

Reliabilism

Juan Comesaña

1 Introduction

The reliabilist theory of epistemic justification is one of the most influential theories in recent epistemology. In this chapter I present reliabilism, starting with its origins as a theory of knowledge and considering the main problems that it faces, together with some of the main solutions to those problems that have been proposed in the literature.

2 A Brief History of Reliabilism

In a short note that didn't attract much attention until *after* reliabilist theories were explicitly defended in the 1970's, Ramsey (1931) wrote:

I have always said that a belief was knowledge if it was (i) true, (ii) certain, (iii) obtained by a reliable process.

We can recognize here perhaps a version of the justified true belief account of knowledge, except that Ramsey talks about the reliability of the process that produced the belief rather than about justification (although it may be that Ramsey is combining the belief and justification conditions in his (ii), and thus appealing to reliability *in addition* to appealing to justification). The next big moment in our brief history involves Edmund Gettier. After Gettier published his counterexamples to the justified true belief account of knowledge (Gettier (1963)) many philosophers embarked in a search for a fourth condition that, added to the traditional three, was to be sufficient for knowledge. The history of that search is well-known (see Shope (1983)). But other philosophers took a different tack: instead of looking for a fourth condition to *add* to the justification condition, they examined the possibility of *replacing* the justification condition. What is missing in subjects in Gettier-style situations, these philosophers thought, is some connection between the belief and the facts that make the belief true. If that connection is present, moreover, then the subject has knowledge, regardless of whether or not

his belief is “justified” as that notion is traditionally understood. Thus, Armstrong (1973) proposed a theory according to which a non-inferential belief amounts to knowledge just in case it has some properties that are nomically connected to its truth. In a similar vein, Goldman (1967) proposed a causal theory of knowledge, according to which (very roughly) a belief constitutes knowledge just in case the fact that p is causally connected (in the right way) with the belief that p . An interesting variant of these theories construes the relation in question not directly as a physical relation, but rather as a certain kind of modal correlation between the belief and the facts. For instance, Dretske (1971) anticipated Nozick (1981) in proposing that a crucial condition for a subject S to know that p is that, if p were true, S wouldn’t believe that p . Other proposals along these lines include Sosa (1999) and Williamson (2000) (for critical examination of these proposals, see Comesaña (2007)).

As said above, what all these proposals have in common is that they jettison the justification condition in favor of a condition that requires that there be some kind of relation between belief and reality. Reliabilism as a theory of epistemic justification is born when that *replacement* thesis is abandoned in favor of an *explanation* thesis—when reliability is seen as what justification *consists in*, as opposed to a condition that is independent of justification.

3 An Initial Formulation of Reliabilism About Justification

That change takes place in the seminal work on reliabilism as a theory of justification: Goldman (1979). In that work, Goldman proposes a recursive definition of justification along the following lines:¹

S is justified in believing that p at t if and only if:

- (i) S ’s belief that p at t results from a belief-independent cognitive belief-forming process that is unconditionally reliable; or
- (ii) S ’s belief that p at t results from a belief-dependent cognitive belief-forming process that is conditionally reliable and all the input-beliefs are justified.

To understand the definition we need to explain what Goldman means by a cognitive belief-forming process (simply “process” in what follows), what the distinction

¹See section 4.2 for an important additional condition added by Goldman (commonly referred to as a “no defeaters” condition).

between belief-dependent and belief-independent processes amounts to, and what it is for a process to be conditionally and unconditionally reliable.

A process is a psychological mechanism that produces beliefs under certain conditions. A process can be represented by a *function* whose outputs are the propositions produced by the mechanism and whose inputs are the conditions under which the mechanism operates (we can, without fear of generating confusion, talk about the inputs and outputs of the processes themselves). Belief-dependent processes will include among its inputs other beliefs of the agent, whereas belief-independent processes will not. A process is unconditionally reliable if and only if it issues in mostly true beliefs, whereas a belief-dependent process is conditionally reliable if and only if it issues in mostly true beliefs if the input beliefs are true.

4 Problems and Refinements

4.1 Justification Without Reliability? The New Evil Demon Problem

Consider the Cartesian skeptical scenario according to which there is a powerful evil demon who is deceiving me with respect to most of the propositions that I believe. According to the traditional description of this skeptical scenario, what the demon does is to arrange things so that the world appears to me to be exactly the same way that it would appear to me if all of my beliefs about the world were true, when they are in fact false. Thus, in the skeptical scenario I am not typing on a computer, but my total experience is still the same as it is right now, and so on and so forth for most of my other beliefs.

Philosophers disagree regarding whether the possibility of skeptical scenarios like those can form the basis for any powerful skeptical argument. But skepticism is not our concern right now. Consider not *our* epistemic predicament, but that of a victim of a skeptical scenario. Regardless of what we think about ourselves, such a victim of course lacks knowledge (if only because most of his beliefs are false, but not only because of that). But ask yourself not whether the victim has knowledge, but whether he has *justified beliefs*. The answer will depend, of course, on the particular epistemic nature of the victim in question. But if the victim is undergoing an experience just like mine right now and, on the basis of that experience, believes that he is typing on a computer, then surely at least that belief of him will be justified. In general, if the victim is a responsible epistemic agent (that is, if he doesn't commit fallacies, takes experience at face value, etc.), then his beliefs will be just as justified as those of his non-victimized counterpart. This is not to say, of course, that there cannot be victims that are unjustified, but it is to say that the mere fact that they are the victims of a skeptical scenario doesn't guarantee that they are unjustified.

So victims of a skeptical scenario may very well have justified beliefs. But those beliefs are, for the most part, false. It seems to follow that the processes that formed those beliefs were not reliable—they certainly didn't tend to produce mostly true beliefs. If those processes are unreliable, then reliabilism has the consequence that those victims can never have justified beliefs, irrespective of the details of how they came to form those beliefs. This is what Sosa (1991) called the “new evil demon” problem. It's *new* because the problem is not Descartes' skeptical argument, but rather the apparent consequence of reliabilism that no victim of skeptical scenarios can ever be justified.

The new evil demon problem has been associated with Cohen (1984), and it has been repeatedly wielded as perhaps the main symptom of everything that's wrong with reliabilism—see, for instance, Pollock (1984), Feldman (1995) and Foley (1985). For friends of reliabilism, the new evil demon problem has proven to be an important reason to modify their theory. Goldman himself, for instance, has reacted to the problem in no less than three different ways (to be considered below).

Sosa (1991) makes an important point in discussing the new evil demon problem. Victims of an evil demon live in worlds where the processes that we use to form our beliefs seem to be unreliable, and they have justified beliefs despite using *those very same processes*. But there may well be creatures who inhabit worlds where the only way to acquire information about the environment is to use processes that would be unreliable if used here (one may argue that the evil demon world is precisely one of those worlds). For instance, there are worlds where when things look round they are square. In those worlds, believing that there is a round thing in front of you whenever something appears round tends to produce false beliefs, while believing that there is a square thing in front of you when it appears round tends to produce true beliefs. Some philosophers (including Sosa (1991) and Bergmann (2006)) think that subjects in those worlds whose beliefs are produced by processes like the latter do have justified beliefs. Some solutions to the new evil demon problem do not have this consequence. Let us now review some of those solutions.

Goldman (1979) actually considered a benevolent-demon world example which is similar in some respects to the kind of cases that Sosa urged us to consider alongside the new evil demon problem. The benevolent-demon arranges things so that wishful-thinking is reliable. What are we to say, then, of a subject who engages in wishful thinking in the benevolent-demon world? Goldman (1979) was ambivalent. One possibility that he considered was that the wishful thinker was not justified, because the reliability of the process in question should be evaluated, not with respect to the world of the wishful thinker, but rather with respect to the actual world. Given that wishful thinking is not actually reliable, the beliefs of the wish-

ful thinker are not justified even if they are produced by processes that are reliable in his world. (Goldman considered that idea not as a theory of justification, but as a theory of justification-attributions. I ignore the complications that this difference raises.) Now, bracketing for the moment any concerns about the consequences that this theory has for the wishful thinker and similar cases, notice that it has the right consequence for the new evil demon problem. For the processes used by those victims of the evil demon that have justified beliefs are actually reliable (or, at least, the objection does not rely on their not being so), and so a theory according to which a belief is justified provided that it was produced by processes that are actually reliable will not succumb to the new evil demon problem.

But let us now go back to the wishful thinker and cases similar to his. Because of obfuscating connotations of wishful thinking, let us concentrate instead on the case of inhabitants of a world where, due to fundamental physical and physiological differences, things that appear round are square (and, maybe, vice-versa). Inhabitants of this world are wired-up so that they believe that there is a square whenever something appears round to them. This appearance-belief transition is for them as basic as the usual transition is for us, and nothing in their other experiences or collective inquiries ever gives them (what would of course be misleading) reasons to think that there is anything wrong with proceeding in this way. Let us also suppose that proceeding like this is the only way for them to know about real squares—if they instantiated instead our usual appearance to belief transitions, then nothing in their other experiences or collective inquiries will ever give them a clue that they are making a fundamental mistake. What are we to say about such creatures? If we insist that justification is proportional to the reliability of the process that produces the belief in the actual world, then we are bound to say that their beliefs are unjustified—but, as mentioned before, many epistemologists believe that there is at least some sense in which they are justified. Call this the “alien cognizers” problem for reliabilism (the name comes from Bergmann (2006)). If, on the other hand, we say that justification is proportional to the reliability of the process that produces the belief in whichever world that belief is held, then we are bound to say that the beliefs of any victim of an evil demon are unjustified. The new evil demon problem and the alien cognizer problem pull reliabilism in opposite directions. We’ll come back to this issue.

Goldman (1986) proposes a different solution to the new evil demon problem, the “normal worlds” approach (in that book Goldman also reconsiders and rejects the actual world proposal). According to this proposal, the reliability of a process should be evaluated with respect to *neither* the actual world *nor* the world where the belief is held, but rather with respect to *normal worlds*, which are worlds that are, in their general features, the way we take the actual world to be. This proposal solves the new evil demon problem insofar as we take the actual world to be one

where the processes used by justified victims are reliable. But it does nothing to alleviate the doubts of friends of the alien cognizer problem. And it also has more serious problems. It seems to offer a radically subjectivist account of justification, for the justificatory status of beliefs is in effect determined by our general beliefs about the world.

Goldman (1988) proposed to solve the new evil demon problem by positing an ambiguity in “justified”. In one sense of the word (the “strong” sense), justification answers to the original reliabilist definition in Goldman (1979), but in another, weak sense, a belief is justified just in case it is non-culpably held. The beliefs of victims of an evil demon can never be strongly justified, but that may well be weakly justified. As Sosa (1991) points out, however, this solution doesn’t distinguish between justified victims of an evil demon and those who have been brain-washed into holding preposterous beliefs. None of them are culpable in holding in their beliefs, but justified victims of an evil demon are more than non-culpable: there is something epistemically positive about how they formed their beliefs, not merely the lack of something epistemically negative.

Elsewhere (Comesaña (2002), building on work by Sosa—Sosa (1993) and Sosa (2008)), I advanced a proposal that aims at taking care of both the new evil demon problem and the alien cognizers problem. Let us first note that, according to many, “actual” is a term whose proper semantics is two-dimensional—thus, to put it in terminology that comes from Chalmers (1996), associated with the term “actual” there is a two-dimensional intension which can be thought of as a function from possible worlds to functions from possible worlds to truth-values. Associated with any utterance of “It is actually true that p ” there is a primary intension, which determines a proposition which is true at a world w just in case p obtains in w , and a secondary intension, which determines a proposition which is true in a world w just in case p obtains in the world of utterance. So, for instance, the primary intension of “It is actually the case that Mount Everest is the mountain with the highest altitude in the world” determines a necessarily true proposition, but its secondary intension determines a contingently true proposition. According to my proposal, the reliability of a process should indeed be evaluated in the actual world, as the position that Goldman (1979) considered suggests, but bearing in mind that “actual” must be given a two-dimensional semantics. Thus, there is a sense in which victims of an evil demon can be justified, although it is also true that, on this view, there is a sense in which they can never be (this captures the sense in which victims of an evil demon cannot have knowledge even if their justified beliefs are true). This view also takes care of the alien cognizers problem, insofar as it has the consequence that there is a sense in which they can be justified (their beliefs may be arrived at by processes which are reliable in their world, and, thus, they can have knowledge), even though it is also true that there is a sense in which they

are not justified (the sense in which victims of an evil demon and ourselves are). These two senses associated with “justification,” however, are not independent of each other (like Goldman’s weak and strong notions of justification), but rather are related to each other in the same way in which, following Kaplan (1989), we can say that the character of an indexical is related to its content.

4.2 Reliability Without Justification?

The new evil demon problem challenges the necessity of reliability for justification. An equally influential problem challenges the sufficiency of reliability for justification. The *locus classicus* for that challenge is BonJour (1980). In that paper, BonJour presents a number of cases involving subjects whose beliefs are formed by processes that, although actually unreliable, are stipulated to be reliable in the subject’s world (so far, this is just like the case of the alien cognizers, but we’ll see that there is a crucial difference). In some cases, the subjects in question have reason to believe that the processes that produce their beliefs are unreliable. BonJour’s most famous case, however, concerns a subject that is not like that. Here is BonJour’s description of the case:

Norman, under certain conditions that usually obtain, is a completely reliable clairvoyant with respect to certain kinds of subject matter. He possesses no evidence or reasons of any kind for or against the general possibility of such a cognitive power, or for or against the thesis that he possess it. One day Norman comes to believe that the President is in New York City, though he has no evidence either for or against this belief. In fact, the belief is true and results from clairvoyant power, under circumstances in which it is completely reliable.

Goldman (1979) had already considered the problems raised by subjects who have reasons to believe that the processes that produced their beliefs are unjustified, and proposed to add a clause to the reliabilist definition of justification to take care of those problems. The proposed clause adds as a necessary condition for a subject’s being justified in believing that p that there be no reliable process available to the subject such that, if the subject were to use that process (in addition to the ones actually used), then he would no longer believe that p . In Goldman (1986) the analogous condition is added that the subject not have a reason to think that his belief is unreliably caused.

The addition of a negative clause may help with the cases where the subjects have reasons to believe that the target belief was unreliably caused. But BonJour’s idea is that the case of Norman is not like this, for Norman does not have any

reason to believe that his belief was unreliably caused. Indeed, BonJour explicitly says that Norman has no evidence one way or the other regarding the reliability of clairvoyance, or even whether he is clairvoyant. Now, this claim by BonJour may well be challenged: after all, Norman suddenly finds himself believing that the President is in New York City, with no clue as to where the belief came from. Presumably, there are reasons to doubt that beliefs whose provenance is completely obscure to us were caused by reliable processes.

So it's not entirely clear that Goldman's original negative clause cannot take care of Norman's case as well. But there is still reason to think that some other modification to reliabilism is needed. For, even if the fact (if it is indeed a fact) that Norman has reasons to believe that his belief about the President was not produced by a reliable process is *one* of the things that is epistemically wrong with Norman, it may well be that there is more than one thing that is epistemically wrong with Norman. Indeed, it is arguably the fact that the clairvoyantly produced beliefs *just pop in Norman's head* that makes them unjustified (this is the crucial difference with the case of the alien cognizers). Elsewhere Comesaña (forthcoming) I have argued that the way to take care of BonJour's objection to reliabilism (and to externalist theories more generally) is to combine reliabilism with evidentialism. Evidentialism, as defended by Earl Conee and Richard Feldman (see, for instance, Conee and Feldman (2004)), is the thesis that adopting a doxastic attitude (for instance, belief) with respect to a proposition is justified for a subject at a time if and only if adopting that attitude fits the evidence that that subject has at that time. If evidentialism is left at that, however, it is a radically incomplete theory, for we would naturally want to know what makes it the case that adopting a certain attitude towards a proposition *fits* a certain body of evidence. I suggest that we should adopt a reliabilist understanding of this notion of fit. The resulting theory looks like this:

S is justified in believing that *p* at *t* if and only if *p* was produced by a process *R* which includes some evidence *e* and:

- (i) *e* doesn't include any beliefs of *S* and *R* is unconditionally reliable;
- or
- (ii) *e* includes beliefs of *S*, all of these beliefs are justified, and *R* is conditionally reliable.

This version of reliabilism has the consequence that Norman's beliefs are not justified, because they are produced by processes that don't include any evidence. If what is motivating the idea that Norman is not justified is the fact that the clairvoyant beliefs *just pop into his head* (and not just that Norman has reasons to believe

that his clairvoyantly acquired beliefs are not justified), then this version of reliabilism takes care of that problem.

4.3 What Is Reliability? The Generality Problem

Only *types* of processes can straightforwardly be said to have a meaningful reliability ratio associated with them. *Token* processes issue in only one belief. If the belief is true, then that token process is perfectly reliable, whereas if it is false, the token process is perfectly *unreliable*. A process-type, by contrast, can be more or less reliable depending on the ratio of true to false beliefs that all the tokens of that type produce (this assumes that we are taking a truth-ratio approach to reliability, as is standardly done in the literature, but see section 4.4 for a problem for this approach, as well as for a suggested solution). When reliabilists say that the justificatory status of a belief depends, ultimately, on the reliability of the processes that led to that belief, therefore, they must mean to refer to process-types. But *any token-process belongs to indefinitely many types of processes, each one of which may have a different reliability ratio*. For instance, the process that created my belief that I am facing a computer screen belongs to all of these types: visual process, visual process while I am not using glasses, visual process while I am not using or needing glasses, perceptual process, process that issues in a true belief, process that occurs in Wisconsin, etc. The reliability problem for the reliabilism is best understood as a challenge: the challenge of finding a principled way of selecting, for each token-process of belief formation, the type whose reliability ratio must be high enough in order for the belief to be justified. Goldman (1979) was aware of the problem, and it has been forcefully pushed by Conee and Feldman (1998).

Many different solutions have been proposed for the generality problem, which has proven to be resilient. Heller (1995) proposed a contextualist “dissolution” of the problem according to which reliabilists should resist trying to come up with a principled way of singling out one of the indefinitely many types as the one that has to be evaluated for reliability, for the context of attribution of justification takes care of that. Among the proposals that tackle the problem head on (as opposed to denying that there is one), one of the most influential ones has been that of Alston (1995). According to Alston, the type of process whose reliability determines whether the target belief is justified or not is the type describing the function that is actually operative in the formation of the belief. Thus, if the psychological process that created the belief is insensitive to the exact geographical location where the process takes place, then the reliability of the type that mentions that kind of information is irrelevant to the justificatory status of the belief. The main question that this proposal raises is whether it singles out a unique process-type or whether, as Conee and Feldman (1998) argue, it still leaves us with a multitude of types of

differing reliability.

As I said, many other solutions have been proposed, but most of them share with Alston's the characteristic of leaving it open that the strictures that they put in place fail to single out a unique process-type. I have argued (see Comesaña (2006)) that *any* plausible epistemological theory will face an analogue of the generality problem. Therefore, either the problem is worse than we thought (because it shows that not only reliabilism is an inappropriate epistemological theory, but more generally that there are no appropriate epistemological theories) or there is a solution to the generality problem. In either case, there is no *special* problem for reliabilism.

Consider a theory that is usually presented as an internalist alternative to reliabilism: evidentialism. As Conee and Feldman (1985) present this theory, a subject *S* is justified in adopting doxastic attitude *D* towards proposition *p* at time *t* if and only if adopting *D* towards *p* fits the evidence that *S* has at *t*. Now, given this definition of justification, it may well happen that a subject has a justified true belief, is not in a Gettier-type situation and yet still fails to have knowledge—because, despite the fact that the evidence that he has justifies believing the proposition in question, he doesn't believe it on the basis of that evidence but, say, out of wishful thinking. Conee and Feldman recognize this, and add a definition of well-foundedness to their definition of justification:

S's doxastic attitude *D* towards *p* at *t* is well-founded if and only if:

1. having *D* toward *p* is justified for *S* at *t*; and
2. *S* has *D* toward *p* on the basis of some body of evidence *e*, such that
 - (a) *S* has evidence *e* at *t*; and
 - (b) having *D* towards *p* fits *e*; and
 - (c) there is no more inclusive body of evidence *e'* had by *S* at *t* such that having *D* toward *p* does not fit *e'*.

Now, once we recognize the necessity of relying on something like the notion of basing one's belief on some evidence (as is done in the evidentialist definition of well-foundedness) we can appeal to that very notion in solving the generality problem for reliabilism. Thus, incorporating the modifications made in sections 4.1 and 4.3, I propose the following as a reliabilist definition of justification that is able to deal with the generality problem (I called the resulting position "Evidentialist reliabilism"):

Evidentialist reliabilism: A belief that *p* by *S* is justified if and only if:

1. S has evidence e ;
2. the belief that p by S is based on e ; and either
 - (a) e doesn't include any beliefs of the subject and the type *producing a belief that p based on evidence e is actually reliable*; or
 - (b) e includes other beliefs of S , all of those beliefs are justified and the type *producing a belief that p based on evidence e is actually conditionally reliable*.

Evidentialist reliabilism solves the generality problem insofar as it provides a process-type whose reliability determines the justification of the belief produced by the token of that type that produced the belief in question. It remains to be seen, of course, whether the resulting theory is adequate in other respects, but objectors cannot now complain that no single type has been selected in a non-principled way.

4.4 What Is Reliability? The Lottery Problem

Somewhat surprisingly, an issue that has not received much attention in the literature on reliabilism is the question of what it means, exactly, for a process to be reliable. Insofar as this question is taken up, a truth-ratio answer is generally given. More precisely, the following seems to be the characterization of reliability that most reliabilist take for granted:

A type T is reliable if and only if most of the actual and close counterfactual tokens of T produce true beliefs.

However, if that is what it takes for a type to be reliable, then there are a range of cases that would constitute counterexamples to reliabilism (Adler (2005) is dedicated to this problem, and it was raised also in Cohen (1988)). Suppose that both Smith and Jones have bought one ticket each in a large fair lottery. The day the result is announced, Smith believes that his ticket lost based solely on the basis that it is one ticket in, say, a million. Jones, on the other hand, believes that his ticket lost based on reading about the result in the newspaper. If we are going by the above definition of reliability, then the relevant type associated with Smith's belief is more reliable than the relevant type associated with Jones's belief (newspapers print typos more often than once in a million cases). It is not part of the official characterizations of reliabilism that we have offered so far that the higher the reliability of the relevant type the higher the justification of the belief, but this is a natural way for a reliabilist to deal with the fact that justification comes in degrees.

But the problem is that we judge Jones to be the one that is better justified than Smith, not the other way around (indeed, some philosophers go as far as believing that Smith cannot know that his ticket lost based on the evidence on which he believes it).

If we adopt the modifications to reliabilism offered in the previous section, there is a different but natural way of explaining what it means for a type to be reliable, in terms of conditional probability (Alston (1988) proposed to understand reliability in terms of conditional probability, and Comesaña (2009) develops this idea):

A type of the form *believing that p based on e* is reliable if and only if the conditional probability of p given e is sufficiently high.

Let s be the proposition that Smith's ticket is a loser, j the proposition that Jones' ticket is a loser, m the proposition that Smith's ticket is one in a million in a fair lottery, and n the proposition that the newspaper announced that a ticket other than Jones' won the lottery. It is easy to verify that (given plausible assumptions about the newspaper's error rate) the conditional probability of j given n is *higher* than the conditional probability of s given m , which gives us the comparative result that we were looking for.

Conceiving of the relevant type as always including some evidence and of reliability as defined in terms of conditional probability also allows us to bypass the distinction between unconditional and conditional reliability that Goldman relied on in order to capture the difference between derived and basic beliefs. The reliabilist definition of justification that we arrive at is therefore the following:

S is justified in believing that p if and only if:

1. S has evidence e ;
2. S bases his belief that p on e ; and
3. the actual conditional reliability of p given e is high enough.

5 Conclusion

In this chapter we have presented a reliabilist theory of epistemic justification. We have also seen that such a theory is subject to some serious problems. But reliabilism is a resourceful theory, and plausible answers can be given to the most pressing problems.

References

- Adler, J. (2005), "Reliabilist Justification (or Knowledge) as a Good Truth-Ratio," *Pacific Philosophical Quarterly* 86(4), pp. 445–58.
- Alston, W. (1988), "An Internalist Externalism," *Synthese* 74, pp. 265–283.
- Alston, W. (1995), "How to Think About Reliability," *Philosophical Topics* 23, pp. 1–29.
- Armstrong, D. (1973), *Belief, Truth and Knowledge*, Cambridge University Press.
- Bergmann, M. (2006), *Justification without Awareness*, Oxford University Press.
- BonJour, L. (1980), "Externalist Theories of Empirical Knowledge," *Midwest Studies in Philosophy* V, pp. 53–73.
- Chalmers, D. (1996), *The Conscious Mind*, Oxford University Press.
- Cohen, S. (1984), "Justification and Truth," *Philosophical Studies* 46, pp. 279–295.
- Cohen, S. (1988), "How To Be A Fallibilist," *Philosophical Perspectives* (2), pp. 91–123.
- Comesaña, J. (2002), "The Diagonal and the Demon," *Philosophical Studies* 110, pp. 249–266.
- Comesaña, J. (2006), "A Well-Founded Solution to the Generality Problem," *Philosophical Studies* 129(127-47).
- Comesaña, J. (2007), "Knowledge and Subjunctive Conditionals," *Philosophy Compass* 2(6), pp. 781–791.
- Comesaña, J. (2009), "What Lottery Problem for Reliabilism?" *Pacific Philosophical Quarterly* 90(1), pp. 1–20.
- Comesaña, J. (forthcoming), "Evidentialist Reliabilism," *Noûs* .
- Conee, E. and Feldman, R. (1985), "Evidentialism," *Philosophical Studies* 48, reprinted with an "Afterword" in Conee and Feldman (2004), pp. 83-107.
- Conee, E. and Feldman, R. (1998), "The Generality Problem for Reliabilism," *Philosophical Studies* (89), pp. 1–29.
- Conee, E. and Feldman, R. (2004), *Evidentialism*, Oxford University Press, New York.

- Dretske, F. (1971), "Conclusive Reasons," *Australasian Journal of Philosophy* 49, pp. 1–22.
- Feldman, R. (1995), "Reliability and Justification," *The Monist* 68, pp. 159–174.
- Foley, R. (1985), "What's Wrong With Reliabilism?" *The Monist* 68, pp. 188–202.
- Gettier, E. (1963), "Is Justified True Belief Knowledge?" *Analysis* 23, pp. 121–123.
- Goldman, A. (1967), "A Causal Theory of Knowing," *Journal of Philosophy* 64(12), pp. 357–372.
- Goldman, A. (1979), "What is Justified Belief?" in G. Pappas (ed.), *Justification and Knowledge*, Reidel, Dordrecht, pp. 1–23.
- Goldman, A. (1986), *Epistemology and Cognition*, Harvard University Press, Cambridge, MA.
- Goldman, A. (1988), "Strong and Weak Justification," *Philosophical Perspectives* 2, pp. 51–69.
- Heller, M. (1995), "The Simple Solution to the Generality Problem," *Noûs* 29(4), pp. 501–515.
- Kaplan, D. (1989), "Demonstratives," in H. W. J. Almog, J. Perry (ed.), *Themes from Kaplan*, Oxford University Press, Oxford.
- Nozick, R. (1981), *Philosophical Explanations*, Harvard University Press.
- Pollock, J. (1984), "Reliability and Justified Belief," *Canadian Journal of Philosophy* 14(103-114).
- Ramsey, F. (1931), "Knowledge," in *The Foundations of Mathematics and Other Logical Essays*, The Foundations of Mathematics and Other Logical Essays.
- Shope, R. (1983), *The Analysis of Knowing*, Princeton University Press.
- Sosa, E. (1991), *Knowledge in Perspective*, Cambridge University Press, New York.
- Sosa, E. (1993), "Proper Functionalism and Virtue Epistemology," *Noûs* 27(1), pp. 51–65.
- Sosa, E. (1999), "How Must Knowledge Be Modally Related to What Is Known?" *Philosophical Topics* 26(1/2), pp. 373–84.

Sosa, E. (2008), "Godlman's Reliabilism and Virtue Epistemology," *Philosophical Topics* 29(1/2).

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