Philosophers' Imprint

VOLUME 26, NO. 5
FEBRUARY 2026

THE NORMATIVE FUNCTION OF INTENTIONAL ACTION

Mikayla Kelley

University of Chicago

doi.org/10.3998/phimp.7981

1. Introduction

This essay will explicate and defend one aspect of the following idea: there are some seemingly non-normative concepts that function in distinctive ways to support our capacities for normative thought, talk, and practice. Some possible such concepts include HAPPINESS, PERSON, HUMAN, FREEDOM, and INTENTIONAL ACTION.[1] These concepts are not normative in the obvious sense that GOOD, OUGHT, and EVIL are normative. Yet they are candidates for concepts that are intimately connected to the space of normative concepts in the following way: they function to make our application of normative concepts possible. Or so I'll argue in one particular case here.

The particular case of concern is the concept of intentional action, a concept that has been of great interest in contemporary philosophy of action. I will argue that INTENTIONAL ACTION functions to support normative life by focusing our evaluative concern on some doings rather than others. The concept acts as a proxy for evaluative priority.

The argument that I offer is methodologically inspired by Craig's (1990) practical explication of KNOWLEDGE.[2] Like others working in the broadly pragmatist tradition, Craig approaches a concept of philosophical significance by first thinking about what it is doing for us—what needs it is satisfying. His work has led to a thriving "function-first" research program in epistemology where concepts and phenomena of epistemological interest are studied primarily through studying their functions in a human life.[3] Here, I enlist this function-first methodology to study intentional action.[4] Accordingly, I will locate a need that we

1. I use all caps to refer to concepts. I follow Queloz (2021) in viewing concepts as norm-governed thinking techniques characterized by distinctive bundles of dispositions to partition and draw inferences about the world in certain ways, ways that are disposed to influence behavior (p. 23). Like Queloz, I don't take this to be a real definition. Whether these bundles of dispositions are best explained in terms of, e.g., abilities, mental representations, or abstract entities, I take no stand.
2. Queloz 2022a also offers methodological inspiration.
3. See, e.g., Dogramaci 2012, Hannon 2015, Reynolds 2017, Hannon 2019, Lawlor 2023, and O'Sullivan and Mace 2024.
4. Most obviously, the function-first methodology that I enlist is a way of

have in being the sorts of normative creatures that we are—a need to focus evaluative concern (Section 2); and I will argue that we have good grounds to infer that INTENTIONAL ACTION functions to satisfy that need (Sections 3 and 4).

Philosophers of action have tended to study intentional action with focus on the intension of INTENTIONAL ACTION or the metaphysical joint picked out by INTENTIONAL ACTION. The outputs of such investigations tend to take the form of (arguments for) necessary and/or sufficient conditions on intentional action. I take neither approach and instead study intentional action with focus on the function of INTENTIONAL ACTION in our lives. The output of the investigation will not take the form of (arguments for) necessary and/or sufficient conditions but rather an argument for an articulation of a function of INTENTIONAL ACTION.[5]

It's surprising that there's been so little interest in the functions of the concept of intentional action among philosophers of action, given the prevalence of functionalist approaches throughout philosophy.[6] I conjecture that this lack of interest stems from a widespread implicit assumption that the concept of intentional action has a single function, namely to track a distinctive sort of control.[7] Thus, once we've understood the contours of the distinctive sort of control tracked by INTENTIONAL ACTION—a project that philosophers of action have been disproportionately engaged in—we've understood all there is to understand about the function of the concept.

I show that this implicit assumption is false: INTENTIONAL ACTION also has a normative function, a function which shapes its control-tracking function.[8] Moreover, I will show that with this additional function in view, we can explain in a unified way a number of features of intentional action that are otherwise puzzling (Section 5). First, it explains a widely documented phenomenon—the *Knobe effect*—where moral considerations affect the folk attribution of INTENTIONAL ACTION. Second, being an intentional action is a binary property of doings but being controlled is gradable; thus even granting that INTENTIONAL ACTION functions to track a distinctive sort of control, there remains a question of what settles whether some *degree* of that distinctive sort of control suffices for intentional action. This is the *threshold problem for intentional action* (Kelley 2025). By identifying a normative function of INTENTIONAL ACTION, we are in position to solve the threshold problem for intentional action. Third, there are a number of disagreements over the scope of intentional action, among philosophers and at a cultural level, which the normative function of INTENTIONAL ACTION can illuminate. These include disagreement over whether non-human animals possess intentional agency and whether intentional action requires knowledge of what one is doing (Anscombe 1957). Finally, a striking feature of intentional action is the radical pluralism of ways that it manifests (O'Brien 2017). The normative function of INTENTIONAL ACTION can explain this too.

studying the *concept* of intentional action. But studying the concept of intentional action is a way of studying the phenomenon of intentional action assuming the following plausible principle: INTENTIONAL ACTION applies to x iff x is an intentional action, and C are conditions in virtue of which INTENTIONAL ACTION applies to x iff C are conditions in virtue of which x is an intentional action (cf. Hannon 2019, Sec. 1.5).

5. An articulation of a function of INTENTIONAL ACTION might put us in a better position to articulate necessary and/or sufficient conditions, but I say little about such an upshot here.
6. There's been more interest in the functions of the concept of intentional action in cognitive science and, specifically, in the literature on the Knobe effect (see, e.g., Knobe 2008). This literature is large, and I can't respond to it fully (see Section 5 for a brief discussion of the Knobe effect). However, my project differs methodologically in a crucial respect: while the literature on the Knobe effect relies heavily on experimental data about the folk usage of 'intentional action' and its cognates, such data plays a small role in my argument.
7. In effect, we've largely heeded Anscombe's (1958) direction to *banish ethics totally* from our minds in studying intentional action. There also may be general methodological commitments supporting such an implicit assumption. For example, Cappelen (2018) claims that the only function of a concept C is denotational, that is, to track the Cs (p. 181); see Thomasson 2024 for a reply.
8. Nothing I say will preclude that INTENTIONAL ACTION has even more functions than these two. If so, the question of how all the functions of INTENTIONAL ACTION "hang together" is a pressing topic for further investigation.

2. A Human Need

We are normative creatures, meaning that normative thought, talk, and practice are deeply ingrained in the way we lead our lives. But there is a basic challenge we face in being normative creatures, on the one hand, and perceptive, social, curious, and cognitively limited creatures, on the other. The challenge is that doings—a central class of targets for practical forms of normative evaluation—permeate our experience of the world in a way that would, without help, overwhelm our cognitive capacities.[9]

To see this, notice how many doings we are acquainted with at a given moment. Imagine, for example, sitting outside a coffee shop on a busy street. In a matter of seconds you might perceive a couple walking, a family conversing, a door slamming, a dog barking, a leaf falling, a person tripping, and a child laughing. We are acquainted with doings not just via perception but via more indirect methods, as well. You might read a story in the newspaper about a local robbery, see a photograph of someone playing fetch with their dog, or your friend might tell you a story featuring their child saying something hurtful to them. I mean for 'acquaintance' to encompass all of these modes of being in contact with a doing.

Notice next that doings of all kinds could be subject to various kinds of normative evaluation.[10] The following list exemplifies the diversity among the proper targets of normative evaluation according to our actual practices of normative evaluation:[11]

- You ought not have accidentally insulted me.
- Your fearful reaction was irrational.
- The dog should have obeyed his owner.
- It is bad that the acid eroded the metal on your bike.
- The lion killing the zebra in that way was cruel.
- I blame you for carelessly leaving the milk out to spoil.

So doings can be assessed with evaluative and deontic concepts—such as GOOD, BAD, IMPERMISSIBLE, RATIONAL, and OUGHT—as well as with thick ethical concepts like CRUEL and GENEROUS; and they can be the target of practices such as praise, blame, and other mechanisms of holding responsible. Essentially every doing we are acquainted with (of which there are many) is a proper target of normative evaluation of some sort.

On the other hand, we are cognitively limited. We can only store, process, and attend to so much information at a time. Moreover, we are interested in spending our time doing more than storing, processing, and attending to doings for the purpose of normative evaluation. An important question follows: given one's cognitive and temporal limitations, on which of the many normatively evaluable doings to which one is acquainted shall one focus evaluative concern? Answering this question requires *focusing evaluative concern*.[12] Focusing evaluative concern is similar to focusing philosophical concern while reading challenging philosophy. We are limited in time and cognitive resources and so can't take in all the information conveyed in a philosophical work with perfect understanding and accuracy; thus, we have methods

9. I assume that a doing is an event of A φing, for some entity A and some act type φ (for further discussion of doings, see Baron-Schmitt 2024). I use a fine-grained individuation of doings, such that there is a distinct doing for each entity-act-type pair. So let d be a doing which is an event of A φing and d' be a doing which is an event of B ψing; then $d \neq d'$ if $A \neq B$ or $\varphi \neq \psi$. This choice of individuation isn't strictly necessary, but it makes the discussion cleaner.
10. Throughout, I use 'normative evaluation' to refer to practical forms of normative evaluation in particular, where the practical forms of normative evaluation are unified in their being centrally about doings.
11. More precisely, a kind of normative evaluation E has a *domain of application* defined by the set of x for which the question "is x E?" is felicitous according to our actual practices of normative evaluation. To say that d is a proper target of normative evaluation is to say that d falls within E's domain of application for some kind of normative evaluation E. I emphasize our *actual* practices of normative evaluation because it is the nature of our actual practices that generate conceptual needs, not our normative practices as they ought to be upon sophisticated philosophical reflection.
12. Prioritization structures play a central role in models of perceptual attention (Watzl 2017). The thought here is that prioritization structures also play a central role in "evaluative attention" (which is connected to, though not identical to, perceptual attention).

of focusing our cognitive resources in order to ensure accurate and thorough understanding of the most important parts of the text. In the evaluative case, when you focus evaluative concern on a doing, you take a special interest in that doing for the purpose of normatively evaluating it. You will be disposed to perceptually attend to the doing, commit it to memory, store more details about the doing than others, begin inquiry into its normative status, flag it as playing a distinctive role in your normative reasoning, and cognitively process it in relation to your normative aims.

In sum, we have a basic need in being the sorts of normative creatures that we are: we need to focus evaluative concern on some of the doings to which we're acquainted rather than others. We should infer that we have a cognitive mechanism up and running that satisfies this basic need. After all, we manage to engage in complex forms of normative evaluation seamlessly and with great skill and frequency. A task for us as theorists is to identify this cognitive mechanism. As we'll see, in taking on this task, we come to understand a central function of the concept of intentional action.

3. Direct vs. Indirect Prioritizing

There are two natural possibilities for how we focus evaluative concern. The first option is that we track a prioritized set of doings directly via a shared concept: the concept of a doing of evaluative priority. The second option is that we track the prioritized doings indirectly via a shared *prioritizing proxy* that acts as a proxy for the concept of a doing of evaluative priority. A prioritizing proxy would be a shared concept (or suite of shared concepts) with three features: (i) its extension would correspond roughly to the doings of evaluative priority, (ii) its intension would *not* include a clause about evaluative priority, and (iii) it would tend to trigger in us the suite of cognitive activities involved in focusing evaluative concern when applied.[13] Such a concept would function to focus evaluative concern but would do so from sideways on. There are three reasons to pursue the second option.

A preliminary point is that in introspecting on the categories with which we divide up the world, we just don't seem to regularly operate with a concept of a doing of evaluative priority. Instead, we regularly operate with concepts of agency. But perhaps the concepts with which we regularly operate are not fully transparent to us, so this point is only preliminary. The following two points show that a concept of a doing of evaluative priority couldn't generally satisfy our need to focus evaluative concern given contingent features of our situation.

First, tracking the prioritized doings directly requires access to information we often lack. Indeed, whether a doing is of priority for normative evaluation depends on its context, in particular, the extrinsic features of the doing that determine whether the doing calls for normative evaluation, which forms of normative evaluation are called for, and the urgency of the normative evaluations called for.[14] These extrinsic features concern how the doing is situated with respect to broader social-causal-historical structures. They include, for instance, the doing's upstream and downstream causal relations and the doing's social meanings that are fixed by the rule-based social structures within which the doing is embedded. For instance, you might witness A greeting B. Knowing only the doing's intrinsic features, you can't tell whether it calls for any sort of evaluation and thus whether it is of priority. But say you find out that the greeting caused B to be offended. This is an extrinsic feature of the doing which, in many contexts, would entail that the greeting calls for, e.g., deontic assessment and is of evaluative priority. Or say the greeting was performed within a community where greetings are faux pas. Being a faux pas would be an extrinsic feature of the doing that, in such contexts, would entail that the doing calls for some sort of assessment and is of evaluative priority.

13. Might we all operate with different prioritizing proxies? Maybe, but I'll identify a shared concept that looks extremely well-suited to be a prioritizing proxy. Once this is established, I see little reason to prefer the alternative that we each use our own.

14. A doing *d calls for* a normative evaluation E if the question 'is *d* E?' is not only felicitous but of some practical significance.

Often, when we are first acquainted with a doing, we know very little about its broader social meaning or its downstream and upstream causal relations. In fact, we often never learn much about a doing's context or, if we do, we often learn about it long after we are acquainted with it (and thus long after we must begin the suite of cognitive activities involved in focusing evaluative concern). For example, I might hear A greet B but not find out until a dinner party weeks later that B was offended by A's greeting. Or consider again sitting outside a coffee shop on a busy street—for the vast majority of the doings you perceive, you'll know very little about their context. In short, we usually only have access to snapshots in the lives of the agents we're surrounded by. Yet to determine whether a doing has the property of being of priority for normative evaluation requires access to much more information than this. So the direct method of prioritizing is too *informationally costly*.

The second problem for the direct method of prioritization is that even when we have access to such contextual information, it can take substantial cognitive resources to use that contextual information to determine whether the given doing is of evaluative priority. For instance, I might notice that A is frowning at B's movement. Determining what this fact means for the evaluative priority of B's movement may be no simple task, given the many other contextual factors that must be balanced. A direct method of prioritizing would require us to move through the world continuously engaging in such relatively difficult "calculations" of what doings are of evaluative priority given the highly specific features of the current context. But our need to prioritize stems from, in part, our cognitive limitations, and so a solution to it should respect these limitations. Thus, the direct method of prioritizing is too *cognitively costly* to generally satisfy our human need to prioritize.[15]

Taken together, these three considerations motivate looking for a prioritizing proxy that can be identified with little access to contextual information and with cognitive ease—a concept that would, in an indirect manner, function to satisfy our need to focus evaluative concern.

15. For all I've said, we may *sometimes* directly flag doings as of evaluative priority (I'll return to this point). But for the reasons discussed, this is not likely to be our only or usual method of focusing evaluative concern—we need a prioritizing proxy.

4. Intentional Action as a Prioritizing Proxy

My strategy for identifying a prioritizing proxy is two-fold. First, I lay out the features we'd expect a prioritizing proxy to have and attempt to match them with one of our existing folk concepts. If we identify a unique match, this gives us strong grounds to infer that the identified concept is a prioritizing proxy. I then strengthen the grounds for the inference by showing its explanatory power. I begin with the first part of the argument in this section.

So we ask: if some concept were to be a good prioritizing proxy, what would it need to be like? In fact, we've already established enough to identify three conditions on a prioritizing proxy that uniquely identify INTENTIONAL ACTION as a strong candidate. Indeed, for some concept to be a good proxy for another, the concept needs to be roughly co-extensive with that which it is a proxy for. Applying this general principle to our case, we have:

> **Extensional Accuracy** (EA): a prioritizing proxy picks out roughly the doings that are of priority for evaluative concern.

A prioritizing proxy must also be sensitive to our informational and cognitive limits discussed in the previous section. This gives rise to the following two conditions:

> **Informational Efficiency** (IE): in the typical case, whether a prioritizing proxy applies to a doing is largely independent of its context.

> **Cognitive Efficiency** (CE): in the typical case, we can apply a prioritizing proxy with cognitive ease.

The qualifiers in the three conditions reflect a tension between the need for an *accurate* proxy (captured by EA) and the need for an *efficient* proxy (captured by IE and CE). Any proxy needs to be both easy to use and

accurate to that which it is a proxy for, but these two constraints tend to pull in opposite directions. In our case, evaluative priority is often fixed by the context of the doing—by its upstream and downstream causal relations, as well as its broader social meaning; so if we wanted to ensure a correct determination of whether a doing is worthy of evaluative focus, we would often need to do relatively difficult calculations using this rich web of contextual information. Insofar as a prioritizing proxy bypasses this demanding work in being informationally and cognitively efficient, we should expect it to get the wrong result sometimes. How exactly the demands for accuracy and efficiency are traded off will depend on a number of factors. By understanding the nuances of these trade-offs, we'd come to a fuller understanding of the nature of a prioritizing proxy (as illustrated in Section 5).

INTENTIONAL ACTION satisfies EA, IE, and CE. Start with IE. A striking feature of INTENTIONAL ACTION is that (in the usual case) its application is largely independent of context: to reliably tell whether some φing is intentional, one needs to know very little about the context of the φing. You don't need to know what the φing caused downstream, and you need to know little about its upstream causal relations—you only need to know that the doing stands in the right sort of (causal) relationship to (psychological states of) the doer. You don't need to know anything about the broader social or normative implications of the φing. You don't need to know if it was right, wrong, cruel, kind, or impolite. You don't need to know historical facts about the source of the agent's tendency to φ or whether the φing was coerced or negligent. So in the typical case, INTENTIONAL ACTION is extremely informationally efficient.

Turn next to CE. First-person reflection suggests that we're highly skilled at telling what others are doing intentionally (and non-intentionally). Anscombe (1957) astutely observes this:

> "I am sitting in a chair writing, and anyone grown to the age of reason in the same world would know this as soon as he saw me, and in general it would be his first account of what I was doing; if this were something he arrived at with difficulty, and what he knew straight off were precisely how I was affecting the acoustic properties of the room (to me a very recondite piece of information), then communication between us would be rather severely impaired." (§4)

Anscombe claims that this "first account" of what another is doing will usually be an account of what the other is doing *intentionally*. Moreover, we are able to come up with such accounts "straight off". As she puts it, if I were sitting in a chair writing, "anyone grown to the age of reason in the same world would know this *as soon as he saw me*" (emphasis added). Even more, if it *were* difficult to know what others are doing intentionally, this would cause great strain on our capacity to communicate—and more generally, cooperate—with one another. Anscombe's observations are well-grounded in first-person reflection: identifying what another is doing intentionally is just not something that we very often find difficult.[16] Of course, there are cases where the intentionality of some doing comes into question and we are forced to think hard about it—for example, in the court room.[17] But these cases are notably rare.

Lastly turn to EA. To get a grip on which doings are of evaluative priority, we can start from the functions of (practical forms of) normative evaluation, for the doings of evaluative priority are roughly those whose prioritization best supports the functions of normative evaluation. Two central functions are (i) behavioral modification and (ii) generating

16. For an empirical perspective on these points, see the extensive literature on "mind-reading".
17. Relatedly, as Anscombe notes in §4, sometimes there is something done intentionally and we do not know it is being done at all. If you see someone intentionally walking across the street, you may not know she is also (intentionally) walking to her grandmother's house. But in such cases, you aren't faced with the task of determining whether to prioritize the doing, for you lack access to its being done at all. In other words, the task we face is prioritizing among the doings we are acquainted with *in a sense that requires some awareness of what one is acquainted with*; and Anscombe seems right that among those, we do not struggle to tell which are intentional.

information about the individual whose behavior is being evaluated. In short, we normatively evaluate in order to control and understand the things in our environment. These two functions would be supported by prioritizing doings that are especially modifiable in light of normative evaluation and those which are representative of the doer as a unified, temporally extended entity. Note now that intentional actions are consciously controlled and so especially modifiable; moreover, intentional actions speak for and are attributable to the agent given the way they express an aim which is attributable to the agent.[18] Doings that are not intentional actions, on the other hand, will tend to fail at least one of these criteria. So prioritizing precisely the intentional actions would tend to support the functions of normative evaluation. From this we can conclude that INTENTIONAL ACTION is roughly co-extensive with the class of doings of evaluative priority.

In fact, we can say more. As noted, the reason that a prioritizing proxy should be expected to only *roughly* pick out doings that are of priority for evaluative concern is that there is a tension between being an accurate proxy and being an efficient proxy. Thus, say that a concept is *maximally extensionally accurate* if mismatches between the extension of the concept and the class of doings of evaluative priority can be explained by appeal to advantages in efficiency. A concept that is maximally extensionally accurate, in addition to satisfying EA, has a more robust case for being the prioritizing proxy. Moreover, by adding this maximality condition, proposals for the prioritizing proxy are more easily falsifiable.

INTENTIONAL ACTION is also maximally extensionally accurate. Indeed, there are two kinds of mismatches between the extension of INTENTIONAL ACTION and the class of doings of evaluative priority. First are doings which are non-intentional because they are not consciously controlled but nonetheless could and should have been consciously controlled. Negligent unintentional actions are examples—for instance, falling down the stairs onto someone because you intentionally omitted to tie your shoelaces. Second are doings which are intentional (resp. non-intentional) but whose evaluation would be practically insignificant (resp. significant). But both whether a non-intentional action could and should have been consciously controlled and the extent to which an evaluation of a doing is of practical significance are features of doings that we often lack immediate access to at the time of acquaintance with the doing. Thus, ignoring such features in prioritizing would be advantageous from an efficiency perspective.[19]

So INTENTIONAL ACTION is a cognitive mechanism that looks extremely well-designed to satisfy our need to focus evaluative concern. I will not argue that we use only a single tool to focus evaluative concern, for it would not be surprising if INTENTIONAL ACTION were supplemented by other cognitive tools (recall Footnotes 15 and 19). That said, if such cognitive tools looked redundant rather than supplementary relative to INTENTIONAL ACTION, this might put pressure on my proposal. A useful heuristic to tell the difference is as follows: consider a scenario where creatures otherwise very much like us lacked a concept of intentional action; then ask: would such creatures struggle to fulfill their need to focus evaluative concern, or would some combination of other cognitive tools serve them just as well? I will consider a number of

18. I take a theory of intentional action to be a theory of what this conscious control and attributability more precisely consist in.

19. I've been emphasizing that we should expect a prioritizing proxy to not *require* significant access to the contextual information which determines the evaluative significance of the doings to which we are acquainted. But when we *do* have access to such information, it could play a role in prioritizing. Generally speaking, a judgment on the basis of a proxy is often made in explicit recognition that it is subject to modification, for one recognizes that one might gain direct access to the property that the proxy is a proxy for. For instance, I might use strength of writing sample as a proxy for readiness to pursue a philosophy Ph.D. But if I received information that immediately entailed a lack of readiness, then I could update my judgment accordingly. Still, the proxy determines my initial, default judgment which is then subject to revision. INTENTIONAL ACTION may work in a similar way. It constructs an initial, default sorting of one's experience into doings which are to be the focus of one's evaluative concern and those which are not. The resulting evaluative focus may then be further refined using whatever information one has access to with whatever cognitive resources one has available to put toward this task.

possible replacements and show that each either fails to meet one of EA, IE, or CE, relies on INTENTIONAL ACTION functioning as a prioritizing proxy, or is in some other way limited in its support of our need to focus evaluative concern.

Take first the concept of acting of one's own free will, which I'll denote with 'FREE ACTION'. Determining whether an agent acted of their own free will is informationally and cognitively taxing. You have to check for all sorts of external influences on the agent's behavior, including coercion, addiction, sleep deprivation, manipulation, and brainwashing. Of course, what exactly you have to check depends on the correct theory of free will, but many theories incorporate elements that are not transparent in day-to-day interactions (e.g., on Frankfurt's (1971) theory of free will, one would need to be able to tell the difference between a willing and an unwilling addict). So FREE ACTION fails IE. There is another problem with FREE ACTION: our grasp of the concept is notably obscure. As the longstanding debate over free will illustrates, it is far from clear what is required to act of one's own free will and which cases count as cases of acting of one's own free will. Thus, even with all possible information about a doing, applying the concept to particular cases will often be no easy task. A concept as vexed as FREE ACTION is too difficult to apply in any regularity and thus fails CE.

The first reason that FREE ACTION fails to be a plausible candidate prioritizing proxy is also a reason that VOLUNTARY ACTION is not a plausible candidate. According to Williams (1997) and Queloz (2022a), a voluntary action is an intentional action that is done in a "normal state of mind"; and Queloz argues that VOLUNTARY ACTION functions to ensure fairness in the attribution of responsibility and freedom to determine the course of our own lives (p. 1605). However, VOLUNTARY ACTION wouldn't fare well as a prioritizing proxy because it, like FREE ACTION, is too informationally rich. To determine whether something is done in a normal state of mind requires one to determine what counts as a normal state of mind for a given agent and what sort of state of mind the agent is in (e.g., whether they're being manipulated, coerced, or sleep deprived)—information that we very often lack when engaging with others. So VOLUNTARY ACTION fails IE.[20]

A third proposal is that we have a set of *prioritized act types*, and we prioritize the token doings that fall under at least one of those act types. An initial problem for this proposal is that presumably the set of act types whose tokens are (roughly) the ones of priority is largely context dependent; yet continually constructing a set of prioritized act types based on the context would be cognitively taxing. Thus, to satisfy EA, this method must fail CE. More importantly, the proposal appears to rely crucially on INTENTIONAL ACTION functioning as a prioritizing proxy and thus would not operate well for creatures very much like us who lacked a concept of intentional action. Indeed, how do we construct the set of prioritized act types (for a given context)? More specifically, could we come to a feasible set of prioritized act types satisfying EA without a concept of intentional action? I'm skeptical. More naturally, the reason that we prioritize the act types we do is at least partly because their tokens tend to be flagged as of priority by a prioritizing proxy. I prioritize the act type 'walk' because walkings tend to be intentional and thus represented as of priority; and I don't prioritize the act type 'sneeze' because sneezings tend to be non-intentional and thus not represented as of priority.

A final proposal is that we have some (perhaps non-conceptual) mechanism which focuses evaluative concern on the doings whose evaluation would serve an immediate goal, and we can tell which doings have this property with little information and cognitive effort.

20. A similar problem arises for Hyman's (2013; 2015) conception of the voluntary which requires, roughly, that the thing done is not done out of ignorance or compulsion. Anscombe (1957) has a conception of the voluntary that is wider than the concept of intentional action but fails maximal extensional accuracy. Anscombe's permissive conception of the voluntary includes doings that "happen to one's delight so that one consents and does not protest or take steps against them", e.g., when one voluntarily heads out to sea by being pushed while on a boat (p. 89). In many such cases, one will not be in a position to modify their doing that they welcome—at least not in the direct and highly effective way that one can modify one's intentional actions. Therefore, such voluntary doings are not generally of evaluative priority. Yet I see no benefit to efficiency in prioritizing them.

Whether the latter is true, something would undoubtedly be lost if we only used such a mechanism to prioritize, for we need to be prepared to navigate our continually changing environment and goals. Thus, we have a strong standing interest in understanding entities in our environment and in being prepared to make interventions on their behavior in a manner that is *flexible* and *open-ended*. To satisfy this interest, we need a method that prioritizes in a (more or less) task-neutral manner, as INTENTIONAL ACTION does.

When all of these arguments are taken together, we have strong grounds to infer the *Prioritizing Thesis*: INTENTIONAL ACTION is a prioritizing proxy and functions to satisfy our need to focus evaluative concern. The case for the Prioritizing Thesis so far has two parts: INTENTIONAL ACTION is well-suited to prioritize and there is no other mechanism which could support prioritization in the distinctive way INTENTIONAL ACTION would.

It is important to note that the Prioritizing Thesis is a thesis about what INTENTIONAL ACTION is *actually* doing for us in our evaluative practices. So, the proposed inference is a defeasible inference from a "how-possibly" explanation to a "how-actually" explanation. Like all such inferences, it is empirically accountable. In the next section, I use the Prioritizing Thesis to systematically explain some features of how INTENTIONAL ACTION actually works in human cognition. This will begin an empirically grounded case for the thesis, but it will be left far from complete.

Before moving on, note that the second part of the Prioritizing Thesis—the claim that INTENTIONAL ACTION has focusing evaluative concern as a function—is not a substantive addition to the first part of the thesis—the claim that INTENTIONAL ACTION is a prioritizing proxy. Rather, it follows immediately from the definition of being a prioritizing proxy and two plausible accounts of conceptual function.

First, conceptual function can be understood as a *system function*. The system function of a concept is the role that the concept plays in promoting the functioning of some conceptual system of which it is a part (Cummins 1975; Thomasson 2024). System functions are therefore relativized to some larger system and a capacity of that larger system. For example, the system function of a heart relative to the body is to disperse oxygen throughout the body, and the relevant capacity of the body is its sustaining life. Focusing evaluative concern is a system function of a concept relative to "the normative conceptual scheme"—the web of interrelated concepts with which we tend to normatively evaluate practical life—and the relevant capacity of the normative conceptual scheme is its guiding normative thought, talk, and practice.

Alternatively, Queloz (2021; 2022b) defends a *need-satisfaction account* of conceptual function. On this account of conceptual function, to say that a concept C has the function of serving some need N is to say that living by C tends to cause effect B and B tends to cause the satisfaction of N. Living by a prioritizing proxy, by definition, tends to cause the effect that we focus evaluative concern on a manageable class of doings of evaluative priority. This causal effect trivially causes the satisfaction of our need to focus evaluative concern. So INTENTIONAL ACTION has the function of focusing evaluative concern on Queloz's need-satisfaction account of conceptual function too.

5. Explanatory Power of the Prioritizing Thesis

In this section, I strengthen the case for the Prioritizing Thesis considerably by illustrating its explanatory power. Consider the following five facts about intentional action:

1. The moral valence of a doing affects whether the folk attribute INTENTIONAL ACTION to it.
2. There is a distinctive threshold of sufficient control which governs the scope of intentional action.
3. There is widespread disagreement about the extension of INTENTIONAL ACTION, especially with respect to uncertain actions, side-effect cases, creative agency, and animal agency.
4. There are cultural differences in how the folk ascribe INTENTIONAL ACTION.
5. There is a radical pluralism of ways to act intentionally.

With the Prioritizing Thesis in hand, we can explain all of these facts in a unified way. While I can offer only sketches of each of the five explanations, I hope to show that they are sufficiently illuminating to constitute genuine support for the Prioritizing Thesis and the broader project of identifying the functions of our concepts of agency.[21]

Start with the first fact, standardly called the *Knobe effect*. The Knobe effect is a well-documented phenomenon where aspects of a doing's moral context affect ascription of INTENTIONAL ACTION. In particular, a doing with a negative moral valence is more likely to be described as intentional by the folk than a doing with a positive moral valence, even when intuitively the agent has the same level and kind of control over the doings. Moreover, a doing with either positive or negative moral valence is more likely to be described as intentional by the folk than a doing with neutral moral valence. Part of what is so puzzling about the Knobe effect is that there is a strong competing intuition that being an intentional action is independent of moral context——an intuition that aligns with the Prioritizing Thesis. Squaring the empirical results with this intuition has led to a large research program. Here, I offer a sketch of an explanation of the Knobe effect using the Prioritizing Thesis. I do so by returning to the tension between the accuracy and efficiency constraints on a prioritizing proxy.

Recall that in order to be an efficient proxy, a prioritizing proxy must typically be (largely) context independent, since it is cognitively and informationally costly to pay close attention to contextual facts in focusing evaluative concern. So assuming the Prioritizing Thesis, INTENTIONAL ACTION must typically be context independent. But there is pressure to have an accurate proxy too, one which roughly tracks the doings of evaluative priority. Thus, when useful contextual information that entails a doing is of evaluative priority is easily accessible, there will be pressure to utilize that information in applying the prioritizing proxy. Notably, the Knobe effect occurs in precisely these sorts of cases, where contextual information which makes clear the moral valence of the doing and thus its evaluative priority is offered up for free. So the Prioritizing Thesis and EA predict the Knobe effect: INTENTIONAL ACTION is ascribed differently in morally charged and morally neutral cases when such moral valences are easily accessible.[22] In this way, exceptional cases where ascription of INTENTIONAL ACTION is context dependent can be explained by appeal to the specific trade-offs that INTENTIONAL ACTION makes between being an accurate proxy and being an efficient proxy for evaluative priority.

One might wonder why some doings—for example, someone forgetting their partner's birthday—strike us as unintentional even when there is easy access to contextual information which clearly entails they are of evaluative priority. Wouldn't my proposed explanation of the Knobe effect using the Prioritizing Thesis predict that one would ascribe INTENTIONAL ACTION to *any* doing which one could easily see is of evaluative priority? The answer is 'no'. INTENTIONAL ACTION manages to be an efficient proxy by functioning to track a distinctive kind of control—a kind of control which is cognitively and informationally easy to track.[23] Call it *intentional control*. Thus, if a doing fails to exhibit any intentional control (as is the case in the birthday example), it will not be flagged as intentional even if it is of evaluative priority. But control is gradable, and there are cases where a doing exhibits a *middling* degree of intentional control. Notably, the Knobe effect shows up in precisely these sorts of cases.[24] Where there is a middling degree

21. Due to space, I will also do little to argue that the explanations I offer are better than all possible alternatives.

22. Moreover, insofar as doings with negative moral valence are of greater evaluative priority than doings with positive moral valence, which seems plausible, the Prioritizing Thesis and EA would also predict the greater likelihood of ascribing INTENTIONAL ACTION in the former case.

23. It's worth making two observations here. First, we see that the widely recognized control-tracking function of INTENTIONAL ACTION works in service of, rather than competes with, its prioritizing function. Second, the practice of ascribing INTENTIONAL ACTION has what Queloz (2021) calls "self-effacing functionality" (p. 57). It is precisely by not using INTENTIONAL ACTION with the "aim" of prioritizing—but rather with the "aim" of tracking a certain kind of control—that INTENTIONAL ACTION is able to prioritize.

24. For instance, in the original chairman example (Knobe 2003), there are many

of intentional control, there is room for the ascription of INTENTIONAL ACTION to incorporate easily accessible contextual information while still tracking intentional control—specifically by tracking a *degree* of intentional control which is influenced by contextual factors. In this way, the Prioritizing Thesis can explain the Knobe effect *and* its limits.[25]

This discussion of the Knobe effect also shows that the Prioritizing Thesis is poised to help us solve the threshold problem for intentional action (Kelley 2025). INTENTIONAL ACTION functions to track a distinctive sort of control, but control is gradable and so there is a question of what sets the degree of that distinctive sort of control which suffices for intentional action in any given case. If the only function of INTENTIONAL ACTION is to track a distinctive sort of control, there is no solution to this problem in sight. However, with the Prioritizing Thesis in hand, a preliminary solution is immediate: the threshold of sufficient control for intentional action is set by the further functions of INTENTIONAL ACTION, including its prioritizing function. In particular, the Prioritizing Thesis entails the following characterization of the control threshold: the control threshold is set in such a way that INTENTIONAL ACTION can play its prioritizing function.

The Prioritizing Thesis can also explain widespread disagreement about the extension of INTENTIONAL ACTION. For instance, there is widespread disagreement over the *Knowledge Thesis* which claims that one must know that one is φing to be intentionally φing (Anscombe 1957). Consequently, there is disagreement over whether INTENTIONAL ACTION applies to purported counterexamples to the Knowledge Thesis where the agent fails to know what they're doing, e.g., Davidson's (1980) much discussed ten carbon copies example. This disagreement can be explained using the Prioritizing Thesis.

On the one hand, the most typical cases of intentional action do involve knowledge of what one is doing, so where one lacks this knowledge, there is a middling degree of intentional control.[26] Since tracking complete (or near complete) intentional control would be cognitively efficient for a prioritizing proxy—given the costs of applying a concept to what differs significantly from a paradigm case—CE supports not ascribing INTENTIONAL ACTION where knowledge is missing. But knowledge of what one is doing isn't required for the doing to be of evaluative priority; often, mere belief or intention suffices. For instance, a behavior guided by mere belief or intention can be highly modifiable and expressive of the agent as a temporally extended entity (cf. Kelley 2025, p. 974). Thus, EA supports sometimes ascribing intentionality even where knowledge is missing. So disagreement over the Knowledge Thesis stems from users of INTENTIONAL ACTION making different implicit trade-offs between using an accurate proxy and using an efficient proxy for evaluative priority. A similar story can be told about disagreement over whether bringing about an anticipated but unintended side-effect is intentional and about whether creative agency amounts to intentional agency.

features of intentional control manifest in the chairman's bringing about the environmental effect as a result of intentionally starting the new program. Stipulatively, he doesn't intend to bring about the environmental effect, but his attitude toward bringing it about is *intention-like*. For instance, he in some sense *prefers* that he bring it about rather than not; moreover, he settles on bringing about the environmental effect and knows that he's bringing it about while doing so. A similar point applies to the cases discussed in Nadelhoffer 2004 and Nadelhoffer 2005.

25. A fuller discussion of the Knobe effect would address a wider range of data and compare the proposal to existing alternatives. I've also said nothing about the psychological mechanism through which contextual factors related to evaluative priority affect one's sense of where the control threshold is set. As far as I can tell, what I've said is consistent with Knobe's own proposal (Pettit and Knobe 2009; Knobe 2010). Indeed, it could be that contextual factors related to evaluative priority affect the strength of pro-attitude that is treated as *default*, and this default partly determines one's sense of where the control threshold is set. The Prioritizing Thesis could also supplement Knobe's proposal by offering a specific explanation of why application of INTENTIONAL ACTION utilizes the normativity-infused psychological mechanism he describes.

26. Accordingly, I propose an alternative to the Knowledge Thesis: exhibiting the kind of control distinctive of intentional action to a complete degree requires knowledge of what one is doing. This is consistent with the claim that a middling degree of intentional control does not require such knowledge, nor does acting intentionally.

There is also widespread disagreement over whether non-human animals can act intentionally. Denial of non-human intentional agency is sometimes explicit (e.g., Stoecker 2009), but often it is implicit—it is seen in the way that theories of intentional action are developed with hardly any interest in non-human animal examples or in illuminating the distinctive forms of intentional agency exhibited by non-human animals. For instance, within a dominant tradition, intentional action is identified with acting for reasons in a sense that implicates language-use (Davidson 1963). On such a view, animals without language can't intentionally act.

But setting aside rationality, many animals clearly exhibit other typical features of intentional control over their behaviors. They act in conscious, skilled, thoughtful, knowledgeable, flexible, goal-directed ways.[27] So consideration of efficiency would support categorizing their behaviors as intentional, given they exhibit near complete intentional control. On the other hand, hesitation to categorize non-human animal behavior as intentional can be explained by the pressure for an accurate proxy: non-human animal behaviors are not generally of evaluative priority. So again, disagreement over animal agency stems from users of INTENTIONAL ACTION making different trade-offs between the pressure to use an accurate proxy and the pressure to use an efficient proxy.

A special form of disagreement over the extension of INTENTIONAL ACTION happens at the cultural level. The Prioritizing Thesis offers a different sort of explanation for cultural differences in ascribing INTENTIONAL ACTION: they reflect differences in what sorts of doings are taken to be of evaluative priority. For instance, studies have shown that in certain cultures, the Knobe effect is reversed: INTENTIONAL ACTION is ascribed more readily in positive moral valence cases than negative moral valence cases (Clark *et al.* 2017; Robbins, Shepard, and Rochat 2017). This could reflect cultural norms that, in certain cases, prioritize praise and other positive forms of evaluation over blame and other negative forms of evaluation.

Finally, turn to our last explanandum. A striking fact about our inquiry into the nature of intentional action is that we've been unable to identify conditions on acting intentionally that are both precise and necessary. Many proposed precise conditions on intentional action have fallen prey to systematic classes of counterexamples, including conditions in terms of knowledge, intention, motivating reasons, event-causation, causal guidance, belief-desire pairs, and counterfactuals.[28] It is, of course, difficult to give precise and necessary conditions on any concept or phenomenon of philosophical interest. But the case of INTENTIONAL ACTION is particularly striking. Proposals for precise necessary conditions are not subject simply to one-off, bizarre counterexamples, but systematic classes of counterexamples that are widely compelling. Note that I add 'precise' because we *are* able to give rough and ready characterizations of intentional action. For instance, I suggested above that intentional action is in some sense consciously controlled and in some sense attributable to the agent. This much is unobjectionable. It is when you begin precisifying the sense in which intentional actions are consciously controlled or the sense in which they are attributable to the agent that you begin to fall prey to counterexample.[29]

27. One might suggest that all disagreement over the extent to which non-human animals intentionally act boils down to disagreement over whether animals actually exhibit these features of intentional control over their behaviors. In my view, this puts at least some deniers of non-human animal intentional agency in a worse light than they deserve: there is just far too much empirical work supporting that capacities for consciousness, skill, thought, knowledge, flexibility and goal-direction in behavior extend far beyond the human species. The explanation that I offer puts the deniers in a more charitable light.

28. Against knowledge conditions, see Shepherd and Carter 2023; on intentional action without intention, see Bratman 1987, p. 126; pp. 137-138; on cases that challenge the identification of acting for a reason and acting intentionally, see Alvarez 2009; against the necessity of causal guidance, see Kelley 2024; against the necessity of belief-desire pairs, see Hursthouse 1991; against proposed counterfactual conditions, see Frankfurt 1969. Perhaps a necessary condition on acting intentionally is that one's action is event-caused. But this condition is hardly precise. One should like to know how and by what exactly the action must be event-caused, and identifying this precise sort of causal connection has proved extremely difficult.

29. This has led some to claim that intentional action should be taken as prime

Another way to state this explanandum is that INTENTIONAL ACTION is radically pluralistic. There are just many ways to perform an intentional action, in the sense that intentional actions exhibit many different causal histories and many different psychological profiles. To use an example from O'Brien (2017), one might intentionally drink a beer out of love for beer, out of habit, from addiction, skillfully, clumsily, on the basis of a longstanding plan to drink beer on this day, overwhelmed by a craving for sugar, just because one feels like it, out of joy or embarrassment, in order to calm social nerves, or out of a sense of duty (pp. 275-276). This radical pluralism in the manifestations of intentional action can be explained by the Prioritizing Thesis: in order to be both an efficient and accurate proxy for evaluative priority, INTENTIONAL ACTION needs to be roughly characterized by some easily identifiable form of control that corresponds well with being of evaluative priority, but there also needs to be room for INTENTIONAL ACTION to weigh accuracy over efficiency. This will lead to systematic intuitive counterexamples to any proposed necessary condition on INTENTIONAL ACTION. Indeed, there will be cases where a doing is clearly of evaluative priority despite failing to meet the proposed necessary condition. When many other usual features of intentional control are present, accuracy will be privileged over efficiency and INTENTIONAL ACTION will intuitively apply.

6. Conclusion

Concepts are psychological structures through which we navigate the world. These structures *do things* for us—they have functions relative to the larger psychological structures within which they're embedded. This is not a new idea. Many throughout the history of philosophy have started their philosophical inquiry from a pragmatic vantage point. But this idea has not been sufficiently appreciated within the philosophy of action. Here I contributed to an effort to make amends. I identified a function of the psychological structure that is the concept of intentional action: INTENTIONAL ACTION functions to prioritize doings for evaluative concern, thereby solving a basic problem we face in being the sorts of normative creatures that we are. Our next question should be: what other needs do we have as the sorts of creatures we are, and how do our concepts of agency function to support those needs? The hope is to eventually come to a more or less complete picture of the way that our lives, in all of their humanity, are propped up by our distinctive ideas of agency.[30]

or primitive (Levy 2013; O'Brien 2017). While I won't explore this further here, the Prioritizing Thesis may offer a new argument for some form of anti-reductionism.

30. For helpful feedback, thanks to Zach Barnett, Nathaniel Baron-Schmitt, Olle Blomberg, Michael Bratman, Greg Brown, Ben Conroy, Bridget Copely, David Finkelstein, Wesley Holliday, Tim Kearl, Josh Knobe, Gus Law, Jed Lewinsohn, Taylor Madigan, Tristram McPherson, Daniel Muñoz, Sven Neth, Antonia Peacocke, Kate Petroff, Juan Piñeros Glasscock, David Plunkett, Kieran Setiya, Ginger Schultheis, Sharon Street, Juanhe (TJ) Tan, Barbara Vetter, and a number of anonymous referees. I benefited greatly from presenting versions of this work to audiences at Ohio State University, the Stanford Graduate Student Workshop, UNC-Chapel Hill's PPE Speaker Series, Agency and Intentions in Language 3, University of Chicago's Practical Philosophy Workshop, and Normativity and Reasoning X. I am also indebted to Matthieu Queloz's excellent book *The Practical Origins of Ideas* which spurred many important changes to the essay.

References

Alvarez, Maria (2009). Acting intentionally and acting for a reason. *Inquiry: An Interdisciplinary Journal of Philosophy*, 52(3):293–305. doi: 10.1080/00201740902917168.

Anscombe, G. E. M. (1957). *Intention*. Harvard University Press.

Anscombe, G. E. M. (1958). Modern moral philosophy. *Philosophy*, 33(124):1–19. doi:10.1017/S0031819100037943.

Baron-Schmitt, Nathaniel (2024). Thing causation. *Noûs*, 58(4):1050–1072. doi:10.1111/nous.12494.

Bratman, Michael (1987). *Intention, Plans, and Practical Reason*. Harvard University Press.

Cappelen, Herman (2018). *Fixing Language: An Essay on Conceptual*

Engineering. Oxford University Press.

Clark, Cory J., Christopher W. Bauman, Shanmukh V. Kamble, and Eric D. Knowles (2017). Intentional sin and accidental virtue? Cultural differences in moral systems influence perceived intentionality. *Social Psychological and Personality Science*, 8(1):74–82. doi:10.1177/1948550616663802.

Craig, Edward (1990). *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Clarendon Press.

Cummins, Robert E. (1975). Functional analysis. *Journal of Philosophy*, 72(20):741–765. doi:10.2307/2024640.

Davidson, Donald (1963). Actions, reasons, and causes. *Journal of Philosophy*, 60(23):685–700. doi:10.2307/2023177.

Davidson, Donald (1980). Agency. In *Essays on Actions and Events*, 43–62. Clarendon Press.

Dogramaci, Sinan (2012). Reverse engineering epistemic evaluations. *Philosophy and Phenomenological Research*, 84(3):513–530. doi:10.1111/j.1933-1592.2011.00566.x.

Frankfurt, Harry G. (1969). Alternate possibilities and moral responsibility. *The Journal of Philosophy*, 66(23):829–839. doi:10.2307/2023833.

Frankfurt, Harry G. (1971). Freedom of the will and the concept of a person. *The Journal of Philosophy*, 68(1):5–20. doi:10.2307/2024717.

Hannon, Michael (2015). The universal core of knowledge. *Synthese*, 192(3):769–786. doi:10.1007/s11229-014-0587-y.

Hannon, Michael (2019). *What's the Point of Knowledge? A Function-First Epistemology*. Oxford University Press.

Hursthouse, Rosalind (1991). Arational actions. *The Journal of Philosophy*, 88(2):57–68. doi:10.2307/2027147.

Hyman, John (2013). Voluntariness and choice. *The Philosophical Quarterly*, 63(253):683–708. doi:10.1111/1467-9213.12074.

Hyman, John (2015). *Action, Knowledge, & Will*. Oxford University Press.

Kelley, Mikayla (2024). How to perform a nonbasic action. *Noûs*, 58(1):175–195. doi:10.1111/nous.12440.

Kelley, Mikayla (2025). Separating action and knowledge. *The Philosophical Quarterly*, 75(3):956–977. doi:10.1093/pq/pqae050.

Knobe, Joshua (2003). Intentional action and side effects in ordinary language. *Analysis*, 63(279):190–194. doi:10.1093/analys/63.3.190.

Knobe, Joshua (2008). The concept of intentional action: A case study in the uses of folk psychology. In Joshua Knobe and Shaun Nichols (Eds.), *Experimental Philosophy*, 129–148. Oxford University Press.

Knobe, Joshua (2010). Person as scientist, person as moralist. *Behavioral and Brain Sciences*, 33(4):315–329. doi:10.1017/S0140525X10000907.

Lawlor, Krista (2023). A genealogy of reasonableness. *Mind*, 132(525):113–135. doi:10.1093/mind/fzac036.

Levy, Yair (2013). Intentional action first. *Australasian Journal of Philosophy*, 91(4):705–718. doi:10.1080/00048402.2012.762028.

Nadelhoffer, Thomas (2004). The butler problem revisited. *Analysis*, 64(3):277–284. doi:10.1093/analys/64.3.277.

Nadelhoffer, Thomas (2005). Skill, luck, control, and intentional action. *Philosophical Psychology*, 18(3):341–352. doi:10.1080/09515080500177309.

O'Brien, Lucy (2017). Actions as prime. *Royal Institute of Philosophy Supplement*, 80:265–285. doi:10.1017/S135824611700011X.

O'Sullivan, Angela and Lilith Mace (2024). Reverse-engineering risk. *Erkenntnis*, 90(5):2021–2046. doi:10.1007/s10670-024-00788-6.

Pettit, Dean and Joshua Knobe (2009). The pervasive impact of moral judgment. *Mind and Language*, 24(5):586–604. doi:10.1111/j.1468-0017.2009.01375.x.

Queloz, Matthieu (2021). *The Practical Origins of Ideas: Genealogy as Conceptual Reverse-Engineering*. Oxford University Press.

Queloz, Matthieu (2022a). The essential superficiality of the voluntary and the moralization of psychology. *Philosophical Studies*, 179(5):1591–1620. doi:10.1007/s11098-021-01720-2.

Queloz, Matthieu (2022b). Function-based conceptual engineering and the authority problem. *Mind*, 131(524):1247–1278. doi:10.1093/mind/fzac028.

Reynolds, Steven (2017). *Knowledge as Acceptable Testimony*. Cambridge University Press.

Robbins, Erin, Jason Shepard, and Philippe Rochat (2017). Variations in

judgments of intentional action and moral evaluation across eight cultures. *Cognition*, 164:22–30. doi:10.1016/j.cognition.2017.02.012.

Shepherd, Joshua and J. Adam Carter (2023). Knowledge, practical knowledge, and intentional action. *Ergo: An Open Access Journal of Philosophy*, 9:556–583. doi:10.3998/ergo.2277.

Stoecker, Ralf (2009). Why animals can't act. *Inquiry: An Interdisciplinary Journal of Philosophy*, 52(3):255–271. doi:10.1080/00201740902917135.

Thomasson, Amie L. (2024). How should we think about linguistic function? *Inquiry: An Interdisciplinary Journal of Philosophy*, 67(3):840–871. doi:10.1080/0020174X.2022.2074886.

Watzl, Sebastian (2017). *Structuring Mind: The Nature of Attention and How It Shapes Consciousness*. Oxford University Press.

Williams, Bernard (1997). Moral responsibility and political freedom. *The Cambridge Law Journal*, 56(1):96–102. doi:10.1017/S0008197300017761.