

The Value of Evidence

Bernhard Salow (Magdalen College, Oxford)

Getting more evidence is good. After all, more informed opinions are more accurate, more informed decisions more successful.

Except, of course, when they're not. Sometimes, additional evidence misleads: it suggests believing P when P is false, suggests ϕ -ing when ϕ -ing is disastrous. This needn't be anyone's fault; it could just be bad luck, as when a randomly selected sample just happens to be unrepresentative.

So getting more evidence is not always *objectively* good. But maybe it is always *subjectively* good: maybe we should always think that, on balance, receiving additional evidence would improve our opinions and decisions.

Even this turns out to be too much to hope for: there are clear cases where we should think no such things. But a restricted version of this claim may be true – it is intuitively attractive, and there is an appealing, if controversial, explanation for why it would be true.

The value of evidence (or 'learning', 'knowledge', or 'information') has received substantial discussion in decision theory. Those discussions usually take place against I.J. Good's (1967) proof that, in a restricted set of circumstances, the expected utility of gathering additional evidence before deciding always has higher expected utility than deciding in ignorance; the interesting question is then whether similar results are available outside of Good's 'classical' setting, or whether it carries over to the epistemic case.¹

By contrast, I won't explain Good's proof here. Instead, I will present a non-mathematical argument for the same claim, which is only loosely connected to Good's proof. One ambition is to make the issues accessible to a wider audience. Another is to gain more clarity on which philosophical assumptions underlie Good's result, and which are extraneous packaging. The hope is that the informal argument will show us not just *whether* this intuitively attractive idea is true, but also *why* it is (or isn't).²

I will start by classifying some clear and obvious counterexamples, leaving us with potentially defensible principles in §1. I will then lay out the argument for these principles in §2. §3 and §4 discuss controversial objections to two different steps in the informal argument and show how they could support

¹ See footnotes 8, 13, and 14 for references.

² Ahmed and Salow (2019, §4) do something similar, but the explanation here is slightly different.

counterexamples to the principles themselves – counterexamples which seem difficult to deal with by just restricting the principle further. §5 concludes.

§1 Qualifications

It's tempting to think that additional evidence is always subjectively good: that we should always think that, on balance, we would be better off – epistemically and/or practically – gathering the additional evidence than avoiding it. Unfortunately, there are clear counterexamples.

Perhaps the most obvious involve anticipated failures of rationality. I might know that, in my belief-formation and decision-making, I attach weight to certain types of evidence that is either excessive (e.g. to my general impression of someone's collegiality when I meet them in the highly artificial environment of a job interview) or completely misplaced (e.g. the weight I am apparently inclined to give to someone's race or gender when evaluating the strength of their CV).³ Or I might know that certain kinds of evidence (e.g. intricate details about the pros and cons of various mobile phones) tend to overwhelm me, leaving me unable to make up my mind about what to do or think. In those cases, I can reasonably think I'm better off (both practically and epistemically) shielding myself from the evidence.

A less common style of counterexamples involves loss of background information. I don't always have to think that *other people's* situation would be improved by receiving additional evidence – I might know that the relevant evidence is unreliable or problematically selective, so that it would mislead them even if they (not knowing about these problems) handled it with perfect rationality. This worry doesn't usually apply in my own case: after all, if I know (or have reason to suspect) that some evidence is unreliable or selective, I will be required to take that background information into account when I handle it – and so, if I know that I will be rational, I shouldn't expect to be misled. But I might worry that I will lose the background information between now and the time I would receive the evidence; the analogue to the third-personal worry then makes sense, and could rationalize thinking that the additional evidence would leave me worse off.

A third source of counterexamples is the fact that additional evidence is usually costly: if nothing else, it requires time, effort and attention to collect and evaluate, all of which could be usefully spent elsewhere. So gathering

³ See e.g. Kunda (1999) on the 'interview illusion', and the "CV studies" of Steinpreis et al. (1999), Bertrand and Mullainathan (2004), and Moss-Racusin et al. (2012). See also Gendler's (2011) worry that I might know about myself that, while I process e.g. statistical evidence about crime rates and race to form reasonable statistical beliefs, I'm disposed to then draw further unreasonable inferences, which also manifest in my actions – and that I might prefer not to have this information for that reason.

additional evidence isn't always worth it. One might respond that *having* the evidence would still be good – if only there was a way to *have* it without *collecting* it. But that's no way out. For there are also cases where simply having the evidence carries costs. The evidence might rationalize beliefs which are painful; more controversially, it might rationalize beliefs that are morally problematic, even when held on good evidence.⁴ Or the world might be out to punish those who know, depriving them of opportunities, including opportunities to learn other things. All of these can generate costs to having evidence; and, in some cases, the benefits won't outweigh these costs.

There is, however, an intuitive distinction between these 'extraneous' effects, and those we consider when we're tempted to think that, on balance, more evidence is always better. In the epistemic case, where the primary potential costs involve missing out on other, more valuable, ways of improving our beliefs, it's tempting to say "sure, getting the additional evidence might not leave my beliefs more accurate than they would have been, since without it they would have become more accurate in other ways – but it still leaves them more accurate than they are right now." In the practical case, where the potential costs are more varied, one still feels "OK, getting the evidence might not make things better overall – but it would still improve my decisions, so that if we could somehow focus only on how it affects what choices I make in the decisions I actually face, we would see an improvement."⁵ In what follows, I will appeal to this intuitive distinction, without attempting to make it precise.⁶

The three issues raised so far can all generate examples where we should think that, on balance, the additional evidence will actively make things worse. Not so with our final one. Sometimes, evidence is useless: it just doesn't bear on any decisions we'll ever have to make, or any issue worth getting right.⁷ When we

⁴ Basu (forthcoming) argues that racist beliefs – e.g. that the black diner at table 8 is likely to be a bad-tipper – wrong their targets, even if they are true, supported by your statistical evidence, and have no effects on your behaviour. I should note, however, that some authors – including Basu herself – argue that these moral costs also raise the evidential standards required for such beliefs to be rational. See e.g. Moss (2018), Gardiner (2018), and Basu and Schroeder (2019). If this is so, examples involving morally costly beliefs will often fall into the first class of counterexamples.

⁵ It's tempting to bring this out by replacing the question "would I be more successful with the additional evidence?" with the question "supposing I could defer my decisions to a hypothetical version of me who has the additional evidence – should I?".

Unfortunately, this doesn't work – after all, I'd still find out what I end up doing, hence what the more informed version of me recommends, and hence (in many cases) what the relevant evidence is; and so I would still bear the costs of knowing.

⁶ I doubt the distinction can be made precise without begging important questions in decision theory. For example, causalists like Gibbard and Harper (1978) or Lewis (1981) and evidentialists like Jeffrey (1965) might disagree about whether facts which give you evidence that an unfavourable scenario obtains without causally bringing it about should count as 'costs'.

⁷ The second of these assumes that not all topics are worth getting right. It's needed because all evidence bears on some question or other – so if every question is worth getting right, there can be no completely useless evidence in the epistemic case.

know for sure that some evidence will be useless, we shouldn't think it will improve matters. But we shouldn't think it will make things *worse* (unless, of course, it comes with costs): useless evidence might not make our decisions more successful or our (significant) opinions more accurate, but it won't make them less successful or less accurate either.

We thus have four reasons why one might sometimes think that gathering additional evidence won't result in an improvement. But it would still be interesting if these were the only reasons. Let an *ideal case* be one where we know for sure that (i) we will handle the evidence with perfect rationality, (ii) we won't lose relevant background information in the interim, and (iii) gathering the evidence incurs no extraneous costs; and where, furthermore, (iv) we leave open that the evidence is relevant. Then the problems discussed here are still compatible with:

Practical Value: In ideal cases, we should always think that, on balance, gathering the additional evidence will lead to greater success in pursuing our goals.

Epistemic Value: In ideal cases, we should always think that, on balance, gathering the additional evidence will make our opinions more accurate.⁸

It would be interesting if one or both of these were defensible. This would be interesting even if ideal cases never occur; for it would show that there are certain reasons, those stemming from the 'direct' effects of evidence on rational and informed belief-formation and decision-making, that always favour getting more evidence.

§2 The Argument

There is an appealing line of argument supporting **Practical** and **Epistemic Value**. I'll first run through a specific case, and then show how to generalize that reasoning.

Suppose you're considering buying a house. Before you decide, you can find out, without cost, whether the house contains asbestos. You think the information might affect whether you buy: if you had to decide without finding out, you'd take the risk and buy the house anyways; but if you found out that it definitely contains asbestos, you would not.

⁸ The literature has mostly focused on **Practical Value**; Oddie (1997) was the first to isolate (and defend) **Epistemic Value**. Buchak (2010) and Campbell-Moore and Salow (forthcoming) discuss a related distinction between the claims that gathering additional evidence is always required by instrumental or epistemic rationality.

Should you find out? The effect of doing so depends on what you find, i.e. on whether the house contains asbestos. If it doesn't contain asbestos, you'll still buy, so the outcome is the same. If it does contain asbestos, you won't buy when you otherwise would have done. Moreover, conditional on the house containing asbestos, you think (even now) that it's better not to buy it – this is both plausible independently, and a consequence of the idea that (a) your future decision not to buy would be the rational result of discovering that the house contains asbestos and (b) the rational reaction to some new evidence E (at least in ideal cases) is to adopt the views you previously held on the condition that E is true.⁹ So from your current perspective, finding out either makes no difference, or else results in a decision that you already think is better, conditional on the circumstances that would lead you to make it. So finding out has no downside and may well have an upside – you'd be silly not to do it.

(It may be worth saying more about 'think better', 'think more accurate' and the like. My use of these is, I think, perfectly ordinary; but making philosophical sense of it is slightly tricky. It does not mean 'think that it is objectively better': you might think it better to insure your house against flood damage despite taking flood damage to be unlikely. Nor does it mean 'think that it is subjectively better': even conditional on the house containing asbestos, you do not think that your current evidence, which doesn't contain this information, tells against buying it. A natural analysis is that to think something better/more accurate) is to assign it higher expected utility/ higher expected accuracy; those who reject credences, utilities, or accuracy measures, or think that expected values aren't central to subjective evaluations, will need some other way to make sense of it.)

This reasoning can be generalized. First, we need some terminology. You face a decision between some options (e.g. *Buy* and *Don't Buy*); let's use " O " to pick out the one you would choose if you were to decide without investigating further (*Buy*, in our example).¹⁰ If you want, you can gather some additional evidence before deciding; let E_1, E_2, \dots be the various (total) things you might receive if you do (e.g. $E_1 = \textit{Asbestos}$ and $E_2 = \textit{No Asbestos}$). Each of those would lead you to choose a particular option, so let's use " O_n " to pick out the option you would choose if you were to receive E_n (so, in our example, $O_1 =$

⁹ I assume, here and throughout, that evidence consists of propositions. Some (such as ????) deny this because they think evidence consists of e.g. mental states, but still allow that there are propositions (namely that you are in such-and-such a mental state) which play a central role in determining what that evidence supports – such theorists may be able to adapt the argument given here. Others insist that there are no such propositions, e.g. because evidence gives rise "directly" to a new assignment of probabilities to at least some propositions; as discussed in Graves (1989), theories like this require additional constraints to validate **Practical** and **Epistemic Value**.

¹⁰ For simplicity, I assume that options are never tied, so that there is always a particular option you'd definitely choose.

Don't Buy and $O_2 = \textit{Buy}$). We can then break the argument into a series of steps:

1. For each n , you think O_n is at least as good as O , conditional on E_n being true; and for some n , you think O_n is better than O , conditional on E_n being true.
2. So, for each n , you think O_n is at least as good as O , conditional on E_n being the evidence you would get if you were to investigate; and for some n , you think O_n is better than O , conditional on E_n being the evidence you would get if you were to investigate.
3. So, you should think the complex action *For each n : do O_n if E_n is the evidence you would get if you were to investigate* is better than O .
4. You know for sure that gathering the evidence is equivalent to performing this complex action, and that deciding without investigating is equivalent to performing O .
5. So, you should think gathering the evidence is better than deciding without investigating.

4 is true by stipulation, and the inference from 4 and 5 to the conclusion seems good. 1 also looks highly plausible, in light of the thoughts that (a) your future decision is rational, (b) in ideal cases, the rational reaction to some new evidence E is to form the beliefs you previously held on the condition that E , and (c) you consider the evidence to be useful so that, for some E_n , O_n is different from O . I will argue shortly that the moves from 1 to 2 and from 2 to 3 ought to be more controversial. But it's worth noting that they, too, have strong *prima facie* appeal. The move from 1 to 2 is supported by the natural thought that " E_n is true" and " E_n is the evidence that you would get if you were to investigate" seem pretty much equivalent. And the move from 2 to 3 is underwritten by the plausible

Sure Things: if you think that A isn't worse than B conditional on any of a finite number of mutually exclusive and jointly exhaustive hypotheses, and better conditional on at least one of these, then you should think that A is better overall.¹¹

We can run a structurally similar argument for **Epistemic Value**. In fact, we can even adapt the example. Suppose that, instead of wondering whether to buy the house, you're wondering how much it will be worth in 10 years. You're fairly confident the house doesn't contain Asbestos, and will hence be worth at least £X. Finding out for sure that it doesn't contain Asbestos would only confirm that, so wouldn't change your belief that this is true; finding out that it

¹¹ To be plausible as a general principle, **Sure Things** should also require that which of the hypotheses obtains is independent of whether you choose A or B. How exactly to spell out the "suitable independence" is controversial; but which E_n is the evidence you would receive if you were to investigate seems like it will be suitably independent of whether you investigate in nearly all situations. For discussion of **Practical Value** in cases where it might not be, see Adams and Rosenkrantz (1980) and Maher (1990).

does contain Asbestos would: you'd then only be willing to say that it will be worth at least $\epsilon X - 1,000$. But, conditional on the house containing Asbestos, you're already unsure that it will be worth at least ϵX – so, presumably, you'd regard this change as an improvement in accuracy. So you think that gathering the evidence will either change nothing or improve your accuracy; and so you should think that, on balance, it will be an improvement.

Again, we can generalize. Let's call your current belief state " B " and the belief state you would end up in if you were to learn E_n , " B_n ". Here, a belief state could be either a complete assignment of attitudes to propositions; or an assignment of attitudes to a particular proposition; this does not affect the structure of the argument, which goes as follows:

- 1e. For each n , you think B_n is at least as accurate as B , conditional on E_n being true; and for some n , you think B_n is more accurate than B , conditional on E_n being true.
- 2e. So, for each n , you think B_n is at least as accurate as B , conditional on E_n being the evidence you would get if you were to investigate; and for some n , you think B_n is more accurate than B , conditional on E_n being the evidence you would get if you were to investigate.
- 3e. So you should think that the result of the complex action *For each n : adopt B_n if E_n is the evidence you would get if you were to investigate* is more accurate than B .
- 4e. You know for sure that gathering the evidence is equivalent (in its consequences for the accuracy of your belief state) to performing this complex action.
- 5e. So, you should think that gathering the evidence will lead you to have more accurate beliefs.

The steps are analogous to 1-4, and mostly plausible for the same reasons. The main exception is premise 1e. The thoughts motivating 1 above get us that B_n just are the beliefs you currently have, conditional on E_n being true, and that, in at least one case, these are different from B . It's a substantial further claim that these beliefs always rate themselves as at least as accurate as B , and as more accurate when they differ from B . This further claim, however, is well-motivated: it's just the requirement that accuracy be measured in such a way that rational states are *immodest*, thinking themselves more accurate than any specific alternative – a requirement that is extensively discussed and defended in the literature on accuracy.¹² So this difference does not render the epistemic argument substantially less plausible than the practical one.

¹² Oddie (1997), Greaves and Wallace (2008), Joyce (2009), and Horowitz (forthcoming) all defend immodesty. For criticism, see Maher (2004), Gibbard (2008), and Myrvold (2012). Oddie's (1997) argument for **Epistemic Value** explicitly invokes immodesty; Myrvold (2012) shows how failures of immodesty can generate counterexamples.

(Another difference between the practical and epistemic cases is that it's rather unnatural to speak of "actions" in the latter. But it should be clear on reflection that such talk in 3e and 4e is ultimately just a useful fiction – for example, we nowhere assume that you can *choose* which belief state to adopt.)

We now have arguments for both **Practical** and **Epistemic Value**. These arguments make assumptions: that some form of "conditionalization" is the rational way to revise your beliefs in ideal cases; that " E_n is true" and " E_n is the evidence you would get" are, for our purposes, equivalent claims; that **Sure Things** is true; that rational belief states are immodest. While plausible, these assumptions aren't obviously *more* plausible than **Practical** and **Epistemic Value** themselves. For this reason, we should think of these arguments as attempting to *explain* (rather than *establish* or *support*) **Practical** and **Epistemic Value**. They articulate the intuitive idea that we should favour gathering the evidence because, no matter what evidence we get, gathering it will (in ideal cases) lead us to make decisions or adopt beliefs that we *already* regard as better, on the supposition that these are the decisions or beliefs which the evidence will encourage.

Note, moreover, that this explanation – unlike Good's own argument – builds in relatively little background theory; in particular, it seems equally intelligible irrespective of how one represents doxastic (e.g. as full beliefs, sharp credences, imprecise credences) or what decision theory one favours (i.e. of how one thinks that beliefs about descriptive matters and evaluations of complete state of affairs determine which things one should "think are better" than which others). Admittedly, *some* structure is taken for granted – most notably that people have beliefs about which things are better than which others conditional on various hypotheses. But this seems a datum about our psychology, something any theory should accommodate.

Of course, as we will see shortly, various decision-theoretic or doxastic frameworks might reject one of our premises, and **Practical** and **Epistemic Value** with them. But, importantly, they do not reject the terms within which the explanation is developed. We can thus pinpoint quite precisely where they think the explanation goes wrong, and thereby get some insight into *how* **Practical** and **Epistemic Value** might fail.

To illustrate this point, I will focus on two potential weaknesses: worries about **Sure Things** which challenge the transition from 2 to 3, and worries about the equivalence between " E_n is true" and " E_n is the evidence you would get if you were to investigate" that challenge the transitions from 1 to 2.

§3 First Objection: Against Sure Things

The transition from 2 to 3 (and from 2e to 3e) relies on **Sure Things**; this is required for us to “put together” your various conditional opinions about what’s best (or most accurate) given the different hypotheses about what evidence you would get into a single, unconditional, opinion. But, while attractive, the **Sure Things** is not uncontroversial. Here are two widely-discussed reasons to think it might fail.

Many people are risk-averse:¹³ they think better of a guaranteed £3,000 than a 4/5 shot at £4,000. But this effect often reduces when the “safer” choice is still quite risky: many of the same people think better of a 1/5 (=4/20) shot at £4,000 than a 1/4 (=5/20) shot at £3,000. If this combination is rational, **Sure Things** is false. For we can divide the 20 possibilities from the second choice into two sets: the 15 in which you’d get nothing no matter what, and the 5 in which you’d get £3,000 from the safer choice and in 4 of which you’d get £4,000 from the riskier one. Obviously, the safe and the risky choice are equally good in the first set; and, by hypothesis, the people in question think the safe choice is better on the condition that they are in the second set. In spite of this, these people think that the risky choice is better overall.

A related, but distinct, fact is that many people are ambiguity-averse:¹⁴ they think better of options with clearer odds. Consider, for example, an issue on which you have complex but reasonably balanced evidence; say whether Federer or Nadal will win the next time they play. If you are ambiguity-averse, you’ll require more favourable odds to bet on either answer than you would to bet on the outcome of a fair coin toss, since it’s less clear how risky such a bet is. But then you violate **Sure Things**. For consider some odds you’d take on the coin toss, but not on the tennis match; and then consider the condition that *either the coin lands heads and Federer wins, or the coin lands tails and Nadal wins*. On that condition, you should think it quite unclear how likely the coin is to land heads; after all, all the complex evidence suggesting Federer wins suggests that it will, and all the complex evidence suggesting Nadal wins suggests that it won’t. So, being ambiguity averse, you should think that, on that condition, it is better not to bet at those odds on the coin. And, for the same reason, you should also think it better not to bet on the coin on the condition that *either the coin lands heads and Nadal wins, or the coin lands tails and Federer wins*. But those two conditions are mutually exclusive and jointly

¹³ The example is due to Allais (1953), the empirical evidence to Kahneman and Tversky (1976: 266f). For discussion of the relevance of risk-aversion to **Practical Value**, see Buchak (2010) and the references therein.

¹⁴ See Ellsberg (1961). Ambiguity is usually thought to require doxastic states more complex than precise credences, e.g. sets of probability distributions representing imprecise credences. The potential relevance of moving to such a framework for **Practical Value** is noted by Good (1974), who attributes the point to Isaac Levi and Teddy Seidenfeld. See Bradley and Steele (2016) and the references therein for discussion.

exhaustive. So **Sure Things** requires you to think unconditionally that it is better not to bet on the coin – which you don't.

It's not hard to see that these counterexamples to **Sure Things** give rise to counterexamples to **Practical Value**. Take the case of risk-aversion: you are faced with the choice between a $1/5$ ($=4/20$) shot at £4,000 and a $1/4$ ($=5/20$) shot at £3,000; you think it better to opt for the first. You are then offered the opportunity to find out whether you are in one of the $15/20$ cases in which you'd lose no matter which you pick. You know that to take this opportunity is, in effect, to opt for the £3,000: for if you find out that you are in the lose-no-matter what cases, it doesn't matter what you pick; and if you find out that you are not in the lose-no-matter-what cases, your risk-aversion will kick in and lead you to choose the safe £3,000. Since you think it's better not to opt for the £3,000, you thus think it's better not to find out. Similarly, in the case of ambiguity aversion, you know that you will refuse to bet on the coin after finding out which of *either the coin lands heads and Federer wins, or the coin lands tails and Nadal wins and either the coin lands heads and Nadal wins, or the coin lands tails and Federer wins* is true; but you think that this is a good bet to take, and so you should avoid finding out which of these is true.¹⁵

We can also try to adapt these examples to generate counterexamples to **Epistemic Value**, but this turns out to be trickier. Consider first the case involving ambiguity. Right now, you believe that the coin is 50% likely to land heads. And you can foresee that, once you find out which of the two disjunctions is true, you will regard it as much less clear how likely the coin is to land heads. Call this latter state "suspended probabilistic judgement". If this is a particular alternative attitude to the coin landing heads, and Immodesty is true, then we can generate an epistemic reason to avoid the evidence. After all, you currently regard this claim as 50% likely; so, by Immodesty, you think regarding it as 50% likely is, on balance, more accurate than any alternative attitude you could take, including suspending probabilistic judgement. But you

¹⁵ A complication is that many decision theories dealing with ambiguity *permit* rather than *require* ambiguity aversion. Some, like Seidenfeld (2004) and especially Bradley and Steele (2016), suggest that this allows these theories to accept something in the vicinity of **Practical Value**. But what exactly can they accept? They can't accept **Weak Practical Value**: that in ideal cases, we should always expect that gathering the evidence will leave us, on balance, no less successful than we would have been otherwise (and hence that gathering the evidence is always rationally permissible, even if it isn't rationally required). For there can surely be ideal cases where I know about myself that if I gather the evidence, I will, in fact, make the rationally permissible ambiguity-averse decision even though other decisions will also be permissible (just like I might know that, in cases of indifference, I always opt for the option I considered last); and that is all that is required for our counterexample. It thus seems to me that the closest thing they can accept is **Very Weak Practical Value**: that we can rationally be such that **Practical Value** is true of us. But that's fairly unexciting; one reason it is true is that we can rationally fail to exhibit any ambiguity aversion and know this about ourselves. It's also equally open to defenders of risk-aversion who, like Buchak (2013), take risk-neutrality to be rationally permissible.

know that, if you were to gather the evidence, you would end up suspending probabilistic judgement. So you should think that gathering the evidence will reduce your accuracy.

This argument, however, only goes through if “suspended probabilistic judgement” is a single, particular attitude; and some have denied that it is, arguing that you should suspend probabilistic judgement in different ways depending on which of the disjunctions you learn.¹⁶ If there are different attitudes here, it’s not clear that we get a counterexample; for Immodesty only commits you to taking your own attitude to be more accurate than any particular alternative, not more accurate than a compound of different attitudes to be adopted in different situations. On the other hand, it also isn’t clear that we don’t get a counterexample; to determine whether we do, we’d need to develop a theory of these states of suspended probabilistic judgement, including a theory of how accurate they are when things turn out to be one way or the other. And it turns out that such a theory is difficult to obtain, especially if we want to respect constraints like Immodesty.¹⁷

On the face of it, it’s even less clear how to adapt the case involving risk-aversion. This is not because it’s hard to make sense of risk-aversion in the epistemic domain; ever since William James, it’s been a common refrain that believers must balance the benefits of forming true beliefs with the cost of forming false ones, and one can clearly be more or less risk-averse in how one does this. But there seems to be no obvious way to use this idea to develop an epistemic version of this case, at least not without appealing to additional (and controversial) ideas about how to “trade-off” the accuracy of beliefs with different contents. Even if we can’t directly adapt the counterexample, however, once we reject the Sure Thing Principle for this kind of reason, it stops being clear that there couldn’t be one – to settle the matter, we’d need to develop a theory of risk-aversion and its interaction with accuracy, and then use that theory to see if counterexamples can be generated or ruled out.¹⁸

§4 Second Objection: The Structure of Evidence

Another crucial step in our arguments assumes that “ E_n is true” and “ E_n is the evidence you would get if you were to investigate” are, for our purposes, equivalent claims – so that your evaluations conditional on one align with your evaluations conditional on the other. This assumption is quite natural in simple cases: given the set-up of the house-buying example, “the house

¹⁶ See e.g. Joyce (2010); Topey (2012) raises concerns about this move.

¹⁷ See Seidenfeld, Schervish, and Kadane (2012), Schoenfield (2015), and Mayo-Wilson and Wheeler (2016).

¹⁸ Campbell-Moore and Salow (forthcoming) pursue this approach, on the basis of Buchak’s (2013) theory of Risk-Aversion; they argue that there will indeed be cases where it’s epistemically better not to have the additional evidence.

contains asbestos” seems equivalent to “finding out whether the house contains asbestos would reveal that it contains asbestos” and “the house contains no asbestos” equivalent to “finding out whether the house contains asbestos would reveal that it contains no asbestos”, as required.

There are, however, reasons for concern. The claims could, in principle, come apart in two ways: (i) E_n could be true without being the evidence you would get or (ii) E_n could be the evidence you would get without being true. Which, if either, is actually possible depends on one’s theory of evidence.¹⁹

Let’s start with (ii). The main worry here is that investigating might add E_n to one’s evidence despite its falsity.²⁰ Theories of evidence that allow this might include certain “dogmatist” or “reliabilist” view, on which having an experience with content p adds p to one’s evidence as long as one has no ‘defeaters’ and/or experience is generally a reliable indicator of what the world is like.²¹ For such theories will allow that, sometimes, illusions will result in false propositions being added to one’s evidence. So something could be the evidence you would get from investigating without being true.

Theories which allow this threaten not just the move from 1 to 2, but also **Practical and Epistemic Value** themselves. For suppose that there are cases where one can antecedently recognize that further investigating could add a false proposition to one’s evidence. Then it should be possible to adapt these cases to generate ones where one knows that a proposition will be added to one’s evidence whether it’s true or not. One could then anticipate that one will, in the future be more inclined to believe or act on the relevant proposition than one is at the moment: after all, it will then be entailed by one’s evidence, while one currently leaves open that it might be false. So there will be specific decisions or beliefs that one can anticipate one will adopt if one gathers the evidence, but not if one doesn’t. Since one currently thinks better of the decisions and beliefs one would adopt if one doesn’t gather the evidence, one thus has reason to avoid the evidence.

¹⁹ Geanakoplos (1989), Ahmed and Salow (2019), Dorst (forthcoming), and Das (ms) all discuss how these issues arise in something closer to Good’s own framework.

²⁰ Another worry comes from cases where E_n would be true if I were to investigate, but wouldn’t be true if I don’t. ‘ E_n would be true if I were to investigate’ and ‘ E_n is true’ bear on most hypotheses the same way, so this isn’t usually a problem. But there may be cases where it is – e.g. ‘I would investigate if I were to investigate’ and ‘I investigate’ seem to bear very differently on how curious (or epistemically responsible) a person I am. I can’t pursue these complications here; they may interact with the ones raised in footnote 11 above.

²¹ Goldman (2009) suggests something like the reliabilist version of this view; Miller (2016) offers it as a way of understanding the “perceptual dogmatism” defended by Pryor (2000) and others. Miller’s preferred way of understanding that position is more complex; but it predicts counterexamples to both **Practical** and **Epistemic Value** for similar reasons.

Let's work through a toy example. Suppose that the dogmatist/reliabilist view sketched above is true, and that you know you will be shown an object of a certain colour: red if the house doesn't contain asbestos, white if it does. Presumably the dogmatist/reliabilist will want to say that, in a case like this, the content of your experience will become part of your evidence, even if you can't antecedently rule out that there will be an illusion which makes a white object appear red – assuming that there is no positive reason to think that there will be such an illusion. Now suppose that, before deciding to look at the object, an oracle – whose reliability and trustworthiness is beyond doubt – tells you that it will certainly *appear* red. This does not, presumably, add to your evidence the proposition that it *is* red, or anything that entails it; after all, you've not had an experience of it being red, and weren't antecedently able to rule out that it would appear red without being red. Presumably the oracle's testimony also doesn't affect what evidence you get when you look at the object; it didn't, after all, give you positive reason to think that the experience will be illusory. So you can work out that, when you look at the object (and find, as you knew you would, that it appears red), you will strengthen your epistemic position with respect to its being red: at that point, unlike now, your evidence will entail that it's red. Presumably this makes a difference to what you'll be willing to do or think: you'll be willing to stake slightly more on the object being red (pay \$Y for the house, when you'd currently pay at most \$Y-100); and start believing claims non-conclusively supported by it's being red (say, that the house will still be worth at least \$X in ten years), which your current evidence falls just short of supporting enough to justify believing. So you know that looking at the object will definitely result in one set of beliefs (believing the house will still be worth \$X in ten years) or decisions (paying \$Y to buy the house), while not looking at it will result in the ones you currently think best (believing only the weaker claim that it will still be worth \$X-100, paying at most \$Y-100 to buy the house). Since – trivially, or by immodesty – you think that the beliefs and decisions you currently think are best are better than any other particular alternative, this means that you should think it better, where both accuracy and decision-making are concerned, not to look at the object for yourself.

So if (ii) can happen, the argument fails, and **Practical** and **Epistemic Value** are under severe pressure. Similar things are true if (i) can happen.

Cases satisfying (i) are ones where some E_n can be true without being the evidence you would obtain if you were to investigate. At first sight, it seems obvious that there can be such cases. In an only slightly more complicated version of the house-buying scenario, there's a third possibility: instead of revealing that the house does or does not contain Asbestos, the investigation might come out inconclusive. Then it seems that *The house contains Asbestos* is one of the E_n , but could be true without being what you'd learn: there is, after all, a possibility in which the house contains Asbestos but the investigation is inconclusive. But this initial impression is mistaken. For

remember that the E_i are *total* bodies of evidence. And if the investigation reveals that the house contains Asbestos, the total evidence it will provide me with plausibly goes beyond the fact that the house contains Asbestos; it will also include the fact that the investigation yielded a result. The real E_n is thus *The house contains Asbestos and the investigation revealed this* – and *that* claim is not true in the case imagined, in which the investigation is inconclusive.

It thus isn't obvious that there can be cases satisfying (i). But it also isn't obvious that there can't be. It all depends on the structure of total bodies of evidence. Let us say that a body of evidence E is self-identifying if for each p , if E does/doesn't contain p , E entails that one's total evidence after the investigation would/wouldn't contain p . If bodies of evidence are self-identifying, they clearly can't be true without being the evidence you would get. If they are not, they clearly can be.

Are bodies of evidence self-identifying? On some accounts, they are. For example, if someone's evidence consists of an exhaustive description of their phenomenal state (together, perhaps, with a clause saying that there is nothing more to their phenomenal state and/or a clause describing the true theory of evidence), it arguably is. On other accounts, they are not. Since it's uncontroversial that we can fail to know something without knowing anything that entails that we don't know it, a subject's body of knowledge is not self-identifying; so if someone's evidence is just everything they know, then bodies of evidence typically won't be self-identifying.²²

If bodies of evidence are not self-identifying, (i) won't be true, and so our argument/explanation for **Practical** and **Epistemic Value** will break down. Will there also be counterexamples to the claims themselves? Not necessarily, but probably. That is: there are conditions weaker than (i) which still guarantee the value of evidence (though they don't lend themselves to the kind of simple, intuitive explanation (i) does);²³ but many of the ideas which would lead one to reject (i) would also lead one to reject these weaker conditions, and thus do give rise to direct counterexamples.

For example, consider again the view that one's evidence is just one's knowledge. On a standard, anti-sceptical picture, perception can give rise to knowledge when everything is going well, but won't give rise to knowledge when they are not (when e.g. appearances don't match reality). We can then

²² Lewis (1996) endorses something like the first theory; Williamson (2000, ch.9) defends the identification of evidence with knowledge. A potential complication is that what matters is not whether bodies of evidence really are self-identifying but whether they are self-identifying in all the possibilities we leave open. Salow (2019) argues that this difference might matter in related problems, but I will ignore that complication here.

²³ Necessary and sufficient conditions within a Bayesian framework are identified by Geanakoplos (1989); for epistemological discussion and defence, see Dorst (forthcoming).

modify the case given above, to get a counterexample to **Practical** and **Epistemic Value**, by introducing a second potential source of knowledge. Suppose that, in addition to looking at the relevant object, you will also get to touch it to determine its temperature. Things have been set up so that the object will be both red and warm if and only if the house contains no asbestos. Again, you have no prior reason to doubt either of your perceptual capacities, and an oracle who tells you that (i) the object will certainly appear both red and warm and (ii) it is certainly at least one of the two. Now it seems that, in this case, you don't yet know either the object is red or that it's warm – after all, you have neither seen nor touched it. But you do know that you will learn either that it's red, or that it's warm, or both.²⁴ And any of those is additional evidence that the house contains no asbestos: each one only rules out possibilities in which the house contains asbestos (and the object thus fails to either be red or warm).²⁵ So you know that the future evidence – no matter what it is – will support staking strictly more on the house containing no asbestos, in both your actions and your beliefs, than you currently would; and so, by the same reasoning as above, you should now think that the future evidence is better avoided.²⁶

§5 Conclusion

Is it, then, good to get more evidence? More precisely, is gathering additional evidence *subjectively good* in *all ideal cases*? Perhaps. There is an intuitively compelling explanation why it would be. Roughly, it's that no matter what evidence we get, gathering it will (in ideal cases) lead us to make decisions or adopt beliefs that we *already* regard as better, conditional on these being the decisions or beliefs which the evidence will encourage. But, as we have seen, that explanation is not uncontroversial. It assumes that we can combine conditional evaluations into a single, unconditional, opinion in the way articulated by **Sure Things** – an assumption challenged by some decision

²⁴ One might deny this, insisting that one can learn by perception that the object is red (warm) only if one already knows that appearances won't be misleading; in which case one can deduce that it is red (warm) before looking at (touching) it. But consistently taking this line undermines much of the anti-sceptical potential of equating evidence with knowledge; for, on this view, we can only receive evidence about the external world if we already have *a priori* knowledge that appearances are not misleading.

²⁵ Here it is crucial that, when you don't learn that the object is red (or that it's warm), you don't also learn that you didn't learn this. Otherwise, your new knowledge (i.e. your new evidence) would allow you to rule out that you're in possibilities where the object is red and warm and everything is working normally, so that you come to know both those things by interacting with it; but those are possibilities in which the house contains no asbestos, so it's then no longer clear that the evidence overall supports there being no asbestos.

²⁶ For more discussion of cases like this, see Salow (2018, §2.3) and Das (ms); Williamson (2000, §10.5) and (forthcoming, §4) discusses related cases. Dorst (forthcoming, Appendix A.1) argues that such cases should be deemed impossible even if bodies of evidence aren't always self-identifying.

theorists. And it assumes that the hypothesis that some body of evidence is true and that it is the evidence we'd receive are interchangeable – an assumption challenged by some epistemologists. Moreover, if either of these challenges are compelling, it's likely that there will be counterexamples to principles like **Practical** and **Epistemic Value** (though, as we saw, this follows more immediately with some challenges than with others). The question thus seems surprisingly – and excitingly – open.

References

- Adams, E and Rosenkrantz, R. (1980) "Applying the Jeffrey Decision Model to Rational Betting and Information Acquisition" *Theory and Decision* 12: 1-20
- Ahmed, A. and Salow, B. (2019) "Don't Look Now" *British Journal for the Philosophy of Science* 70: 327-350
- Allais, M. (1953) "Le Comportement de l'Homme Rationnel devant le Risque: Critique des Postulats et Axiomes de l'École Américaine" *Econometrica* 21: 503-546.
- Basu, R. (forthcoming) "The Wrongs of Racist Beliefs" *Philosophical Studies*
- Basu, R. and Schroeder, M. (2019) "Doxastic Wronging" in B. Kim and M. McGrath (eds) *Pragmatic Encroachment in Philosophy*. Routledge.
- Bertrand, M. and Mullainathan, S. (2004) "Are Emily and Greg more employable than Lakisha and Jamal?" *American Economic Review* 94: 991-1013.
- Bradley, S. and Steele, K. (2016) "Can Free Evidence be Bad? Value of Information for the Imprecise Probabilist" *Philosophy of Science* 83: 1-28
- Buchak, L. (2010) "Instrumental Rationality, Epistemic Rationality, and Evidence-gathering" *Philosophical Perspectives* 24: 85-120.
- Buchak, L. (2013) *Risk and Rationality*. Oxford: Oxford UP
- Campbell-Moore, C. and Salow, B. (forthcoming) "Avoiding Risk and Avoiding Evidence" *Australasian Journal of Philosophy*
- Das, N. (ms)
- Dorst, K. (forthcoming) "Evidence: a Guide for the Uncertain" *Philosophy and Phenomenological Research*
- Gardiner, G. (2018) "Evidentialism and Moral Encroachment" in K. McCain (ed.) *Believing in Accordance with the Evidence: New Essays on Evidentialism*. Springer.
- Geanakoplos, J. (1989) "Game Theory Without Partitions, and Applications to Speculation and Consensus" *Research in Economics*, Cowles Foundation (914).
- Gibbard, A. and Harper, W. (1978) "Counterfactuals and Two Kinds of Expected Utility" in C. Hooker, J. Leach, and E. McClennan (eds) *Foundations and Applications of Decision Theory* (University of Western Ontario Series in Philosophy of Science, 13a). Dordrecht: D. Reidel.
- Gibbard, A. (2008) "Rational credence and the value of truth" in T. Szabo Gendler and J. Hawthorne (eds.) *Oxford Studies in Epistemology* (vol 2). Oxford: OUP
- Goldman, A. (2009)
- Good, I. J. (1967) "On the Principle of Total Evidence" *British Journal for the Philosophy of Science* 17: 319-21.

- Good, I. J. (1974) "A Little Learning can be Dangerous" *British Journal for the Philosophy of Science* 25: 340-42.
- Graves, P. (1989) "The Total Evidence Principle for Probability Kinematics" *Philosophy of Science* 56: 317-324.
- Greaves, H. and Wallace, D. (2006) "Justifying conditionalization: Conditionalization maximizes expected epistemic utility" *Mind* 115: 607-632
- Horowitz, S. (2019) "Accuracy and educated guesses" in T. Szabo Gendler and J. Hawthorne (eds.) *Oxford Studies in Epistemology (vol 6)*. Oxford: OUP
- Jeffrey, R. (1965) *The Logic of Decision*. Chicago: University of Chicago Press.
- Joyce, J. (2009) "Accuracy and coherence: Prospects for an alethic epistemology of partial beliefs" in F. Huber and C. Schmidt-Petri (eds.) *Degrees of Belief*. Dordrecht: Springer
- Joyce, J. (2010) "A Defence of Imprecise Probabilities in Inference and Decision Making" *Philosophical Perspectives* 24: 281-323
- Kahneman, D. and Tversky, A. (1979) "Prospect Theory: An Analysis of Decision under Risk" *Econometrica* 47: 263-91.
- Kunda, Z. (1999) *Social Cognition: Making Sense of People*. Cambridge MA: MIT Press.
- Lewis, D. (1981) "Causal Decision Theory" *Australasian Journal of Philosophy* 59: 5-30.
- Lewis, D. (1996)
- Maher, P. (1990)
- Maher, P. (2002) "Joyce's argument for probabilism" *Philosophy of Science* 69: 73-81
- Mayo-Wilson, C. & Wheeler, G. (2016) "Scoring Imprecise Credences: A Mildly Immodest Proposal" *Philosophy and Phenomenological Research* 92: 55-78.
- Miller, B. (2016) "How to be a Bayesian Dogmatist" *Australasian Journal of Philosophy* 94: 766-780
- Moss, S. (2018) "Moral Encroachment" *Proceedings of the Aristotelian Society* 118: 177-205.
- Moss-Racusin, C., Dovidio, J., Brescoll, V., Graham, M., and Handelsman, J. (2012) "Science Faculty's Subtle Gender Biases Favor Male Students" *PNAS* 109 (41): 16395-16396.
- Myrvold, W. (2012) "Epistemic Values and the Value of Learning" *Synthese* 187: 547-68
- Oddie, G. (1997) "Conditionalization, cogency, and cognitive value" *British Journal for the Philosophy of Science* 48: 533-541
- Pryor, J. (2000) "The Sceptic and the Dogmatist" ???
- Salow, B. (2018) "The Externalist's Guide to Fishing for Compliments" *Mind* 127: 691-728
- Salow, B. (2019) "Elusive Externalism" *Mind* 128: 397-427
- Schoenfield, M. (2015) "The Accuracy and Rationality of Imprecise Credences" *Noûs* 51: 667-685
- Seidenfeld, T. (2004). "A Contrast between Two Decision Rules for Use with Convex Sets of Probabilities: Gamma-Maximin versus E-Admissibility." *Synthese* 140: 69-88.
- Seidenfeld, T., Schervish, M.J., and Kadane, J.B. (2012) "Forecasting with Imprecise Probabilities" *International Journal of Approximate Reasoning* 53: 1248-1261

- Steinpreis, R., Anders, K., and Ritzke, D. (1999) "The Impact of Gender on the Review of the Curricula Vitae of Job Applicants and Tenure Candidates: A National Empirical Study" *Sex Roles* 41 (7/8): 509–528.
- Topey, B. (2012) "Coin Flips, Credences, and the Reflection Principle" *Analysis* 72: 478–88
- Williamson, T. (2000) *Knowledge and its Limits*. Oxford: OUP.
- Williamson, T. (forthcoming) "Evidence of Evidence in Epistemic Logic" in M. Skipper and A. Steglich-Petersen (eds.) *Higher-Order Evidence: New Essays*. Oxford: Oxford University Press.