim about the conditions under which an agent identifies with (and so authors) her actions is not an account of what makes something an action, nor of the abilities that make something a practical agent. It is, rather, an account of a particularly important feature of robust forms of practical agency, oriented around what is distinctive of *persons* (who act).⁷

The foregoing discussion has been merely expositional. I explained that the bottom-up, Davidsonian model is an account of what distinguishes action from mere behavior, not what determines the extent to which an agent authors her actions. There are generally two sorts of reactions to the bottom-up model, one being that it simply cannot accommodate any interesting notion of authorship, and the other being that we need to accommodate authorship by supplementation. I'll go on to argue that epistemology admits of more or less the same bottom-up model of inference, the same questions about authorship, and the same menu of solutions.

2. A Corresponding Model for Inference

The discussion of bottom-up theories of action has given us a framework for thinking about explanations of agency and control. In particular, practical agency is characterized by intention-level control, which is to say both that the bottoming out point for manifestations of practical agency is behaviors that are initiated and sustained by intentions, and that anything that counts as actional does so in virtue of a connection to behaviors that are initiated and sustained by intentions.

I'll here argue that epistemic agency is characterized by belief-level control, which is to say that the kind and degree of control that epistemic agents exercise over their

mental behavior is (at least) a matter of their *conditional beliefs* initiating and sustaining it. This way of thinking about thinking also starts at the bottom; thought in general is understood in terms of its connection to conditional beliefs as a bottoming-out point. But to make this idea plausible requires elaboration.

2.1. Belief-level control

Start with a brief account of conditional beliefs, beliefs with content of the form q if p . Frank Ramsey (1929/1990) offered what has come to be known as the “Ramsey test” of conditional commitment: when two people are “arguing ‘if p will q ?’ and are both in doubt as to p , they are adding p hypothetically to their stock of knowledge and are arguing on that basis whether q ” (155). The conditional beliefs that q if p and $not-q$ if p are “in a sense... contradictories”, insofar as they reflect or encode commitments to form contradictory q -attitudes in light of p -information (*ibid*).⁸ On this conception, an agent’s conditional beliefs encode her commitment to accept certain propositions, given that she accepts (or supposes) that others hold. In virtue of this commitment-encoding function, we might think of conditional beliefs as *plans for thinking*, plans encoding what the agent takes to be evidence for what.

Conditional beliefs play a number of distinguished roles in our epistemic lives; believing q if p is a way of ruling out hypotheses in which p and $not-q$, a way of scaffolding the hypotheses one imagines,⁹ and a way of guiding one’s reasoning from p to q .

Consider what happens in a somewhat artificial (but not extraordinary) case of inductive inference. An agent, Tia, is asked to pull out the first 99 balls in an urn that contains 100 balls and note their colors. She sees, one by one, that the first 99 balls are black and thereby comes to acquire a great deal of perceptual evidence concerning the colors of the balls in the urn. Tia is then asked about the possible colors of the 100th ball.

On the basis of her belief that each of the first 99 balls were black, she becomes confident that the 100th ball is black too.

What makes this mental behavior an inference is its causal connection to her conditional belief *the 100th ball is black if the first 99 balls were black*. That conditional belief, itself sensitive to the evidence given to her in perceptual experience, guides her thinking about the color of the 100th ball in the right way so as to render her resulting attitude inferential. Tia's confidence that the 100th ball is black manifests belief-level control.

Compare Tia with Ham. Ham is asked to participate in the very same exercise, but after pulling out 99 black balls from the urn, Ham is hit on the head with a hammer at just the right force and at just the right angle so as to become confident that the 100th ball is black. Ham's confidence is in no way causally connected to his conditional beliefs about the color of the 100th ball, given the observed colors of the first 99 balls. His confidence fails to manifest belief-level control; this change in confidence is, so to speak, something that merely happens to Ham.

The difference between Tia and Ham can be explained in terms of the basing relation. In epistemology it is common to accept a distinction between propositional (or *ex ante*) and doxastic (*ex post*) justification. More specifically, an agent (at a time) is propositionally justified in adopting a particular attitude A towards some proposition P just in case the agent's total evidence (at that time) favors taking A towards P, whether or not the agent considers P in the first place, let alone adopts any particular attitude towards it. Doxastic justification is then understood in terms of propositional justification. An agent (at a time) is doxastically justified in adopting A towards P just in case (i) she is propositionally justified in adopting A towards P, and (ii) her adopting A towards P is *properly based* on the evidence that propositionally justifies it.¹⁰

There have been a number of competing views about the nature of the basing relation offered in the last few decades, but few of us seriously think that the basing relation is entirely non-causal.¹¹ To the extent that all of us are at least implicitly

committed to some or another causal theory of basing being the correct one, then there is a bottom-up model of inference that parallels Davidson's bottom-up model of action, and this model can underwrite a difference between manifestations of epistemic agency and mere mental behavior. To wit, according to this parallel bottom-up model of inference, Tia's believing that the 100th ball is black is an inference because it is mental behavior caused in the normal way by her conditional belief that *the 100th ball is black if the first 99 balls were black*, which is itself a reasons-responsive attitude sensitive to the earlier event of Tia's observing and so coming to believe that the first 99 balls were black. The reasons *for which* Tia infers—the complex mental state embedded in that earlier event causally antecedent to her mental behavior—constitute her “inferential basis”.

This bottom-up model of inference inherits some of the same explanatory power as the bottom-up model of action. For instance, we can explain why Ham's mental behavior differs importantly from Tia's: it stands in no interesting, non-deviant causal connection to his other mental states; it is thus mere mental behavior. Also, we can make assessments about the rationality of an inference by comparing one's inferential basis to one's total evidence. If Tia were to instead have let her (strange but deep-seated) desire to pull an array of colorful balls out of the urn cloud her judgment and prevent her from inferring as she did, instead becoming confident that the last ball would be non-black, her inferential basis would have included considerations irrelevant to the truth of the question at hand, rendering her mental behavior both inferential and irrational.

Conditional beliefs, like intentions, are reasons-responsive states. And just as intentions encode an agent's plans for acting, so too do conditional beliefs encode an agent's plans for thinking. Whereas intentions are responsive to a complex of beliefs and desires that constitute an agent's *practical reasons*, conditional beliefs are responsive to a complex of experiences, memories, background information that constitute an agent's *evidence*.¹² So, just as intentions are responsive to practical reasons, and manifestations of practical agency bottom out in their relationship to intentions, conditional beliefs are

responsive to evidence, and manifestations of epistemic agency bottom out in their relationship to conditional beliefs. Our epistemic agency is manifested *within the bounds* of our conditional beliefs, just as our practical agency is manifested *within the bounds* of our intentions.

In this way, epistemic agents can be thought to manifest a kind of control over their mental behavior that is at least as strong as the intention-level control that practical agents exercise over their non-mental behavior. Epistemic agency bottoms out, as does practical agency, in a certain reasons-responsive attitude encoding an agent's plans. This attitude—be it a conditional belief or intention—then plays an ineliminable role in explaining why some behavior is an expression of agency.

2.2. Voluntarism and skepticism about epistemic agency

The previous section laid out, in somewhat sketchy terms, a notion of belief-level control that stood to inference as intention-level control stands to action. But talk of “epistemic control” tends to get people worked up, and this is an important point of disanalogy between practitioners of action theory and epistemology; eliminativism was not and is not generally considered a live option within the realm of theories of practical agency, but it is among the more popular options within the realm of theories of epistemic agency. And agent-causation was and is not generally considered a live option within the realm of theories of epistemic agency, despite its staying power within the realm of theories of practical agency. I'll address these matters in turn.

Talk of “epistemic control” raises the specter of doxastic voluntarism and makes eliminativism about epistemic agency begin to have a certain appeal that it shouldn't and, absent the salience of voluntarism, wouldn't. So, I'll spend some time discussing eliminativism about epistemic agency, doxastic voluntarism, and why those eliminativist sentiments are misguided as responses to belief-level control.

To get a sense of what we're up against: Strawson (2003) remarked that "the role of genuine action in thought is at best indirect. It is entirely *prefatory*, it is essentially—merely—*catalytic*" (231). Kornblith (2012) suggests that epistemic agency is "nothing more than a bit of mythology" (90).¹³ Williams (1973) thinks that voluntarism rests on a conceptual confusion: "...if in full consciousness I could will to acquire a 'belief' irrespective of its truth, it is unclear that before the event I could seriously think of it as a belief, i.e., as something purporting to represent reality" (148). And Alston (1988) argues that it is "psychologically impossible" to adopt any particular propositional attitude at will (the implication being that it is psychologically impossible to adopt beliefs at will, although his argument purports to show more).

The only agency exhibited in judgment on this picture seems to be grounded, indirectly, in manifestations of *practical* agency, wherein we more or less cajole our minds down some path. Strawson admits that one can "[set] one's mind at the problem", "shepherd or dragoon one's wandering mind back to the previous thought-content in order for the train of thought to be restarted or continued", or "[initiate] a kind of actively receptive blanking of the mind in order to give any missing elements a chance to arise". It is about as dim a view of the matter as one might take, the agency manifested in judgment resembling the agency manifested in digestion, which is to say none, apart from avoiding spicy foods, the latter expressing practical agency if anything.

Why the skepticism? The answer has at least two component pieces. The first component of skepticism behind eliminativism may be a reaction to over-intellectualized conceptions of epistemic agency. Proponents of these intellectualized views treat powers of deliberation, reasoning, or higher order thought as the marks of, if not also necessary conditions on, epistemic agency. I will discuss these "intellectual" views in the next section.

The second component of skepticism is a reaction to doxastic voluntarism, so to appreciate why so many are repelled from that view, we've got to do a bit of conceptual spadework.

Now, there's doxastic voluntarism and there's *doxastic voluntarism*, and we shouldn't uncritically lump together more ordinary voluntaristic notions with the heavyweight ones. For instance, Strawson and the eliminativists seem to concede, with the rest of us, that we exercise *indirect* voluntary control over our beliefs, as when we choose to do more research before making up our minds, resign to continue thinking about a problem later and so suspend judgment on it in the meantime, and so on. Thus, we manifest very ordinary indirect voluntary control over our beliefs all the time by performing actions that affect the beliefs we are in a position to form. The idea that we have and exercise this very ordinary sort of indirect voluntary control over our beliefs is undeniable.

What's really at issue is whether we exercise *direct* voluntary control over at least some of our mental life. To exercise direct voluntary control is at least to exercise intention-level control, intentions being a mark of voluntariness. Perhaps it involves stronger notions like a "dual" or "two-way" power, as in cases of choice or decision-making where one has a number of options available. But to have a dual power to realize one course of action or another is to have intention-level control over each. Thus, to ask whether we can *decide* or *choose* to believe is not importantly different from asking whether our beliefs admit of intention-level control.

Two particularly admirable attempts at motivating a limited form of direct doxastic voluntarism come from Steup (2008; 2012; 2018) and Ginet (2001). According to Steup, we exercise direct voluntary control over our doxastic attitudes when we deliberate about what to believe, the result of our deliberation being a *decision* about what to believe. Steup's notion of direct voluntary control is supposed to be a species of compatibilist control; whether or not our deliberative beliefs are determined, we exercise

control over our beliefs because, for instance, *if* we were to reflect on them in such a way that we decided to believe differently, we would believe differently, or *if* we were to attend to available evidence that speaks against the truth of some proposition, we would decide to abandon our commitment to it.¹⁴

Ginet's picture is not too different in letter, also involving decisions to believe, but the spirit of his proposal is more Jamesian. When one has evidence that is ambiguous, or when the stakes are high, one has some leeway in choosing one's commitments and what one takes for granted. For instance, a juror may, upon hearing strong arguments from both the prosecution and defense, decide to believe an exonerating witness's testimony and vote to acquit (2001, p.64).

Seeing voluntarism as the going view for the kind of control manifested in practical agency, one is tempted to shoehorn any legitimate notion of epistemic agency into voluntarism. But this conflates intention-level control with belief-level control. Whether some of our beliefs admit of intention-level control is an interesting question in its own right, both empirically and conceptually, but it is not the right question to ask; if skepticism about epistemic agency is bound up in considerations of intention-level control, skeptics have been using the wrong measuring stick. *Choosing* or *willing* are not the only ways to exercise direct control, even if they are paradigmatic ways to exercise direct intention-level control.

What we need to make sense of epistemic agency is a kind and degree of control over our mental behavior that resembles the kind and degree of control we have over our non-mental behavior. But it does not follow from this that these two kinds of control must be identical. Here is a way to see what I mean. When one acts intentionally, one's behavior is causally initiated and sustained by one's intentions. Part of what renders this relationship one of *control* is that if one had intended differently, one's behavior would have differed correspondingly.¹⁵ But *our intentions themselves* are not necessarily under our intention-level control.

Of course, *some* intentions manifest intention-level control in fairly complex practical agents.¹⁶ For instance, *deliberating* involves the ability to intentionally settle uncertainty about what to do by way of deciding. If, after deliberating, an agent decides to (and thereby intends to) ϕ , her intention to ϕ manifests intention-level control. And *monitoring* involves the ability to intentionally cancel one's intentions prior to or during execution. If, while monitoring, an agent *refrains from canceling* her intention to ϕ , that intention to ϕ manifests intention-level control (at least on the assumption that omissions can sometimes be causes, or even quasi-causes).

That *some* intentions manifest intention-level control is compatible with the idea that intention-level control bottoms out in reasons-responsiveness. Often enough, the intentions characteristically operative in deliberation, or in monitoring, themselves fail to manifest intention-level control; my deliberative intentions, for instance, are typically sensitive to uncertainty about what to do, not some *other*, even-higher-order intention. These complex and iterated intention-level abilities seem important for the control characteristic of *autonomous* or *responsible* agents, but not for the control characteristic of practical agency *per se*.

To sum up: the control characteristic of practical agency is dual. There is the control we have over our actions in virtue of those actions being initiated and sustained by our intentions, and there is the control we have over our intentions in virtue of those intentions being sensitive to our reasons for action.

Similarly, when one infers, one's mental behavior is causally initiated and sustained by one's conditional beliefs. Part of what renders this relationship one of control is that if one had believed differently, one's mental behavior would have differed correspondingly.¹⁷ But *our conditional beliefs themselves* are not necessarily under our belief-level control;¹⁸ there are, here as above, two notions of control operating in tandem. There is the control we have over our inferences in virtue of those mental activities being

guided by our conditional beliefs, and there is the control we have over our conditional beliefs in virtue of those attitudes being guided by our evidence.

I may exercise belief-level control by *coming to terms with* a proposition after deliberating about it, treating the evidence in its favor as sufficiently probative. I may exercise epistemic control by *double-checking* a surprising result, treating the evidence in its favor as though it may be misleading. But neither coming to terms with nor double-checking are manifestations of practical agency; at least they need not be. I may have no stake in whether a question is true or false but nevertheless wonder about it, come to terms with an answer, or double-check because the answer was inconsistent with my priors. These notions are surely not restricted to the outputs of *decisions* about what to think. Intention-level control is not a necessary condition on epistemic agency, even if something that is similar to intention-level control in kind and degree is.

Without claiming that at least some of our beliefs and other attitudes can be formed at will, we have a robust and legitimate notion of control over our inferences in terms of our beliefs, and, in turn, our beliefs in terms of their sensitivity to evidence. But if beliefs stand to epistemic agency as intentions stand to practical agency, it makes no sense to evaluate the prospects of epistemic agency by whether one manifests intention-level control over one's beliefs.

Let's take stock. Skepticism about epistemic agency tends to follow a pattern. First, one understands control in terms of voluntariness or what is done "at will", these notions being tied to intention-level control. Then, one argues that no part of our epistemic lives is governed or guided directly by our intentions; we neither believe nor infer "at will". From these premises one concludes that there is no legitimate and interesting sense of control left to be the one manifested in epistemic agency. But the arguments in Section 2.1 cast doubt on the first premise. Rather than measure the cogency of epistemic agency by intention-level control, I have suggested that we appeal to a different notion that is similar in kind and degree, belief-level control. Belief-level control, like intention-level

control, bottoms out in responsiveness to certain kinds of reasons, and conditional beliefs causally initiate and sustain our mental behavior as intentions do our non-mental behavior.

If I'm correct that, skepticism about epistemic agency looks much less appealing, and perhaps rests on mistaken assumptions about the nature of agency more generally, namely that intention-level control (or something stronger!) is the only game in town. To the extent that voluntarists try to motivate a conception of epistemic agency by way of intention-level control, their views rest on those same mistaken assumptions.

But that only speaks to one half of the debate; the other half buys into epistemic agency but sees it manifested, first and foremost, in the higher altitudes of mental activity. In the next section, I suggest that these accounts of epistemic agency are either over-intellectualized or mistake certain ideals of epistemic agency with the conditions of it.

3. Various Ideals of Epistemic Agency

In the last section I tried to undermine various forms of skepticism about epistemic agency by arguing first that epistemic agency does not need intention-level control, and second that intention- and belief-level control were similar enough in important respects for the latter to underwrite epistemic agency. In particular, I suggested that practical agency has two faces; there is what we do at will, guided by our intentions (manifesting "intention-level control"), and there is the formation of intention itself, which is ultimately guided by practical reasons. Likewise, epistemic agency has two faces; there is what we think because of what we believe (manifesting "belief-level control"), and there is the formation of belief itself, which is ultimately guided by evidence.

This might strike some as an exceedingly deflationary account of epistemic agency, especially those inclined to think about agency in terms of higher order thought, reflection, and so on. If anywhere is the locus of epistemic agency, one might be inclined

to think, it is where mental behavior is complex, creative, or reflects upon itself, not in its rudimentary inferential operations, which seem hardwired and can operate merely implicitly.

Recently, Jenkins (2021) has argued that *reasoning*, by which he means “a *conscious* process with constituent events such as judgments, inferences, and acts of supposition”, is the locus of epistemic agency (my emphasis added). One way to endorse a non-eliminativist conception of epistemic agency is to think that consciously reasoning about whether some proposition p is true—with the aim or purpose of settling the question “ p ?”—is a manifestation of agency, conscious effort and purposiveness being characteristically agentic features.

And Ernest Sosa (2007) has defended a form of virtue reliabilism according to which knowledge requires the manifestation of an epistemic virtue. In order to “know full well” that some proposition p is true, as Sosa once put it, one must not only be connected to the truth of p by a manifestation of virtue or skill (“animal knowledge”), one must also be in a position to know that one’s connection to the truth of p is safe (“reflective knowledge”); full blooded epistemic agency exhibits a kind of mesh between first- and second-order knowledge, or between animal and reflective.

Exegetical points about virtue epistemology aside, being a virtue epistemologist already commits you to *some* interesting notion of epistemic agency; the virtues are manifested by *agents* if they are manifested by anything at all.

I don’t have knock-down arguments against these intellectualized conceptions of epistemic agency; instead, I hope to explain their purported conditions on agency in terms of the ideals or complexities of agency characterized in terms of belief-level control.

Nothing in my account is inconsistent with the idea that conscious reasoning can manifest epistemic agency, but the idea that it must *conscious* reasoning is overly demanding. Belief-level control, like intention-level control, can operate largely under the surface of our conscious attention. For one thing, it would be odd if the demands of

epistemic agency required that the agent make a conscious effort to settle the question of whether p , but the demands of practical agency did not require that the agent make a conscious effort to settle the question of whether to ϕ . Deliberation and reasoning can operate under the hood; this is, in fact, the normal mode of each. For another, to require that reasoning be conscious to manifest epistemic agency seems to exclude exercises of epistemic agency that the agent does more or less “automatically” because they are highly skilled. It might take me hours of conscious, effortful reasoning to reproduce the automatic inferences of gifted young mathematicians at the International Mathematical Olympiad.

With all that said, there is surely special value in being able to sometimes deliberate or engage in reasoning through conscious attention. It permits us to recognize connections between things we might miss in our more automatized moments, it gives us the benefit of critical distance, and so on. Given the complexities of our practical and epistemic lives, perhaps certain courses of action or thought are only available to those whose control rises to the level of consciousness. This seems especially true when the reasons themselves are complex or incommensurable, or when other things hang in the balance.

But if all that is true, then it doesn't follow that manifesting an ability to engage in episodes of conscious reasoning is necessary to be an epistemic agent any more than it follows that manifesting an ability to engage in episodes of conscious deliberation is necessary to be a practical agent. We can surely appreciate the value, *qua* epistemic agency, to be able to consciously deliberate. As Sosa (2003) once remarked, “selective attention is the index finger of the mind”, and to employ this sort of power in reasoning, as in deliberating about what to do, greatly increases the bounds of our agency. But this does not impugn the idea that epistemic agency is characterized, first and foremost, by belief-level control, the operation of which may be merely implicit, just as intention-level control may operate merely implicitly.

Much the same can be said for virtue epistemology, at least of the “mesh” variety Sosa espouses. This sort of picture stands to epistemic agency as Frankfurt’s mesh view of identification stood to practical agency. Recall that Frankfurt’s claim about the conditions under which an agent identifies with her actions is not an account of what makes something an action, nor of the abilities that make something a practical agent. It is, rather, a theory oriented around the nature of persons.

We might think that being an agent put together so as to be able to exhibit a kind of mesh between first- and second-order judgments is, no doubt, extremely valuable *qua* epistemic agent. For instance, we might think that certain inferences are *autonomous* because they flow from such a mesh, or perhaps our attitudes about the world can be held with some more robust modal profile; the mesh might render them safer from error or skeptical challenges. Moreover, some manifestations of epistemic agency flow from a coherent evaluative perspective, and others don’t, and this matters for the ways in which we assess the epistemic rationality of the thinking agent, and for how we hold the thinking agent responsible. All of *that* is compatible with the idea that epistemic agency is characterized by belief-level control.

Again, I want to stress that I am only criticizing these views insofar as one might take them to supply necessary conditions on what it is to be an epistemic agent. If that’s how the authors intend them, then they over-intellectualize epistemic agency, and they mistakenly treat certain ideal or valuable forms of it as conditions on it in the first place. But if these are meant instead to characterize robust or interesting forms of epistemic agency, ones we might value as ideals of the kinds of things we are, then belief-level control is not only consistent with but also seems to reinforce this idea.¹⁹

In fact, one of the virtues of my belief-level control account of epistemic agency is that it untethers questions of epistemic authorship, responsibility, and autonomy from the yokes of conscious reasoning, reflective and higher-order thinking, and voluntarism.

The special value of each of these things can be explained as complex and valuable manifestations of epistemic agency, not conditions on being an epistemic agent.

4. Concluding Remarks on Authorship and Responsibility

The idea that epistemic agency is characterized by belief-level control faced challenges from two sides, one eliminativist, the other intellectualist. The eliminativist trend in debates around epistemic agency sees no room for exercises of agency in the workings of our minds, but those views, I argued, were too caught up in arguing against the idea that we have intention-level control over our attitudes. Whether or not we do, we have belief-level control, which is similar in degree and kind but not identical to intention-level control. This has the nice upshot of not shoehorning the whole topic of epistemic agency into a chapter of the final theory of practical agency. To the extent that anyone was ever attracted to doxastic voluntarism *because it gave shape* to an underlying sentiment that epistemic agency is a real and important phenomenon, I have offered another way to vindicate that sentiment without collapsing the practical and the epistemic.

The intellectualist trend in debates around epistemic agency sees conscious reasoning and reflection as its characteristic features. This, I argued, either over-intellectualizes epistemic agency or describes certain ideals of it. But if they are describing ideals of epistemic agency, then what they say is totally harmonious with the idea that such agency bottoms out in belief-level control. But belief-level control does not, all by itself, settle questions of authorship and responsibility in epistemology, just as intention-level control, all by itself, does not settle questions of authorship and responsibility within the realm of practical agency. I hope to gesture, in some brief concluding remarks, at what the space of options looks like.

What of authorship? As far as I know, no epistemologist has echoed Chisholm's sentiment that unless we have an analysis of "Tia inferred the 100th ball is black" into the

event-causal nexus, then we are claiming the benefits of honest philosophical toil without even having a theory of human inference. Perhaps that means the joke is entirely on us, but I have never found primitivism a live option, and as a sociological fact, epistemology writ large hasn't either.

The prospects for some form of reductionism about epistemic authorship seem, in contrast, pretty bright. We might think that epistemic authorship can be understood in terms of the role an agent's fundamental epistemic perspective plays in the causes of her mental behavior. Not all inferences flow from what the agent really takes to be evidence for what, for instance, and this seems to matter for how we hold her responsible for her irrationality.

There are already a number of philosophers who argue that agents can be *morally* responsible for their attitudes despite lacking voluntary control over them. The so-called "involuntarists", like Smith (2005) and Adams (1985), treat the fact that an agent's attitudes are *properly her own*, or are in-principle answerable to the agent's capacities for reason, as sufficient to ground responsibility for merely having an attitude (of contempt when a friend fairly wins a prize you also sought after, say). It is, however, a thorny question whether we can bear moral responsibility for beliefs and other attitudes, or even whether we can morally wrong someone by, for instance, falsely believing less of them.²⁰ To the extent that questions of moral responsibility are intimately bound up in considerations of intention-level control, the case for this sort of involuntarism faces the same old challenges. But moral responsibility for belief is not at all a part of the product I've tried to sell you; I leave it an open question.

Instead, what seems to carry over to my project from the involuntarists' is that one can be *epistemically responsible* for one's thinking even if one's thinking does not manifest voluntary control.²¹ If I am correct, then involuntarists of this stripe are right that responsibility for attitudes does not require voluntary control, but, depending on the

details of their commitments, they are wrong about which kind of responsibility and why.²²

References

- Adams, E.W. (1965). A logic of conditionals. *Inquiry* 8:166-197.
- Adams, R. (1985). Involuntary sins. *The Philosophical Review* 94(1):3-31.
- Alston, W. (1989). The deontological conception of epistemic justification. In *Essays in the Theory of Knowledge*. Cornell University Press.
- Aronowitz, S. (2021). Exploring by believing. *The Philosophical Review* 130(3):339-383.
- Arpaly, N. (2003). *Unprincipled Virtue: An Inquiry into Moral Agency*. Oxford University Press.
- Arpaly, N. and Schroeder, T. (2013). *In Praise of Desire*. Oxford University Press.
- Ayer, A. (1954). Freedom and necessity. In his *Philosophical Essays*. London: MacMillan Publishers.
- Basu, R. (2018). Can beliefs wrong? *Philosophical Topics* 46(1):1-17.
- Boult, C. (2021). Epistemic blame. *Philosophy Compass* 16(8):1-12.
- Boult, C. (2020). There is a distinctively epistemic kind of blame. *Philosophy and Phenomenological Research*, online first.
- Bratman, M. (1987). *Intentions, Plans, and Practical Reason*. Harvard University Press.
- Chisholm, R. (1978). Comments and replies. *Philosophia* 7:597-636.
- Clarke, R. (1993). Toward a credible agent-causal account of free will. *Noûs* 27(2):191-203.
- Clarke, R. (2018). Free will, agent causation, and 'disappearing agents'. *Noûs* 52(1):76-96.
- Davidson, D. (1963/1980). Actions, reasons, and causes. Reprinted in *Essays on Actions and Events*. Oxford University Press.
- Edgington, D. (1995). On conditionals. *Mind* 104:235-329.
- Fara, M. (2008). Masked abilities and compatibilism. *Mind* 117(468): 843-65.

- Frankfurt, H. (1987). Identification and wholeheartedness. In *Responsibility, Character, and the Emotions: New Essays in Moral Psychology*. Cambridge University Press.
- Franklin, C. (2016). If anyone should be an agent-causalist, everyone should be an agent-causalist. *Mind* 500:1101-1131.
- Gärdenfors, P. (1986). Belief revision and the Ramsey test for conditionals. *The Philosophical Review* 95:81-93.
- Ginet, C. (2001). Deciding to believe. In M. Steup (ed.) *Knowledge, Truth, and Duty*. Oxford University Press.
- Goldman, A. (1970). *A Theory of Human Action*. Princeton: Princeton University Press.
- Goldman, A. (1979). What is justified belief? In G. Pappas (ed.) *Justification and Knowledge*. Dordrecht: D. Reidel.
- Grice, H.P. (1989). *Studies in the Ways of Words*. Harvard University Press.
- Harman, G. (1976). Practical reasoning. In A. Mele (ed.) *The Philosophy of Action*. Oxford University Press.
- Hitchcock, C. (2007). Prevention, preemption, and the principle of sufficient reason. *The Philosophical Review* 116(4):495-532.
- Jackson, F. (1987). *Conditionals*. Oxford: Basil Blackwell.
- Jenkins, D. (2021). The activity of reasoning: How reasoning can constitute epistemic agency. *Pacific Philosophical Quarterly* 102(3):413-428.
- Kauppinen, A. (2018). Epistemic norms and epistemic accountability. *Philosophers' Imprint* 18(8).
- Kearl, T. (2022). A plea for exemptions. *Erkenntnis*. <https://doi.org/10.1007/s10670-022-00615-w>.
- Kearl, T. (Unpublished manuscript). Quality of thought.
- Kearl, T. and Wallace, R. (Forthcoming). Agentive modals and agentive modality: A cautionary tale. *American Philosophical Quarterly*.

- Kornblith, H. (2012). *On Reflection*. Oxford University Press.
- Lehrer, K. (1968). Cans without ifs. *Analysis* 29(1): 29-32.
- Lehrer, K. (1971). How reasons give us knowledge, or the case of the gypsy lawyer. *The Journal of Philosophy* 68(10):311-313.
- Lewis, D. (1976). Probabilities of conditionals and conditional probabilities. *The Philosophical Review* 85(3):297-315.
- McCain, K. (2012). The interventionist account of causation and the basing relation. *Philosophical Studies* 159: 357-382.
- Mele, A. (1992a). *Springs of Action: Understanding Intentional Behavior*. Oxford University Press.
- Mele, A. (1992b). Recent work on intentional action. *American Philosophical Quarterly* 29(3):199-217.
- Mele, A. (2017). *Aspects of Agency: Decisions, Abilities, Explanations, and Free Will*. Oxford University Press.
- Mele, A. and Moser, P. (1994). Intentional Action. *Noûs* 28(1):39-68.
- Moore, G.E. (1912). *Ethics*. Oxford University Press.
- Moran, R. (2001). *Authority and Estrangement: Essays on Self-Knowledge*. Princeton University Press.
- Moser, P. (1989). *Knowledge and Evidence*. Cambridge University Press.
- O'Connor, T. (2000). *Persons and Causes: The Metaphysics of Free Will*. Oxford University Press.
- Pereboom, D. (2014). *Free Will, Agency, and the Meaning of Life*. Oxford University Press.
- Ramsey, F. (1929/1990). General propositions and causality. In D.H. Mellor (ed.) *Philosophical Papers*. Cambridge University Press.
- Schlick, M. (1939). *Problems of Ethics*. Prentice-Hall.
- Schroder, M. (2021). *Reasons First*. Oxford University Press.

- Smith, A. (2005). Responsibility for attitudes: Activity and passivity in mental life. *Ethics* 115(2):236–271.
- Smith, M. (2003). Rational capacities, or: How to distinguish recklessness, weakness, and compulsion. In S. Stroud and C. Tappolet (eds.) *Weakness of Will and Practical Irrationality*. Oxford University Press.
- Sosa, E. (2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume I*. Oxford University Press.
- Sosa, E. (2003). Reply to BonJour. In L. BonJour and E. Sosa (eds.) *Epistemic Justification: Internalism vs. Externalism, Foundations vs. Virtues*. Blackwell.
- Sripada, C. (2016). Self-expression: a deep self theory of moral responsibility. *Philosophical Studies* 173 (5): 1203-1232.
- Steup, M. (2008). Doxastic freedom. *Synthese* 161:375-392.
- Steup, M. (2012). Belief control and intentionality. *Synthese* 188:145-163.
- Steup, M. (2018). Doxastic voluntarism and up-to-me-ness. *The International Journal of Philosophical Studies* 26:611-618.
- Strawson, G. (2003). Mental ballistics or: The involuntariness of spontaneity. *Proceedings of the Aristotelean Society* 103:227-256.
- Sturgeon, S. (2020). *The Rational Mind*. Oxford University Press.
- Swain, M. (1985). Justification, reasons, and reliability. *Synthese* 64(1):69-92.
- Turri, J. (2011). Believing for a reason. *Erkenntnis* 74(3):383-397.
- Velleman, D. (1992). What happens when someone acts? *Mind* 101(403):461-481.
- Vihvelin, K. (2004). Free will demystified: A dispositionalist account. *Philosophical Topics* 32(1):427-450.
- Vihvelin, K. (2013). *Causes, Laws, and Free Will: Why Determinism Doesn't Matter*. Oxford University Press.
- Watson, G. (2004). *Agency and Answerability: Selected Essays*. Oxford University Press.

Williams, B. (1973). Deciding to believe. In *Problems of the Self*. Cambridge University Press.

Williamson, T. (2016). Knowing by imagining. In A. Kind and P. Kung (eds.) *Knowledge Through Imagination*. Oxford University Press.

Williamson, T. (Forthcoming). Justification, excuses, and sceptical scenarios. In J. Dutant and F. Dorsch (eds.) *The New Evil Demon*. Oxford University Press.

Woodward, J. (2003). A theory of singular causal explanation. *Erkenntnis* 21(3):231–262.

Yablo, S. (2002). De facto dependence. *The Journal of Philosophy* 99(3):130-148.

¹ This at least applies to creatures with the sort of complexity we find in normally functioning human beings, but my own view is that it applies to much simpler creatures too.

² Davidson (1963)

³ On a coarse-grained view of events (Davidson (1963)), the unintentionality of my alerting the burglars is explained by the fact that that event is identical to the event of my intentionally illuminating the room. On finer-grained views of events (Goldman (1970)), the events are non-identical but nevertheless appropriately connected so as to exhibit the dependence I mention. Nothing in the text hangs on this choice.

⁴ Here's Chisholm (1978) reacting to what he sees as the poverty of event-causal explanations of action: ...Now if you can analyze such statements as "Jones killed his uncle" into event-causation statements, then you may have earned the right to make jokes about the agent as cause. But if you haven't done this, and if all the same you do believe such things as that I raised my arm and that Jones killed his uncle, and if moreover you still think it's a joke to talk about the agent as cause, then, I'm afraid, the joke is entirely on you. You are claiming the benefits of honest philosophical toil without even having a theory of human action (622-623).

⁵ One can see the same sentiment, more recently, in O'Connor (2000), Pereboom (2014), Clarke (1993, 2018) and Franklin (2016), although the latter two views have a somewhat deflationary stance about the nature of agent-causation.

⁶ Note that it's not obvious that Velleman is carrying Frankfurt's torch, at least if we look past the project of making good on some notion of "identification" within the event-causal nexus. That's because Velleman's account of identification is wholeheartedly event-causal, while Frankfurt's picture of authorship is non-causal; for an agent to be the author of her actions is for her effective desires (the ones that are causes of her behavior) to be such that she wishes them to be the effective ones. But her wishing that certain desires be effective need not be a cause of her behavior.

⁷ One might see the work of Bratman (1987), Arpaly (2003), Watson (2004), Arpaly and Schroeder (2013), and Sripada (2016) in this vein.

⁸ A great deal of work has been done on the related topics of whether and how conditional beliefs are true and the relationship between *q if p*—the contents of a conditional belief—and the material conditional $p \supset q$ (Adams 1965, Grice 1967, Lewis 1976, Jackson 1987, Gärdenfors 1986, and Edgington 1995). I won't stick

my neck out about the correct semantics and logic of indicative conditionals. Instead, I'm only concerned with the normative function of conditional beliefs in inference.

⁹ Aronowitz (2021) argues that background beliefs place side-constraints on the exercises of our imagination; these beliefs guide what and how we imagine, and so whether our imaginative searches result in knowledge. See also Williamson (2016).

¹⁰ Some prefer to explain propositional justification in terms of doxastic justification; on this inverted picture, a subject's being propositionally justified in believing P is a matter of her having available some doxastic justification-conferring route to believing P. The discussion below does not hinge on whether we take propositional justification to be explanatorily prior or posterior to doxastic justification.

¹¹ Moser (1989), McCain (2012), and Turri (2011) give causal accounts; Swain (1985) offers a counterfactual account, but given the tight conceptual connection between counterfactual and causal relations, this is a different species within the genus. Even doxastic accounts of basing, which require that the agent have a "meta belief" about what her reasons support, tend not to eschew a causal requirement entirely. (Though Lehrer's (1971) lawyer case purports to do just that; I follow Lehrer's critics but lack the space to offer a refutation here.)

¹² I'm using 'evidence' here very broadly, so as to capture both sides of many debates within epistemology about evidentialism, the claim that only evidence rationalizes belief-like attitudes.

¹³ Kornblith's problem seems to revolve around claims made by the likes of Moran (2001) and Sosa (2007), which involve the idea that *capacities for reflection* and *reflectively formed beliefs* are the marks of epistemic agency. See Jenkins (2021) for a recent defense of a similar, reasoning-based view of epistemic agency.

¹⁴ This picture resembles the compatibilist-friendly, conditional analyses of free action due to Moore (1912), Schlick (1939), and Ayer (1954), and more recent dispositional analyses due to Smith (2003), Vihvelin (2004, 2013), and Fara (2008).

¹⁵ I don't here mean to buy into a simple, counterfactual account of causal control à la Moore (1912), Schlick (1939), and Ayer (1954). Various problems with that account have been widely discussed (see, e.g., Lehrer (1968), Shope (1978), and Kearl and Wallace (forthcoming)). I only mean to point out that we can get a grip on the role of intentions in guiding intentional action by thinking about how the latter, at least in a range of paradigmatic cases, counterfactually depends on the former. This is true even in light of so-called "prevention" and "intervention" cases (Hitchcock (2007); Woodward (2003); Yablo (2002)). But if even that suggestion is objectionable, I invite the reader to slot in their preferred account of non-deviant causation instead.

¹⁶ See also Mele's *Aspects of Agency*, especially chapter 2.

¹⁷ As above, the appeal to counterfactual dependence is meant to be illustrative rather than reductive.

¹⁸ Though we sometimes reason about and update our conditional beliefs; see Sturgeon's *The Rational Mind*, chapter 4, for helpful discussion.

¹⁹ Jenkins also says that conscious reasoning manifests agency in virtue of its being aim-directed, but he goes on to say that "I will hold back from trying to say what in general the aim of reasoning is (and could even refrain from insisting that all reasoning is performed with an aim)" (10). Perhaps this suggests that Jenkins is better understood as articulating the features of *paradigmatic expressions of epistemic agency*, rather than as articulating *necessary conditions on epistemic agency*.

²⁰ Basu (2018); Schroder (2021)

²¹ See Boulton (2020, 2021), Kauppinen (2018), Kearl (2022; Ms.), and Williamson (forthcoming).

²² I am grateful to Rhys Borchert, Juan Comesaña, Michael McKenna, Carolina Sartorio, Robert Wallace, and Chris Willard-Kyle for comments on earlier drafts of this paper. This research was made possible with funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No 948356, KnowledgeLab: Knowledge-First Social Epistemology, PI: Mona Simion).
