

## Believing Well

Knowing is believing well. Or so I will argue.

When I say that knowing is believing well, I do not mean only that when you know, you believe well, or conversely, or both. I mean that the nature of knowledge is believing well – that knowledge is, with respect to belief, what acting well is, with respect to action. It is a corollary of this thesis that knowledge is a normative property, through and through. It is not merely a psychological state on a par with belief, as Williamson [2000] would have it, or a conjunction of psychological and modal properties, as Dretske [1981] or Nozick [1981] would have it.<sup>1</sup> So it is a corollary of my thesis that epistemology is a fundamentally normative discipline, by its very nature and at its very heart.

When I say that knowing is believing well, I do not mean that knowing is just a matter of having the right belief – a belief that satisfies some norm or rule. Indeed, it is hard to see what norm or rule could be such that you know just in case your belief satisfies it, unless it is the rule ‘believe only what you know’.<sup>2</sup> And obviously this gives us no informative insight into the nature of knowledge at all. So that is not what I mean. Instead, when I say that knowing is believing well, what I mean is that knowing is to having the right belief as *acting well* is to *doing the right thing*.

There is much circumstantial evidence to be had that knowing is believing well. Knowers believe the right thing, and those who act well, do the right thing. But knowers do more than just believe the right thing, and similarly for those who act well. Knowers hold beliefs that are rational to hold, in their situation, and those who act well make choices that are rational, in their situation. Knowers generally do not know by accident, except for the accident of opportunity, and similarly, people who act well do not act well merely by accident, again except for the accident of opportunity. One could make a long and interesting paper out of collecting this circumstantial evidence and systematizing it. But that is not what I am interested in doing in this paper. The evidence is all there for the taking, and in any case it can all or nearly all be coopted by the

---

<sup>1</sup> According to some (compare Wedgwood [2002], Shah [2003], and discussion in McHugh and Whiting [2014]), belief itself is a normative kind. Still, even if this is correct, it follows from my thesis that knowledge is normative in a way that belief is not.

<sup>2</sup> Compare Williamson [2000], [2013], Engel [2004], Smithies [2012].

contrary thesis that acting (or believing) well requires knowing what it is right to do (or believe) – a task I leave as an exercise for the reader.<sup>3</sup> My argument, in contrast, cannot be so coopted. More on this, later.

My argument, though it is ambitious in aims and scope, will be simple. After clarifying the distinction that underlies my thesis – the distinction between acting well and doing the right thing – I will distinguish between two contrary theses about the analytic priority of these two normative concepts. Then I will offer a new argument that I believe should resolve this priority question. It is a corollary of this general argument about the priority of the concepts of acting well and doing the right thing that something like the following principle must be true:

**Well Principle**      Every normative standard N which specifies right and wrong ways of  $\phi$ -ing *gives rise to a corresponding standard of what it takes to  $\phi$  well.*

The penultimate section of this paper will use the Well Principle as the key premise in my argument that knowledge is believing well, and in the final section I will return to address some objections.

## I      Acting Well vs Doing the Right Thing

It is important, since one of the issues at stake in this paper is the priority relationship between acting rightly and acting well, not to build in too much theory up front into our understanding of this distinction. So I will focus on examples. I take it that it will be agreed on all sides that acting well, in the sense in which I am interested, entails doing the right thing, and so the helpful examples for distinguishing the two will all be cases in which someone does the right thing but does not act well.<sup>4</sup> The most famous such example is Kant's prudent shopkeeper, who gives his customers correct change because, in the twenty-first century version of the example, he is concerned about his Yelp ratings. Kant's shopkeeper definitely does the right thing – it is not controversial whether shopkeepers ought to give correct change. But there is something that his action lacks. What he lacks, is that his actions do not give him moral credit. He does not, as I will put it, *act well*.

It is easy to over-theorize this distinction up front. For example, it is natural to want to contrast Kant's shopkeeper with the example of a *conscientious* shopkeeper, who does act well in giving correct change. But because it is controversial what is required in order to act well, it rapidly becomes controversial exactly

---

<sup>3</sup> Compare Sliwa [2016], Johnson King [unpublished].

<sup>4</sup> As will emerge shortly, my view is that there are different senses of or dimensions of rightness, and hence that it can be denied that acting well entails acting rightly, if we cross dimensions. For example, those who think that moral worth is equivalent to praiseworthiness may accept that moral worth is acting well *subjectively*, but does not entail doing what is *objectively* right. This, I believe, is the view of Markovits [2010], for example.

how best to describe the conscientious foil. And conversely, descriptions of the conscientious foil lend themselves to interpolated theories about what, exactly, is involved with acting well. For example, according to an Aristotelian view, the shopkeeper who acts well will be motivated out of appropriate concern for his customers, or out of a sense of justice. But according to Kant, desires for justice or for his customers' well-being are not enough for moral merit – a sense of duty is required, instead. What is important for us, therefore, is not exactly how to *describe* the foil of the conscientious shopkeeper, but that we know him when we see him, and he has something important that Kant's prudent shopkeeper lacks.

I will take Kant's shopkeeper to be a paradigm case in which someone does the morally right thing but does not act morally well. But there are other paradigms of morally right action without acting well. For another moral case, consider an agent who, through a combination of errors, comes to believe the truth about what she ought to do, but for the wrong reasons. Just to make the example concrete, suppose that she comes to believe that she ought to reveal to her friend her knowledge that his wife is cheating on him, reasoning that his wife is simply bringing him down and the only way that he will divorce her is if he comes to believe that she is having an affair. In fact, this is what she should do, but not for the reasons that she believes – in fact the reason why she should reveal her knowledge is simply out of duty to her friend, so he can make his own decision in light of it, or open up to her about his open marriage. If the agent in this case does what she believes she ought to do, I believe, she may do the right thing, but she does not act well, in the relevant sense.

Non-moral cases are possible, as well. To focus on an example similar to one that I will rely on later, in many situations in chess there is a unique best available move – the right move to make in that situation. But there is a difference between someone who makes the right move in understanding of what makes it the right move, and someone who makes it without such understanding. For example, someone could make a move that happens to be the right one because they are trying to illustrate the movements of the pieces to a novice player. Or they could make a move that happens to be the right one because they have been bribed to capture both bishops at the earliest opportunity. Or they could make the move that happens to be the right one as the result of a miscalculation or combination of miscalculations about possible future states of the board. In each of these cases, the player makes the right move, but in each of them, their move lacks merit – it does not redound to their credit.

The distinction between acting rightly and acting well is matched by similar distinctions between taking the rational action and acting rationally, between believing what it is rational to believe and believing rationally, and between fearing what it is rational to fear (e.g., the man chasing you with a knife) and fearing rationally (e.g., because he is chasing you with a knife, rather than because he is wearing clown makeup).

Epistemologists often distinguish, similarly, between *propositional* and *doxastic* justification – someone who believes what she has a propositional justification to believe may yet, for all that, fail to be doxastically justified.<sup>5</sup> So the distinction between propositional and doxastic justification has much in common with the distinction between acting rightly and acting well, and correspondingly with the distinction between doing (or believing, or fearing) what it is rational for one to do, and acting (or believing, or fearing) rationally.

As should be clear from the foregoing discussion, the distinction between doing well and doing the right thing can be applied to belief, fear, and other attitudes as well as action. It also arises for different things that we might mean by ‘the right thing’. An action might be the right thing to do from an objective perspective that is not filtered by the information an agent happens to have about her situation, without being a rational thing for her to do – that is, the right thing to do from a subjective perspective that is so filtered.<sup>6</sup> Or conversely, an action might be the rational thing for an agent to do, and hence the right thing to do from her subjective perspective, but not the right thing to do from a more objective perspective.

There are corresponding distinctions between acting well and doing the right thing. An agent may do the right thing, objectively speaking, without acting objectively well. For example, she may do it only by the coincidence of her motives with the correct motives, as with Kant’s shopkeeper. An agent may also do the right thing, subjectively speaking, without acting subjectively well. This can also be illustrated with Kant’s shopkeeper, provided that we are clear to stipulate that the shopkeeper has the information required to easily determine that giving correct change is what is morally required. Similar points go for fear and other attitudes. An agent may fear the person who it is objectively correct to fear – for example, someone who unbeknownst to her in fact intends to kill her – but only by coincidence, as where his clown makeup triggers one of her phobias. Or she may fear the person who it is subjectively rational for her to fear, but again for the wrong reasons – as where she knows that he intends to kill her but underestimates his abilities, but is still phobic about his clown makeup.

Because there are dimensions of acting well corresponding to both objective and subjective dimensions of acting rightly, and the objective dimensions of acting rightly can come apart, it follows that the objective and subjective dimensions of acting well can come apart as well. For example, all it takes to act well along the subjective dimension but not act well along the objective dimension is to have incomplete information. Many examples fit this profile. In these cases an agent does the rational thing – for example, gives their spouse the pills from the bottle carefully marked as their pain medication – and does it well, in

---

<sup>5</sup> Compare, for example, Pollock and Cruz [1999]. More on this later.

<sup>6</sup> Exactly how to understand this more objective point of view is a somewhat vexed question that I’ll ignore here. See, for example, Schroeder [unpublished].

the careful way that they have always done it – but still fails to do the objectively right thing or act well along the objective dimension, because unbeknownst to her, the pills have been undetectably swapped out for poison pills by an assassin.

Cases of acting well along the objective dimension while acting poorly along the subjective dimension are also plausible in cases in which agents possess misleading information, but require a little bit more setup. Some examples of blameworthy right action plausibly fit this profile. In these cases an agent does the correct thing – for example, delivering their spouse’s pain medication – and does it for the right reason – it is the time of day for the pain medication and the pills are in the box correctly marked ‘pain medication’ – but still deserves blame for her action, because she ignores her nagging but false apparent memory of swapping out the pain medication for poison pills.

But the highest standard to which we can hold action is that it be acting well along *both* dimensions. Cases of acting well subjectively but not objectively miss something important. And cases of acting well objectively but not subjectively miss something important. The best actions – the ones worth aiming at – are both. We may refer to such actions – ones which involve acting well along *both* the objective and subjective dimensions – as *acting well full-stop*. When I say that knowledge is believing well, what I mean is that it is *believing well full-stop* – believing well along both the objective and subjective dimensions. Since objective rightness for belief is truth and subjective rightness for belief is rationality, that is why it follows (on this view) that knowledge entails truth, and knowledge entails rationality of belief. Similarly, when Kant says in the first *Critique* that knowledge happens “when the holding of a thing to be true is sufficient both subjectively and objectively”, I take him to be endorsing a version of the view that knowledge is believing well full-stop.<sup>7</sup>

## 2 The Priority Thesis

So far, I have distinguished between several kinds of right action, noted corresponding notions of acting well for each, and introduced the notion of acting well full-stop. I now turn to the question of the analytical and explanatory priority between right action and acting well – and similarly for the other, corresponding distinctions between believing the rational thing and believing rationally, fearing the thing it is rational to fear and fearing rationally, and so on.

We should start by being clear that this is a real and important question. Even among those who distinguish between propositional and doxastic justification, for example, and who agree that doxastic justification entails propositional justification, and who even agree that this entailment is evidence of a close

---

<sup>7</sup> Kant [1999, A822/B850]. See, in particular, Chignell [2007a], [2007b].

analytic or constitutive connection between these two properties, there is much disagreement about the order of this analytic (or constitutive) connection. Some, including Alston [1985], Pollock and Cruz [1999], Swain [1982], and Feldman [2002], say that doxastic justification simply consists in propositional justification plus some further, differentiating, condition. Whereas others, including Chisholm [1966] and Turri [2010], say that propositional justification is a kind of modal shadow cast by facts about doxastic justification.

The question of the order of priority of acting rightly and acting well is also important in moral philosophy. Aristotle famously clearly distinguishes between acting from virtue and acting in accordance with virtue, and on some interpretations of Aristotelian ethics, such as that of Rosalind Hursthouse [1999], mere right action is a kind of modal shadow of what is done by someone who acts well. Similarly, on some interpretations of Kant [2002], the reason why acts that can only be done for non-universalizable principles are wrong is that, as the case of the shopkeeper is supposed to illustrate, acting well (with moral worth) requires acting according to principles. And so on this interpretation, Kant's view is also that what is right or wrong is just a modal shadow of what could possibly be done well.

In contrast to these interpretations of both Aristotle [2009] and Kant [2002], most contemporary normative ethical theories, including those of theorists like Thomson [1990], Scanlon [1998], Parfit [2011], Kamm [2007], and McMahan [2002], provide rich and illuminating explanations of which actions are right or wrong without adverting in any way to considerations of what it takes to act well. All of these theories are committed, therefore, to the reverse order of explanation – since acting well is not prior to acting rightly, it must be posterior, if either priority thesis is true. Similarly, contemporary discussions of moral worth, such as those of Arpaly [2002], Hills [2009], Markovits [2010], Sliwa [2013], [2016], Howard [unpublished], and Johnson King [unpublished], take for granted that moral worth is to be explained in terms of right action, rather than conversely, and merely disagree about how this is to be done. Since moral worth is a kind of acting well, these theorists are also committed to the order of explanation from rightness to acting well, in my terms.

In a moment I will argue that the standards for acting well must be explained in terms of the standards for acting rightly. But before giving that argument, I first want to emphasize why this view is not obvious, and requires argument.

Notably, each of these disputes has much in common with other 'common factor' disputes in philosophy. For example, some say that veridical perceptual experience is a matter of something that is shared with non-veridical experience plus some further condition, but others contend that all that veridical and non-veridical perceptual experiences have in common is that they are shadows of or approximations to veridical

perceptual experience.<sup>8</sup> Similarly, the orthodox traditional view has been that knowledge is belief plus some further condition, but Williamson [2000] claims that belief is just a shadow or approximation of knowledge. And legal positivists claim that legitimate legal regimes share a common nature – that of being legal regimes – with illegitimate legal regimes but satisfy a further condition – legitimacy – but some of their critics contend that legitimate legal authority is prior to and explanatory of legal authority as such, and that the broader notion of legal authority, which encompasses illegitimate as well as legitimate regimes, is simply its shadow or an approximation.

In all common factor debates, there are two important conditions, one of which entails the other. In all of these debates, the naïve view is that the stronger of the conditions can be defined in terms of the weaker one, together with some further condition. And in all of these debates, critics of the naïve position object that we can't assume without argument that the order of analysis or explanation must go in this way, because there are perfectly intelligible possibilities on which it goes the other way around, while it is not hard to find defenders of the naïve view saying no more in defense of their position than that one condition is stronger than the other.<sup>9</sup> So as with each of these disputes, we should tread carefully and not jump to conclusions about which must be prior, but should look for arguments.

So I'm going to try to give an argument. In contrast to the common-factor inference, my argument will *not* establish that acting well is just a conjunctive condition of acting rightly plus satisfying some further, differentiating condition. But it *will* establish that the conditions on acting well exist and have the particular shape that they do *because* there are conditions on acting rightly, and because of the particular shape of those conditions.

### 3 The Argument From Games

My argument is simple. It works by carefully choosing a special case of the acting rightly/acting well distinction where we can *control* for whether the conditions of acting rightly or acting well are explanatorily prior. In this case we can prove that there is only one possible answer as to whether the conditions on acting rightly come from the conditions on acting well or conversely. And this case gives us both direct and indirect arguments that other cases of the acting rightly/acting well distinction must work in the same way.

---

<sup>8</sup> For example, McDowell [1982].

<sup>9</sup> And of course, some critics of the naïve view in each of these disputes rejects the idea that there even *is* any unified account of the weaker condition, maintaining that it is just the disjunction of the stronger condition and something else. This is prominently the position of disjunctivists in the philosophy of perception, but disjunctivism has analogues for each of these other debates.

The special case of the acting rightly/acting well distinction on which we will focus is the case of artificial games. I will use card games to illustrate, since card games come with a wide variety of rules and often exhibit minor variants. Sheepshead, for example, the case that I will focus on, is a trick-taking card game played mostly in Wisconsin and parts of southern Indiana populated by the descendants of German immigrants. The rules are complicated – a deck consists of four suits of cards 7 through 10, jacks, queens, kings, and aces, but all queens and jacks are trump along with diamonds, jacks take priority over aces for taking tricks, cards have point values that don't correspond to their priority for taking tricks, the objective of each hand, including whether players have partners, who their partners are, and whether the objective is to take points or leave them on the table can change based on the bidding phase of the hand. And in the version that I played growing up, the game dynamics are significantly different depending on whether it is played with three, four, or five players. Wikipedia's page for Sheepshead describes many other variations.

The rules of Sheepshead, as with other card games, dictate what players are allowed to do at any point in the game, and they dictate the victory conditions for the game. These rules are conventional, so we know where they come from – they come from the conventional stipulations that we set at the beginning of each game, in which we clarify which form of the game we intend to be playing, specifying things like which suit is trump, the order of the cards, the bidding process, and the number of players.

The rules about what players are allowed to do at any point during the game constitute a kind of standard of rightness, but not a very interesting one. Because these rules structure what counts as playing this particular form of the game, if you break these rules, you are not really playing this variant of the game, any more than you are playing chess if you move your rook diagonally. But together with the victory conditions, the rules of play also determine which moves are the right ones to make at any given point during the game. In any given version of Sheepshead, as in most other well-designed games, there are uniquely best moves to make in many different situations in the game – right moves to make – and in many other situations, there may be a small range of equally good moves to make. And part of the pleasure of playing card games that admit of multiple variants or in which the objectives of play are situational, like Sheepshead, is precisely that of working out for oneself what the appropriate move is in the particular variant that one is playing.

The moves that are best in any given situation for achieving the victory condition of the game count as satisfying a different standard of rightness. They are right not just in the sense of being permissible moves of the game, but in the sense of being the *right move* to make in that situation within the game. This is a more interesting standard of rightness, but it is still determined *by* the underlying rules of the game. Which moves are right in this sense is a mathematical or game-theoretical consequence of the rules of admissible play and the victory conditions. The way in which the rules of admissible play and the victory conditions determine



which moves are the right ones to make is illustrated well by the case of simple games like tic-tac-toe, for which all players except for very young children are familiar with the full calculation. But the same principles apply in the study of chess endgames, for example, and generalize to other well-defined games – including Sheepshead.

Finally, it is possible, when playing a card game like Sheepshead, to make the right move, but not play well. You might, for example, just like leading trump and do it because you like it, and it turns out that this is what it is right to do in your situation in the game. Or you might get confused by the rules and miscalculate in two ways that cancel each other out. So there are, without question, clear standards on what it takes to play Sheepshead well, and in particular, to play well in making any particular move. So there are both standards on making the right move and on playing well, for Sheepshead. Indeed, there are such standards for *every version* of Sheepshead – versions in which clubs are trump instead of diamonds, versions which change the order of the cards, versions which change the point values of the cards, versions which change the rules for revealing partners, and so on, for all of the many variants of the game.

But again, because we invent each version of Sheepshead, we know exactly where the standards on playing Sheepshead well must come from. When we invent a new variant of the game, we do not invent a standard for what it takes to play well; rather, we invent new rules of admissible play. Since what we control when we vary the game are the rules of admissible play and the victory conditions, any changes in the standards on what it takes to play well must come *from* the rules of admissible play and the victory conditions. But the standards on playing well can't come directly from the rules of admissible play and the victory conditions, either, because all that those give us are the tree of possible ways the game could play out, together with probabilities attached to each node, and a classification of end-points of the game as wins or losses (or by how much money is won or lost, for many card games). The facts about the right moves to make in each situation fall directly out of the structure of these trees for possible play. So the facts about the right moves to make do not fall out of standards on playing well. On the contrary, the only way that rules of admissible play and victory conditions could possibly determine standards of playing well, is *by* determining the standards for which moves are the right ones to make, in each situation.

So the case of artificial games, such as card games, is controlled in a way that allows us to see that at least in their case, there *is* a distinction between playing well and making the right move, but that the standards on playing well must come *from* the standards on making the right move. And this gives us the means to develop two distinct arguments that the standards on acting rightly are prior to and explanatory of the conditions on acting well more generally – even outside the case of artificial games.

The first argument is indirect, and non-deductive. If we take seriously the analogy between the distinctions between acting rightly and acting well across the many examples that I gave in the last section – examples involving action, belief, and emotion, examples involving both moral and non-moral standards for right action, and examples involving both objective and subjective standards of rightness – then we may draw the conclusion *by analogy* that since in the case of artificial games, the conditions on acting well come from and are explained by the standards for making the right move, the same goes for every other analogous distinction.

I'm friendly to this way of thinking, but it doesn't strictly rule out the possibility that these distinctions are merely analogous in some other respects that doesn't go this deep. But fortunately, we also have the tools for a more direct, deductive, argument that the conditions on acting well derive from, and are explained by, the standards for acting rightly. And this is because we can see from the case of artificial games that *something* about the existence of a standard on right play is enough, all by itself and in the absence of further conventions, to establish conditions on playing well. When we invent a new variant of Sheepshead or another card game, we do not also decide what will count as playing it well, and it may be that no one has yet played it well. Yet the condition on playing well is there, already. When we look back, later, once we have all figured out the consequences of these new rules, we may observe that Wen was playing well already on our very first hand, but it took Sahil a few hands to get the hang of it.

So it is clear that conditions on playing well somehow get automatically generated by the standards for right play. In other words, something like the following principle must be true:

**Well Principle**      Every normative standard  $N$  which specifies right and wrong ways of  $\phi$ -ing *gives rise* to a corresponding standard of what it takes to  $\phi$  *well*.

Only if there is some background explanatory principle like the Well Principle could it be that establishing a standard for right play will automatically establish conditions for playing well.

But if anything like the Well Principle is true, then wherever there are standards of rightness, there will be corresponding conditions on acting well that are explained by them. And so it would be redundant to think that there is something worth calling 'acting well' which explains a standard on right action, which then gives rise, through the Well Principle, to a further corresponding condition of acting well. There never needs to be such a thing, because the Well Principle guarantees that there will *always* be conditions of acting well.

Finally, the Well Principle predicts and explains something exciting: it predicts that standards of rightness will *always* come paired with conditions of acting well. But that is precisely what we already observed in the last section. For any norm that we could think of – on action, belief, emotions, or other attitudes, moral or otherwise, and objective or subjective – there were corresponding distinctions worth being made between satisfying that norm and doing well. The Well Principle explains why there are always such pairings, where a condition for which an agent gets some kind of credit but which constrains the *way* that she does something, and not just *what* she does, gets paired with a less-restrictive condition that is also of normative significance.

In contrast, if the normative significance of doing the right thing is just some kind of shadow or approximation of the normative significance of acting well, there is no good explanation of why the more fundamentally normative significant thing, of which other normatively significant things may be shadows, always concerns *how* you act, believe, or feel, and not just *what* you do, believe, or feel. There is no good – certainly no obvious – explanation of why the fundamental norms governing action, belief, and the emotions, governing morality, prudence, games, and the law, and along both objective and subjective dimensions would always need to specify *how*, and not just *what*. But on the contrary, if norms may come from pretty much anywhere, and govern pretty much anything, but whenever a norm exists, the Well Principle generates the existence of a corresponding condition of acting well that constrains *how* the underlying norm is satisfied, then that is exactly what would elegantly explain this unity.

#### 4 Knowledge is Believing Well

So here is where we are. In section 1 I distinguished between acting rightly and acting well, and in section 2 I argued that the property of acting rightly is prior to and explanatory of the property of acting well, rather than conversely, and we encountered the Well Principle – a principle whose truth we deduced but have not yet explained, which guarantees that whenever there is a standard of acting rightly, there is a corresponding property of acting well. But recall that my goal in this paper is to argue that knowledge is believing well – by which I mean, believing well full-stop, along both objective and subjective dimensions. And this thesis is important in part because, if it is true, then knowledge is a normative property in its own right – that epistemology is a fundamentally normative discipline, through and through.

Let us start with this consequence of the thesis that knowledge is believing well. What would it take for it to turn out that knowledge is *not* a normative property in its own right? This thesis is a consequence, I take it, of Timothy Williamson's [2000] view that knowledge is a mental state on a par with belief, as well

as of traditional forms of epistemological externalism such as those espoused by Armstrong [1973], Goldman [1967], [1988], Dretske [1981], and Nozick [1981]. On each of these views, knowledge has no intrinsically normative nature – the essence of knowledge can be completely unpacked without encountering anything about norms, reasons, obligations, values, duties, or the like.

This claim seems striking because knowledge certainly *seems* to be important or valuable. When we say that someone knows, we certain *communicate* something that seems to matter for the knower’s standing or authority, for how we take her to be entitled to act or reason, and for the credit that she is due. These are all transparently normative properties. So knowledge certainly seems to be normatively *significant*. But of course, some things are normatively significant without having normative natures. Death is one such example. When we learn that something would result in our imminent death, that matters in an obvious way – not because death’s nature is normative, but just because there are important normative facts *about death* – in particular, that death is *bad*.

So those who deny that knowledge has a normative nature can similarly explain away its seeming normative import by insisting that this is simply because there are important normative facts *about knowledge*. This is, for example, what Williamson [2000] is doing when he says that belief ‘aims at’ knowledge, or endorses the principle (Williamson [2013]) that a belief is rational just in case it is knowledge. These are two different ways of trying to formulate the principle that belief is *normed* for knowledge, or in other words, that there is a *rule* governing belief, to believe only what you know.<sup>10</sup> There could also be other ways of defending rules governing belief that mention knowledge. For example, the rule could be to believe only what you don’t know that you don’t know, or only what you don’t know that you are not in a position to know, or only what you are not in a position to know that you are not in a position to know.<sup>11</sup> All of these are possible normative standards for belief – possible standards of rightness – that mention knowledge. And if any of these is correct, then knowledge will turn out to be highly *significant* for the norms governing belief, even though knowledge itself does not have a normative nature. The normative significance will come from the fact that the norms mention knowledge.

There is a kind of picture against which this idea makes sense. The picture goes like this: the *highest* standard for belief to meet – the very best kind of belief – is knowledge. And then there are lower standards for belief to meet – truth, or justification. Being true, or being justified, are nice properties for belief to have, but they are consequences of knowing. So if the fundamental norm governing belief is to know, then we can

---

<sup>10</sup> Though contrast Smithies [2012], who distinguishes between the claim that belief ‘aims’ at knowledge from the claim that knowledge is the norm for belief.

<sup>11</sup> Compare Schroeder [2015b].

get a tidy explanation of the values of truth and of justification – they are a kind of second-best norm that are explained by the existence of a more fundamental rule that says to believe only what you know.

But the Well Principle shows that this picture cannot be right. For it predicts that wherever there is a standard of rightness, there is a corresponding standard of acting well that entails it but further constrains *how* you meet the standard of rightness. So if the standard of rightness for belief is to believe only what you know, then there must be some further, more demanding standard of what counts as believing well – a standard that you meet only if you know (since that is believing rightly), but which goes over and above knowing, because it requires knowing *well*. But I don't know what such a standard would be.

It can't, for example, be that you know well when you know that you know. This view would be the extension of the principle that you act well when you take the action that you know is right. But this condition on acting well is too weak – you can act poorly even in knowledge that what you do is the right thing – indeed, this is plausibly the plight of Kant's selfish shopkeeper, one of the classic examples introducing the distinction between acting well and doing the right thing. This is why advocates of the view that knowledge of the right action plays a central role in an account of acting well require the agent to be motivated *by* her knowledge of which action is right. But the corresponding requirement for knowledge would be that you know *because* you know that you know. But this condition is too strong – you can't know that you know that P prior to and independently of knowing that P, in order to base your knowledge on that knowledge.

Indeed, the point generalizes – there is no such standard. To see why, it suffices to pay attention to the kinds of cases in which someone does the right thing, but fails to act well. For example, you can do the right thing, but only by coincidence. But you cannot know, but only by coincidence – if it is a coincidence, then you don't know. Another kind of example of doing the right thing without acting well is if the reasons for which you do what you do have nothing to do with what makes it the right thing to do. But again, if your reasons for believing have nothing to do with what makes it knowledge, then it is not knowledge, after all. And more generally, it is possible to do the right thing without acting well, because standards of rightness only affect *what* you do, and not *how* you do it. But knowing is not just a matter of what you do – whether you know depends intimately on how you do it.

So this is my argument: knowledge is – at least – normatively important. But if the Well Principle is true, then if there is any standard of rightness *about* knowledge, then there must be a corresponding standard of what it takes to believe *well*, which is more demanding. But there is no such standard, and as I have argued, the Well Principle is true. So there is no standard of rightness *about* knowledge. It follows that the best explanation of why knowledge is – at least – normatively important is that it is *intrinsically* normative. The

thesis that knowledge is believing well, in contrast, provides an elegant explanation of all of these things. It explains why knowledge is intrinsically normative, because acting well is intrinsically normative and entails doing the right thing, which is intrinsically normative. It explains why it is normative without generating a further standard of acting well, because it *is* the standard of acting well that is so generated. And it explains why knowledge is not subject to coincidence – far from being a special fact about knowledge, this is a highly general fact about acting well, of which knowledge is merely a special case.

Of course this argument doesn't settle exactly what sort of standard of rightness knowledge is believing well with respect to. But we can do a bit to triangulate on an answer to this question, as well. For one, knowledge is widely assumed to entail truth, which is intuitively an objective standard. And it is also widely entailed to entail propositional rationality of belief, which is intuitively a subjective standard. Correspondingly, knowledge can be defeated by both objective and subjective factors – both by psychological facts about what else the agent believes, and by facts about her situation of which she is unaware. So it is natural to conjecture, since knowledge entails satisfying two different standards for belief – one objective, and one subjective – that knowledge is believing well full-stop – that is, that it is believing well with respect to both the objective and subjective standards of rightness governing belief.

The thesis that knowledge is believing well full-stop doesn't just predict that knowledge entails belief that is both objectively right (i.e., true) and subjectively right (i.e., propositionally rational); it also predicts that even someone who believes well subjectively and believes the right thing objectively may fail to know. But of course this prediction is correct – this is what Gettier cases are – cases in which someone rationally believes the truth, but fails to know.

Similarly, the thesis that knowledge is believing well full-stop predicts that even someone who believes the thing that it is propositionally rational to believe and believes well by the objective standard can fail to know. And I believe that these cases are possible, as well. For example, the literature on defeasibility analyses of knowledge is full of examples like this:<sup>12</sup> Carl has ample direct evidence that P. But Carla has told him that  $\sim P$ . But Carl also believes – truly – that Carla is trying to trick him. In this case, the fact that he possesses the defeater also ensures that it is propositionally rational for Carl to believe that P, despite his misleading evidence from Carla's testimony. And similarly, the fact that Carl has misleading evidence about P doesn't prevent him from knowing all by itself, because he also possesses a defeater for this misleading evidence. But despite all of this, Carl might mis-weigh his evidence, treating his independent evidence about P as outweighing Carla's testimony out of disrespect for Carla, rather than on the basis of his appreciation

---

<sup>12</sup> Starting with Lehrer and Paxson [1969].

that she is trying to trick him. In this version of the case, I claim, Carl believes well objectively and believes the thing that it is propositionally rational for him to believe, but his belief is not doxastically rational – he does not believe well subjectively. And the right thing to say about such a case, I believe, is that Carl does not know.

So the best version of the view that knowledge is believing well, I believe, is the version which says that it is believing well full-stop – believing well along both the objective dimension, which is truth, and along the subjective dimension, with which we are familiar under the guise of propositional rationality. When Kant claims, in the first *Critique*, that “when the holding of a thing to be true is sufficient both subjectively and objectively, it is *knowledge*,” I believe that this is what he meant.<sup>13</sup>

## 5 Corollaries

I close with a few parting observations.

In setting up my aims for this paper, I noted that there are more direct ways of arguing that knowledge is believing well, simply on the basis of collecting observations about parallels between the conditions that defeat knowledge and those that defeat acting well. For example, accidentality seems like a defeating condition on knowledge, and accidentality is a defeating condition on acting well. My argument, I suggested, could be expected to improve over this strategy, by better avoiding the possibility of the alternative explanation according to which doing the right thing requires knowledge.

I can now make good on that claim. If acting well requires knowledge of what is right, as claimed by, for example, Sliwa [2013], [2016] and Johnson King [unpublished], then acting well will inherit some of the features of knowledge – for example, non-accidentality – *from* knowledge. But if my thesis is correct, in contrast, then knowledge and acting well are on a par, with respect to where these features come from – both are special cases of  $\phi$ -ing well, simply for different values of  $\phi$ . If we only look at the circumstantial evidence, it may be hard to adjudicate between these two views, or the points of adjudication may turn on subtle judgments about marginal cases. But my argument provides a way of telling. Since the Well Principle cannot be applied to knowledge, knowledge *must* be a special case of  $\phi$ -ing well.

On the flip side, my claim that knowledge is believing well full-stop raises *prima facie* troubles for the main contemporary theories of acting well that compete with those of Sliwa and King. According to Julia Markovits [2010], for example, a prominent representative of this alternative camp of theories, moral worth

---

<sup>13</sup> Kant [1999, A822/B850]. See also Chignell [2007a], [2007b], and Schroeder [2015a], [2015c].

consists in doing the right thing *for the reasons that make it right*.<sup>14</sup> If this is correct, then knowledge will require believing the right thing *for the reasons that make it right*. But intuitively, truth is what makes a belief right. So the reasons that make a belief right must be reasons that make it true. Not surprisingly, this is a not-uncommon idea at various places throughout the Gettier literature – that what Gettier cases miss is that the reasons for which someone believes something are not the reasons why it is true. But when it comes to knowledge, this thought is too general. Testimony can be an adequate source of knowledge, but in most cases, testimony that P does not make it true that P.

So if the *right reasons* approach to theorizing about  $\phi$ -ing well is on the right track, the fact that knowledge is just a special case of  $\phi$ -ing well constrains how we can reasonably think about what sorts of considerations an agent is allowed to base her action or belief on, in order to count as  $\phi$ -ing well. The restriction to things that *make*  $\phi$ -ing right cannot be exactly the right restriction. It turns out that I think this is the right conclusion about how to develop right reasons accounts of acting well, for independent reasons, but that would take us too far afield, here.

In this paper, I've argued that knowledge is believing well, full-stop. Epistemology is, therefore, a normative discipline through and through, in content as well as in import.

## References

- Alston, William [1985]. 'Concepts of Epistemic Justification.' *Monist* 68(I): 57-89.
- Aristotle [2009]. *Nicomachean Ethics*. Trans. W.D. Ross. Revised and notes, Lesley Brown. Oxford: Oxford University Press.
- Armstrong, David [1973]. *Belief, Truth, and Knowledge*. Cambridge: Cambridge University Press.
- Arpaly, Nomy [2002]. *Unprincipled Virtue: An Inquiry into Moral Agency*. Oxford: Oxford University Press.
- Arpaly, Nomy, and Timothy Schroeder [2014]. *In Praise of Desire*. New York: Oxford University Press.
- Chignell, Andrew [2007a]. 'Belief in Kant.' *Philosophical Review* 116(3): 323-360.
- \_\_\_\_\_ [2007b]. 'Kant's concepts of justification.' *Notus* 41(I): 33-63.

---

<sup>14</sup> The *right reasons* account of acting well has the virtue of offering an elegant explanation of the Well Principle – if what it is to act well in some respect *is* to act for certain reasons, then wherever there are reasons, there will be a corresponding standard of acting well. So all it takes to predict the Well Principle is the assumption that wherever there are normative standards, there are corresponding reasons. It is also worth noting that against the backdrop of this *right reasons* account of acting well, there is much in common between my conclusions in this paper and the observations in Mantel [2013], who applies observations about knowledge to cases of action for reasons – the reverse of what I am trying to do in this paper.



- Chisholm, Roderick [1966]. *Theory of Knowledge*. Englewood Cliffs, NJ: Prentice Hall.
- Dretske, Fred [1981]. *Knowledge and the Flow of Information*. Cambridge: MIT Press.
- Engel, Pascal [2004]. 'Truth and the Aim of Belief.' In D. Gillies, ed., *Laws and Models in Science*. London: King's College, 77-97.
- Feldman, Richard [2002]. *Epistemology*. Upper Saddle River, NJ: Prentice Hall.
- Goldman, Alvin [1967]. 'A Causal Theory of Knowing.' *The Journal of Philosophy* 64(12): 357-372.
- \_\_\_\_\_ [1988]. *Epistemology and Cognition*. Cambridge: Harvard University Press.
- Hills, Alison [2009]. 'Moral Testimony and Moral Epistemology.' *Ethics* 120(1): 94-127.
- Howard, Nathan [unpublished]. Unpublished paper.
- Hursthouse, Rosalind [1999]. *On Virtue Ethics*. Oxford: Oxford University Press.
- Johnson King, Zoe [unpublished]. Unpublished paper.
- Kamm, Francis [2007]. *Intricate Ethics: Rights, Responsibilities, and Permissible Harm*. Oxford: Oxford University Press.
- Kant, Immanuel [2002]. *Groundwork for the Metaphysics of Morals*. Trans. Arnulf Zweig, Ed. Zweig, Arnulf and Thomas E. Hill, Jr. Oxford: Oxford University Press.
- \_\_\_\_\_ [1999]. *Critique of Pure Reason*. Trans. Paul Guyer and Allan Wood. Cambridge: Cambridge University Press.
- Lehrer, Keith, and Thomas Paxson [1969]. 'Knowledge: Undefeated Justified True Belief.' *Journal of Philosophy* 66: 225-237.
- Mantel, Susanne [2013]. 'Acting for Reasons, Apt Action, and Knowledge.' *Synthese* 190(17): 3865-3888.
- Markovits, Julia [2010]. 'Acting for the Right Reasons.' *Philosophical Review* 119(2): 201-242.
- McDowell, John [1982]. 'Criteria, Defeasibility, and Knowledge.' *Proceedings of the British Academy* 68: 455-479.
- McHugh, Conor, and Daniel Whiting [2014]. 'The Normativity of Belief.' *Analysis* 74(4): 698-713.
- McMahan, Jeff [2002]. *The Ethics of Killing: Problems at the Margins of Life*. Oxford: Oxford University Press.
- Nozick, Robert [1981]. *Philosophical Explanations*. Cambridge: Harvard University Press.
- Parfit, Derek [2011]. *On What Matters: volumes I and 2*. Oxford: Oxford University Press.
- Pollock, John, and Joseph Cruz [1999]. *Contemporary Theories of Knowledge, 2<sup>nd</sup> Edition*. New York: Rowman and Littlefield.
- Scanlon, T.M. [1998]. *What We Owe to Each Other*. Cambridge: Harvard University Press.
- Schroeder, Mark [2015a]. 'In Defense of the Kantian Account of Knowledge: Reply to Whiting.' *Logos and Episteme* 6(3): 371-382.

- \_\_\_\_\_ [2015b]. 'Is Knowledge Normative?' *Philosophical Issues* 25: 379-395.
- \_\_\_\_\_ [2015c]. 'Knowledge is Belief for Sufficient (Objective and Subjective) Reasons.' *Oxford Studies in Epistemology* 5: 226-252.
- \_\_\_\_\_ [unpublished]. 'Getting Perspective on Objective Reasons.' Unpublished paper.
- Shah, Nishi [2003]. 'How Truth Governs Belief.' *Philosophical Review* 112(4): 447-482.
- Sliwa, Paulina [2012]. 'In Defense of Moral Testimony.' *Philosophical Studies* 158(2): 175-195.
- \_\_\_\_\_ [2016]. 'Moral Worth and Moral Knowledge.' *Philosophy and Phenomenological Research* 93(2): 393-418.
- Smithies, Declan [2012]. 'The Normative Role of Knowledge.' *Noûs* 46: 265-288.
- Swain, Marshall [1979]. 'Justification and the Basis of Belief.' In Pappas, George, ed., *Justification and Knowledge*. Dordrecht: D. Reidel Publishing Co.
- Thomson, Judith Jarvis [1990]. *The Realm of Rights*. Cambridge: Harvard University Press.
- Turri, John [2010]. 'On the Relationship between Propositional and Doxastic Justification.' *Philosophy and Phenomenological Research* 80(2): 312-326.
- Wedgwood, Ralph [2002]. 'The Aim of Belief.' *Philosophical Perspectives* 16 (Language and Mind): 267-297.
- Williamson, Timothy [2000]. *Knowledge and Its Limits*. Oxford: Oxford University Press.
- \_\_\_\_\_ [2013]. 'Response to Cohen, Comesaña, Goodman, Nagel, and Weatherson on Gettier Cases in Epistemic Logic.' *Inquiry* 56(1): 77-96.