

# Knowledge, Justification and Normative Coincidence<sup>1</sup>

Martin Smith

University of Glasgow

Say that two goals are *normatively coincident* just in case one cannot aim for one goal without automatically aiming for the other. While knowledge and justification are distinct epistemic goals, with distinct achievement conditions, this paper begins from the suggestion that they are nevertheless normatively coincident – aiming for knowledge and aiming for justification are one and the same activity. A number of surprising consequences follow from this – both specific consequences about how we can ascribe knowledge and justification in lottery cases and more general consequences about the nature of justification and the relationship between justification and evidential probability. Many of these consequences turn out to be at variance with conventional, prevailing views.

## I. TWO EPISTEMIC GOALS

Knowledge and justification are both worthy goals to which we might aspire in forming beliefs. These goals are clearly separate in so far as they have distinct achievement conditions. In particular, one could achieve the latter goal without achieving the former – a belief could qualify as justified without qualifying as knowledge. But, if the guiding supposition of this paper is correct, there is another sense in which these goals cannot be

---

<sup>1</sup> Thanks to audiences at the University of Aberdeen, the University of Edinburgh and Queen's University, Belfast. Particular thanks to Adam Carter, Aaron Cotnoir, Dylan Dodd, Davide Fassio, Andreas Fjellstad, Allan Hazlett, Ben Jarvis, Jesper Kallestrup, Federico Luzzi, Aidan McGlynn, Walter Pedriali and Duncan Pritchard.

separated – they are, to use Crispin Wright’s phrase, *normatively coincident* (Wright, 1992, pp18-19). Put simply, two goals are normatively coincident just in case there is no way to *aim* for one without automatically aiming for the other – though it may be possible to *succeed* at one without succeeding at the other<sup>2</sup>.

Suppose I’m poised at the starting line, about to compete in a race. Here are two worthy goals to which I might aspire: (A) To win the race in the fastest time that I can. (B) To complete the race in the fastest time that I can. Evidently, I could accomplish the latter goal without accomplishing the former – and yet, it would appear as though I cannot *aim* to accomplish the latter goal without aiming to accomplish the former. Whatever things I need to do in order to achieve B (get a good start, don’t go out too hard etc.) these are exactly the things that I need to do in order to achieve A – nothing more and nothing less. If I do end up achieving B without achieving A then this will be entirely due to factors beyond my control – namely, the performance of the other runners. While it is possible for me to achieve B without achieving A, I couldn’t aim to bring this about, even if I wanted to. Such a thing is not within my power.

---

<sup>2</sup> For Wright, this characterisation of normative coincidence is offered as a casual gloss on a more formal definition: ‘Say that two predicates *coincide* in (positive) normative force with respect to a practice just in case each is normative within the practice and reason to suppose that either predicate characterises a move is reason to suppose that the other characterises it too’ (Wright, 1992, pp18) The specifics of this definition very much reflect the overall dialectic in which it is embedded in the opening chapter of *Truth and Objectivity* – and are less of a natural fit with my concerns here. As such, I shall make do with the former characterisation.

Contrast B with the following goal: (C) To complete the race. C, like B, is a more modest or achievable goal than A. Unlike B, though, it seems as though I *could* aim for C without also aiming for A. It's relatively easy to imagine a course of action directed at C that would effectively scupper my chances of achieving A – I could, for instance, take it deliberately easy, set a slow steady pace etc. While B and C are both more modest goals than A, B is normatively coincident with A while C is not. According to the guiding supposition of this paper, the relationship between knowledge and justification is more like the relationship between A and B than it is like the relationship between A and C.

It is, admittedly, somewhat unusual to think about knowledge and justification in these terms. And the claim that knowledge and justification are normatively coincident has never, as far as I'm aware, been made – at least not in so many words<sup>3</sup>. But, once we do conceptualise knowledge and justification as epistemic goals, and the issue of normative coincidence is raised, the suggestion that knowledge and justification are normatively

---

<sup>3</sup> A number of epistemologists have advanced claims in the near vicinity. Reynolds (2002) claims that knowledge and justified belief are *indiscriminable* from the perspective of the believer – as Reynolds puts it 'Justification is what knowledge looks like from inside' (Reynolds, 2002, pp151). This claim has been echoed by others (Smithies, forthcoming, see also Bird, 2007). From this it seems a short step to the conclusion that knowledge and justification are normatively coincident. If I am unable in principle to discriminate the conditions under which two distinct goals are achieved then, plausibly, I am unable to take aim at one without taking aim at both. What of the converse entailment? It is not generally the case that normatively coincident goals are indiscriminable to the aspirant. For instance, I can easily discriminate the conditions under which I win the race in the fastest time I can and the conditions under which I finish the race in the fastest time I can – but the goals are no less normatively coincident for that. There may, however, be special reason to think that this entailment does hold in the case of knowledge and justification, in which case Reynolds' claim and my own may amount to much the same thing. I won't explore this further here.

coincident strikes me as quite an attractive one, at least as a working hypothesis. If we do take this suggestion seriously, then a number of issues in epistemology get cast in a somewhat unfamiliar light. In particular, as I shall argue, this suggestion has surprising consequences for how we ought to ascribe knowledge and justification in lottery cases and concomitant consequences regarding the relationship between justification and evidential probability.

The plan for this paper is as follows: In the next section I will attempt to clarify and develop the claim that knowledge and justification are normatively coincident goals. In section III, I will briefly outline two possible arguments in favour of this claim. These arguments are admittedly sketchy. Although I'm inclined to think that knowledge and justification are indeed normatively coincident I am, in a way, unsure of how best to argue for this. What *can* be clearly argued, though, is that this claim has a number of surprising implications, as suggested above. The implications for the treatment of lottery cases and for the relationship between justification and evidential probability, will be explored in sections IV and V. I'll conclude with some more speculative thoughts about how epistemology might end up looking if we were to take the normative coincidence of knowledge and justification seriously as a starting point.

Before moving on, it's worth making a few further observations about the goals A and B as described above. When thinking about the relationship between these two goals, it's quite natural to regard B as a kind of *fallback* or *consolation* that one could potentially salvage in the event that one fails to achieve A. If I'm simply outclassed in the race, then I could still take some comfort in the fact that I ran as fast as I physically could and got the best

time of which I was capable – better this, surely, than losing as a result of a poor start, poor strategy etc. No doubt I could declare B to be my ultimate objective in running the race and renounce any interest in A – but, given that the purpose of a race is, in effect, to measure the relative performance of the competitors, this may seem a little disingenuous.

While this will play no particular role in what is to come, I'm inclined to think that the analogy between the A/B relationship and the knowledge/justification relationship may hold up, even in these respects. It may be that the role of justification is precisely that of a fallback status that one can potentially salvage if one's belief falls short of knowledge. This is not to deny that one could, in principle, regard justification as one's ultimate epistemic goal and disavow any interest in knowledge – one can take this kind of attitude to a fallback or consolation goal. If the present suggestion is correct, though, then this attitude won't be the natural or default one. In any case, if knowledge and justification are normatively coincident then such an attitude would be largely academic and make no difference as to how one conducts inquiry or forms beliefs. More on this in due course.

On the suggested conception of justification, justified belief is not so much an *ingredient* or constituent of knowledge in the way that epistemologists have conventionally presumed. Rather, it is a kind of *residuum* that can linger when knowledge is blocked or precluded by factors that are beyond one's control or responsibility. This conception of justification may be conducive to a 'knowledge first' approach to epistemology in which the notion of knowledge is used as a resource for the explication of other epistemic notions. On the present view, understanding the value and the nature of justification in terms of the value and nature of knowledge will be the right order of explanation. In any case, as I've said, this

conception of justification won't play any essential role for what follows and I won't be taking it for granted. The claim that knowledge and justification are normatively coincident does not, in and of itself, imply that the former enjoys any particular priority over the latter.

## II. THE NORMATIVE COINCIDENCE CONSTRAINT

To claim that two goals are normatively coincident, as we've seen, is to claim that one cannot aim at one without automatically aiming at the other. This, in turn, implies that a certain kind of project is impossible. If two goals are normatively coincident, then there is no possible project of aiming for a situation in which one of the goals is achieved but the other is not. While normatively coincident goals can come apart, there is no method by which one could *cleave* them apart. If we suppose, as most epistemologists do, that justified belief is necessary for knowledge, it follows at once that one cannot aim for a situation in which one forms a belief that qualifies as knowledge but not as justified – for the situation itself could never obtain. If knowledge and justification are normatively coincident, however, then the reverse project will also be impossible. If knowledge and justification are normatively coincident then it is not possible for one to aim for the (possible) situation in which one forms a belief that qualifies as justified but not as knowledge. Call this the *normative coincidence constraint*.

A few points of clarification are in order. First, it is standard to distinguish between two kinds of justification that a belief might possess – namely, *epistemic* justification and *practical* or *prudential* justification (see, for instance, Jenkins, 2007). The distinction is

typically motivated by examples such as the following: Suppose I'm about to take an aptitude test and firmly believe that I will pass. In fact, my belief is overly optimistic – I've taken several practice tests and I haven't yet managed to achieve a passing mark. My belief is not, then, epistemically justified – I have little or no evidence to back it up. Nevertheless, there is a sense in which my holding the belief may be in my overall best interests – it will enable me to approach the test more calmly and with more confidence and may even increase the chances of my passing. For this reason, the belief might be described as prudentially justified.

It is, of course, epistemic, and not prudential, justification that is my concern in this paper. As such, it is epistemic, and not prudential, justification that is supposed to enter into the normative coincidence constraint. For what it's worth, it seems clear that prudential justification is *not* something that normatively coincides with knowledge. It seems clear that I *could* form beliefs in a way that is aimed at prudential justification and not at knowledge (more on this in the next section). If we accept that epistemic justification is normatively coincident with knowledge, then this could perhaps provide a principled way of distinguishing between justification of the two varieties<sup>4</sup>.

---

<sup>4</sup> While the distinction between epistemic and prudential justification is widely accepted, there is substantial disagreement as to how it should best be analysed. Some have proposed that what sets epistemic justification apart from prudential justification is the fact that it is conditioned by evidence (Conee and Feldman, 1985, Williamson, 2000, section 9.8, Conee, 2004). Others have claimed that what makes justification epistemic is the fact that it is 'truth-linked' in a way that prudential justification is not (Alston, 1985, Lemos, 2007, pp13-14). Still others have suggested that epistemic justification serves to promote a specialised set of 'epistemic' interests, while prudential justification serves to promote one's overall interests (Hazlett, 2006). Characterising epistemic justification in terms of its normative coincidence with knowledge should, I think, be regarded as a

Second, it is also standard to distinguish between two kinds of epistemic justification – namely, *prima facie* justification and *ultima facie* or *all things considered* justification (see, for instance, Pollock, 1970, pp72-78, Alston, 1985, pp104-105 and Señor,1996). Roughly speaking, a belief is taken to be justified *prima facie* if it is based upon evidence that provides adequate support and a belief is taken to be justified *all things considered* if it is, in addition, adequately supported by one's *total* body of evidence. A *prima facie* justified belief will qualify as justified *all things considered*, provided that one does not possess any contrary or otherwise defeating evidence.

Suppose I have an experience as of someone calling my name and come to believe that I've just been called. If this is my only relevant evidence then, plausibly, my belief is justified both *prima facie* and *all things considered*. Now suppose that I acquire some additional evidence that I fail to take into account – suppose I have been undergoing auditory hallucinations and have been experiencing a series of bizarre voices and sounds. In this case, my belief would continue to be justified *prima facie*, but it would no longer be justified *all things considered* – my total body of evidence would include a defeater.

When I have used the term 'justification' thus far, it is *all things considered* justification that I have had in mind. When I have claimed that knowledge and justification

---

further viable alternative. Of course the claim that epistemic justification is normatively coincident with knowledge might be thought contentious – but so too are the claims that epistemic justification is always conditioned by evidence or always truth-linked or always guaranteed by a maximisation of one's epistemic interests.

are normatively coincident, it is all things considered justification that was intended. If I were to believe a proposition on the basis of discredited or defeated evidence then this could be a way of aiming for prima facie justification without aiming for knowledge – this would, in effect, be a cleaving method with respect to the two goals. But this would not, of course, be a way of aiming for all things considered justification and, thus, could not serve to cleave knowledge from it.

The very fact that one could achieve prima facie justification in the way described strongly suggests that it shouldn't be regarded as a worthy epistemic goal in and of itself. It's role, rather, is that of a necessary condition for the attainment of worthy epistemic goals such as all things considered justification and knowledge. For the remainder of the paper I shall continue to use 'justification' to mean all things considered justification, unless noted.

If the normative coincidence constraint is correct then it is impossible for *me* to aim for a situation in which *I* justifiably believe a proposition without knowing it. Another point that should be clarified is that, for all that the constraint says, it may well be possible for me to aim for a situation in which *someone else* justifiably believes a proposition without knowing it. Indeed, this would seem to be possible – telling someone a lie might be a simple way of bringing such a situation about. I can treat others' beliefs as states in the world to be controlled for various ends. The normative coincidence constraint may go some way towards capturing the idea that I cannot treat my *own* beliefs in this way.

This observation does, however, point us towards certain, somewhat contrived, counterexamples to the normative coincidence constraint as stated – for there are certain, somewhat contrived, situations in which one’s attitude towards one’s future self can take on the character of an attitude to another. Suppose I know that I’m about to have my memory wiped and I deliberately plant, for my future self, some evidence that misleadingly supports a proposition I know to be false. This would seem to be a viable way of aiming for a situation in which my future self, suitably estranged from my present self, forms a belief that qualifies as justified but not as knowledge. As I’ve suggested, these sorts of cases are better assimilated to the class of cases in which one attempts to manipulate the beliefs of another – in a way it is incidental that the target of the manipulation happens to be a disconnected future self. In any case, I shall set these cases to one side for the purposes of the present discussion.

Finally, it’s worth pointing out that the normative coincidence constraint does not prevent one from *desiring* to form a belief that qualifies as justified but not as knowledge. As already discussed, it is possible, in general, for one’s desires to differentiate between normatively coincident goals. Furthermore, it may also be possible for one to *intend* to form a belief that qualifies as justified but not as knowledge. As strange as such an intention would seem, the normative coincidence constraint need not prevent one from forming it – but merely from realising or enacting it. If the normative coincidence constraint holds, then there is no way of forming beliefs that is, in actual fact, directed towards justification and not knowledge.

### III. TWO ARGUMENTS

It is not my intention to mount a detailed defence of the normative coincidence constraint here. As I mentioned in the first section, the claim that knowledge and justification are normatively coincident makes for a plausible working hypothesis – and it is largely in this spirit that the claim is adopted. Having said this, though, I will briefly outline two possible motivations for accepting the constraint – one that comes from reflection on the nature of doxastic deliberation and one that comes from considering the epistemic version of Moore’s paradox. By ‘doxastic deliberation’ I simply mean deliberation over what, if anything, to believe about a given issue. Belief formation may, in general, be automatic and unreflective – but doxastic deliberation characterises cases in which it is more self-conscious.

When I’m deliberating, say, over what to eat or what car to buy, one thing that I might do is to consider various different goals and weigh them up with respect to one another. When buying a car, for instance, I could easily have a number of different goals in mind – to buy a powerful car, to buy an environmentally friendly car, to buy a car with good fuel economy, to buy an inexpensive car etc. To reach a decision about which car to buy I may have to figure out just how important or weighty each of these different goals is with respect to the others. Deliberation over what to *believe* however doesn’t seem to be like this. In deliberating over what to believe about a given issue I may evaluate available reasons and evidence, but I don’t evaluate different *goals* – this extra element seems to be missing. If I ask a friend what car I should buy, it would be quite reasonable for him to reply ‘Well, what are you looking for in a car?’ But if I ask a friend what I should believe about a given issue he could hardly respond by asking ‘Well, what are you looking for in a belief?’

What is distinctive about the activity of buying a car is not just that there are a plurality of different goals that one might pursue, it is the fact that these goals can pull us in *different directions* – the fact that they don't normatively coincide. It's for this reason that the evaluation of these different goals can assume a tangible practical importance. Evaluating normatively coincident goals, in contrast, has no practical import. If two goals are normatively coincident then, for the purposes of deliberation, they may as well be treated as one and the same. In particular, it makes no difference how important I take them to be in relation to one another or, indeed, whether I take any view about this at all.

Returning to the race example from the previous section, when crouching at the starting line and turning over possible running strategies in my mind, I don't need to weigh up the relative importance of A and B, but I may need to weigh up the relative importance of A and C – if I take them both seriously as potential goals that is. In general, deliberation over an activity  $\phi$  can take on a goal evaluating layer whenever there are some worthy  $\phi$ -goals that do not normatively coincide. If doxastic deliberation really lacks this extra layer, then this can only be because all the worthy goals of such deliberation *do*. If knowledge and justification are worthy doxastic goals, then they must be normatively coincident as required.

This argument for the normative coincidence constraint proceeds, then, via an ostensibly stronger claim – namely, that *all* worthy epistemic goals normatively coincide. But even if one thought that there were examples of non-normatively coincident epistemic goals, there might still be reason to deny that knowledge and justification constitute an

example. Even if doxastic deliberation could take on the same character as deliberation over what car to buy, we might still deny that the weighing up of *knowledge* and *justification* could ever be a feature of such deliberation. Some epistemologists have posited epistemic goals that are autonomous from knowledge and make different claims and demands upon us. Jonathan Kvanvig, for instance, seems to have such a role in mind for the notion of *understanding* (Kvanvig, 2003, 2009). Kvanvig sometimes writes as though we really would inquire differently if understanding, rather than knowledge, were our ultimate goal (see, for instance Kvanvig, 2009, section 2). If this is right, then the weighing up of understanding and knowledge could become a real practical issue.

In addition, there are, perhaps, some cases of doxastic deliberation in which the practical consequences of holding a certain belief can seem to compete with purely epistemic considerations. Suppose, as in the example from the previous section, that I'm about to face an aptitude test and I have strong evidence to the effect that I won't pass (I've taken a number of practice tests and never achieved a passing mark etc.). Should I believe that I'm not going to pass? I could imagine considering different goals as I agonise over this question – believing that I won't pass is realistic given my evidence and will help me avoid disappointment, but, on the other hand, if I believe that I won't pass before I even sit the test, this will likely adversely affect my performance, so maybe I should keep an open mind or even try to convince myself that I will pass etc. I think it's far from clear how best to understand what is going on in an example like this – but one interpretation, certainly, is that I'm weighing up the relative importance of prudential vs. epistemic justification in deliberating over what to believe.

In any case, I needn't take a view as to whether understanding and/or prudential justification can, in the right circumstances, serve as doxastic goals that compete with goals such as knowledge and epistemic justification. Provided that knowledge and justification cannot compete *with one another*, everything that follows will be on a secure footing. And neither of the preceding considerations give us any reason to doubt this.

The second consideration I shall adduce in favour of the normative coincidence constraint derives from the epistemic version of Moore's paradox which is exemplified by assertions such as the following: 'It's raining but I don't know that it's raining'. There is something absurd – almost self-conflicted – about this assertion, and yet it could perfectly well be true (see Moore, 1962, pp277). It's not essential to this puzzle, though, that the content in question be asserted. After all, if I were to *believe* this, then there would be something absurd and almost self-conflicted about the belief, even if I were never to voice it<sup>5</sup>.

The normative coincidence constraint straightforwardly predicts that no belief of this kind could ever be justified. Suppose I believe that: P but I don't know that P. On pain of contradiction, this Moore's paradoxical belief could never qualify as knowledge. If I know that: P but I don't know that P, it follows that I know that P and I know that I don't know that P. Given the factivity of knowledge, it follows, further, that I know that P and I don't know that P. If there were a method for justifiably holding a Moore's paradoxical belief, this would

---

<sup>5</sup> While most philosophers accept that epistemic Moore's paradoxical assertions and beliefs are always absurd, this has been denied (see, for instance, McGlynn, forthcoming). If one held this view, then the forthcoming considerations would seem to offer little support for the normative coincidence constraint (and might even be thought to constitute an objection to it). I won't pursue this further here.

be a guaranteed method for aiming for justification but not for knowledge. The normative coincidence constraint predicts, then, that there is no method for justifiably forming a Moore's paradoxical belief.

I don't mean to suggest, of course, that this is the *only* possible explanation for the strangeness of Moore's paradoxical beliefs or that this strangeness *couldn't* be explained without the normative coincidence constraint. This is not intended as any sort of decisive argument in favour of the constraint. But the fact that the constraint does offer a fairly simple explanation of the strangeness of Moore's paradoxical beliefs should I think be regarded as a (tentative) point in its favour. Furthermore, this explanation may enjoy a certain advantage over some of the alternative explanations that have been offered.

The strangeness of Moore's paradoxical beliefs is often thought to show that knowledge serves as the *norm* of belief – that belief is governed by the rule that one ought to believe only what one knows (Williamson, 2000, section 11.4, Adler, 2002, Sutton, 2007, Huemer, 2007, Bach, 2008). Importantly, though, such a rule fails to predict that Moore's paradoxical beliefs are unjustified – it predicts merely that one should not hold such beliefs, in the same sense that one should not hold justified beliefs that happen to fall short of knowledge. But drawing attention to a feature of Moore's paradoxical beliefs that they *share* with certain justified beliefs offers no explanation of what is distinctively amiss about them. At the very least, a treatment of Moore's paradoxical beliefs using the knowledge norm of belief requires supplementation in a way that a treatment of Moore's paradoxical beliefs using the normative coincidence constraint does not. I won't explore these issues any further here.

#### IV. LOTTERY CASES

In this section I shall begin exploring some of the consequences of the normative coincidence constraint as I see them. Suppose I hold a single ticket – ticket #542 say – in a fair 1000 ticket lottery with a single guaranteed winner. Suppose the lottery has been drawn and ticket #457 is the winner, but I've not yet heard the result. Both of the following claims are widely held amongst epistemologists:

- (i) I cannot know, purely on the basis of the odds involved, that ticket #542 has lost.
- (ii) I can justifiably believe, purely on the basis of the odds involved, that ticket #542 has lost.

Defenders of (i) include Harman (1968), Dretske (1971), DeRose (1996), Nelkin (2000), Williamson (2000, chap. 11), Hawthorne (2003, partic. chapter 1) and Pritchard (2005, pp162-163, 2007) amongst others and defenders of (ii) include Kyburg (1970), Klein (1985), Foley (1993, chap. ), Lewis (1996), Hawthorne (2003, pp8, 9) and Pritchard (2007) amongst others. Both claims do, undoubtedly, have a certain intuitive appeal. Accepting (ii) in conjunction with other intuitively plausible principles leads, as is well known, to the lottery paradox. Many philosophers, though, have attempted to resolve the paradox in such a way as to hold on to (ii).

If the normative coincidence constraint is correct, however, then (i) and (ii) cannot *both* be maintained. The reason, in a way, is simple: If (i) and (ii) are both true, then believing that ticket #542 has lost, on the basis of the odds involved, would be a way of

aiming for justification but not for knowledge – it would be a ‘cleaving’ method with respect to the two goals. If knowledge and justification are normatively coincident then there can be no such method.

In drawing this conclusion, though, we do need to tread with some care. Cases in which a justified, true belief falls short of knowledge are very familiar and, presumably, their existence poses no general threat to the normative coincidence constraint. Why, then, should it make any difference if we choose to include *lottery* cases amongst their number? The answer, I think, is that there are significant differences between lottery cases and Gettier cases of a traditional sort. Consider, for contrast, the following standard Gettier case (originally described by Skyrms, 1967): Suppose I’m holding a ‘Sure Fire’ brand match, ready to strike it. I know that the match is dry and the environment is normal. Furthermore, I have used any number of Sure Fire matches in the past and they have always lit on the first attempt. I strike the match and form the belief that it’s about to light. In actual fact, the match I’m holding, unlike almost all Sure Fire matches, has a tip with a high proportion of impurities – so much so that it could never be lit just by friction. Nevertheless, the match *is* about to light, due to a coincidental burst of anomalous Q-radiation. In this case my belief that the match is about to light is justified and true and yet, intuitively, I don’t know that the match is about to light.

In both the lottery case and the Gettier case, my belief fails to qualify as knowledge. In the Gettier case, however, this failing is clearly attributable to *extenuating circumstances*. What prevents my belief from qualifying as knowledge, in the Gettier case, is the abnormal condition of the match that I’m holding. Had the match been a normal one, my belief *would*

have qualified as knowledge. Although knowledge and justification come apart in the Gettier case, they are driven apart by circumstances of which I have no awareness and over which I have no control. It is for this reason that the Gettier case supplies no cleaving method. My situation in the Gettier case is one that could *befall me* – but it's not one that I could aim for<sup>6</sup>.

In the lottery case, in contrast, there don't appear to be any extenuating circumstances that prevent my belief from qualifying as knowledge. What would such circumstances be? The lottery was drawn as expected, the proceedings were fair, everything was above board, my ticket did in fact lose etc. How exactly could the situation be made any *more* conducive for the belief that I form? But if the failure of my belief to qualify as knowledge is not attributable to extenuating circumstances of any kind then the buck, as it were, must stop with me and the way that I formed the belief. If we insist that my belief in the lottery case is justified, then we have before us the blueprint for a cleaving method. If knowledge and justification come apart in the lottery case, they are driven apart by *me* and not by circumstance. My situation in the lottery case *is* one that I could aim for.

The lottery case and the Gettier case, then, are significantly different. And, needless to say, there is nothing particularly special about the Sure Fire match case that I chose – any

---

<sup>6</sup> What if I were aware of the impurities in the match but stubbornly continued to believe that the match was about to light? Would this be a way of aiming for the Gettier situation? It would not, since the new information would crucially alter the case – and, indeed, alter it in such a way that it would no longer count as a Gettier case. In the new case my belief that the match is about to light will still fail to qualify as knowledge – but it will also fail to qualify as justified. After all, there is relevant defeating evidence that I neglect. My belief will still qualify as justified *prima facie* but, as already discussed, the existence of a method that cleaves apart knowledge and *prima facie* justification should come as no surprise.

standard Gettier case could have been used to illustrate the point. All standard Gettier cases involve extenuating circumstances of one kind or another that lie beyond the ken and control of the protagonist. As such, the divergence of knowledge and justification in such cases is in no tension with the normative coincidence constraint. If the normative coincidence constraint is correct, though, then our attributions of knowledge and justification in the lottery case must converge – this is simply not the sort of case in which they could come apart. Either I don't know that ticket #542 has lost, and I'm not justified in believing it, or I am justified in believing that ticket #542 has lost, and I know it.

There is, I think, a general and pervasive picture of justification that lies behind the acceptance of (ii). The picture is something like this: Securing justification is a matter of minimising one's risk of error. More precisely, if a belief is based upon evidence that makes it highly likely to be true and it is highly likely to be true given one's total evidence, then this is sufficient for the belief to be justified. This sort of view is widely held amongst epistemologists. Indeed, it seems to be shared even by epistemologists who otherwise disagree quite profoundly about the nature of justification (see, for instance, Russell, 1948, chap. VI, Chisholm, 1957, pp28, Goldman, 1986, section 5.5, Alston, 1988, Fumerton, 1995, pp18-19, Pryor, 2004, pp350-351, 2005, BonJour, 2010)<sup>7</sup> In believing that ticket #542 has

---

<sup>7</sup> Fumerton (1995) portrays the very debate between internalists and externalists about justification as driven by an underlying dispute about the nature of epistemic probability. As he writes: '...the present debate between inferential internalism and inferential externalism was foreshadowed over forty years ago by another controversy about the way to understand the concept of *probability* relevant to epistemology' (Fumerton, 1995, pp109, emphasis mine). For Fumerton, externalism was foreshadowed by the view that epistemic probabilistic relations are contingent and worldly while internalism was foreshadowed by the view that epistemic probabilistic relations are necessary. On this portrayal, both internalism and externalism about justification are

lost the lottery I am clearly running a very low risk of error. My only relevant evidence is that the lottery is fair and has 1000 tickets. The probability that ticket #542 has lost, given this evidence, is 0.999. My risk of error in believing this proposition is only 0.001. If risk minimisation is sufficient for justification, then the belief seems entirely beyond reproach.

There are a number of different ideas about knowledge that could potentially motivate the acceptance of (i). First, and most simply, there is the thought that knowledge requires the complete *elimination* of error risk. On this sort of view, in order to know a proposition P, P must be *certain* given one's evidence (Klein, 1995, Williamson, 2000, chap. 10). Clearly, my evidence in the lottery case doesn't eliminate the possibility that ticket #542 has won. If we combine the claim that risk minimisation is sufficient for justification with the claim that risk elimination is necessary for knowledge, we derive an overall picture on which knowledge and justification are not normatively coincident goals. Lottery cases can be used to make this vivid.

One needn't, though, think that knowledge requires evidential certainty in order to motivate (i) – (i) could equally be motivated by the imposition of a *sensitivity* or *safety* requirement upon knowledge (Dretske, 1971, Nozick, 1981, chap. 3, Williamson, 2000, chap. 5, Sosa, 1999, Pritchard, 2005, chap. 6). Suppose I truly believe that P based upon evidence E. My belief might be described as sensitive iff had P been false then I would not have believed that P based on E. This, in turn, is sometimes cashed out in terms of the following

---

wings of the risk minimisation conception, separated by a disagreement over how the relevant notion of risk is to be understood. This portrayal is, I think, quite faithful to the self conception that many internalists and externalists have.

possible worlds condition: In none of the closest or most similar possible worlds in which P is false do I believe that P based on E. My belief might be described as safe iff I could not easily have falsely believed that P based on E. This, too, can be cashed out in terms of possible worlds: In all close or similar possible worlds in which I believe that P based on E, P is true. The sensitivity and safety conditions offer, I think, slightly different ways of capturing the requirement that a belief be attuned or responsive to the condition that makes it true.

My belief that ticket #542 has lost the lottery, based upon the odds involved, will satisfy neither condition. Consider sensitivity first. If ticket #542 had actually won the lottery then I would still have believed erroneously that it lost on the basis of the odds involved. In the closest worlds in which ticket #542 wins, my belief persists undisturbed. By stipulation, I have not heard anything about the lottery draw. As such, if ticket #542 had won this would not have impinged upon my evidence or beliefs at all.

My belief, then, is not sensitive. Neither is it safe. Ticket #457 may have actually won the lottery, but it could have easily lost – the lottery was not rigged and the ticket was not ordained to win etc. Furthermore, ticket #542 could have just as easily won as any other ticket – there were no special impediments or obstacles to ticket #542 winning. The circumstance in which ticket #542 won, then, could have easily come about and, as we've seen, this would have been a circumstance in which I falsely held my belief. As long as there are close possible worlds in which ticket #457 loses, there must be close possible worlds in which other tickets win. But in this case, given that all the tickets are on a par, there must be

close possible worlds in which each ticket wins – including ticket #542. In these worlds, I falsely believe that ticket #542 lost, on the basis of the odds involved<sup>8</sup>.

If we combine the claim that risk minimisation is sufficient for justification with the claim that either sensitivity or safety is necessary for knowledge, we derive an overall picture on which knowledge and justification are not normatively coincident. Once again, lottery cases can be used to make this vivid. Whether or not the normative coincidence constraint makes any predictions about the nature of justification or the nature of knowledge, it does make predictions about how the two natures must be *coordinated* with each other. With the normative coincidence constraint in place, certain widespread views of justification and of knowledge cannot be combined.

## V. ANOTHER LOOK AT LOTTERY CASES

Consider again verdicts (i) and (ii) about the lottery case described in the previous section:

- (i) I cannot know, purely on the basis of the odds involved, that ticket #542 has lost.

---

<sup>8</sup> Some epistemologists have suggested a watered down safety condition that requires only that P be true in *most* close worlds in which one believes that P on the basis of E. This condition could presumably be met by my lottery belief in which case this alone will not be enough to motivate (i). In *Epistemic Luck* Duncan Pritchard describes such a principle (see Pritchard, 2005, pp156) before opting for a strengthened safety condition that he takes to preclude lottery beliefs from qualifying as knowledge.

- (ii) I can justifiably believe, purely on the basis of the odds involved, that ticket #542 has lost.

Clearly, there are four combined attitudes that one could take towards these claims: One could deny (i) and accept (ii), one could deny (ii) and accept (i), one could accept both or one could deny both. The final option can be ruled out on the grounds that justification is necessary for knowledge. As argued in the last section, the normative coincidence constraint allows us to rule out the penultimate option. This leaves us with two remaining alternatives – either I cannot know that ticket #542 has lost and cannot justifiably believe it or I can justifiably believe that ticket #542 has lost and can know it too.

It's tempting to think that this is where the normative coincidence constraint will leave things. If knowledge and justification are normatively coincident then, as discussed, the lottery case is not the kind of case in which they could diverge. But all that follows from this is that we must be even-handed in the way that we attribute knowledge and justification. We must either be maximally *generous* – granting that lottery beliefs are both justified and constitute knowledge (Hill and Schechter, 2007, Lycan, 2006, footnote 23, Reed, 2010, see also Weatherson, forthcoming, section 4) – or maximally *stingy* – insisting that lottery beliefs fall short of both knowledge and of justification (Ryan, 1996, Nelkin, 2000, Smith, 2010a, Smithies, forthcoming). Presumably, the normative coincidence constraint won't allow us to adjudicate between these two remaining options. It would be surprising if this constraint alone allowed us to draw conclusions about the justificatory status of lottery beliefs (or any beliefs for that matter).

In this final section I shall argue that there is a sense in which this impression is correct and also a sense in which it is wrong. The normative coincidence constraint, as formulated in section II, is crucially ambiguous – it permits of a weaker and a stronger construal. While the weaker construal is indeed neutral between the generous and stingy stances toward lottery beliefs, the stronger construal is not – it weighs against the generous stance, leaving the stingy stance as the last option standing.

In section II, the normative coincidence constraint was stated like this: It is not possible for one to aim for a situation in which one forms a belief that qualifies as justified but not as knowledge. The ambiguity here is simply one of scope – more precisely, it concerns the relative scope of the modal operator and the existential quantifier. Let A be read ‘It is possible for one to aim for a situation in which...’, let J be read ‘One justifiably believes that ...’, let K be read ‘One knows that...’ and let X be a variable ranging over propositions. The normative coincidence constraint could, I suggest, be permissibly formalised in either of the following two ways:

$$(NCC1) \quad \sim\exists X A(JX \wedge \sim KX)$$

$$(NCC2) \quad \sim A(\exists X (JX \wedge \sim KX))$$

According to NCC1, there is no proposition such that one can aim for a situation in which one justifiably believes it but does not know it. According to NCC2, one cannot aim for a situation in which there is some proposition that one justifiably believes but doesn’t know. We might say that NCC1 forbids us from aiming *de re* to form a justified belief that is not knowledge, while NCC2 forbids us from aiming for this outcome *de dicto*.

NCC2 is, I suspect, the more natural formalisation of the informal constraint – but it is important to distinguish NCC1, as it is the weaker, less committal, of the two. If there is some proposition that I can aim to justifiably believe without knowing, then I can clearly aim for a situation in which I justifiably believe some proposition without knowing it. But the converse does not hold. Just because I can aim for a situation in which I justifiably believe some proposition without knowing it, it doesn't follow that there is some proposition that I can aim to justifiably believe without knowing.

In section III, I sketched two possible motivations for accepting the normative coincidence constraint – one that drew upon the nature of doxastic deliberation and one that drew upon the epistemic version of Moore's paradox. The support provided by the second of these motivations would seem to extend only as far as NCC1. Suppose it is possible for one to aim to justifiably hold a Moore's paradoxical belief –  $A(J(P \wedge \sim KP))$ . Since  $\sim K(P \wedge \sim KP)$  is derivable just from the fact that  $K$  distributes over conjunction and is factive, it plausibly follows that  $A(J(P \wedge \sim KP) \wedge \sim K(P \wedge \sim KP))$ . This conflicts with NCC1 (as well as NCC2). NCC1, then, is enough to generate the prediction that one cannot justifiably hold a Moore's paradoxical belief.

The first motivation for the normative coincidence constraint, however, would appear to apply equally to NCC1 and NCC2. The motivation, in effect, was this: If there are worthy doxastic goals that do not normatively coincide, then we would expect doxastic deliberation to incorporate an element of goal evaluation. But such an element seems conspicuously absent. Suppose that, in violation of NCC1, I have a method of aiming for a situation in which I justifiably believe a particular proposition but don't know it. Deliberation about

whether to implement such a method could clearly implicate deliberation over the relative importance of knowledge and justification.

Now suppose instead that I have a method of aiming for a situation in which some *unspecified* proposition is justifiably believed but not known. The existence of such a method is in violation of NCC2 but is perfectly consistent with NCC1. And yet this method, no less than the former, sacrifices knowledge for the sake of justification. Deliberation over this method, just like deliberation over the former, could naturally lead to deliberation about the relative importance of the two goals. And this, according to the suggestion made in section III, is a place to which epistemic deliberation never leads.

While one of the arguments that I've offered in favour of the normative coincidence constraint may bear the weight of NCC2, NCC1 is enough to support the consequences that I've drawn from the normative coincidence constraint thus far. In particular, NCC1 will be enough to block the combination of claims (i) and (ii). If (i) and (ii) hold then there will be a proposition – namely, the proposition that ticket #542 has lost – that I can aim to believe without knowing. More formally, if L is the proposition that ticket #542 has lost, we have it that  $A(JL \wedge \sim KL)$  – which is inconsistent with both NCC1 and NCC2. But this is as far as NCC1 will take us – it holds no further lessons as to how the lottery case should be described. But, as suggested above, there is more yet that might be extracted from NCC2.

Suppose I am justified in believing that ticket #542 has lost the lottery, based purely on the odds involved. Given that every ticket has an equally slim chance of winning, I would

presumably be justified in believing the same thing about each of them. This might be thought of as a kind of *symmetry assumption*. Such an assumption is familiar from certain formulations of the lottery paradox<sup>9</sup>. But, however generous we wish to be in appraising these 1000 beliefs, it's clear that they could not *all* constitute knowledge. One of these beliefs, after all, would have to be false. Given that knowledge is factive, the maximum number of these beliefs that could possibly constitute knowledge will be 999.

By forming these 1000 beliefs I would, in effect, be guaranteed to form some belief that is justified but falls short of knowledge – naturally, though, I could not specify in advance which belief this would be. While this does not describe a method of justifiably believing any *particular* proposition without knowing it, it does describe a method for justifiably believing *some proposition or other* without knowing it. More formally, if  $L_n$  is the proposition that ticket #n has lost, we have it that  $A(JL_1 \wedge \dots \wedge JL_{1000} \wedge \sim KL_1 \vee \dots \vee \sim KL_{1000})$  which is inconsistent with NCC2, though consistent with NCC1. If I can aim for a situation in which each of a range of propositions is justifiably believed, but some amongst their number are not known, then I can aim for a situation in which I justifiably believe some proposition without knowing it. If NCC2 is to be maintained, our original assumption must

---

<sup>9</sup> Some – such as Harman (1986, pp70-72) – have attempted to resolve the lottery paradox by questioning this assumption. On the picture that Harman defends, I can justifiably believe that ticket #542 has lost and that ticket #712 has lost etc. provided I don't also believe this about too many other tickets. More precisely, I can justifiably believe, of each of a certain *number* of tickets, that it has lost – but any lottery beliefs that I form in excess of this quota will not be justified. On Harman's view, the *order* in which I form lottery beliefs can, then, have an impact on their justificatory status. This picture, I suspect, is also in conflict with NCC2 – but I won't argue this here.

be rejected – I am not justified in believing, on the basis of the odds involved, that ticket #542 has lost.

In order to elicit the result that lottery beliefs are not justified, NCC1 must be combined with certain substantial claims about knowledge – such as the claims that knowledge requires safety or sensitivity or evidential certainty. NCC2, however, seems capable of generating this conclusion almost by itself. The only assumption about knowledge that is needed is the assumption of factivity – that one can only know truths. And the only other assumption used is the symmetry assumption flagged above – if one can justifiably believe, of a given ticket, that it will lose the lottery on the basis of the odds involved, then one can justifiably believe the same about each ticket.

As discussed in the last section, the claim that lottery beliefs are justified is underwritten by a certain widespread conception of justification – namely, the risk minimisation conception. As such, while NCC1 weighs against combining the risk minimisation conception of justification with certain requirements on knowledge – such as safety, sensitivity or evidential certainty – NCC2, surprisingly enough, weighs against the risk minimisation conception all by itself. As the lottery case illustrates, the interests of risk minimisation can, if the setting is right, be well served by believing each of a set of propositions that is known to be inconsistent. In general, if my evidence suggests that some member of a set of propositions is false but a high proportion of its members are true, then this is quite compatible with each member having a high evidential probability. Indeed, if this is my only relevant evidence about the propositions then, as the proportion of truths tends towards one, so too will the individual evidential probabilities. If my epistemic goal is to

ensure that my beliefs, taken individually, are very likely to be true, then I could scarcely do better than to believe every member of the set.

Clearly, though, by believing every member of such a set I effectively ensure that some of the beliefs I form will be false and thus will fall short of knowledge. Knowledge and high evidential probability are not normatively coincident – at least not in the sense captured by NCC2. It is possible to aim *de dicto* for a situation in which I believe a proposition with a high evidential probability without knowing it. If high evidential probability is sufficient for justification, as the risk minimisation conception suggests, then it must be possible to aim *de dicto* for a situation in which I justifiably believe a proposition without knowing it.

If NCC2 is accepted, the risk minimisation conception of justification must be rejected. But where exactly does the rejection of this entrenched picture leave us? If it's wrong to think about justification in terms of risk minimisation or high evidential probability, then how exactly *should* we think about it? Speculation about these questions is, for the most part, beyond the scope of this paper. Suffice it to say, though, the 'consolation prize' conception of justification briefly sketched in the first section may make for the beginnings of an alternative picture. As it stands, of course, this conception is little more than a collection of somewhat suggestive slogans. And this contrasts sharply with my formulation of the risk minimisation conception as a relatively precise thesis. Until the commitments of the former conception are stated with more precision, any comparison between the two is bound to be somewhat unfair. I won't pursue this further here.

It is also important to emphasise that NCC2 does not force us to sever *any* link between justification and evidential probability. While NCC2 is in conflict with the claim that high evidential probability is *sufficient* for justification, it is perfectly compatible with the claim that high evidential probability is *necessary* for justification. It could well be that NCC2 allows us to salvage something recognisably similar to the original risk minimisation picture – a kind of refinement of the picture rather than something completely different<sup>10</sup>.

In this paper I have experimented with the normative coincidence constraint as a possible starting point for theorising about knowledge and justification. I've attempted to demonstrate just how much can be wrung from this seemingly innocuous principle – including general conclusions about the nature of justification and of knowledge as well as particular conclusions about the justificatory and knowledge status of test-case beliefs. Relatedly, I've attempted to demonstrate just how far this principle can take us from certain sorts of conventional thinking. If I'm right, then the normative coincidence constraint forces us to abandon a well entrenched picture of justification – a picture on which justification is linked to the minimisation of error risk. However, just because the normative coincidence constraint offers a new perspective on a range of familiar topics it does not follow, of course,

---

<sup>10</sup> I'm inclined to think that no condition defined exclusively in terms of evidential probabilities could suffice for justification compatibly with NCC2 – so refinement in this direction, at the very least, is closed off. If C is a condition defined exclusively in terms of evidential probability and satisfied by some propositions that are less than evidentially certain then, given appropriate structural constraints, it can be shown that C can be simultaneously satisfied by each member of an inconsistent set of propositions (see Douven and Williamson, 2006, Smith, 2010b). In this case, it will be possible to aim de dicto to believe, but not know, a proposition that satisfies C.

that it's correct or ought to be accepted. And, in spite of a somewhat provisional motivation offered in its favour, this is admittedly an issue that is still to be fully addressed.

## References

- Adler, J. (2002) *Belief's Own Ethics* (Cambridge, MA: MIT Press)
- Alston, W. (1985) 'Concepts of epistemic justification' *The Monist* v68, pp57-89
- Alston, W. (1988) 'An internalist externalism' *Synthese* v74(3), pp265-283
- Bach, K. (2008) 'Applying pragmatics to epistemology' *Philosophical Issues* v18 pp69-88
- Bird, A. (2007) 'Justified judging' *Philosophy and Phenomenological Research* v74(1), pp81-110
- BonJour, L. (2010) 'The myth of knowledge' *Philosophical Perspectives* v24(1), pp57-83
- Chisholm, R. (1957) *Perceiving* (Ithaca, NY: Cornell University Press)
- Conee, E. and Feldman, R. (1985) 'Evidentialism' *Philosophical Studies* v48, pp15-35
- Conee, E. (2004) 'The truth connection' in Conee, E. and Feldman, R. eds. *Evidentialism: Essays in Epistemology* (Oxford: Oxford University Press)
- DeRose, K. (1996) 'Knowledge, assertion and lotteries' *Australasian Journal of Philosophy* v74, pp568-58
- Douven, I. and Williamson, T. (2006) 'Generalising the lottery paradox' *British Journal for the Philosophy of Science* v57(4), pp755-779
- Dretske, F. (1971) 'Conclusive reasons' *Australasian Journal of Philosophy* v49(1), pp1-22

- Foley, R. (1993) *Working Without a Net* (Oxford: Oxford University Press)
- Fumerton, R. (1995) *Metaepistemology and Scepticism* (London: Rowman and Littlefield)
- Goldman, A. (1986) *Epistemology and Cognition* (Cambridge, MA: Harvard University Press)
- Harman, G. (1968) 'Knowledge, inference and explanation' *American Philosophical Quarterly* v5(3), pp164-173
- Harman, G. (1986) *Change in View* (Cambridge, MA: MIT Press)
- Hawthorne, J. (2003) *Knowledge and Lotteries* (Oxford: Oxford University Press)
- Hazlett, A. (2006) 'How to defeat belief in the external world' *Pacific Philosophical Quarterly* v87, pp198-212
- Hill, C. and Schechter, J. (2007) 'Hawthorne's lottery puzzle and the nature of belief' *Philosophical Issues* v17(1), pp102-122
- Huemer, M. (2007) 'Moore's paradox and the norm of belief' in Nuccetelli, S. And Seay, G. eds. *Themes from G. E. Moore: New Essays in Epistemology and Ethics* (Oxford: Oxford University Press)
- Jenkins, C. (2007) 'Entitlement and rationality' *Synthese* v157(1), pp25-45
- Klein, P. (1985) 'The virtues of inconsistency' *The Monist* v68, pp105-135
- Klein, P. (1995) 'Scepticism and closure: Why the evil genius argument fails' *Philosophical Topics* v23, pp213-236
- Kvanvig, J. (2003) *The Value of Knowledge and the Pursuit of Understanding* (Cambridge: Cambridge University press)

- Kvanvig, J. (2009) 'The value of understanding' in Haddock, A., Millar, A. and Pritchard, D. *Epistemic Value* (Oxford: Oxford University Press)
- Kyburg, H. (1970) 'Conjunctivitis' in Swain, M. ed. *Induction, Acceptance and Rational Belief* (Dordrecht: Reidel)
- Lemos, N. (2007) *An Introduction to the Theory of Knowledge* (Cambridge: Cambridge University Press)
- Lewis, D. (1996) 'Elusive knowledge' *Australasian Journal of Philosophy* v74(4), pp549-567
- Lycan, W. 'On the Gettier problem problem' in Hetherington, S. ed. *Epistemology Futures* (Oxford: Oxford University Press)
- McGlynn, A. (forthcoming) 'Believing things unknown' *Noûs*
- Moore, G. (1962) *Commonplace Book 1913-1953* (London: Allen and Unwin)
- Nelkin, D. (2000) 'The lottery paradox, knowledge and rationality' *Philosophical Review* v109, pp373-409
- Pollock, J. (1970) 'The structure of epistemic justification' in Rescher, N. ed. *Studies in the Theory of Knowledge* (Oxford: Basil Blackwell)
- Pritchard, D. (2005) *Epistemic Luck* (Oxford: Clarendon Press)
- Pritchard, D. (2007) 'Sensitivity, safety and anti-luck epistemology' in Greco, J. ed. *The Oxford Handbook of Scepticism* (Oxford: Oxford University Press)
- Pryor, J. (2004) 'What's wrong with Moore's argument?' *Philosophical Issues* 14: *Epistemology*

- Pryor, J. (2005) 'Is there immediate justification?' in Sosa, E. and Steup, M. eds. *Contemporary Debates in Epistemology* (Oxford: Blackwell)
- Reed, B. (2010) 'A defence of stable invariantism' *Noûs* v44(2), pp224-244
- Reynolds, S. (2002) 'Testimony, knowledge and epistemic goals' *Philosophical Studies* v110, pp139-161
- Russell, B. (1948) *Human Knowledge: Its Scope and Limits* (London: Unwin)
- Ryan, S. (1996) 'The epistemic virtues of consistency' *Synthese* v109, pp121-141
- Señor, T. (1996) 'The prima/ultima facie justification distinction in epistemology' *Philosophy and Phenomenological Research* v56, pp551-566
- Skyrms, B. (1967) 'The explication of 'X knows that p'' *Journal of Philosophy* v64, pp373-389
- Smith, M. (2010a) 'What else justification could be' *Noûs* v44(1), pp10-31
- Smith, M. (2010b) 'A generalised lottery paradox for infinite probability spaces' *British Journal for the Philosophy of Science* v61(4), pp821-831
- Smithies, D (forthcoming) 'The normative role of knowledge' *Noûs*
- Sosa, E. (1999) 'How to defeat opposition to Moore' in Tomberlin, J. ed. *Philosophical Perspectives* v13 (Atascadero: Ridgeview)
- Vogel, J. (1990) 'Are there counterexamples to the closure principle?' in M. Roth and G. Ross (eds.), *Doubting: Contemporary Perspectives on Skepticism* (Dordrecht: Kluwer)

Weatherson, B. (forthcoming) 'Probability and scepticism' in Dodd, D. and Zardini, E. eds.

*Contemporary Perspectives on Perceptual Justification and Scepticism* (Oxford:

Oxford University Press)

Williamson, T. (2000) *Knowledge and Its Limits* (Oxford: Oxford University Press)

Wright, C. (1992) *Truth and Objectivity* (Cambridge, MA: Harvard University Press)