

Refitting the Mirrors:

On Structural Analogies in Epistemology and Action Theory

Abstract: Structural analogies connect Williamson's (2000; 2017) epistemology and action theory: for example, action is the direction-of-fit mirror image of knowledge, and knowledge stands to belief as action stands to intention. These structural analogies, for Williamson, are meant to illuminate more generally how 'mirrors' reversing direction of fit should be understood as connecting the spectrum of our cognitive and practically oriented mental states. This paper has two central aims, one negative and the other positive. The negative aim is to highlight some intractable problems with Williamson's preferred analogical picture, which links the cognitive and the practical through the nexus of direction-of-fit mirroring. The positive aim of the paper is to propose a better alternative. In particular, we show that an achievement-theoretic proposal captures what is in common across the range of attitudes that exhibit the kind of structure that knowledge-belief, action-desire/intention do, while at the same time avoiding the problems shown to face Williamson's proposed picture. Moreover, we draw attention to several key theoretical benefits of embracing our proposed achievement-theoretic picture, including some of the key benefits of the knowledge-first programme that Williamson's own analogies were designed to secure.

1 Introduction

Timothy Williamson (2000) begins *Knowledge and its Limits* with a structural analogy between what he takes to be the two key relations between mind and world: knowledge and action.

In action, world is adapted to mind. In knowledge, mind is adapted to world. When world is maladapted to mind, there is a residue of desire. When mind is maladapted to world, there is a residue of belief. Desire aspires to action; belief aspires to knowledge. The point of desire is action; the point of belief is knowledge (2000, 1).

As Williamson characterizes this analogy in a bit more detail, action is the direction-of-fit mirror image (i.e., DFMI) of knowledge in a mirror that reverses direction of fit. On this picture, knowledge and belief go on the mind-to-world

side of the proposed analogy (i.e., the left side), and action and desire on the world-to-mind side (i.e., the right side), and knowledge stands to belief as action stands to desire.

Fitting mind-to-world		Fitting world-to-mind
knowledge	mirror	action
belief	mirror	desire

Table 1 : The *KAIL* analogy (2000)

More recently (e.g., 2017), Williamson’s knowledge-action analogy has been tweaked and refined, and these tweaks principally concern the place of desire in the analogy—in particular, how desire should be understood as interfacing with belief, intention and action, respectively.¹ We will argue here that Williamson’s insightful analysis concerning the general structural relations between items *within* these columns—what we will call an aptness, or achievement-theoretic structure—can and should be divorced from the more problematic account of his proposed relationship *between* these columns in terms of direction of fit.

Here is the plan. In §2, we briefly outline the key contours of the kind of direction-of-fit theory that undergirds Williamson’s analogy. In §3, we discuss Williamson’s reasoning for updating and refining the original analogy from *Knowledge and its Limits*, which largely stems from the contrasts between a knowledge-first approach and traditional belief-desire psychology. We then raise and develop two objections to his preferred analogical picture. In §4 we argue that the primary focus on knowledge and action as the two principal attitudes that instantiate this symmetry closes-off many other cognitively important attitudes, such as objectual understanding, from satisfactorily entering into the analogy (no matter how we tweak it); and, by Williamson’s own lights, an expansion of the analogy should offer a place for such attitudes. Moreover, these attitudes seem to exhibit parallel “column” structure to knowledge and action—we explain how many attitudes exhibit the kind of achievement-theoretic structure that parallels the knowledge-belief, action-intention structure, and suggest that a central driver of the structural analogy is not an “axis” of symmetry, but rather structure related to achievement.

In §5, we provide a more general argument against the view that knowledge and action can be illuminated as having reverse directions of fit. We argue that *many* attitudes and purposeful activities can be understood in relation to achievement, and that they all have both mind-to-world and world-to-mind normative constraints, a point Williamson and others have problematically overlooked (but, perhaps surprisingly, Anscombe did not). It is an issue not just of theoretical but also of ethical importance that we do not conceive of action as a matter of bending the world to one’s mind. The normativity governing action is more complex than that, as we will explain.

In §6, we conclude, showing that our achievement-theoretic proposal captures what is in common across the range of attitudes that exhibit the kind of struc-

¹A further update to his analogies concerns their placement in his wider view of what is involved in a cognitive-practical system’s functioning well—viz., when one acts on what one knows, and thus where the input to practical reasoning is knowledge (not mere true belief) and the output is action (not mere successful intention)(p. 174).

ture that knowledge-belief, action-desire/intention do, without committing to problematic normative views. We summarize several key theoretical benefits, including some of the insights of the knowledge-first programme that Williamson's own analogies were designed to secure.

2 Mind to world, world to mind

In G.E.M. Anscombe's (1957, 56) famous vignette, a man is going around town with a shopping list in his hand, while a detective follows him about, hoping to find out what he buys. In this situation, Anscombe maintains that the shopper tries to fit his purchases to his own shopping list. But the detective is in a sense doing just the opposite: she is trying to fit her list to the man's purchases. If the list and purchases do not agree, the error is to be located in different places for the shopper and detective: for the shopper the error would be in the purchases, for the detective in the list. Anscombe is often read as endorsing the idea that 'success' in the former case involves the world-fitting the mind—viz., when the shopper's desire is satisfied; in the latter, the mind fitting the world—viz., when what the detective believes about the shopper's purchases is actually what's so (we shall further discuss this interpretation in §5).

Anscombe's case has been taken to suggest a more general point: that there are two kinds of attitudes: those that aim at representing things as they are, and those that aim at getting things done. These two aims are realized, as the shopping list vignette indicates, in opposite directions, and correspond more generally to the cognitive and the practical aspects of intelligent life, respectively.

Searle (1983), Humberstone (1992), Smith (1994), and Velleman (2000), among others, have been impressed with Anscombe's core idea, and have attempted to develop it further in different ways.² Common to their proposals is a core commitment to the following ideas which direction-of-fit (DOF) theorists accept: *exclusivity*, *realization* and *normativity*.

Exclusivity: There are two directions of fit that characterize intentional mental states: mind-to-world and world-to mind.

Realization: Realization (i.e., success), for a cognitive (or theoretical) intentional mental state, involves fitting mind-to-world; realization for a practical mental state (e.g., desire, intention, etc.) involves fitting world-to-mind.

Normativity: Intentional mental states with a mind-to-world direction of fit are such that they ought to fit the world; intentional mental states with a world-to-mind direction of fit are such that the world ought to fit them.

²See, however, Frost (2014) for a recent critique of DOF theory, and for an argument to the effect that the conclusions typically drawn from Anscombe's vignette are the wrong ones.

As Kim Frost (2014) succinctly puts it, these core DOF theses lead DOF theorists to embrace a connected idea concerning symmetry, and which features prominently in Williamson’s own analogy. As he puts it, the core components of DOF theory:

[...] present an image of symmetry at work in the thetic and telic DOFs: whatever the thetic relation of mind to world is, the telic relation is somehow the *mirror image* of the same. (2014, 430)

Call this fourth idea ‘mirroring’:

Mirroring: For a given thetic (mind-to-world) intentional mental state, there is a telic mirror—viz., a symmetrical world-to-mind state; for a given telic (world-to-mind) intentional mental state, there is a thetic mirror (viz., a symmetrical mind-to-world state).

The symmetry that features in mirroring will presumably be a matter of (a given state’s) standing in connection with realization. For example, on the cognitive side, some intentional attitudes aim haplessly at realization—viz., a guess; a candidate symmetrical world-to-mind DFMI state will stand in an analogously hapless relationship to realization (e.g. a hail Mary).³

The Williamsonian claim that knowledge is the DFMI of action, accordingly, may be understood as a special case of mirroring, one where not only the corresponding paradigm good states (knowledge, action), mirror each other, but also corresponding attempts—or “aspirations” (Williamson, 2000). In the cases of knowledge and action, each mirror image is maximally realized with respect to its own direction-of-fit.

3 The updated analogy: desire and intention

Even if we grant Williamson that knowledge stands to the cognitive as action to the practical (and correspondingly, that knowledge is the direction-of-fit mirror image of action), it wouldn’t simply follow that *desire* is what analogously mirrors belief by aiming at action as belief aims at knowledge. Williamson (2017) accordingly revisits the matter of how desire originally earned its spot (opposite belief) in the *KAIL* analogy.

3.1 Desire’s place in the *KAIL* analogy

Why is it that desire, rather than something else, was initially taken to stand to action as belief to knowledge? In short, the explanation is as follows: it’s a hallmark of belief-desire psychology that beliefs and desires are propositional

³We set aside for the present purposes the question of whether all or most mirrors of common intentional mental states have single-term mirrors in English, though we submit that it would be surprising if they did. That said, the DOF theorist’s commitment to mirroring is in the first instance metaphysical, rather than semantic; correspondingly she needn’t be read as committed to anything like ‘single-term’ mirroring.

attitudes that have characteristic success conditions with *opposite* directions of fit (e.g., Humberstone 1992)—viz., the thetic and the telic.⁴ In the case of belief, the content is supposed to represent the world, and thus “fit” the mind to the world. With desire, the content is supposed to be actualized—and so the world is “fit” to the mind. In this respect, belief and desire are supposed to be mirror images of one another, in a mirror that reverses direction of fit.⁵

Now, if belief-desire psychology were true—that is, if belief and desire were the core explanatory mental attitudes—this would recommend taking desire as the ‘pre-selected’ direction-of-fit mirror image (DFMI) of belief in a mirror reversing direction of fit and then solving for the blanks *above* belief and desire by asking what success involves in each case. Williamson’s original answer, of course, was ‘knowledge’ and ‘action’—as he puts it “The point of desire is action; the point of belief is knowledge” (Williamson, 2000, 1). However, *even if* knowledge and action are the ‘right answers’, we should—as Williamson (2017) now thinks—consider more critically *how* those answers were arrived at. In broad outline, the strategy Williamson (2000) originally used in reasoning that knowledge is the DFMI of action simply took for granted that desire is the DFMI of belief. But this deserves scrutiny for two reasons: first, it leaves the idea that *belief is to mind-to-world realization as desire is to world-to-mind realization* as explanatorily central. Second, and relatedly, such a move stands in tension with a picture of intelligent life on which knowledge is supposed to be the cognitively central attitude, more explanatorily central than belief.

If we “start with success” i.e., that knowledge is to mind-to-world realization as action is to world-to-mind realization, and thus, that knowledge is the DFMI of action rather than that belief is the DFMI of desire, we abandon any inadvertent tie to belief-desire psychology as explanatorily central, and in doing so, *leave open* whether desire is *what it is* that stands to action as belief to knowledge. On Williamson’s updated thinking, we find a principled “success first” starting point for establishing structural analogies between cognitive and practical attitudes. Such an approach must go beyond the bare fact that in other contexts the knowledge-firster prioritises knowledge over belief and action over desire.

Williamson’s idea is as follows. Knowledge is cognition working well; but what characterises a *cognitive-practical* system working well, one that not only represents the world but aims to manipulate it through good practical reasoning? Whether any given cognitive-practical system is working well will be a matter of the quality of its *inputs* and *outputs*, quality that is for Williamson maximal if and only if one acts (the *output* to practical reasoning) and not merely desires or intends, on what one *knows* (the *input* to practical reasoning) and not merely on what one believes. (Williamson, 2017, §4).

With this starting point—viz., that knowledge and action, *the inputs and outputs of practical reasoning working well*, are our pre-selected DFMI—rather than that belief and desire are the preselected DFMI—we are now solving for *different blanks* (i.e., solving *downward*); and so we *might* now see belief’s DFMI as something *other* than desire. After all, with this new starting point, that belief and desire are DFMI is not simply *given*, as (among others) Humberstone took them to be.

⁴These terms are originally due to Humberstone (1992).

⁵For discussion, see Williamson (2017, 163).

Fitting mind-to-world		Fitting world-to-mind
knowledge	mirror	(intentional) action
belief	mirror	?

Table 2

3.2 Knowledge is to belief as action is to intention (2017)

What stands to action as belief stands to knowledge? On closer inspection, *desire* is not a satisfying answer. For one thing, consider that desire is not more closely connected to action than belief is, even on belief-desire psychology, where the premises of practical reasoning are supplied by beliefs and desires taken *together*.

As Williamson sees it, the obvious place to look is at *conclusion* of practical reasoning, i.e., its *output*, not at the *input* (beliefs/desires) to the premises. But, as he notes, '[...]the conclusion of a piece of practical reasoning should *not be just another desire*; it should be an *intention* to do or be something' (2017, 164).

Accordingly, on the updated view, we get:⁶

Fitting mind-to-world		Fitting world-to-mind
knowledge	mirror	action
belief	mirror	desire intention

Table 3 : The 2017 analogy

3.3 Assimilating desire to belief

Where does desire go, then, now that intention stands to action as belief to knowledge? One natural option would be to try something like the following: since desire is further from action than intention, we should (i) place desire on the right-hand side of the analogy, below intention, and then place opposite desire on the left-hand-side a mind-to-world attitude that's further from knowledge than belief as desire is further than action to intention. One such candidate might be: an 'appearance'⁷. The idea would go as follows:

Fitting mind-to-world		Fitting world-to-mind
knowledge	mirror	action
belief	mirror	intention
appearance?	mirror	desire

Table 4

⁶One wonders, on this updated view, why intention should stand as the mirror to belief. After all, typically knowledge is output as well as input to theoretical reasoning. This suggests that really the contrast between knowledge and belief is that of achievement and attempt, not anything inherent to do with (theoretical or practical) reasoning. We shall return to this below.

⁷Williamson (2017, 164) also canvasses the related possibilities of suspicion or a level of credence short or outright belief.

But Williamson forecloses any such possibility, and he does so regardless of whether we replace appearance with something slightly further or closer to knowledge on the left-hand-side. His concern is this: such a strategy would leave desire’s DFMI further than desire itself from action, since (i) it is further than outright belief from action, and (ii) outright belief is at least as far as desire from action. Therefore, desire can’t simply be added to the right-hand-side below intention.⁸

Instead, Williamson’s idea is to (perhaps, counterintuitively) shift desire from the right- to the left-hand-side of the analogy, via the strategy of assimilating desire to belief.⁹ On such an approach, to desire something is to believe that it is good. More carefully, the idea is that to desire that p requires fitting world to mind with respect to the proposition that p , but fitting mind-to-world with respect to the proposition *that it is good that p* (Williamson 2017, 165).

This move, is, of course, a controversial one. Though for the present purposes, we are happy to grant that it is viable, as our criticisms will lie elsewhere.¹⁰ The expanded 2017 picture Williamson offers fills in the remaining details as follows:¹¹

Fitting mind-to-world		Fitting world-to-mind
knowledge	mirror	action
belief	mirror	intention
truth	mirror	success
falsity	mirror	failure
input to practical reasoning	mirror	output from practical reasoning

Table 5 : The 2017 analogy (completed)

4 The Exclusion Problem

In this section, we raise the first of two objections to Williamson’s analogies. Both, we will suggest, reveal that what drives the structural analogy between knowledge and action is not a “mirror”-reflecting different directions of fit, but rather that knowledge and action both are achievements, in the sense of

⁸See Williamson (2017, 166) for further discussion. As the reader will see below, our ultimate view does not depend on accepting this line of reasoning, because it is achievement-theoretic structure, rather than practical reasoning, that motivates the analogies represented in table rows.

⁹For a notable line of resistance to this suggestion, see Lewis (1988). Lewis’s objection to assimilating desire to belief is that such a view cannot be squared with a plausible Bayesian theory of belief conditionalization. For a criticism of Lewis’s Bayesian challenge to desire-belief assimilation, see Weintraub (2007).

¹⁰Apart from concerns about over-intellectualization, one wonders why it is legitimate to characterize desire in this way and not intention, perhaps, say, as a belief about what to do.

¹¹The idea that direction of fit is to be understood in terms its role in reasoning echoes Avery Archer (2015)’s claim that direction of fit is to be understood in terms of inferential role—states with indicative content can play roles in reasoning that states with practical content cannot—and this is to be distinguished from the *revision conditions* for the attitude (p. 177). This enables him to apply direction-of-fit analysis to a wider range of attitudes, including assuming and fantasizing, which have a mind-to-world direction of fit because of their indicative contents but have different revision conditions than belief. We are obviously sympathetic to this expansion of applicability of analysis, as we shall now turn to.

being paradigmatic *good cases* of doing what one aims to do. They and related attitudes, and success and failure conditions, therefore exhibit structurally analogous achievement-theoretic structure.

Once we have abandoned traditional belief-desire psychology as driving the symmetry regarding knowledge and action, there are many other attitudes one might consider.¹² How, for example, does *understanding* fit in the above picture? Objectual understanding takes not a proposition, expressible by a that-clause, as an object, but rather an object that just is, or can be treated as, subject matter or body of information (see, e.g., Kvanvig (2003)). Such attributions take the form “A understands X” as opposed to “S understands that x” (propositional understanding) or “S understandings why-x” (e.g., understanding-why). The following are some representative examples where objectual understanding is attributed:

1. "Akira understands organic chemistry; I don't."
2. "We were watching Australian rules football for nearly an hour, and we still didn't understand it."
3. "Jon understands the Comanche's dominance of the southern U.S. plains during the 19th century." (cf., Kvanvig 2003, 197-98).

According to an increasingly popular view in epistemic axiology, this kind of understanding is particularly valuable to possess (perhaps more so than knowledge), and moreover, it is a candidate epistemic goal of inquiry.¹³ Given the direction of fit – with some outlying exceptions¹⁴, most epistemologists think objectual understanding (hereafter, understanding) in some sense answers to the facts – we should expect understanding to fall on the left side of the analogy.

Two knowledge-first approaches to understanding defended by Christoph Kelp (2015; 2016) and Paulina Sliwa (2015) place it *de facto* on the left side of the analogy by assimilating understanding to knowledge. This offers an elegant way to envisage how understanding should enter the analogy.

Fitting mind-to-world		Fitting world-to-mind
knowledge/ understanding	mirror	action
belief	mirror	intention

Table 6

Kelp and Sliwa agree that understanding a subject matter involves possessing (perhaps, a significant amount of) propositional knowledge about the subject

¹²While we focus here on objectual understanding, one could make analogous points with other attitudes. See, e.g., §4.2 for some such examples.

¹³While we are sympathetic to such views in epistemic axiology (e.g., Kvanvig (2003), Pritchard (2008; 2010b)), we will not here presuppose or argue for any such claims. It suffices for our starting point to register that aim at this kind of understanding, we are attempting to fit mind to world in a way that succeeds just when what we get is understanding.

¹⁴For an influential recent challenge to factivist construals of objectual understanding, see Elgin (2017).

matter. If understanding is assimilated to knowledge, understanding's DFMI is knowledge's DFMI, which is action; and since knowledge stands to belief as action stands to intention, so understanding stands to belief as action to intention.¹⁵

But these analogies run in to several hitches. Understanding's relationships to both belief (on the mind-to-world side) and action (on the world-to-mind side) are, in important respects, disanalogous with those of knowledge.

Firstly, consider the problem that arises once understanding's DFMI is taken to be action. Practical reasoning is, on Williamson's wider picture of intelligent life, functioning well just when one acts on what one knows.¹⁶ But it's not obvious that acting on objectual understanding is necessary for practical reasoning's functioning well. For example, practical reasoning is surely functioning well when one attempts to understand something one does *not* understand, by acting on what one knows until one does.

This suggests that the structural analogies related to direction of fit come apart from those motivated by reflection on practical reasoning. If objectual understanding has a mind-to-world DFMI but plays a different role in practical reasoning, then either the direction-of-fit mirror analogy should be revised or we should abandon the view that the analogies are motivated by the roles of the attitudes in theoretical and practical reasoning, respectively.

The second, and perhaps more serious set of concerns, has to do with understanding's connection on the left-hand-side with belief.

Fitting mind-to-world		Fitting world-to-mind
knowledge/ understanding	mirror	action
belief	mirror	intention

Table 7

When knowledge is maladapted to world, there is residual belief. When understanding is maladapted to world, there is residual X, where X should be a kind of botched attempt at understanding. Is such a botched attempt simply *belief* (or, more specifically, a kind of belief)? Plausibly not. For one thing, it's hard to see how botched understanding could be belief given that botched understanding can't be a (mere) propositional attitude relation.

This point is secured through two observations; first, understanding is an objectual-attitude relation, an attempt to fit mind-to-world that takes a body of information as an object. Second, a botched attempt at understanding will be some form of a non-factive objectual attitude relation, one that understanding entails, but which does not entail understanding. It will be the non-factive objectual attitude that is residual when (in attempting to understand by taking up such an objectual attitude) mind is not fit to world.

¹⁵Of course, even if understanding is assimilated to knowledge, understanding (which asymmetrically entails knowledge) is a *type* or species of knowledge (just like, on the 2017 analogy, desire is a type or species of belief); accordingly, if understanding's DFMI is knowledge's DFMI, understanding (qua a type of knowledge) would stand to a (type) of belief, as action stands to intention.

¹⁶For an even more recent discussion of this picture, specifically in the context of epistemic dilemmas, see Williamson (2021, §3).

Relatedly, it is worth considering that you can maladapt the mind to world by believing, even without making *any attempt* at objectual understanding—viz., without even putting yourself in the market for it when things go right. Imagine here the algebra student who comes to believe various memorized propositions about the axioms of geometry (with the aim of recalling just enough of them to pass an exam). Such a person, though a believer, has not attempted to fit mind to world in the way one does when one understands—viz., by fitting mind to world through an objectual attitude. Such an individual has not *taken up* an objectual attitude, non-factive or otherwise. Belief is accordingly not botched understanding (even if people who believe things about a subject may fail to understand that subject).

The point here is not whether objectual understanding can be reductively explained in terms of propositional knowledge, but that *even if it can be*, we are still owed an account of what is residual when understanding is botched. This cannot just be belief. Understanding and knowledge cannot both be *THE* mind-to-world attitudes because they—and correspondingly their botched attempts—have different functional and normative roles.

A final problem concerning the analogy with belief is that the *normative* relationship between belief and understanding is very different from the normative relationship between knowledge and belief. To appreciate this point, it will be helpful to draw from Williamson's recent paper 'Justification, Excuses and Sceptical Scenarios' (2018). A key claim of that discussion is that primary truth-related epistemic norms have compliance conditions (like promising) that we may fail to satisfy even if we do our best and do what the very best would do in our shoes.

We may, however, blamelessly violate a primary truth-related epistemic norm (i.e., believe only what you know) provided we satisfy derivative norms: e.g., do what the person who usually conforms to the primary norm would do—viz., in the case of belief, proportion her beliefs to the evidence, etc.¹⁷

On the assumption that truth-related norms constrain other attitudes, such as objectual understanding, we should expect there to be blameless violations of these norms provided one satisfies derivative norms—viz., do what one who what one who understands would do. But here a disanalogy emerges. It's not clear that doing what constitutes a blameless violation of knowledge norm for belief would also be a blameless violation of the derivative norm: do what one who understands would do (as the latter will at least involve taking up an objectual attitude, something not entailed by simply forming a belief and proportioning it to the evidence). What this all suggests is that excuses for violating primary (truth-related) belief norms will not always give us coverage for violations of truth-related norms governing our attempts at understanding. In this respect, the normative relationship between knowledge and belief differs importantly from that between belief and understanding.

These considerations are problematic for Williamson's analogies, analogies that purport to connect the cognitive and the practical. Within these analogies, there

¹⁷See also Simion et al. (2016) for a development of this picture, and which distinguishes further between evaluative and prescriptive norms.

seems to be no room for understanding, a cognitively valuable way of fitting mind to world.

4.1 The Problem with Mirrors

The case of objectual understanding suggests that things are not as simple as the *mirror analogy* suggests. There are multiple attitudes that can be construed as fitting mind-to-world. Regardless of whether it might be possible to analyze understanding in terms of knowledge, the functional and normative roles of understanding differ substantially from those of knowledge, and thus bear different relations to action, belief, truth, falsity, and practical reasoning.

We can begin to see the root of the difficulty by noting that implicit in the foregoing knowledge-first attempt to assimilate objectual understanding to knowledge is an acknowledgement that objectual understanding and knowledge seem to have the same place in the hierarchy (namely, the top). That is, objectual understanding seems to be an *achievement*.¹⁸ Achievements, in the sense of Miracchi (2014, 2017) are attitudes or activities that do not merely entail that the world is a certain way (e.g., are not merely factive), but also that the subject is related in a certain kind of characteristic way to that entailed condition (ruling out, e.g. Gettiered beliefs and deviant causal chains). This characteristic relation is one that reflects well on the subject as bearer of that attitude or agent of that activity, thus the term “achievement”.¹⁹ To say that an attitude, activity, or action is an achievement is not to say that it is difficult or the result of effort.²⁰ We are using the term in a technical sense.

Objectual understanding plausibly has these characteristics. It involves not only substantial propositional knowledge, and so multiple entailments about how the world is, but also the kind of coherent grasp of such knowledge and competence to deploy that knowledge in the right circumstances and connect it up with other things one learns in real time. (One might know all the rules of Australian soccer, and still not understand it.) It thus also plausibly involves a characteristic kind of relation to those entailments (see, e.g., Grimm (2012) and Kvanvig (2003)).²¹

Beginning with the relationship between knowledge and belief, we note that corresponding to each achievement is a corresponding attempt: belief in the case of knowledge, intentional attempt into the case of action. For our purposes

¹⁸See Miracchi (2017) for discussion. See also Carter et al. (2015) for discussion of achievements in epistemology in connection with reliability and risk.

¹⁹Some may want to fill out this picture in a way that commits to virtue epistemology—see, e.g., Greco (2010)—but we’re trying to stay as neutral on this here as possible.

²⁰For a different philosophical development of *achievements* on which being difficult is necessary for something’s counting as an achievement, see Bradford (2015).

²¹Whether either the entailments or the characteristic kind of relation are reducible to knowledge—that is an interesting question, which we do not need to take a stand on here. Even if this were possible, understanding would still be an achievement with a different functional and so normative profile from knowledge—that.

here, one need not read too much into “attempt.”²² Williamson (2000) uses the term “aspires” (see quote above), and one might use related terms such as “aims”. The important thing for our purposes—which is a subject of wide agreement between knowledge-first and belief-first epistemologists alike—is that knowledge is the paradigm good case of what happens when belief goes well (the existence of Gettier cases cries out for explanation).²³ Similarly with deviant causal chains in the case of action (Miracchi, 2014; Sosa, 2015). Knowledge and action stand as cases, not only of the relevant content being satisfied, but its being satisfied “in the right way.”

We can now include in our table objectual understanding, and explicit characterization of the rows in relation to achievement.

Fitting mind-to-world			Fitting world-to-mind	
achievement	understanding	knowledge	<i>mirror?</i>	action
attempt	?	belief	<i>mirror?</i>	intention

Table 8 : Adding Columns

“Intention” in our table thus deserves to be clarified. Sometimes when people use the term they talk about a state that may or may not cause action when the time is right. In this sense, an intention is something like a decision to *A*, rather than (part of) an attempted *A*-ing. In order to maintain parity with belief, as attempted knowledge, or the state indiscriminable from knowledge, we should clarify that “intention” in this table should refer to *intention-in-action* or *intention-with-which* in Anscombe’s sense, not an independent decision.²⁴

We can then ask what attempted understanding is. Let us stipulatively use the term “conception” for attempted understanding. The subject, through forming her conception, is making an attempt to be related in the competent way characteristic of understanding to her knowledge of the subject matter that such understanding entails.²⁵ And the kind of relatedness characteristic of understanding involves (perhaps, among other things) grasping some significant subset of the coherence and explanatory relations that hold between the knowledge she has of the relevant subject matter.²⁶ Of course, not all such attempts one makes by forming a conception are achievements of understanding. In some cases, this attempt is botched simply because one lacks the knowledge

²²See, however, Sosa (2021) for a new version of virtue epistemology which takes the kind of normativity that features in knowledge related evaluations in epistemology (that is, in what Sosa calls ‘gnoseology’) to be the normativity of attempts *as* attempts. This newer view marks a transition between Sosa’s thinking of epistemic normativity as a species of performance normativity to his updated thinking on which epistemic normativity is a species of telic normativity. See here especially (2021, Ch. 2).

²³See Hetherington (1999, 2001) for rare attempts to assimilate Gettier cases and knowledge as being of the same epistemic kind.

²⁴This interpretation does not challenge the view that intention is the output of practical reasoning, although we do not have to accept such a view.

²⁵In the case of a conception, this attempt is made by means of taking up an objectual rather than propositional attitude relation.

²⁶For discussion on this point, see, Kvanvig (2003) and Grimm (2012), and Gordon (2019).

understanding entails.²⁷

But, more importantly, such an attempt may fall short even when the relevant knowledge of the subject matter entailed by understanding is present. Consider that a skilled amateur chess player's conception of a complex chess position might include ample knowledge about the key facts of the position—about what pieces are threatening what other pieces directly and even indirectly—while at the same time the player fails to appreciate how all of this sophisticated knowledge of the position is connected, and related to other (chess-relevant) things she knows, e.g., about positions of that type, about how such positions usually transform into endgames, etc. As a result, we can imagine such a player's conception of the position will be less effective for her in assessing what move to make next than the grandmaster's, who understands the position and thus easily sees exactly what move should be made next. In this case, then, our novice (unlike, say, a casual spectator who makes no serious attempt at understanding the position) has a conception of the position—and even has some knowledge of exactly the sort that understanding the position entails—but her conception nonetheless falls short. She thus lacks the kind of achievement to which her conception was an attempt. There is surely more to say here, but we take the idea that objectual understanding is an achievement, with a corresponding characteristic attempt, to be plausible.²⁸

Are there are more achievement attitudes that we should consider? [Miracchi \(2014, 2017\)](#), and [Sosa \(2015\)](#) argue that perception is an achievement, analogous to knowledge and action, with perceptual experience as the corresponding attempt and veridical hallucinations as analogous to Gettier cases. Perception, as well as understanding, not only has corresponding attempts, it also has what we will call corresponding *entailment conditions*, conditions whose obtaining is necessary but not sufficient for the attempt to be an achievement. We use the term “entailment” in order to leave open whether all achievements have propositional contents. Regardless of whether, e.g., perception has propositional content, it is widely agreed that perception has *accuracy conditions*, and so entails certain conditions that can be specified propositionally.

In addition to objectual understanding and perception, there are other attitudes that we might think of as achievements. For example, the kind of objectual knowledge one has in virtue of the capacity for singular thought plausibly entails not only the existence of the object and its having some of the features attributed to it (entailment) but also facts about the subject's justification for beliefs about its features (see Imogen [Dickie \(2015\)](#) for defense of this claim).

This is by no means a comprehensive list, but we already see that several columns must be added to our table. Below is a table that replicates the first

²⁷For example, consider Aristotle's physics, which was deeply flawed—not because of any lack of competence Aristotle had in assimilating and piecing together the information he had, but rather, because the scientific inputs he had available were bad (or incomplete) inputs. For discussion on this point, see, e.g., [Montmarquet \(1993\)](#). We want to clarify that we are not suggesting Aristotle did not advance understanding. Rather, we focus only on those narrow physical subject matters that Aristotle failed to understand on account of bad inputs, as opposed to on account of any defect in his capacity to piece together the inputs he had.

²⁸For a defence of the view that objectual understanding is a valuable achievement, see, along with [Kvanvig \(2003, Ch. 8\)](#), also [Carter and Gordon \(2014\)](#) and [Pritchard \(2010a\)](#).

three rows of Williamson’s 2017 table (Table 5), but includes columns for perception, understanding, and objectual knowledge, as well as a left-hand column generalizing features. In each case, we can include an attitude that corresponds the attempt (either to the achievement or to the purported achievement), and the entailment condition (with “truth +” representing the additional coherence relations that are required for understanding.

				Fitting mind-to-world knowledge		Fitting world-to-mind action
achievement	perception	understanding	objectual knowledge		<i>mirror?</i>	
attempt	perceptual experience	conception	attempted thoughts	belief	<i>mirror?</i>	intention
entailment	accuracy	truth +	existence	truth	<i>mirror?</i>	success

Table 9 : Adding Columns

There are now several questions to ask. First, can more achievements be specified on the world-to-mind side as well? Indeed, we can make similar points about other manifestations of agency than intentional action. Some forms of agency, such as language production, are not always well characterized as intentional actions, understood as potential outputs of practical reasoning. Language production is often automatic, effortless, and although we sometimes have to choose and plan particular phrasing or careful articulation, this is plausibly best understood as intentional modulation of the process of speech production. Nevertheless, language production is an achievement in the sense at issue. It has corresponding attempts which may or may not be botched, and it requires not only the production of certain linguistic items such as sentences (entailment), but it requires that those sentences be related to one’s mental life in the right kind of way (as expressions of beliefs, orders, requests, etc).

Or, consider *improvisation*, which is an action that does not aim at the production of particular notes or note patterns, but rather more abstract musical features, such as rhythmic flow. Improvisation inherently involves unplanned aspects, and so is not properly viewed as the output of practical reasoning. (*That* one improvises can be an output of practical reasoning, but not *what* one improvises.) Improvisation is nevertheless an achievement in our sense—viz., it has attempts that may or may not be botched, it does entail that certain conditions obtain (one can fail), and it does entail that the subject bear a certain relation to those conditions that reflects upon the subject *qua* bearer of that relation.

Sometimes, the achievement is not just intentional action, but rather expert action, for example, of the kind exhibited by professional athletes, dancers, or musicians. This kind of expert action involves not only the kind of satisfaction of intention involved in intentional action, but also grace, coordination, excellence of technique, mastery of alternatives, etc. Our already long table is now getting very long:

				Fitting mind-to- world		Fitting world- to- mind			
achievement	perception	understand- ing	objectual know- ledge	knowledge	<i>mirror?</i>	action	speech produc- tion	improvis- ation	skilled ac- tion
attempt	perceptual experi- ence	conception	attempted thoughts	belief	<i>mirror?</i>	intention	attempted speech	attempted improv	attempted skilled action
entailment	accuracy	truth +	existence	truth	<i>mirror?</i>	success	attitude expres- sion	rhythmic flow, etc.	success +

Table 10 : Adding Columns

A few things should be noted about this table. First, we lack natural language terms for many of these attempts and entailments. What we have filled in here is a rough approximation. Rather than being a difficulty with the proposal, we hope to have said enough to make it highly plausible that many more kinds of attitudes are achievements in our sense, and thus have corresponding attempts, entailments, and relational structure. That many more kinds of attitudes and activities are achievements in this sense, and that the corresponding aspects of their relational structure are less conceptually clear or easy to articulate, generalizes the Williamsonian idea regarding knowledge and action: we should treat achievements in all cases as conceptually and metaphysically prior to their corresponding attempts.²⁹

Second, a healthier diet of examples raises even more concerns about the supposed centrality of practical reasoning into the analogies between knowledge and action. By increasing our examples on the action side, we make reasoning less central to the discussion, and resist reifying forms of reasoning into kinds: theoretical and practical. This allows for wider flexibility in considering the role of reasoning for these various achievements: for example, dance often involves reasoning (see [Montero \(2018\)](#)), but the point of such reasoning may not be the formation of an intention.

We are approaching the heart of the problem that objectual understanding and these other achievements raise with the supposed mirroring of knowledge and action, namely, direction of fit as the axis of symmetry. If there were compelling reason to think that each achievement on the mind-to-world side corresponded to an achievement on the world-to-mind side (and vice versa) in a deep way, then maybe we could preserve the symmetry. But what, for example, corresponds to improvisation?

Moreover, the theoretical motivation for this symmetry claim is now highly suspect. Once we have given up on belief-desire psychology, we should give up on the idea that there need to be two attitudes, practical and theoretical³⁰, that serve a deep explanatory purpose. Attempts to do this with the relation between knowledge and action in practical reasoning fail to illuminate the other attitudes that seem to deserve similar analysis. Instead, we now put forward the hypothesis that the structural relationships are explained, not by symmetry of the columns, but by the structure of achievements. It is the *rows* of the columns that do the explanatory work.

²⁹See Ichikawa and Jenkins (2017) for discussion of various senses of "knowledge-first."

³⁰Or, as Williamson (2017) puts it "cognitive", as contrasted with "practical" (2017, §4). In a bit more detail, as Williamson writes, "Call 'cognitive' those aspects of intelligent life which concern fitting mind to world, and 'practical' those which concern fitting world to mind (2017, 164).

5 Against Direction of Fit

One point in favor of the mirrors metaphor still remains. What about the idea that the functional and normative features of belief and action seem to be governed by different *directions of fit*? If it still makes sense to think in these terms, then even if knowledge and action are not the only achievements, and the fact that they have these directions of fit doesn't explain the rows, it still might be a deep and illuminating feature of the mind that our achievements cleave into these two functional and normative kinds.

We will now argue that there is no such cleavage in mental kinds, and that this itself is an illuminating and important consequence of knowledge-first epistemology. Let us return to the famous passage of Anscombe's that is supposed to motivate the idea that there are two kinds of attitudes with reverse directions of fit. She writes:

Let us consider a man going round a town with a shopping list in his hand. Now it is clear that there is a different relation when a list is made by a detective following him about. If he made the list itself, it was the expression of an intention; if his wife gave it him, it has the role of an order. What then is the identical relation to what happens, in the order and the intention, which is not shared by the record? It is precisely this: if the list and the things the man actually buys do not agree, and if this and this alone constitutes a *mistake*, then the mistake is not in the list but in the man's performance (if his wife were to say: 'Look, it says butter and you have bought margarine', he would hardly reply: 'What a mistake! we must put that right' and alter the world on the list to 'margarine'); whereas if the detectives record and what the man actually buys do not agree, then the mistake is in the record (1957, 56).

Here we draw your attention to the caveat in this passage that has been almost universally overlooked (though see Frost (2014)): if the *only* mistake made by the man is that the list and the purchases do not agree, the mistake is in the purchasing, not in the list. This strongly suggests that Anscombe thinks there can be other kinds of mistakes we can make in performance. Without this caveat, the DOF theory has some reprehensible predictions. As Frost points out, it is clear that merely having an intention is not sufficient to make it that one should make the world fit one's mind. One should not intentionally do morally reprehensible things, and so one should not carry out one's intentions to do such things. Of course, there is a sense in which a person who maintains intentions and does not execute them behaves irrationally or incoherently, but this does not speak to whether it is better to revise one's intentions or to carry them out. The mistake may be in the *failure to revise*, rather than the performance.

So, we should ask, under what conditions does the failure to satisfy the intention *alone* constitute a mistake, and does this motivate the idea that knowledge and intention (or action) have reverse directions of fit? At the very least, this is only the case when the intention is an intention to do something that is otherwise permissible or possible. The action might be a poor choice either

because it is wrong or cannot be accomplished.³¹ In such cases one should revise one's intention. But what makes an action permissible or possible are—typically—facts that go beyond an agent's psychology, viz., facts about the world. So intentions are governed, it would seem, both by norms governing the satisfaction of one's aim and norms governing the possession of the aim in the first place. The latter norms are plausibly construed as mind-to-world norms. So in what sense can we say that desire, intention, or action have a world-to-mind direction of fit?³²

Let us consider an example. A professor is strongly attracted to a graduate student, and decides to convince her to have a sexual relationship with him. Let us suppose that she rejects his advances repeatedly. What should the professor do? Clearly he should revise his intention. If he continues to try to bend the world to his mind, for example by issuing threats that he will not support her career if she refuses, or retaliating against her by mocking her in front of other students in efforts to get her to agree, the error is clearly *not* that the world fails to fit his intention. The professor is making a mistake of the utmost importance in continuing to try to satisfy his intention. Here we see that Anscombe's caveat is crucial—only once we have ruled out other mistakes can the difference of error be brought into such relief.³³

Think for a moment about what the idea that action inherently has only a world-to-mind direction of fit actually entails. It entails that action aims to exert a kind of utter control over what is outside of us. But this isn't the right way to think about most endeavors at all. Consider the case of ballroom dancing. When one is the follower, one's job is to dance in a way that aligns with the intentions and actions of the leader. When the follower acts out of line with the leader's intentions, her action is defective for this reason. Should we say that the actions of the follower have a mind-to-world direction of fit, while those of the leader have a world-to-mind direction of fit? This distorts the phenomenon. The follower is performing an action, the kind of which involves coming into a certain kind of harmony with another dancer's actions. It is no less agential, and no less subject to performance norms, for that. When the follower fails to conform to the leader's intentions, she fails in her own aim: her action is defective because she has not done what she intended. The leader's actions also involve a mind-to-world aspect, since if the leader is insensitive to the actualities of their dance partner's abilities and bodily sensitivities, they will not perform well. We are in this predicament for all our endeavors—any performance with an aim is subject not just to norms regarding the satisfaction of the aim, but whether pursuing the aim, and pursuing it in the way one is, is adequately responsive to the facts.

What lessons should we draw for a discussion of direction of fit, and sub-

³¹ Exactly what we should say here re: feasibility, practicality, possibility, etc. we can leave open.

³² Remark on Sosa and his attempt to circumscribe epistemic normativity as solely the normativity of achieving aims one has. When we look at other domains it's not like this.

³³ In case the reader is skeptical about introducing such a charged example, consider that a main proponent of direction-of-fit analyses, John Searle, was stripped of his emeritus status at U.C. Berkeley for violation of its policy on sexual harassment (Weinberg, 2019). In contrast with the way our philosophical pursuits are often framed, they are not a matter of disinterested reflection, but rather a manifestation of our broader reasoning and engaging with the world. When the mind-to-world norms governing action are not adequately attended to, they can have terrible results.

sequently for a mirror analogy for the attitudes, from the observation that in attempting to satisfy one's intentions one is subject to norms not just involving the satisfaction of one's intentions, but also their permissibility and possibility? We have multiple norms at play. One is indeed a performance norm that is internal to the possession of intention and attempted intentional action. Failures to satisfy one's intention are inherently failures of performance, in that they are defective qua intentions if they fail to be satisfied. This could be considered a world-to-mind norm (although as the ballroom dancing case makes clear, we must take care not to conceptualize this norm as the imposition of will on the world to make the world "fit" the mind). But, additionally, there are action-relevant mind-to-world norms as well. One must only have permissible intentions; one must have intentions that *can* be satisfied. And so on.

Now, when we reexamine the purportedly mind-to-world attitudes, we find the same thing. Rather than being a point of disanalogy with knowledge and other purportedly mind-to-world attitudes, *all* attempts at achievement are inherently subject to this kind of performance norm. From a belief-first perspective, the norms governing belief are characterized in terms of truth, where the agent should strive to be in a state that accurately represents the world. Here we might see how thinking of these norms in terms of a mind-to-world direction of fit would be alluring, as success in belief can be considered independently of the attitude, viz. truth of content. However, from a knowledge-first perspective things are not so simple. Ignorance is sometimes the result of false belief, but can arise for other reasons besides false belief: Gettier cases, paradigmatically, are cases of ignorance even where one is justified and one's belief is true. When knowledge—being in right relation to the truth—is the aim of belief, questions about performance become inseparable from the norms governing them. Sometimes, when one fails to know, the error is not in the world but in the performance. What is an error in performance for knowledge? Failing to secure the right relation to the facts.

This suggests that the mind-to-world / world-to-mind division falls apart once we take an achievement-first perspective. Because achievements not only have entailment conditions as conditions of success, but also performance conditions (bringing about the success "in the right way", they are inherently subject to performance norms, even if the aims are primarily *thetic*. Conversely, even when aims are primarily *telic*, because we are dealing with performances instead of the imposition of satisfaction conditions, there will be norms of permissibility and possibility (and perhaps other norms as well) that are applicable to the performances as such.

We are now in position to reconceive of what the mistake consists in for the detective. The detective aims at knowledge of what the shopper buys, whereas the shopper aims at buying what is on his list. In the detective's case, his list is an expression of *his performance*; whereas in the shopper's case, the list is an expression of *his aim*. It is not that we have two different kinds of attitudes here characterized by different directions of fit, but rather that the list plays a different role with respect to each endeavor. If the list fails to be an accurate expression, the normative assessment will be different for this reason, not a deep mirroring structure between knowledge and action.

At this point, one might be concerned that we have conflated different ways in

which normativity can be applicable to knowledge and action. While perhaps mind-to-world and world-to-mind norms are applicable to both knowledge and action, only one of the two is *internal* to each of these kinds of performances, respectively. The thought is this: while we can criticize action from many different perspectives, including being out of touch with the facts, only the applicability of world-to-mind normativity is inherent to what action *is*. In intending to Φ , one is aiming to bring about the satisfaction of the intention one has in mind. Conversely, in aiming to know, one is aiming to conform one's beliefs to the facts. This is why, the traditionalist line goes, reasons to believe that p must be reasons that p , not reasons why it would be practically or morally good to believe that p .³⁴

The concern here misunderstands the claim we are making about the practical normativity that is applicable to knowledge. The claim is not that knowledge is subject to norms regarding *other* aims (such as recovering from an illness, or winning a contest), but rather inherent in aiming to know the epistemic agent is aiming to bring about a state of the world: one in which she is in the right relation to the facts, and as such is subject to performance aims. Similarly in the case of action, one's performances will be governed by norms for certain kinds of actions, and these will include making sure your action is adequately sensitive to the relevant facts about the world.³⁵

Lastly, our approach makes better sense of what Anscombe calls *practical knowledge*, knowledge of what one does. Here the idea that in aiming to know we aim to fit the world is entirely out of place, as we are the cause of what we understand (1957, 87)). Of course, knowledge requires being in the right relation to the facts, but in this case the right relation is not plausibly understood as "fitting". One's action and one's corresponding knowledge of that action are both impacts one makes on the world, but the latter is no less epistemic for that.

6 Conclusion

In drawing deep structural analogies between the cognitive and the practical dimensions of intelligent life, Williamson recommends that we, put simply, "begin with things going well" in practical reasoning, viz., when one acts on what one knows. This methodological choice is "knowledge-first" in that it takes as a starting point that knowledge stands to the cognitive as action stands to the practical. The resulting picture is meant to represent a stark alternative to the popular and broadly internalist picture of belief-desire psychology as the core of intelligent life, a picture that places our *aspirations* conceptually and metaphysically prior to that *to which we are aspiring*, both cognitively and practically. We have argued that Williamson's attempt to preserve direction-of-fit within

³⁴This is what Pamela Hieronymi (2005) has termed the Wrong kind of reasons problem.

³⁵There is an interesting question, in both cases, whether it is inherent to knowledge or action that it is subject to moral normativity—some might argue that questions of moral (im)permissibility are inherently applicable to action but not to knowledge. The foregoing suggests parity of treatment in both cases, as both are inherently subject to mind-to-world and world-to-mind norms; however the case we make here does not depend on answering this question. The applicability of practical (world-to-mind) norms does not entail the applicability of moral norms. Conversely, if moral norms are inherently applicable to action in a way they are not to knowledge, the case could be made on another basis than that practical normativity inherently applies only to action.

a knowledge-first paradigm contains holdovers from belief-desire psychology that should be abandoned. Instead, a more thoroughgoing adoption of the idea that we should take “things going well” as conceptually and metaphysically prior requires us to refit the “mirrors” entirely, dropping direction-of-fit as importantly explanatory of different kinds of attitudes and instead focusing on the structural relationships that emerge when we take achievements first.

Williamson is entirely right to begin as he does with things going well in a cognitive-practical system. Where he errs is in how he proceeds from this promising starting point, which is to then attempt to connect the spectrum of our cognitive and practically oriented attitudes through the nexus of direction-of-fit mirroring.

The problems we’ve raised in §4 and §5 for such a strategy are intractable ones. They reveal not only how certain attitudes are *de facto* screened off from entering into the mirror-theoretic analogies (even when we begin with success rather than attempted success), but also how such a proposal commits itself to problematic normative views.

That said, we have *not* recommended—in response to these intractable problems—a return to belief-desire psychology as the center of intelligent life, from which we might then attempt to fit, through analogical reasoning, *different* mirrors between the cognitive and the practical than those mirrors which Williamson favors. And this is the case, even though, as we have argued in the negative part of this paper, Williamson’s mirrors do not all fit.

Rather, our proposal is, more fundamentally, to abandon from the very start the idea that knowledge and action (and their corresponding attempts) are “mirrors” of each other—mirrors reversing direction of fit—while at the same time retaining a view of intelligent life that prioritizes things working well (in a wide range of cognitive and practical endeavors) as explanatorily central and prior to their corresponding mere attempts.

On the “achievement-first” view we favour, achievements have relatively basic explanatory force, and so corresponding attempts, entailment conditions, etc. are to be understood in terms of what it takes to execute the achievement. Any performance aimed at achievement is subject to world-to-mind and mind-to-world directions of fit (although the cases of practical knowledge discussed above suggest that we not take the term “fit” too seriously), and relationships between knowledge, action, and other achievements can be straightforwardly explained as structural analogies in achievement-theoretic terms. This proposal succeeds in capturing—as Williamson helpfully aims to capture—what is in common across the range of attitudes that exhibit the kind of structure that knowledge-belief, action-desire/intention do, albeit, without any of the baggage that comes with direction-of-fit mirrors and what they tend to distort.

References

- Anscombe, G. E. M. (1957). *Intention*. Harvard University Press.
- Archer, A. (2015). Reconceiving direction of fit. *Thought: A Journal of Philosophy*, 4(2):171–180.

- Bradford, G. (2015). *Achievement*. Oxford University Press, USA.
- Carter, J. A. and Gordon, E. C. (2014). Objectual understanding and the value problem. *American Philosophical Quarterly*, 51(1):1–13.
- Carter, J. A., Jarvis, B. W., and Rubin, K. (2015). Varieties of cognitive achievement. *Philosophical Studies*, 172(6):1603–1623.
- Dickie, I. (2015). *Fixing Reference*. Oxford: Oxford University Press.
- Elgin, C. Z. (2017). *True Enough*. Cambridge, MA: MIT Press.
- Frost, K. (2014). On the very idea of direction of fit. *Philosophical Review*, 123(4):429–484.
- Gordon, E. C. (2019). Elgin on understanding: How does it involve know-how, endorsement and factivity? *Synthese*, pages 1–18.
- Greco, J. et al. (2010). *Achieving knowledge: A virtue-theoretic account of epistemic normativity*. Cambridge University Press.
- Grimm, S. (2012). The value of understanding. *Philosophy Compass*, 7(2):103–117.
- Hetherington, S. (1999). Knowing failably. *Journal of Philosophy*, 96(565–587).
- Hetherington, S. (2001). *Good Knowledge, Bad Knowledge*. Oxford: Oxford University Press.
- Hieronymi, P. (2005). The wrong kind of reason. *Journal of Philosophy*, 102(9):437–457.
- Humberstone, I. L. (1992). Direction of fit. *Mind*, 101(401):59–83.
- Ichikawa, J. and Jenkins, C. S. I. (2017). On putting knowledge ‘first’. In Carter, J. A., Gordon, E. C., and Jarvis, B., editors, *Knowledge First: Approaches in Epistemology and Mind*. Oxford University Press.
- Kelp, C. (2015). Understanding phenomena. *Synthese*, 192(12):3799–3816.
- Kelp, C. (2016). Towards a knowledge-based account of understanding. In Grimm, S., Baumberger, C., and Ammon, S., editors, *Explaining Understanding: New Perspectives from Epistemology and Philosophy of Science*, pages 251–269. Taylor & Francis.
- Kvanvig, J. L. (2003). *The value of knowledge and the pursuit of understanding*. Cambridge University Press.
- Lewis, D. (1988). Desire as belief. *Mind*, 97(387):323–332.
- Miracchi, L. (2014). *Getting Things Done*. PhD thesis, Rutgers University, New Brunswick.
- Miracchi, L. (2017). Perception first. *Journal of Philosophy*, 114(12):629–677.
- Montero, B. (2018). *Thought in Action*. Oxford: Oxford University Press.

- Montmarquet, J. A. (1993). *Epistemic Virtue and Doxastic Responsibility*. Rowman & Littlefield.
- Pritchard, D. (2008). Knowing the answer, understanding and epistemic value. *Grazer Philosophische Studien*, 77(1):325–339.
- Pritchard, D. (2010a). Knowledge and final value. In *The nature and value of knowledge: Three investigations*, pages 1–30. Oxford University Press.
- Pritchard, D. (2010b). *The Nature and Value of Knowledge: Three Investigations*. Oxford University Press.
- Searle, J. R. (1983). *Intentionality: An essay in the philosophy of mind*. Cambridge: Cambridge University Press.
- Simion, M., Kelp, C., and Ghijsen, H. (2016). Norms of belief. *Philosophical issues*, 26(1):374–392.
- Sliwa, P. (2015). Iv—understanding and knowing. In *Proceedings of the Aristotelian Society*, volume 115, pages 57–74.
- Smith, M. (1994). *The Moral Problem*. London: Blackwell.
- Sosa, E. (2015). *Judgment and Agency*. Oxford: Oxford University Press.
- Sosa, E. (2021). *Epistemic Explanations: A Theory of Telic Normativity, and What it Explains*. Oxford: Oxford University Press.
- Velleman, D. (2000). *The Possibility of Practical Reason*. Oxford University Press.
- Weinberg, J. (2019). Searle found to have violated sexual harassment policies (updated with further details and statement from berkeley). *Daily Nous*.
- Weintraub, R. (2007). Desire as belief, lewis notwithstanding. *Analysis*, 67(2):116–122.
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.
- Williamson, T. (2017). Acting on Knowledge. In Carter, J. A., Gordon, E. C., and Jarvis, B., editors, *Knowledge First: Approaches in Epistemology and Mind*, pages 163–181. Oxford: Oxford University Press.
- Williamson, T. (2018). Justifications, excuses, and sceptical scenarios. In Dutant, J. and Dorsch, F., editors, *The New Evil Demon*. Oxford: Oxford University Press.
- Williamson, T. (2021). Epistemological ambivalence. In Hughes, N., editor, *Epistemic Dilemmas*, pages 1–20. Oxford: OUP.