

Swiss Philosophical Preprint Series

87

Julien Dutant

Two Notions Of Safety

Added 15/01/2010

ISSN 1662-937X

© Julien Dutant

Two notions of safety*

Julien Dutant

January 14, 2010

Abstract

Timothy Williamson (1992, 224–5) and Ernest Sosa (1996) have argued that knowledge requires one to be *safe from error*. Something is said to be safe from happening iff it does not happen at “close” worlds. I expand here on a puzzle noted by John Hawthorne (2004, 56n) that suggests the need for two notions of closeness. *Counterfactual closeness* is a matter of what could *in fact* have happened, given the specific circumstances at hand. The notion is involved in the semantics for counterfactuals and is the one epistemologists have typically assumed. *Normalized closeness* is rather a matter of what could *typically* have happened, that is, what would go on in a class of normal alternatives to actuality, irrespectively of whether or not they could have happened in the circumstances at hand. The distinction sheds light on recent apparent counterexamples to safety (Szabó Gendler and Hawthorne, 2005, 315n14, Comesaña, 2005, Sosa,

*Part of this material has been presented at a Virtue Epistemology workshop in Edinburgh in april 2009, and the paper was presented in the SoPhA congress (sept 2009) and in the PhilEAs talk series (oct 2010) in Geneva and at the Graduate Philosophy Conference in Oxford (oct 2010). I thank the audience on each occasion for helpful discussion, Pascal Engel, John Hawthorne, Jesper Kallestrup, Anne Meylan, Duncan Pritchard, Ernest Sosa and Timothy Williamson for detailed discussions, and the later for his comments at the Oxford Graduate Philosophy Conference. The present version essentially dates from august 2010 and barely addresses the problems and questions that were raised on those occasions. A newer version is in preparation.

2007, 31): they involve failure of counterfactual safety but not of normalized safety. This suggests that normalized safety is the one that guides intuitive judgements about knowledge. However, it is plausible to think that counterfactual safety is the substantive property the intuitions are tracking; if so, the intuitive judgements should be rejected.

Peter Unger (1968) argued that one fails to know that p if it is a mere accident that one's belief that p is true and diagnosed that Gettier and lottery cases involve a violation of that condition. These views are now widely shared, though there is still disagreement over the appropriate theoretical framework for non-accidentality. Three main options have been explored: something is accidental iff it is, in some relevant sense (a) unexplained, (b) improbable, or (c) highly contingent. The topic of the present paper is a promising and influential version of the modal approach (c): the *safety* condition put forward by Timothy Williamson (1992, 224–5) and Ernest Sosa (1996) according to which one knows only if one could not easily have been wrong in a similar case.

The weak necessity modal in the safety condition is commonly represented in terms of “close” possible worlds. Here I expand on a puzzle noted by John Hawthorne (2004, 56n) that suggests that we should distinguish two notions of “closeness”. *Counterfactual closeness* is a matter of what could easily have happened in the specific circumstances at hand, and can be thought of as proximity on a tree of branching time. The notion is involved in the semantics for counterfactuals and is the one epistemologists have typically assumed. The second, *normalized closeness*, is a matter of what would go on in a normalized class of alternatives to actuality, that is, alternatives that are generated from actuality by some constrained recombination procedure, irrespectively of whether or not they could have happened in the circumstances at hand. While counterfactual safety is a matter of what could *in fact* have happened, normalized safety is rather a

matter of what could *typically* have happened.

The distinction between counterfactual and normalized safety sheds light on recent apparent counterexamples to safety (Szabó Gendler and Hawthorne, 2005, 315n14, Comesaña, 2005, Sosa, 2007, 31). Such apparent cases of unsafe knowledge involve failure of counterfactual safety but not of normalized safety. This suggests that normalized safety is the notion that guides intuitive judgements. However, it is plausible to think that counterfactual safety is the substantive property the intuitions are tracking; if so, the intuitive judgements should be rejected.

1 Safety

In its most basic form, the safety condition states:

(BS) *S* knows that *p* only if *S*'s belief that *p* could not easily have been false.

The condition is unsatisfactory on several counts, but it allows a more perspicuous exposition of our puzzle. (Refinements are postponed to section 4.)

The safety condition straightforwardly explains (1) why a lucky guess is not knowledge; (2) why one cannot know in advance that a ticket in a fair lottery is a loser, (3) why one fails to know in a range of Gettier cases.

1. If Leo guessed rightly that the coin would fall heads, his guess could easily have turned out wrong.
2. If Max holds a ticket in a standard fair lottery, my belief that he will loose could easily be false, no matter how low the odds of Max's winning are.
3. Consider Chisholm's (1966) sheep case:

Hidden Sheep Seeing a rock that resembles a sheep in the distance, Sharon comes to believe that there is a sheep in a field. There happens to be one, but it is hidden behind the rock.

The sheep could easily have failed to be there, in which case Sharon’s belief would have been false.

The safety condition reaps these benefits without fueling scepticism, at least at first sight. It is metaphysically possible, but could hardly be the case, that all your experiences are artificially produced by a supercomputer. One could not easily be wrong on such matters (Sainsbury, 1997, 908). Thus safety does not support sceptical intuitions, and that is welcome to its proponents. (Of course that does not by itself advert scepticism since safety is not claimed to be sufficient for knowledge.)

The consequences of safety are best seen when it is represented in terms of the kind of ordering of possible worlds that is familiar from Stalnaker-Lewis’ semantics for counterfactuals (Stalnaker, 1968; Lewis, 1973). We assume a preorder of possible worlds in terms of closeness to actuality. We further assume a threshold that divides worlds into those that are close enough to actuality (henceforth, “close worlds”) and those that are not (henceforth, “distant worlds”). We say:

Easy-possibility as closeness p could easily happen iff p is true at a some close world.¹

A state of affairs p is said to be an *easy possibility* iff it obtains at some close world and a *weak necessity* iff it obtains at all close worlds. The notions are

¹More formally: \preceq_x is a triadic relation over a set W of worlds such that for all $w \in W$, \preceq_w is a preorder over W (\preceq_w allows ties but no incomparabilities) with w as a minimal element. When $w' \preceq_w w''$ we say that w' is closer to w than w'' . R is a reflexive accessibility relation over worlds such that for any $w, w', w'' \in W$, if wRw' and $w'' \preceq_w w'$ then wRw'' . When wRw' we say that w' is close to w , and the condition ensures that any close world is closer than any non-close world. Finally, we introduce a weak necessity operator \Box such that $\Box p$ is true at w iff p is true at any world w' such that wRw' . We say that p could easily happen iff $\Box p$.

intended to capture ordinary claims about what could or could not easily have happen.

According to (BS), knowledge is weakly necessarily true belief. Safety is thus a restricted form of infallibility: one's belief cannot be false at close worlds. By requiring infallibility over *all* close worlds, the condition rules out Gettier cases and lottery cases insofar as these cases involve a close possibility in which one's belief would be false. By requiring infallibility across close worlds *only*, safety avoids fueling sceptical arguments insofar as those rely on distant possibilities of error. By contrast, a condition that only requires avoidance of error at *most* worlds fails to block Gettier cases and the lottery, and a condition that requires avoidance of error at *all* worlds indiscriminately implies scepticism. Safety deals with both issues by requiring avoidance of error at *all* worlds *within a certain range*.²

2 A puzzle for safety

Now the puzzle. For almost any Gettier case the safety condition handles, one can build a variant where specific circumstances ensure that the belief is safe yet no less accidentally true:

Shy Sheep As before, except that the sheep always hides behind this rock — precisely because it has a sheep shape.

The circumstances ensure that there could not easily have failed to be sheep behind the rock.³ Similarly (1) and (2) are intuitively true:

²Williamson (2000, 124) explicitly formulates safety as truth over *all* worlds within a range. Some authors have instead required avoidance of error at *most* or *nearly all* close worlds (Pritchard, 2005, 71, 163). That is unadvisable because the resulting condition fails to account for Gettier and lottery cases (Greco, 2007, 301-2), and plausibly entails failure of closure (Hawthorne, 2004, 145). Consequently Pritchard (2007, 292) has amended his condition to require that one avoid error in "*all very close*" worlds. But the distinction between "*close*" and "*very close*" seems simply to relabel the "*close*" / "*distant*" one.

³If needed, stipulate that local sheep are genetically bound to behave like that and have moved to the area or survived in it precisely because of that trait, of which Sharon

- (1) There was hardly any chance that there would be no sheep behind the rock.
- (2) One could not easily have won a bet that there was no sheep behind the rock.

Since there could not easily have failed to be a sheep there, Sharon could not easily have been wrong. However, the intuition that she does not know remains unaffected.

That is not a counterexample since safety was only claimed to be necessary for knowledge. It is in principle possible that some other condition explains Sharon's failure to know. However, the similarity between the cases (and the fact that one can build a continuum between them) puts pressure on the safety condition: if some other condition explains the Shy Sheep case, it will plausibly explain the initial case as well.

The reader will easily check that analogous variants can be built for a wide range of Gettier cases. Anticipating on my diagnosis, I call them *abnormal safety cases*.

3 Two notions of closeness

In abnormal safety cases, relevant possibilities of error are not "close". Christopher Peacocke (1999, 310, 321-3) provides a plausible characterization of the relevant notion of "closeness": at a given time t , p is close iff p would have resulted from a slight variation of the initial conditions.⁴ There are two sources of vagueness here: which variations are sufficiently

is unaware.

⁴See also Williamson (2000, 123-4). Peacocke (1999, 315) objects to using the same closer-than order that underlies counterfactuals for closeness, but for reasons that seem to me mistaken: he worries that the Lewis-Stalnaker ordering conjoined with determinism will result in close worlds that make backtracking counterfactuals true, while the ordering counts as closer worlds where small miracles have happened precisely in order to avoid that result.

small, and how far back in time are the “initial” conditions. These appear to be sensitive to subtle contextual matters and I can only give rough guidelines. (1) as Peacocke argues (1999, 321), variations should hold *ceteris paribus* laws fixed and defeasibly keep fixed “robust” properties of objects; (2) in the knowledge case, the relevant initial conditions are typically those obtaining not too long before the time of belief-formation or the time relevant for the target proposition, depending which one is earlier.⁵ In the Shy Sheep case, these constraints deliver the verdict that there is no close possibility in which Sharon’s belief is false.

The resulting structure can be laid out on a tree of branching time: at each time t we open parallel branches corresponding to slight variations of the conditions at t . A possibility p is close to w at t iff there is some sufficiently close prior time t' where a p branch opens. More generally, the later the last common node with a p branch is, the closer p is.⁶ The resulting closeness measure appears to be the relevant one for evaluating counterfactuals.⁷ Call it *counterfactual closeness*.

John Hawthorne has already used an abnormal-safety case to point out that counterfactual closeness is not the right notion for a safety-like condition on knowledge:

“For the purpose of epistemic evaluation, the epistemically relevant notions of danger aren’t to be cashed out in terms of any straightforward metric on objective risk, or any general-purpose notion of metaphysical closeness at play in the discus-

⁵This comports well with Sturgeon’s (1993) point that whether some belief is knowledge depends only on facts concerning times prior to belief formation — leaving cases of multiple support aside. But when considering an historian’s knowledge we consider variants of times long past.

⁶Note that while “ w' is close to w ” is time- and possibly context-sensitive, “closer-than” claims are not.

⁷For instance Lewis’ (1973) semantics (“if p would q ” is true iff the closest p worlds are q worlds), or von Fintel-Gillies’ one (“if p would q ” is true in context c iff $\neg p \vee q$ holds across all worlds that count as “close” in c) (von Fintel, 2001; Gillies, 2007). I remain neutral on these.

sion of counterfactuals. Suppose a demon protects X in such a way that whenever X is in danger of forming a false belief, the demon intervenes and stops the belief from being formed. X's belief-forming process up to that point might exactly resemble that of Y, who does form a false belief. Here we are tempted to claim that X does not know when, but for the demon, he would have formed a false belief in nearby worlds and, correlatively, to speak of a 'danger' or error, even though there is also a very good sense in which the objective risk of X falling into error was vanishingly low."(Hawthorne, 2004, 56n)

However, most authors have (more often implicitly than not) relied on counterfactual closeness.⁸ Williamson's (2009) recent distinction between "branching" and "close" worlds may be seen as giving up the counterfactual-closeness account, though it is not clear how he would apply it to abnormal safety cases.

An obvious reply on behalf of safety is to argue that the no-sheep possibility *is* close in the Shy Sheep case even though it is not *counterfactually* close. That requires the safety theorist to provide an alternative notion of closeness.⁹ (She may or may not further hold that the alternative notion also provide the right semantics for "could easily" claims.)

Here is the suggestion I want to examine here: the no-sheep possibility is relevant because it is a *normal* one. It would be perfectly normal for there *not* to be a sheep behind the rock, even though that is weakly impossible in the circumstances. Rather, it is the *abnormality* of the circumstances that make some normal possibilities counterfactually distant. More generally,

⁸See Williamson's (2000, 124) reference to Peacocke, Comesaña (2005, 399) and (2007), Pritchard (2005, 47, 71), Sosa (2007, 25).

⁹By contrast, if one ends up with a set of relevant possibilities of error that cannot be naturally seen as a "sphere" of possible worlds but rather as a set of unconnected "islands", we do not have a safety view in any recognizable sense. See Schaffer (2005, 125).

the idea is that an alternative is *normally* close to actuality iff it is a sufficiently similar variant of actuality that is *at least as normal* as actuality. I call such variants “normalized” rather than “normal”, for they may be relatively abnormal. *Normalized* safety is thus avoidance of error at *normally close* worlds.¹⁰

One expected feature of normalized is that it will be much less sensitive to the specific background set-up of cases than counterfactual safety, because normalized alternatives to various cases are likely to be similar to each other. That seems to be closer to the way our intuitions about knowledge behave.

Yet normality is a notably elusive notion, and I will attempt to put some flesh on the bones of the suggestion (section 5). But let me first defuse the potential worry that the abnormal safety puzzle calls for a more straightforward solution.

4 Sophisticated versions of safety and the puzzle

Basic safety is known to be unsatisfactory on several counts. (a) It fails to exclude irrelevant possibilities of error in which one would have a false belief on a different basis (Williamson, 2000, 128). (b) It is trivially satisfied by beliefs in necessary truths (Sosa, 2002, 275). (c) It cannot deal with the fake-barn style of Gettier cases, in which relevant possibilities of error do not involve belief in the target proposition p but in a different one

¹⁰The idea that the relevance of possibilities of error depends on their normality can also be found in Alvin Goldman (1986, 107) and John Greco (2003, 129–31). On Greco’s explanationist view, abnormalities in the set-up trump the default salience of our faculties in explaining why a belief is true. On Goldman’s view, the reliability of a process depend on its success in normal (not *normalized*) worlds. For reasons of space I will not compare the present view with theirs.

(Hawthorne, 2004, 56n).¹¹ The abnormal-safety puzzle may appear to be an instance of (b)-(c); in this section I argue that it is not.

The three problems (a)-(c) can be adverted by replacing the safe-belief condition with a safe-method one:

(MS) *S* knows that *p* only if one could not easily have formed a false belief on the basis of the same method.¹²

To illustrate:

Prime numbers Primo has a mistaken method for identifying prime numbers: he adds the digits and declares the number prime if the resulting sum is prime. I write “47” on a piece of paper.

Variant A: I ask him whether that number is prime.

Variant B: I ask him whether I have written a prime number.

Given that $4+7=11$ is prime, Primo answers (rightly) “yes” in both cases.

I could easily have written 49, and Primo would have falsely believed that I had written a prime number. So Primo’s belief in variant B violates (BS). However 47 could not easily have been non-prime, so Primo’s belief in variant A *satisfies* (BS). (BS) thus imposes a surprising asymmetry between the two variants. For intuitively, Primo fails to know for the same reason in both cases: he uses a wrong method. (MS) captures the intuition: in variant A, Primo could easily have formed the false belief that 49 is prime,

¹¹A fourth problem is that possibilities of error include truth-valueless (instead of false) belief (Hawthorne, 2002). That is solved by substituting non-truth to falsity; we leave the amendment aside here.

¹²I will not argue the point here, but I think that safety should require that *no one* could easily have formed a false belief on the same method. That deals for instance with Gendler and Hawthorne’s Fake Bar case (2005, 338), in which three subjects use the same method to detect whether their drinks are gin, but due to their respective habits only one of them is confronted with the undetectable fake gin their favourite bar serves on sundays.

even though his actual belief could not easily have been false.¹³

Now it should be easy to see that the abnormal-safety puzzle arises for method-safety just as it does for belief safety, only with slightly more complex cases:

Safe prime numbers As before except that I pick up the numbers from a list that, as a matter of weak necessity, contains only prime numbers whose sums of digits is prime.

Now there is no counterfactually close possibility where Primo forms a false belief on the same basis; yet he still fails to know.¹⁴ But here as well normality could provide a solution: an alternative where the list would contain other numbers is at least as normal as Primo's actuality.

The diagnosis sheds light on two types of cases recently put forward against safety. First, a range of cases straightforwardly fall into the abnormal-safety category:

Protective angels A hidden agent or demon ensures that the subject cannot form false beliefs *via* a certain method.

(See Hawthorne, 2004, 56n, Pritchard's Temp case in Haddock et al., 2010, chap. 3, Sosa, 2007, 29.)

While the angel prevents error at counterfactually close worlds, normalized variants of actuality include worlds without angels.¹⁵

Second, if the normalized safety view is right, one expects cases in which some possibilities of error are *counterfactually* close but not *normally*

¹³Sainsbury (1997, 908–9) proposes instead to construe “one’s belief” in a way that does not keep the belief content fixed. Williamson (2000, 124) and Hawthorne (2004, 56) rather formulate safety as avoidance of error in “similar-enough” cases; their condition is not specific enough to distinguish respects pertaining to counterfactual closeness, normalized closeness and methods. Sosa (2002, 279–80) adds a “guidance” condition such that one must be guided by the condition under which one’s basis is safe for a given belief, but it is not clear to me what the condition amount to.

¹⁴See also Sosa’s tomato-ripeness case (2002, 280).

¹⁵The diagnosis might also apply to cases involving “strange and fleeting” but safe processes (Greco, 1999, 285–286) commonly used against reliabilism.

close and in which we still ascribe knowledge. That is indeed what seems to go on in apparent counterexamples to the safety conditions:

Lost supporter Wandering in an Italian city, Andy has just removed his Man U cap because of the heat. He asks an apparently knowledgeable woman if she will tell him where the station is. She answers “yes” and proceeds to do so. However, she is an Intern Milan fan, and if Andy had wore his cap, she would have answered the same but sent him in the wrong direction. (Variant of [Comesaña 2005](#), 397.)

Hallucinogenic cocktails Picking up a glass at a cocktail party, you go to the next room and recognize your mother. However, most of the other glasses on the tray contained an hallucinogenic drug that would have made anybody looked to you like your mother.

(See also Gendler and Hawthorne’s ([2005](#), 351n14) real daisies case and Sosa’s ([2007](#), 31) kaleidoscope case.)

Many are disposed to ascribe knowledge to both subjects, even though their methods could easily have produced false beliefs. But the possible worlds in which they do so are intuitively *less normal* than the actual one. On the normalized safety account, such worlds are distant. Call such cases *abnormal unsafety cases*.

5 Normally close possibilities

Can we flesh out an account of normalized alternatives that delivers the advertised goods? Here is an attempt. Let the “area of interest” be the space-time region of the world relevant to a certain knowledge ascription we’re focusing on:

1. Normalized alternatives to w are *variants* of w : they are furnished with the same kind of things at broadly the same locations. Give or

take some sheep, mountains or blades of grass, but do not rearrange the molecules so that there is no sheep, grass nor mountains left. Similarly, do not move things across wide distances in time or space: no normalized variants of actuality have dinosaurs in Central Park nor cheese on the Moon.

2. Normalized alternatives to w satisfy the basic laws of w , but for a few exceptions that are kept as much as possible outside of the area of interest.
3. Normalized alternatives to w do not breach *ceteris paribus* laws of w more than w does, especially within the area of interest. If the woman does not normally lie, do not add a world where she does. (The rule should probably be reinforced to apply to each law individually: an actual abnormal dog does not allow you to introduce a normal-dog-but-abnormal-liar alternative.)
4. Normalized alternatives to w do not have less natural properties that are highly likely (by the lights of the laws of w) than w does, especially within the area of interest.¹⁶ If the world contains just 100 fair coin tosses, do not add a world with a series of 100 heads; but if it contains 2^{100} tosses, do so.

This is admittedly rough.¹⁷ Let me sketch how the account should be applied to abnormal unsafety cases. Here is a pair:

Absent Fake barns Oscar drives by the countryside, spots a barn in a field and believes that it is a barn. There is similar-looking building on the

¹⁶The rule draws on Adam Elga's (2004) account of typicality, itself based on Gaifman and Snir's Gaifman and Snir (1982) account of randomness. See also Robbie Williams (2008, 408-9).

¹⁷And I do not recommend testing the account on worlds where what is normal differs greatly from what is normal in ours. It is plausible that our intuitions are rigidly tuned to what we take to be actually normal, and that they cannot be swiftly recalibrated.

other side of the hill which Oscar cannot see from the road.

Variant A: the other building is a real one too. But earlier on, villagers had tossed a coin to decide whether to build two fakes instead.

Variant B: the other building is a fake. Earlier on, villagers had tossed a coin to decide on which side of the hill they would put it.

The possibility of error is equally counterfactually close in both variants: namely, it would have resulted from an alternative outcome of the coin toss. However, the error possibility in A involves a breach of a *ceteris paribus* law (namely, that there are no fake buildings) that the actuality in A does not breach. By contrast, in B, actuality itself breaches this law. The fake barn possibility is thus a normalized alternative of B but not of A. Correspondingly, there is more pressure to deny knowledge in B than in A.¹⁸

Here is another test pair of cases:

Series of lotteries *Variant A.* I hold a ticket in each of nine small ten-ticket lotteries.

Variant B. I hold a ticket in one big one-million-tickets lottery.

Assuming that you did not win *each* lottery in A, the possibility that you did is less normal than the actual result by rule (4), even though it was equally likely. By contrast, in variant B, the possibility that you win is as normal as the actual result. This predicts our tendency to say that we know that I won't win *every* lottery in A, but we do not know that I won't in B.¹⁹

¹⁸See Szabó Gendler and Hawthorne (2005, 351n14) for such a verdict on a structurally similar case.

¹⁹See Jonathan Vogel's (1999, 165) Heartbreaker case.

6 Should we opt for normalized safety?

Granting that normalized safety does guide our intuitive ascriptions of knowledge, should we adopt it as a condition on knowledge? That is far from obvious.

Counterfactual closeness goes with objective risk and chance. If a danger is counterfactually close, then however abnormal it may be, one is not in a safe situation; and if a possibility has a high chance of obtaining, it is counterfactually close, no matter how abnormal it may be. Now objective risk, rather than normal risk, seems to provide adequate norm for actions. For instance, driving one's car after having had a drink at the hallucinogenic cocktail party is wrong — though it may be excusable if one was unaware of the risk. This is so even though one did not in fact take an hallucinogenic drink and that hallucinations are not a normal possibility. Aligning knowledge with normal safety rather than counterfactual safety thus threatens to sever the intuitive role of knowledge as a norm for action (Hawthorne and Stanley, 2008).

One can hardly avoid noticing that normalized safety looks very much like a heuristic. It allows one to build a quick and dirty space of possibilities around actuality to evaluate the reliability of a given method, without caring much about specific details of the set-up. That is most useful when we are left in the dark, as we often are, about specific factors in the set-up that affect the space of counterfactual possibilities. It also plausibly makes it simpler to assess than counterfactual safety.²⁰ Moreover, normalized safety will most of the time coincide with counterfactual

²⁰Compare Lewis (1996, 563): “Why have a notion of knowledge that works in the way I described? [...] I venture the guess that it is one of the messy short-cuts - like satisficing, like having indeterminate degrees of belief- that we resort to because we are not smart enough to live up to really high, perfectly Bayesian, standards of rationality. You cannot maintain a record of exactly which possibilities you have eliminated so far, much as you might like to. It is easier to keep track of which possibilities you have eliminated if you - Psst! - ignore many of all the possibilities there are.”

safety — normalized safety is safety over the kind of counterfactual space one has in normal (most frequent) cases. This makes it plausible to consider normal safety as a heuristic we use to track counterfactual safety. The heuristic breaks down when the counterfactual close possibilities do not match the normal ones.

It both ideas are right, we should maintain the counterfactual construal of safety, and discard instead the intuitions that knowledge is incompatible with abnormal safety and compatible with abnormal unsafety.

References

- Roderick M. Chisholm. *Theory of Knowledge*. Prentice Hall, Englewood Cliffs, NJ, 1966. ISBN 0139141502.
- Juan Comesaña. Knowledge and subjunctive conditionals. *Philosophy Compass*, 2(6):781–791, 2007. doi: <http://dx.doi.org/10.1111/j.1747-9991.2007.00076.x>. URL <http://dx.doi.org/10.1111/j.1747-9991.2007.00076.x>.
- Juan Comesaña. Unsafe knowledge. *Synthese*, 146(3):395–404, 2005. doi: <http://dx.doi.org/10.1007/s11229-004-6213-7>. URL <http://dx.doi.org/10.1007/s11229-004-6213-7>.
- Adam Elga. Infinitesimal chances and the laws of nature. *Australasian Journal of Philosophy*, 82(1):67, 2004. doi: 10.1080/713659804. URL <http://www.informaworld.com/10.1080/713659804>.
- H. Gaifman and M. Snir. Probabilities over rich languages, testing and randomness. *The Journal of Symbolic Logic*, 47(3):195–548, 1982.
- Anthony Gillies. Counterfactual scorekeeping. *Linguistics and Philosophy*, 30:329–360, 2007. doi: 10.1007/s10988-007-9018-6.

- Alvin I. Goldman. *Epistemology and Cognition*. Harvard University Press, 1986.
- John Greco. Agent reliabilism. *Philosophical Perspectives*, 13, 1999. URL <http://www.jstor.org/stable/2676106>.
- John Greco. Knowledge as credit for true belief. In Michael Depaul, editor, *Intellectual Virtue: Perspectives from Ethics and Epistemology*. Oxford University Press, 2003.
- John Greco. Worries about pritchard's safety. *Synthese*, 158(3):299–302, October 2007. doi: <http://dx.doi.org/10.1007/s11229-006-9040-1>. URL <http://dx.doi.org/10.1007/s11229-006-9040-1>.
- Adrian Haddock, Alan Millar, and Duncan Pritchard. *The Nature and Value of Knowledge*. Oxford University Press, 2010.
- John Hawthorne. Deeply contingent a priori knowledge. *Philosophy and Phenomenological Research*, 65(2):247–269, 2002. doi: <http://dx.doi.org/10.1111/j.1933-1592.2002.tb00201.x>. URL <http://dx.doi.org/10.1111/j.1933-1592.2002.tb00201.x>.
- John Hawthorne. *Knowledge and Lotteries*. Oxford University Press, 2004. ISBN 0199287139.
- John Hawthorne and Jason Stanley. Knowledge and action. *Journal of Philosophy*, 105(10):571–590, 2008.
- David Lewis. Elusive knowledge. *Australasian Journal of Philosophy*, 74: 549–567, 1996.
- David K. Lewis. *Counterfactuals*. Harvard University Press, 1973. URL <http://www.worldcat.org/oclc/795075>.
- Christopher Peacocke. *Being Known*. Oxford University Press, 1999.

- Duncan Pritchard. Anti-luck epistemology. *Synthese*, 158:277–297, 2007. doi: 10.1007/s11229-006-9039-7.
- Duncan Pritchard. *Epistemic luck*. Oxford University Press, 2005. URL <http://www.worldcat.org/oclc/57134424>.
- Richard M. Sainsbury. Easy possibilities. *Philosophy and Phenomenological Research*, 57(4):907–919, 1997.
- Jonathan Schaffer. What shifts? thresholds, standards, or alternatives? In Gerhard Pryer and Georg Peter, editors, *Contextualism in Philosophy*, pages 115–130. Oxford University Press, 2005.
- Ernest Sosa. Postscript to “proper functionalism and virtue epistemology”. In John L. Kvanvig, editor, *Warrant in Contemporary Epistemology*. Rowman & Littlefield, Lanham, Md, 1996.
- Ernest Sosa. Tracking, competence, and knowledge. In Paul Moser, editor, *The Oxford Handbook of Epistemology*. Oxford University Press, 2002.
- Ernest Sosa. *A Virtue Epistemology*. Oxford University Press, 2007. URL <http://www.worldcat.org/oclc/85814007>.
- Robert Stalnaker. A theory of conditionals. In Nicholas Rescher, editor, *Studies in Logical Theory*. Blackwell, 1968.
- Scott Sturgeon. The gettier problem. *Analysis*, 53(3):156–164, 1993. URL <http://www.jstor.org/stable/3328464>.
- Tamar Szabó Gendler and John Hawthorne. The real guide to fake barns: A catalogue of gifts for your epistemic enemies. *Philosophical Studies*, 124:331–352, 2005.
- Peter Unger. An analysis of factual knowledge. *Journal of Philosophy*, 65(6): 157–170, 1968. URL <http://www.jstor.org/stable/2024203>.

- Jonathan Vogel. The new relevant alternatives theory. *Noûs*, 33 (s13):155–80, 1999. doi: 10.1111/0029-4624.33.s13.8. URL [10.1111/0029-4624.33.s13.8](https://doi.org/10.1111/0029-4624.33.s13.8).
- Kai von Fintel. Counterfactuals in a dynamic context. In M. Kenstowicz, editor, *Ken Hale: A life in language*. MIT Press, 2nd edition, 2001.
- J. Robert G. Williams. Chances, counterfactuals, and similarity. *Philosophy and Phenomenological Research*, 77(2):385–420, 2008. doi: 10.1111/j.1933-1592.2008.00196.x. URL <http://dx.doi.org/10.1111/j.1933-1592.2008.00196.x>.
- Timothy Williamson. Inexact knowledge. *Mind*, 101:218–242, 1992. URL <http://www.jstor.org/stable/2254332>.
- Timothy Williamson. *Knowledge and its Limits*. Oxford University Press, 2000.
- Timothy Williamson. Reply to john hawthorne and maria lasonen-aarnio. In Patrick Greenough and Duncan Pritchard, editors, *Williamson on Knowledge*, pages 313–329. Oxford University Press, 2009.