

This is a preprint of an article whose final and definitive form will be published in the Australasian Journal of Philosophy. The Australasian Journal of Philosophy is available online at: <http://www.tandf.co.uk/journals/>.

Avoiding Risk and Avoiding Evidence

Catrin Campbell-Moore and Bernhard Salow*

July 31, 2019

Abstract

It is natural to think that there is something epistemically objectionable about avoiding evidence, at least in ideal cases. We argue that this natural thought is inconsistent with a kind of risk-avoidance that is both widespread and intuitively rational. More specifically, we argue that if the kind of risk-avoidance recently defended by Lara Buchak is rational, avoiding evidence can be epistemically commendable.

In the course of our argument we also lay some foundations for studying epistemic value, or accuracy, when considering risk-avoidant agents.

Is it ever reasonable not to gather evidence? Sure it is. Gathering and processing evidence is almost never free, costing time and energy if nothing else. Even if it were free, you might know that the evidence doesn't bear on anything important. And even if the evidence were both free and relevant, you might still be worried that you'll miscalculate it.

But what about cases in which none of these worries arise? *Ideal cases*, in which (i) gathering the evidence incurs no cost whatsoever, (ii) the evidence is potentially relevant, and (iii) you're certain to process it rationally? Here, it seems wrong to ignore evidence. Good [1967] shows that classical decision theory agrees: it says that not gathering the evidence is always *instrumentally irrational*, a bad way of pursuing your goals. However, Wakker [1988] shows that decision theories that allow for *risk-avoidance* (a form of risk-aversion not permitted by classical decision theory) do not share this verdict. The two thus disagree about:

Look-I In ideal cases, one is *instrumentally* required to gather the evidence.

*Authors in alphabetical order. For helpful comments and discussion, we would like to thank Arif Ahmed, Richard Bradley, Adam Bjorndahl, Kevin Dorst, Kenny Easwaran, Branden Fitelson, Alan Hajek, Matt Hewson, Jason Konek, Ben Levinstein, Samir Okasha, Richard Pettigrew, Jack Spencer, Chris Stephens, and several anonymous referees, as well as audiences at the Universities of Bristol, Cambridge, and Glasgow, the London School of Economics, the Australian National University, and the 2017 Formal Epistemology Workshop.

So it looks as though evidence-avoidance can be rational (in ideal cases) if and only if risk-avoidance is.

Recently, however, Lara Buchak [2010] has pointed out that this reasoning ignores an important distinction. For, even if avoiding evidence is rational from an instrumental perspective, it might nonetheless be irrational from an epistemic one. This would mean that instrumental and epistemic rationality sometimes conflict; but that would hardly be surprising. Perhaps, then, defenders of risk-avoidance can accept:

Look-E In ideal cases, one is *epistemically* required to gather the evidence.

If so, they can preserve an important sense in which avoiding evidence is unreasonable.

In this paper, we close this gap: focusing on the picture of risk-avoidance developed by Buchak [2013, 2014], we argue that if risk-avoidance is rational, then Look-E fails. We leave open whether the correct reaction is to reject the rationality of risk-avoidance, or to conclude that, even from an epistemic point of view, there needn't be anything wrong with avoiding evidence.¹

The plan is as follows. Sections 1 and 2 present relevant background, explaining the connection between risk-avoidance and Look-I, as well as Buchak's risk-sensitive decision theory. Section 3 begins the main argument, making the case that agents are epistemically required to gather (or avoid) evidence if doing so is conducive towards securing epistemic goods; and hence that risk-sensitive agents are required to gather (or avoid) evidence if doing so maximizes risk-weighted expected epistemic value. Section 4 discusses how to measure epistemic value when working with risk-sensitive agents. Section 5 presents the core result showing that Look-E fails for risk-avoidant agents. Section 6 relaxes an assumption about how epistemic value is measured; and section 7 relaxes the assumption that rational agents revise their beliefs by conditionalization. Section 8 sums up.

1 Risk Avoidance and Look-I

Classical decision theory requires agents to maximize expected utility. This means that all rational aversion to risky options must be reflected in the agent's utility function: if it's rational to prefer a sure \$5 over a bet that pays \$10 if a fair coin lands heads (and nothing otherwise), this is only because the utility of gaining \$10 is less than twice that of gaining \$5. But many have argued that this does not capture all the ways in which rational agents can be risk-averse. Perhaps the best-known argument is that adjusting the utilities won't let us capture the *Allais Preferences*.²

Consider the following four lotteries

¹Strictly speaking, another option is to maintain that some risk-avoidance is rational, but not the kind described by Buchak's theory.

²See Allais [1953]. The preferences usually involve \$1m where we have \$1,000 and \$5m where we have \$2,000. These changes affect nothing of substance, but simplify the calculations in section 2.

	Ticket 1	Ticket 2-11	Ticket 12-100
L_1	\$1,000	\$1,000	\$1,000
L_2	\$0	\$2,000	\$1,000
L_3	\$1,000	\$1,000	\$0
L_4	\$0	\$2,000	\$0

Expected utility theory tells us that, regardless of how you value money, if you prefer L_1 to L_2 , you should prefer L_3 to L_4 : the sure outcome of the tickets 12-100 should not influence your preferences. However, many people favour L_1 over L_2 but L_4 over L_3 . Moreover, this preference seems to make sense: L_1 feels preferable to L_2 partly because it is safe; L_3 offers no such advantage over L_4 .

Cases like this motivate decision theories that incorporate risk-avoidance: a form of risk-aversion not reflected in the utilities assigned to the outcomes.³ Revising classical decision theory, however, means that we need to reexamine evidence-gathering. I.J. Good [1967] famously proved that, when faced with a choice between either (i) choosing an option now or (ii) gathering some cost-free evidence, conditionalizing on it, and then choosing the option that maximizes expected utility relative to the updated credences, (ii) will have a strictly higher expected utility than (i) whenever the new evidence might lead one to choose a different option. Given classical decision theory, which says that agents are instrumentally required to maximize expected utility, this establishes Look-I. But given a different theory, it obviously does not.

Moreover, no theory which permits the Allais preferences can recover an analogous result:⁴ Either when choosing between L_1 and L_2 , or when choosing between L_3 and L_4 , someone with the Allais preferences will want to avoid learning whether their ticket is one of 1–11. The reason, roughly, is that you are bound to ‘switch’ preferences in one of the two cases upon finding out whether your ticket is 1–11 or 12–100. For the choice between L_1 and L_2 on the one hand, and between L_3 and L_4 on the other, is the same if you know that your ticket is 1–11 and is inconsequential if you know that your ticket is 12-100; and so you can’t preserve your original preferences after receiving the information in both cases. But this means that, in the case in which you would switch, you know you would pick the lottery you currently consider worse if you were to find out whether your ticket is 1–11 before choosing.

This result is somewhat strange, but Buchak [2010] offers a helpful diagnosis. We can call evidence *instrumentally misleading* if it rationalizes an action that turns out to be worse than the one you would have performed if you hadn’t received that evidence. Evidence which is relevant, i.e. which rationalizes an action different from the one you would have performed otherwise, almost

³Following [Buchak, 2013], we use ‘risk-aversion’ to refer to the general phenomenon, and ‘risk-avoidance’ for any aspect of it not permitted by classical decision theory.

⁴Cf Buchak [2013, p.171-173]. Wakker [1988] establishes the even stronger claim that no theory which allows for any violations of the ‘Sure-Thing Principle’ (a principle of classical decision theory violated by the Allais preferences) can recover Look-I. Both assume, as will we, that rational agents make sequential decisions through *sophisticated choice* – a theory which, like that defended by McClennen [1990], instead has agents employ *resolute choice* can vindicate Look-I (albeit somewhat trivially).

always has a risk of being instrumentally misleading. Of course, this risk must be weighed against the possibility that the evidence is *instrumentally ‘truth-guiding’*, rationalizing an action that is better than the one you would have performed otherwise. For agents who follow expected utility theory this potential benefit always outweighs the risk; but for agents who prefer safe options, it may not.

Non-classical decision theories that allow for intuitively rational forms of risk-avoidance thus reject Look-I. But what about Look-E? To answer this, we focus on a particular theory of risk-avoidance: the one defended by Buchak herself.

2 Buchak’s Theory

Buchak’s decision theory departs from expected utility theory by allowing agents to engage in worst-case-scenario style reasoning. Moreover, it allows different agents to have different tendencies to engage in such reasoning. So, in addition to credences c and utilities U , agents have a risk profile r describing that tendency. Following Quiggin [1982],⁵ Buchak’s theory uses this risk profile to modify the weight given to an outcome depending on how it ranks relative to the others.

To see how this works in detail, note that ordinary expected utilities can be written, slightly non-standardly, as follows:

Expected utility Suppose an act, A , leads to outcomes o_1, \dots, o_n in states s_1, \dots, s_n , with $U(o_1) \leq \dots \leq U(o_n)$. Then

$$\begin{aligned} \text{Exp}_c U(A) &= U(o_1) \\ &\quad + (c(s_2) + \dots + c(s_n)) \cdot (U(o_2) - U(o_1)) \\ &\quad \dots \\ &\quad + \qquad \qquad \qquad c(s_n) \cdot (U(o_n) - U(o_{n-1})) \end{aligned}$$

Intuitively, the expected utility is here calculated by first taking the utility of the worst-case scenario; adding the improvement over the worst-case scenario secured in the second-worst-case scenario, weighted by the probability of securing at least that improvement; adding the improvement over the second-worst-case scenario secured in the third-worst-case scenario, weighted by the probability of securing at least *that* improvement; and continuing like this until all the possible improvements have been taken into account.

Buchak’s theory replaces *expected utility* calculations with *risk-weighted expected utility* calculations, where the risk profile, r , modifies the weight which possible improvements receive. Mathematically, r is an increasing function from $[0, 1]$ to $[0, 1]$ with $r(0) = 0$, $r(1) = 1$. It is used as follows:

Risk-weighted expected utility Suppose an act, A , leads to outcomes

⁵Cumulative prospect theory [Tversky and Wakker, 1995, Wakker, 2010] is a related theory.

o_1, \dots, o_n in states s_1, \dots, s_n with $U(o_1) \leq \dots \leq U(o_n)$. Then:

$$\begin{aligned} \text{RExp}_c^r U(A) &= U(o_1) \\ &\quad + r(c(s_2) + c(s_3) + \dots + c(s_n)) \cdot (U(o_2) - U(o_1)) \\ &\quad + \quad \quad r(c(s_3) + \dots + c(s_n)) \cdot (U(o_3) - U(o_2)) \\ &\quad \dots \\ &\quad + \quad \quad \quad \quad \quad r(c(s_n)) \cdot (U(o_n) - U(o_{n-1})) \end{aligned}$$

If an agent has the risk-neutral risk profile $r(x) = x$, this is equivalent to the expected utility formula given above. But if r is more interesting, it yields different results.

Suppose, for example, that $r(x) < x$. Then the weight given to the potential improvements will be less than it is in the expected utility calculation; and so the relative weight given to the worst-case scenario will be greater. This means, for example, that one can have utilities that are linear with money and still prefer the sure \$5 (which has an $\text{RExp}_c^r U$ of 5) over the possible \$10 (which has an $\text{RExp}_c^r U$ of $0 + r(0.5) \cdot (10 - 0) < 5$) – one thus prefers the safe choice not because one values the ‘second’ \$5 less, but because one gives more weight to the worst-case scenario when evaluating one’s options. We will call such risk profiles ‘risk-avoidant’; following Buchak, $r(x) = x^2$ will be our main example.⁶

The presence of the risk profile allows the theory to rationalize the Allais preferences described earlier. Consider, for example, an agent for whom utility and money are interchangeable and whose credences match the objective probabilities, but who is risk-avoidant in line with the risk profile $r(x) = x^2$. Then it’s easy to check that $\text{RExp}_c^r U(L_1) > \text{RExp}_c^r U(L_2)$ but $\text{RExp}_c^r U(L_3) < \text{RExp}_c^r U(L_4)$. Moreover, if the agent learns that her ticket is in the range 1–11, she will come to prefer L_2 to L_1 , since L_2 has higher risk-weighted expected utility according to the updated credences. For this agent, then, finding out whether her ticket is in the range 1–11 before choosing is equivalent to taking L_2 , and not finding this out is equivalent to taking L_1 . So, since $\text{RExp}_c^r U(L_1) > \text{RExp}_c^r U(L_2)$, she prefers not finding out. So Look-I fails on this theory in exactly the way described earlier.

In what follows we will assume that if risk-avoidance is rationally permissible, then the kind of risk-avoidance described by Buchak’s theory, with some choice of a risk-avoidant r , is rationally permissible.⁷ Part of the motivation is dialectical: Buchak is the one who suggests that advocates of risk-aversion retreat from Look-I to Look-E. Part of the motivation is principled: Buchak’s theory is elegant, intuitive, and well-developed, and thus a leading a contender for what rational risk-aversion might look like. And part of the motivation is pragmatic. An important attraction of Buchak’s theory is that it neatly separates out an agent’s attitude to risk from her beliefs on the one hand, and her utilities on the other.⁸

⁶Wakker [2010, p.174] and Buchak [2013, p. 63] argue convincingly that risk-avoidance should really be identified with a strictly stronger property of risk-profiles called ‘convexity’; but the weaker (and simpler) condition that $r(x) < x$ is sufficient for our result.

⁷We also assume that at least one such r is differentiable; but this looks innocuous.

⁸See Buchak [2013, pp.34-47, 53-56].

This makes her theory especially well-suited to studying the interaction between risk-aversion and *epistemic* rationality: because the risk profile is separated from the utilities, we can easily study the theory’s predictions when ‘what matters’ isn’t fixed by the agent’s own desires; and because the agent’s attitude to risk is separated from her beliefs, we can use normal (i.e. probabilistic) epistemology when we do so.

3 Epistemic Rationality of Actions

Our question is whether a risk-avoidant agent is always epistemically required to gather evidence. This presupposes that epistemic norms apply to actions, such as the gathering or avoiding of evidence. It isn’t obvious that they do: epistemic norms most naturally apply to doxastic states such as beliefs and credences, or to belief-producing procedures such as inference to the best explanation or conditionalization. Equally, however, it isn’t obvious that they do not. When we say that it is irresponsible to consult only one kind of source, or wise to forego a quick but ambiguous test in favour of running a more thorough analysis later, we seem to engage in some form of epistemic evaluation. We will take this appearance at face value and assume that actions can be described as ‘epistemically rational’ (or ‘commendable’ or ‘responsible’ if those sound better to you); if that’s wrong, the retreat from Look-I to Look-E that we are criticizing is equally doomed, albeit for less interesting reasons.

Epistemic rationality, as applied to actions, is plausibly understood in consequentialist terms: an action is epistemically rational if it promotes epistemic goods, such as true beliefs and accurate credences. That’s not to say that epistemic rationality is consequentialist in all domains. Several philosophers have recently argued against the idea that, as applied to *beliefs* or *credences*, epistemic rationality is consequentialist. Suppose that your overall evidence weakly points towards p , but that someone powerful guarantees that you’ll later find out for sure whether p if and only if you now disregard the inconclusive information and adopt a credence of 0.5. Then a naive consequentialist approach to the rationality of beliefs will make the counter-intuitive prediction that it would be epistemically rational for you to disregard your evidence and believe p to degree 0.5; for the loss in accuracy you’re likely to incur by ignoring your current evidence is outweighed by the gain in accuracy when you later conform your credence to the conclusive evidence which will then be available. Perhaps more sophisticated forms of consequentialism can avoid such conclusions.⁹ But it’s worth noting that the analogous consequence of consequentialism about the epistemic rationality of *actions* is actually very intuitive. If I know that running a quick first-pass test into whether p will destroy the only sample and thus prevent me from carrying out a more conclusive analysis in the future, it seems epistemically responsible to refrain. While it may be weird to factor in the consequences on the accuracy of one’s later beliefs, or beliefs in other propositions, when wondering whether

⁹Greaves [2013], Berker [2013], Caie [2013] and Carr [2017] argue that they can’t. Konek and Levinstein [2017] are more optimistic.

to believe p , it is perfectly normal to do so when wondering what evidence to gather. Reasons to doubt consequentialism elsewhere in epistemology thus do not apply when the evaluations concern actions.

We will thus assume that the epistemic rationality of gathering and avoiding evidence can be understood in terms of its anticipated effects on matters of epistemic value, specifically the accuracy of one's credences.¹⁰ However, we will side-step a different controversial issue: how the accuracy of credences regarding different propositions contributes to the overall epistemic value of a credal state.¹¹ It is natural to think that not all propositions contribute equally (that's why we should investigate ambitious scientific theories instead of counting blades of grass); but also that there is a large amount of incommensurability (that's why it's fine to make slow progress on arcane issues in philosophy instead of reading Wikipedia all day). Fortunately, we can side-step these issues by considering only how your action affects your accuracy on a single proposition, which the evidence in question is supposed to bear on. In this simple case, accuracy and value plausibly coincide.

Before saying more about how accuracy is measured, we need to discuss how exactly it should be pursued. If we assume classical decision theory, it is natural to think that an action is epistemically rational only if it maximizes *expected* epistemic value. But if we're working with Buchak's risk-sensitive theory, it's more natural to think that an action is epistemically rational only if it maximizes *risk-weighted expected* epistemic value.

This idea can be developed in different ways depending on which risk profile is used in the risk-weighted expected utility calculations. On a subjective view, it is the agent's: each agent is epistemically required to maximize risk-weighted expected epistemic value relative to their own risk profile. On an alternative view, the agent's actual risk profile is irrelevant – what matters is only which ones are rationally defensible. So an action is epistemically permissible if it maximizes risk-weighted expected epistemic value according to *some* permissible risk profile, and epistemically required if it maximizes risk-weighted expected epistemic value according to *every* permissible risk profile. It mostly won't matter which of these we go with, and so we will presuppose the simpler, subjective picture; when it does matter, we will say so.

There are two reasons why epistemic rationality should depend on *risk-weighted* expected values. The first is that this fits better with the picture painted by risk-sensitive decision theory. This theory departs from classical decision theory precisely because it maintains that a rational agent's attitude towards risk should not be understood as a generalization about her goals (e.g.

¹⁰Buchak seems sympathetic to this assumption; in defending the suggestion that we can retain Look-E while rejecting Look-I, she writes that “what you have to do in connection with maximizing instrumental value is not necessarily constrained by what you have reason to do in connection with maximizing epistemic value” [2010, p.105], which suggests that epistemic rationality is a matter of pursuing epistemic value.

¹¹As the examples we're about to discuss bring out, this issue is particularly pressing if we want to use accuracy evaluations to determine the rationality of actions; Greaves [2013], among others, argues that it does not arise if we evaluate only the rationality of beliefs or belief-revision procedures.

that she has a concave utility function) but as a fact about *how she rationally pursues* her goals.¹² But if that’s true, then we should allow for variation in this feature even when subjective utilities are swapped for epistemic values to determine epistemic obligations. The second reason is that it’s natural to expect epistemic and instrumental norms to align for agents whose goals align with what’s epistemically valuable. But if we tie epistemic obligations to expected rather than risk-weighted expected value, we lose this convergence. So epistemic obligations should depend on risk-weighted expected epistemic value.¹³

Our claim, remember, is that Look-E fails in the risk-sensitive setting. At first sight, this may look like an obvious corollary of the failure of Look-I and the view that epistemic rationality requires us to maximize risk-weighted expected epistemic value, much as the moral permissibility of avoiding evidence is an obvious corollary of the failure of Look-I and the view that morality requires us to maximize risk-weighted expected moral value. A closer look, however, reveals that this isn’t so. For we have significantly fewer ‘degrees of freedom’ when constructing counterexamples to Look-E than when constructing counterexamples to Look-I. This is because the epistemic utilities associated with learning are fixed by the agent’s credences after learning the evidence, which are in turn determined by the credences used in deciding whether to gather the evidence in the first place. So we cannot simply stipulate probabilities and utilities independently. For example, (unlike in the moral case) we cannot simply use the Allais case, stipulating that the utilities in question represent epistemic, rather than instrumental, value.

This explanation also brings us to our final piece of set-up. So far, we have said nothing about how epistemic value, or accuracy, is measured; but without measuring accuracy, we cannot calculate the risk-weighted expected accuracy of gathering or avoiding evidence. Showing how to measure accuracy will be our next task; it deserves its own section, since, when working with risk-weighted expectations, we have to measure accuracy in a somewhat non-standard way.

4 Accuracy

We are interested in measuring the epistemic value, or accuracy, of our agent’s credence in some proposition of interest, X . We can think of this as a measure of how ‘close’ the credence is to the truth-value, i.e. to 1 if X is true and to 0 if X is false. However, there are many ways of measuring such closeness. The obvious absolute-difference measure, on which the distance between her credence, x , and the truth value, v , is simply $|v - x|$, and the closeness between them is hence $-|v - x|$, has various drawbacks; a popular alternative is the Brier Score,

¹²See especially Buchak [2013, p.34-36].

¹³It’s worth noting, however, that as long as one accepts our arguments in the next section that epistemic value should be measured by an r -proper accuracy measure (with the agent’s own risk-sensitive r), Look-E will fail even if epistemic rationality is determined by *risk-neutral* rather than *risk-sensitive* expectations.

which measures distance as the square of the absolute difference:¹⁴

$$\text{BrierScore}(x, v) := -(v - x)^2.$$

Following the literature we will not defend a particular measure but will instead adopt some general constraints and show that our claim holds given any measure \mathcal{A} meeting these constraints. Two of our constraints are uncontroversial:¹⁵

- \mathcal{A} is (weakly) truth directed, i.e.
 - If $x_1 < x_2 < 1$, then $\mathcal{A}(x_1, 1) \leq \mathcal{A}(x_2, 1)$,
 - If $x_1 > x_2 > 0$, then $\mathcal{A}(x_1, 0) \leq \mathcal{A}(x_2, 0)$.
- \mathcal{A} is 0/1 symmetric,¹⁶ i.e.
 - $\mathcal{A}(x, 1) = \mathcal{A}(1 - x, 0)$.

Our final condition, by contrast, requires discussion – not only because it is more controversial, but also because it will need to be spelled out in a non-standard way.

Many authors writing on the connection between credence and truth have argued that rational agents should be immodest: they should regard their credence in X as giving the best shot at the truth, compared to any other particular credence distribution.¹⁷ This thought comes in a weak and a strong version. On the weak version, it merely requires that rational agents should regard their credence in X as giving them *no worse* a shot at the truth than any other; on the strong version, it requires that agents should regard their credence as giving them a *better* shot at the truth than any other.

One standard motivation is that someone who isn't immodest exhibits internal conflict. This clearly motivates weak immodesty: if you think some other credence distribution has a better shot at the truth than yours, you seem divided in much the way as when you believe a contradiction. But it may also motivate strong immodesty: if you think that some other credence distribution has just as good a shot at the truth as yours does, it feels as though, in some sense, you're also taking this other credence, thus bearing rival attitudes towards X .

¹⁴See especially Joyce [2009], Leitgeb and Pettigrew [2010a] and Pettigrew [2016, ch. 4]. To simplify the interaction with Buchak's theory, we consider *accuracy* rather than the more standard *inaccuracy*.

¹⁵See e.g. Joyce [2009]. Continuity will be met by our specific example, but will not be required for the general results. Other standard constraints, e.g. those known as 'Normality' and 'Separability', concern the relationship between the accuracy of particular beliefs and the accuracy of the overall belief state; since we are focusing only on the accuracy of a single belief, they have no bearing on our discussion.

¹⁶In fact, our theorem only requires the much weaker constraint that $\mathcal{A}(x, 0) \leq \mathcal{A}(x, 1)$ when $x > 1/2$.

¹⁷We take talk of 'best shot at the truth' from Horowitz [forthcoming]. Other sympathetic discussions include Lewis [1971], Oddie [1997], Greaves and Wallace [2006], Gibbard [2008], and Joyce [2009].

Another standard motivation is that modest states are problematically unstable. If you fail to be strongly immodest, you have no epistemic reason to stick with your beliefs: if the opportunity arises, you might as well abandon them for one of the alternatives you think equally good. And if you fail to even be weakly immodest, you will not only lack reason to remain, but also have positive reason to switch. Since such changes of mind look epistemically irrational (no new evidence is required to initiate them), this again suggests that rational agents must be immodest.

Most of the literature treats these arguments as establishing strong, rather than merely weak, immodesty;¹⁸ and we will follow this trend in our initial discussion. However, the arguments for strong immodesty are less conclusive than those for weak immodesty. Moreover, there may be special reasons why strong immodesty is too strict a constraint when allowing for risk-avoidance; for while there are ways of measuring the accuracy of *one's attitude to a single proposition* that vindicate strong immodesty even in the risk-avoidant setting, Campbell-Moore and Levinstein [ms] show that there are difficulties to provide accuracy measures of *entire credal state* that do so. Now, for the reasons discussed in section 3, the former kind of measure is more obviously relevant to our project, and so this is not a conclusive reason to reject strong immodesty; but it may, nonetheless, temper one's enthusiasm for this requirement. So, after presenting the initial argument in a way that presupposes strong immodesty, we will, in section 6, consider how it fares given only weak immodesty.

We can leverage immodesty (of either kind) into a precise constraint on \mathcal{A} given two things: a sufficient condition for a distribution being rational, and an account of 'how good a shot' a distribution gives to a particular credence in X .

Plausibly, a sufficient condition for the rationality of a distribution is that it's probabilistically coherent. After all, for any value x , one's evidence might establish that the objective chance of the target proposition X is x ; and an agent with such evidence would be rationally permitted to hold $c(X) = x$ and $c(\neg X) = 1 - x$.¹⁹ So any such distribution should give itself a/the best shot at the truth.

Understanding how good a shot a distribution c gives to credence y is slightly subtler. If we were working in the risk-neutral framework, it would be natural to identify a credence's shot at the truth with its expected accuracy. We could then guarantee that all probabilistic agents are weakly immodest by requiring \mathcal{A} to be *weakly proper*: for all probabilistic c , $\text{Exp}_c \mathcal{A}(c(X)) \geq \text{Exp}_c \mathcal{A}(y)$. And we could guarantee that all probabilistic agents are strongly immodest by requiring \mathcal{A} to be *strictly proper*: for every probabilistic c and $y \neq c(X)$, $\text{Exp}_c \mathcal{A}(c(X)) > \text{Exp}_c \mathcal{A}(y)$.

However, we are assuming that some rational agents are risk-sensitive. And

¹⁸See e.g. Oddie [1997], Greaves and Wallace [2006], and Joyce [2009]. Maher [2002] and Gibbard [2008] raise doubts about using strict immodesty to constrain accuracy measures.

¹⁹ See Joyce [2009] for this argument. Our result actually goes through with a weaker assumption: that there is some range such that if x falls within that range, a distribution with $c(X) = x, c(\neg X) = 1 - x$ can be rational. (Thanks to Jason Konek and Richard Pettigrew for discussion on this point.)

it seems more natural to say that risk-sensitive agents consider those credences to give them a/the ‘best’ shot at the truth that maximize *risk-weighted* expected accuracy. So we define:

- \mathcal{A} is *weakly r-proper* if for all probabilistic c , $\text{RExp}_c^r \mathcal{A}(c(X)) \geq \text{RExp}_c^r \mathcal{A}(y)$;
- \mathcal{A} is *strictly r-proper* if for all probabilistic c , $\text{RExp}_c^r \mathcal{A}(c(X)) > \text{RExp}_c^r \mathcal{A}(y)$ whenever $y \neq c(X)$.

If we are to have a single measure of accuracy that is appropriate for every rational agent, and which also guarantees weak or strong immodesty, it must be weakly or strictly r-proper for every rational risk profile r . Unfortunately, Pettigrew [2016, Section 16.4] shows that for many pairs of risk profiles r_1 and r_2 (including e.g. $r_1(x) = x$ and $r_2(x) = x^2$), no accuracy measure can be both strictly r_1 -proper and weakly r_2 -proper.²⁰ Assuming the members of at least one such pair are both rational, this shows that there is no accuracy measure such that every rational agent comes out as strongly immodest when epistemic value is determined by that measure; it also suggests that it’s even implausibly restrictive to just require the accuracy measure to be such that every rational agent comes out as weakly immodest relative to it.

If we want both immodesty and risk-avoidance, we thus need to reject a picture of epistemic value on which there is a single such quantity, which every rational agent is epistemically required to pursue.²¹ The most similar (‘objectivist’) picture holds that for each risk profile there is a unique value quantity which all agents with that risk profile are epistemically required to maximize. (Compare: we may reject ‘single value quantity’ consequentialism in moral philosophy on the grounds that their histories and personal relationships require different people to prioritize differently the well-being of various individuals; but we might still say that there is a single function, which determines for each kind of history and personal circumstance a unique quantity which people with that history and circumstances are morally required to pursue.) We would then place as a constraint on this function that it associates each risk profile r with an r -proper accuracy measure, thus ensuring immodesty.

An alternative, ‘subjectivist’ picture maintains that agents choose a particular accuracy measure, subject to certain constraints, and are then epistemically required to pursue accuracy as evaluated by that measure. We could then ensure immodesty by imposing as one of the constraints that the chosen accuracy measure be r -proper relative to the agent’s risk profile.

A third alternative is a ‘permissivist’ picture. On this account agents are not required to pursue epistemic value as measured by any particular measure. Rather, given an agent’s risk profile there are a number of legitimate epistemic

²⁰Pettigrew’s argument isn’t as general as he states it, but a slightly fixed-up version shows this for all pairs where there is some $x > 1/2$ such that $r_1(x) \neq r_2(x)$ and $r_1(x), r_2(x) > 1/2$.

²¹As a referee pointed out, this may be a problem for risk-weighted expected utility theory all by itself; after all, part of the attraction of this theory is that it allows us to separate an agent’s attitude to risk from her utilities – an attraction which is somewhat undercut when, as happens here, the former still imposes substantial constraints on the latter.

value measures; and an agent is epistemically permitted to perform any action that maximizes risk-weighted epistemic value on some legitimate accuracy measure or other. We could then impose immodesty by maintaining that \mathcal{A} is a legitimate accuracy measure for an agent with risk profile r only if \mathcal{A} is r -proper.

All three of these accounts hold that an agent’s actual risk profile matters to her epistemic obligations. As mentioned in section 3, one might want to deny this. This leads naturally to a view we will call *super-permissivism*: instead of aggregating verdicts across legitimate accuracy measures like the permissivist, the super-permissivist aggregates them across legitimate *standards of evaluation* consisting of both a risk profile and an accuracy measure. On this picture, Immodesty is preserved by maintaining that $\langle r, \mathcal{A} \rangle$ is a legitimate standard of evaluation only if \mathcal{A} is r -proper.

We have now encountered all the constraints on accuracy measures we will need: weak truth directedness, 0/1 symmetry and strict r -propriety. We think they are all reasonable. We know that all three can be jointly met in the case of a risk-neutral agent (for whom $r(x) = x$), since in that case popular measures such as the Brier Score will do. However, we also know that this measure won’t work for many risk-avoidant agents, including those with risk profile $r(x) = x^2$. Fortunately, there are suitable measures even for risk-avoidant profiles; relative to $r(x) = x^2$, for example, the following measure meets all three conditions:²²

$$\text{AltBS}(x, v) := \frac{-(v - x)^2}{\max\{x, 1 - x\}}$$

With these constraints on accuracy measures in place, we can finally address our central question: whether epistemic rationality ever permits avoiding evidence.

5 Avoiding Evidence

The principle we are interested in is:

Look-E In ideal cases, one is epistemically required to gather the evidence.

We will show that, if Buchak’s risk-avoidant agents are rational and our account of epistemic requirements and epistemic value are accepted, then there are counterexamples to this principle. In fact, we will show that there are even counterexamples to the weaker

Weak Look-E In ideal cases, one is always epistemically *permitted* to gather the evidence.

²²In unpublished work Adam Bjorndahl and Ben Levinstein develop methods which can be applied to almost any risk profile r to yield an inaccuracy measure \mathcal{A} that has all the properties we require (relative to r).

Kothiyal et al. [2010] explain how to find scoring rules that elicit the beliefs of agents who follow risk-weighted expected utility theory. But these rules are not r -proper; there is merely a method for reading off beliefs from the reported values.

Our discussion will use a set of abstract examples. In these, c is a probability function that makes a piece of evidence, E , relevant to our proposition of interest, X . And we assume that it's certain that, were our agent to learn whether E , she would conditionalize: she would adopt $c(X|E)$ as her new credence in X if E is true, and adopt $c(X|\neg E)$ otherwise. (We will show in 7 that we can relax the conditionalization assumption.) Since we identify epistemic value with the accuracy of the agent's credence in X , the value secured by gathering or avoiding the evidence depends on X and on our agent's posterior credence – and hence (in the case of gathering) on E :

	$E \wedge \neg X$	$\neg E \wedge \neg X$	$\neg E \wedge X$	$E \wedge X$
AVOID	$\mathcal{A}(c(X), 0)$	$\mathcal{A}(c(X), 0)$	$\mathcal{A}(c(X), 1)$	$\mathcal{A}(c(X), 1)$
GATHER	$\mathcal{A}(c(X E), 0)$	$\mathcal{A}(c(X \neg E), 0)$	$\mathcal{A}(c(X E), 1)$	$\mathcal{A}(c(X \neg E), 1)$

Since c determines the probability of each of the four possibilities, it's straightforward, given a particular c , r , and \mathcal{A} , to calculate both $\text{RExp}_c^r \mathcal{A}(\text{GATHER})$ and $\text{RExp}_c^r \mathcal{A}(\text{AVOID})$.

In this set-up, Oddie [1997] shows that, in the risk-neutral setting, the expected accuracy of gathering the evidence always exceeds that of avoiding it:

Theorem 5.1. *If \mathcal{A} is strictly proper, then for all c ,*

$$\text{Exp}_c \mathcal{A}(\text{GATHER}) > \text{Exp}_c \mathcal{A}(\text{AVOID}).$$

So if rationality requires risk-neutrality, both Look-E and Weak Look-E are true – regardless of whether we opt for an objectivist, subjectivist, or permissivist picture of the accuracy measures.

When we introduce risk-aversion, by contrast, there will always be cases where the risk-weighted expected accuracy of avoiding the evidence exceeds that of gathering it:

Theorem 5.2. *Suppose that r is differentiable and risk-avoidant. Then there is some probability distribution c such that, for every \mathcal{A} which is weakly truth-directed, 0/1-symmetric, and strictly r -proper,*

$$\text{RExp}_c^r \mathcal{A}(\text{GATHER}) < \text{RExp}_c^r \mathcal{A}(\text{AVOID})$$

We will present a sketch of the proof, but readers should feel free to skip this. For full details see appendix A.

Proof Sketch. For concreteness, let us consider some specific numbers: $c(X) = 0.65$, $c(X|E) = 0.7$ and $c(X|\neg E) = 0.6$. The decision table thus becomes:

credence	$E \wedge \neg X$	$\neg E \wedge \neg X$	$\neg E \wedge X$	$E \wedge X$
	0.15	0.2	0.3	0.35
GATHER	$\mathcal{A}(0.7, 0)$	$\mathcal{A}(0.6, 0)$	$\mathcal{A}(0.6, 1)$	$\mathcal{A}(0.7, 1)$
AVOID	$\mathcal{A}(0.65, 0)$	$\mathcal{A}(0.65, 0)$	$\mathcal{A}(0.65, 1)$	$\mathcal{A}(0.65, 1)$

Note that $\text{RExp}_c^r \mathcal{A}(\text{AVOID}) = \text{RExp}_c^r \mathcal{A}(0.65) > \text{RExp}_c^r \mathcal{A}(0.7)$ by \mathcal{A} 's strict r -propriety. So to show that $\text{RExp}_c^r \mathcal{A}(\text{AVOID}) > \text{RExp}_c^r \mathcal{A}(\text{GATHER})$, it suffices to show that $\text{RExp}_c^r \mathcal{A}(0.7) > \text{RExp}_c^r \mathcal{A}(\text{GATHER})$.

To determine $\text{RExp}_c^r \mathcal{A}(\text{GATHER})$, we need to know how the outcomes are ordered. Given the credences chosen, \mathcal{A} 's truth-directedness and 0/1-symmetry ensure that, in the table, they go from worst to best. So we have

$$\begin{aligned} \text{RExp}_c^r \mathcal{A}(\text{GATHER}) &= \mathcal{A}(0.7, 0) \\ &\quad + r(0.85) \cdot (\mathcal{A}(0.6, 0) - \mathcal{A}(0.7, 0)) \\ &\quad + r(0.65) \cdot (\mathcal{A}(0.6, 1) - \mathcal{A}(0.6, 0)) \\ &\quad + r(0.35) \cdot (\mathcal{A}(0.7, 1) - \mathcal{A}(0.6, 1)) \end{aligned}$$

We also have

$$\text{RExp}_c^r \mathcal{A}(0.7) = \mathcal{A}(0.7, 0) + r(0.65) \cdot (\mathcal{A}(0.7, 1) - \mathcal{A}(0.7, 0))$$

Rearranging these shows that

$$\begin{aligned} \text{RExp}_c^r \mathcal{A}(0.7) &> \text{RExp}_c^r \mathcal{A}(\text{GATHER}) \\ \text{iff } \mathcal{A}(0.7, 0) &+ \frac{r(0.65) - r(0.35)}{r(0.85) - r(0.35)} (\mathcal{A}(0.7, 1) - \mathcal{A}(0.7, 0)) \\ &> \mathcal{A}(0.6, 0) + \frac{r(0.65) - r(0.35)}{r(0.85) - r(0.35)} (\mathcal{A}(0.6, 1) - \mathcal{A}(0.6, 0)) \end{aligned}$$

Note that if we replace $\frac{r(0.65) - r(0.35)}{r(0.85) - r(0.35)}$ by $r(0.7)$, the second inequality becomes $\text{RExp}_c^r \mathcal{A}(0.7) > \text{RExp}_c^r \mathcal{A}(0.6)$ for a probability distribution c with $c(X) = 0.7$, and thus holds by strict r -propriety. So, since by truth-directedness $\mathcal{A}(0.7, 1) - \mathcal{A}(0.7, 0) > \mathcal{A}(0.6, 1) - \mathcal{A}(0.6, 0) > 0$, it suffices to show that

$$\frac{r(0.65) - r(0.35)}{r(0.85) - r(0.35)} > r(0.7).$$

This inequality holds for $r(x) = x^2$. So, for this choice of r and c we have $\text{RExp}_c^r \mathcal{A}(\text{AVOID}) > \text{RExp}_c^r \mathcal{A}(\text{GATHER})$ for any \mathcal{A} . For other r , this inequality might not hold given the particular numbers; however, provided that r is risk-avoidant, one can always find numbers for which the corresponding inequality does hold. So the result holds for any risk-avoidant r .

Moreover, one can show that the inequality will hold whenever the values of $c(X)$, $c(X|E)$ and $c(X|\neg E)$ are close enough (in the sense made precise in corollary A.2); so the range of cases in which $\text{RExp}_c^r \mathcal{A}(\text{AVOID}) > \text{RExp}_c^r \mathcal{A}(\text{GATHER})$ is reasonably wide. \square

Clearly the theorem shows that Weak Look-E fails, given the rationality of risk-avoidance, on both the subjectivist and objectivist picture of accuracy measures. Given its generality, it shows this even if there are constraints on accuracy measures or on risk profiles beyond those discussed here.

Interestingly, the theorem also establishes that Weak Look-E fails on the permissivist picture, which requires us to avoid the evidence only if all the legitimate measures recommend avoiding. For the theorem states not only that for every accuracy measure there will be an example in which it recommends avoiding the evidence (which is sufficient to refute Weak Look-E on the subjectivist and objectivist pictures but only Look-E on the permissivist picture), but also that there will be a single example in which all the relevant accuracy measures recommend against gathering the evidence (which is needed to refute Weak Look-E on the permissivist picture).

On the super-permissivist picture, however, the theorem only establishes failures of Look-E. In fact, Weak Look-E should hold on this picture as risk-neutrality is surely rationally permissible; and by theorem 5.1 standpoints of evaluation using a risk-neutral risk profile will always recommend gathering the evidence. So gathering the evidence is always rationally permissible.

In summary: all four pictures of how accuracy considerations determine epistemic obligations lead to the conclusion that Look-E fails if Buchak-style risk-avoidance is rational; and all but one lead to the conclusion that Weak Look-E fails as well.

Buchak offered an informal explanation of the failure of Look-I in terms of the risk of receiving *instrumentally misleading* evidence. A similar explanation is available here. Evidence can be *epistemically misleading*, moving an agent’s credence away from the actual truth value. Suppose, for example, that E is evidence for and $\neg E$ evidence against X ; then in the state $E \wedge \neg X$, the evidence will lead the agent to increase her credence in the false proposition X , while in $\neg E \wedge X$ it leads her to decrease her credence in the true proposition X . In both cases, the epistemic value decreases by gathering the evidence. Drawn out in our table, we have:

	$E \wedge \neg X$	$\neg E \wedge \neg X$	$\neg E \wedge X$	$E \wedge X$
AVOID	$\mathcal{A}(c(X), 0)$	$\mathcal{A}(c(X), 0)$	$\mathcal{A}(c(X), 1)$	$\mathcal{A}(c(X), 1)$
	\vee	\wedge	\vee	\wedge
GATHER	$\mathcal{A}(c(X E), 0)$	$\mathcal{A}(c(X \neg E), 0)$	$\mathcal{A}(c(X E), 1)$	$\mathcal{A}(c(X \neg E), 1)$

Of course, E and $\neg E$ can also be truth-guiding, as in the two scenarios in which X and E have the same truth value. For the risk-neutral agent, this potential truth-guiding benefit always outweighs the risk of being misled; but for the risk-avoidant agent, we’ve shown that it might not.²³

This completes our main argument. Before closing, however, we will show that much of our conclusion can be established even if we significantly weaken

²³What, in general, characterizes the cases where Look-E fails? An extension of our result shows that it fails whenever $c(X|E)$ and $c(X|\neg E)$ are close enough to $c(X)$, i.e. the evidence in question isn’t too decisive (and $c(X) \neq 1/2$). See corollary A.2 in the Appendix.

This is different to the instrumental case; there Buchak [2010, Appendix C] observes that Look-I fails when the evidence is reasonably informative in addition to not being overly decisive. But this difference makes sense: in the instrumental case, the evidence needs to be reasonably informative to affect the agent’s actions; in the epistemic case, any E that is evidentially relevant at all will affect the agent’s credences.

two of the assumptions: that accuracy measures must be *strictly* r -proper and that rational risk-sensitive agents update by conditionalization.

6 Weak Propriety

The previous section discussed the status of Look-E and Weak Look-E, on the assumption that accuracy measures must be strictly r -proper. However, the arguments for strict r -propriety are less compelling than those for weak r -propriety. And, as we mentioned in section 4, there are specific worries that in the risk-sensitive setting strict r -propriety is too demanding. It's thus worth considering what happens to Look-E and Weak Look-E if we assume only that accuracy measures should be *weakly* r -proper.

Let us begin, again, with the risk-neutral case. The standard result is easily adapted to show

Theorem 6.1. *If \mathcal{A} is weakly proper, then for all c , $\text{Exp}_c \mathcal{A}(\text{GATHER}) \geq \text{Exp}_c \mathcal{A}(\text{AVOID})$.*

This establishes Weak Look-E on any of our four pictures: gathering evidence might not always be required, but it is never forbidden.

Moreover, there is an account on which this, combined with theorem 5.1, can be used to defend Look-E. When discussing the permissivist picture, we aggregated recommendations across different measures by universal agreement: an option is required only when it uniquely maximizes value according to every measure. However, we could instead have aggregated recommendations by a kind of 'pareto dominance': an option is required if it (perhaps non-uniquely) maximizes value according to every measure, and uniquely maximizes value according to at least one measure. Combined with the plausible thought that *some* legitimate measures are strictly proper, theorem 5.1 and theorem 6.1 then show that gathering evidence is always required after all.

So much for the risk-neutral case. What about the risk-sensitive one? Here too it's easy to adapt the result to weakly r -proper measures:

Theorem 6.2. *Suppose that r is differentiable and risk-avoidant. Then there is some probability distribution c such that, for every \mathcal{A} which is weakly truth-directed and 0/1-symmetric,*

- *if \mathcal{A} is strictly r -proper, then $\text{RExp}_c^r \mathcal{A}(\text{GATHER}) < \text{RExp}_c^r \mathcal{A}(\text{AVOID})$*
- *if \mathcal{A} is weakly r -proper, then $\text{RExp}_c^r \mathcal{A}(\text{GATHER}) \leq \text{RExp}_c^r \mathcal{A}(\text{AVOID})$*

Clearly this is enough to show that Look-E fails: we are not always required to gather evidence. However, since, on most pictures, the retreat to weak propriety left us with no argument for Look-E even in the risk-neutral case, this does not establish a connection between risk-avoidance and epistemic norms on evidence gathering.

On some pictures of the accuracy measures, we can still demonstrate such a connection. Most saliently, given the discussion of the risk-neutral case, we can

do so on a version of the permissivist picture on which verdicts are aggregated via pareto dominance. For then Look-E and Weak Look-E both hold in the risk-neutral case. And Look-E and Weak Look-E both fail in the risk-avoidant case on the plausible assumption that for some risk-avoidant risk profile there is at least one legitimate strictly r -proper accuracy measure.²⁴

For similar reasons, we get a connection on a pareto dominance version of the super-permissivist picture. For Look-E holds, on this view, if risk-neutrality is rationally required. But if risk-avoidance is rational, theorem 6.2 shows that Look-E fails (assuming always that *some* strictly r -proper measures are acceptable). So the rationality of risk-avoidance makes a difference to Look-E (though not to Weak Look-E, which holds whether risk-avoidance is rational or not).

We also get a connection on the subjectivist picture, on which agents can pick any accuracy measure meeting certain constraints. For even if we allow that strict r -propriety is not a constraint on how agents make their choice, the constraints should still *allow* some strictly r -proper measures. The original 5.2 shows that those agents who have such strictly r -proper measures will be required to avoid evidence. So the rationality of risk-avoidance refutes Weak Look-E, a principle that is true if we're required to be risk-neutral.

On other pictures, we can only give plausibility arguments. Given the objectivist picture, for example, what we need to get violations of Weak Look-E is that some risk-avoidant profile is associated with an accuracy measure for which the inequality is strict. If any such risk profile is associated with a strictly r -proper measure, that is sufficient for the failure of Weak Look-E. But even if every risk profile is associated only with a weakly r -proper measure, it would still be surprising if in all the cases produced by 6.2, the weak inequality held because GATHER and AVOID have exactly the same expected accuracy. This does not conclusively establish a connection between risk-avoidance and evidence-avoidance; further constraints on accuracy measures would be required to do that. But it does make such a connection overwhelmingly plausible. Similar things are true for a version of the permissivist picture which aggregates by universal verdicts rather than pareto dominance.

There is only one view on which we do not even get a plausibility argument: a version of super-permissivism that aggregates by universal agreement. On this view, there is no connection: Look-E will fail even if risk-neutrality is rationally required; and Weak Look-E will be true even if risk-avoidance is rationally permissible.

Based on these results, summarized in table 1, we conclude that, (almost) regardless of which picture one adopts, our theorem provides at least a plausibility argument for a substantial connection between risk-avoidance and evidence-

²⁴This assumption is plausible even in light of challenges for strictly proper measures that evaluate entire credal states [Campbell-Moore and Levinstein, ms]. For, there can certainly be measures which are (a) strictly proper in their evaluation of credences in a particular proposition X and (b) weakly proper in their evaluations of entire credal states. (Just consider the 'global' measure which cares only about your accuracy with respect to X .) And measures which are strictly proper relative to X are all we need here.

	Weak Look-E			Look-E	
	risk neutral	≠	risk avoidant	risk neutral	risk avoidant
Subjectivist	yes	≠	no	no	no
Objectivist	yes	≠	likely no	no	no
Permissivist	Pareto	yes	≠	no	yes ≠ no
	Universal	yes	≠	likely no	no
Super-Permissivist	Pareto	yes		yes	yes ≠ no
	Universal	yes		yes	no

Table 1: Look Principles given only Weak Propriety.

avoidance, even if *strict* propriety is rejected as a constraint on accuracy measures.

7 Updating by Conditionalization

To calculate the expected accuracy associated with finding out whether E , we need to know what credence the agent will adopt in response to what she learns. Since we are concerned with ideal cases, this means that we need to know what credence the agent should adopt in response to what she learns. In section 5, we assumed that this is the agent’s prior credence conditional on what she learns – i.e. that rational agents update by conditionalization.

Is this assumption plausible when working with risk-sensitive agents? Several standard arguments for conditionalization (e.g. that it yields intuitively plausible verdicts, or, following van Fraassen’s Muddy Venn Diagram model [1989], that to update agents should simply renormalize their distribution after the possibilities inconsistent with what they learned have been eliminated) are independent of decision theory, and are thus as compelling as ever. Others however, such as Dutch Book arguments or arguments from expected accuracy, are potentially more problematic, since they usually assume that rational agents maximize expected (rather than risk-weighted expected) values.²⁵ We don’t want to oversell the significance of this, since (1) these arguments are independently controversial (assuming, for example, that pragmatic efficacy or consequentialist reasoning is probative for the epistemic rationality of *updating rules*) and (2) it may be possible to recover these arguments for conditionalization even in a

²⁵By equating credences with betting prices, Lewis’s [1999] Dutch Book argument assumes that agents evaluate bets by their expected value. The accuracy arguments of Greaves and Wallace [2006], Leitgeb and Pettigrew [2010b], Easwaran [2013], and Pettigrew [2016, ch.4] all assume that rationality requires us to update so as to maximize expected accuracy. The argument in Briggs and Pettigrew [ms] assumes only that rational agents don’t use accuracy dominated belief-revision procedures, so is not subject to this problem; however, the formal result assumes that inaccuracy measures are strictly proper, and doesn’t hold for the strictly r -proper measures we associate with risk-avoidant agents.

risk-sensitive setting.²⁶ But even with these reservations in mind, it's worth asking how much our argument relies on the conditionalization assumption.

The answer is: not much. For theorem 5.2 holds given just a few general assumptions about our agent's update procedure.²⁷ Since the only cases we need are the simple ones involving just propositions X and E , we can think of an update procedure as a pair of functions $(f_E, f_{\neg E})$, which take as inputs a probability distribution c over the algebra generated by X and E , and output a real number in $[0, 1]$ which is the credence in X the agent adopts upon receiving E and $\neg E$ respectively.²⁸ Using this notation, the necessary assumptions are:

1. f_E and $f_{\neg E}$ are continuous.
2. If $c(E) = 1$ then $f_E(c) = c(X)$; and if $c(\neg E) = 1$ then $f_{\neg E}(c) = c(X)$.
3. $f_E(c) = c(X)$ only if $c(X) = c(X|E)$ (i.e. c makes X and E independent); and similarly for $f_{\neg E}$.
4. There is some t such that if $c(X)$, $c(X|E)$, and $c(X|\neg E)$ are all $> t$ then $f_E(c)$ and $f_{\neg E}(c)$ are both $> 1/2$.

Conditions 1. and 2. are extremely plausible: they say, respectively, that arbitrarily small changes in the input distribution shouldn't make for large changes in the output, and that learning something of which one is already certain shouldn't change one's views. Conditions 3. and 4. are also quite plausible, saying, respectively, that conditionally relevant evidence is relevant and that when X is highly probable regardless of what's the case with E , learning E shouldn't make one think X improbable. While plausible, however, the latter two both assume that conditional probabilities place substantial constraints on the update function; so someone who rejects conditionalization as fundamentally mistaken might reject one or both of them. So it's worth highlighting what role they play in the proof. Some condition like 3. is needed to guarantee that the update function considers a reasonably wide range of evidence to be 'potentially relevant', and hence doesn't merely validate Look-E in a vacuous way. And condition 4. is needed only to ensure that there is some range of cases in which we know how the various possible outcomes of learning the evidence are to be ordered. It is thus far from clear that denying 3. would allow one to vindicate Look-E in an interesting way – or that denying 4. would allow one to vindicate it at all.

Conditions 1.-4. are thus fair assumptions about the update procedure. This is not to say that they couldn't possibly be rejected.²⁹ But it does mean that

²⁶We defend this claim with respect to accuracy arguments in Campbell-Moore and Salow [ms].

²⁷For the proof, see appendix A.

²⁸We can allow that f_E is undefined when $c(E) = 0$ and $f_{\neg E}$ is undefined when $c(\neg E) = 0$.

²⁹Are they met by the policies motivated by risk-sensitive versions of expected accuracy arguments? In Campbell-Moore and Salow [ms] we argue that the best ways of adapting these arguments still supports conditionalization, so that conditions 1.-4. are met. But on other ways of adapting the arguments, they support a revision procedure which is custom-made to vindicate Look-E, and which violates condition 3.

the burden is not on us to show that risk-avoidant agents should still update by conditionalization, but rather on our opponent to offer some reasons for thinking that risk-avoidant agents should update in a radically unfamiliar – and intuitively odd-looking – way.

8 Conclusion

Can it be rational to avoid evidence, even when gathering it would cost you nothing, you expect it to be relevant, and you know that you would process it rationally? It's well-known that if risk-avoidance is rational then it can be *instrumentally* rational to do so. But, as Buchak [2010] observes, there's more to life than instrumental rationality: there's epistemic rationality as well. What we have shown is that this does not threaten the connection; for if risk-avoidance (of the kind defended by Buchak 2010, 2013) is rational, then evidence avoidance can also be epistemically rational.

Our main argument, presented in section 5, relies on five main assumptions. First, that people are epistemically required to pursue epistemic value. Second, that for rationally risk-avoidant agents, this amounts to maximizing risk-weighted expected epistemic value. Third, that the epistemic value of an action is the accuracy of the anticipated resultant credences. Fourth, that inaccuracy is measured according to a strictly r -proper accuracy measure. And, fifth, that risk-averse agents who know they will be rational expect to conditionalize on the evidence they encounter.

The first three assumptions are natural starting points; in section 3, we motivated them and distinguished them from stronger claims that might be problematic. We defended a strong version of the fourth assumption in section 4 by appeal to the thought that rational agents regard their credences as giving them the best shot at the truth; and we showed, in section 6, that much of our argument can be run from a much weaker version of the fourth assumption. Finally, in section 7, we showed that the fifth assumption, too, was unnecessarily strong: the argument goes through for any update procedure satisfying a few fairly weak constraints.

We thus conclude that there is a serious tension between risk-avoidance and the claim that, in ideal cases, we should always gather the evidence. But we leave open what to ultimately make of that conclusion: whether it is a reason to reject the rationality of risk-avoidance, or a reason to embrace the epistemic rationality of avoiding relevant and freely available information.

References

Maurice Allais. Fondements d'une théorie positive des choix comportant un risque et critique des postulats et axiomes de l'Ecole Americaine. *Econométrie*, pages 257–332, 1953.

- Selim Berker. Epistemic teleology and the separateness of propositions. *Philosophical Review*, 122:337–93, 2013.
- Rachael Briggs and Richard Pettigrew. An accuracy-dominance argument for conditionalization. May 27, 2015., ms.
- Lara Buchak. Instrumental rationality, epistemic rationality, and evidence-gathering. *Philosophical Perspectives*, 24(1):85–120, 2010.
- Lara Buchak. *Risk and Rationality*. Oxford University Press Oxford, 2013.
- Lara Buchak. Risk and tradeoffs. *Erkenntnis*, 79(6):1091–1117, 2014.
- Michael Caie. Rational probabilistic incoherence. *Philosophical Review*, 122:527–575, 2013.
- Catrin Campbell-Moore and Ben Levinstein. Accuracy and risk-sensitivity. ms.
- Catrin Campbell-Moore and Bernhard Salow. Accurate updating for the risk aware. ms.
- Jennifer Carr. Epistemic utility theory and the aim of belief. *Philosophy and Phenomenological Research*, 95:511–534, 2017.
- Kenny Easwaran. Expected accuracy supports conditionalization and conglomerability and reflection. *Philosophy of Science*, 80(1):119–142, 2013.
- Bas C. Van Fraassen. *Laws and Symmetry*. Oxford University Press, 1989.
- Allen Gibbard. Rational credence and the value of truth. Oxford University Press, 2008.
- Irving J Good. On the principle of total evidence. *The British Journal for the Philosophy of Science*, 17(4):319–321, 1967.
- Hilary Greaves. Epistemic decision theory. *Mind*, 122(488):915–952, 2013.
- Hilary Greaves and David Wallace. Justifying conditionalization: Conditionalization maximizes expected epistemic utility. *Mind*, 115(459):607–632, 2006.
- Sophie Horowitz. Accuracy and educated guesses. Oxford University Press, forthcoming.
- James Joyce. Accuracy and coherence: Prospects for an alethic epistemology of partial beliefs. Springer, 2009.
- Jason Konek and Benjamin A Levinstein. The foundations of epistemic decision theory. *Mind*, page fzw044, 2017.
- Amit Kothiyal, Vitalie Spinu, and Peter P Wakker. Comonotonic proper scoring rules to measure ambiguity and subjective beliefs. *Journal of Multi-Criteria Decision Analysis*, 17(3-4):101–113, 2010.

- Hannes Leitgeb and Richard Pettigrew. An objective justification of bayesianism I: Measuring inaccuracy. *Philosophy of Science*, 77(2):201–235, 2010a.
- Hannes Leitgeb and Richard Pettigrew. An objective justification of bayesianism II: The consequences of minimizing inaccuracy. *Philosophy of Science*, 77(2): 236–272, 2010b.
- David Lewis. Immodest inductive methods. *Philosophy of Science*, 38:54–63, 1971.
- David Lewis. Why conditionalize? Cambridge University Press, 1999.
- Patrick Maher. Joyce’s argument for probabilism. *Philosophy of Science*, 69: 73–81, 2002.
- Edward F McClennen. *Rationality and Dynamic Choice*. Cambridge University Press, 1990.
- Graham Oddie. Conditionalization, cogency, and cognitive value. *British Journal for the Philosophy of Science*, 48:533–541, 1997.
- Richard Pettigrew. *Accuracy and the Laws of Credence*. Oxford University Press Oxford, 2016.
- John Quiggin. A theory of anticipated utility. *Journal of Economic Behavior & Organization*, 3(4):323–343, 1982.
- Amos Tversky and Peter Wakker. Risk attitudes and decision weights. *Econometrica: Journal of the Econometric Society*, pages 1255–1280, 1995.
- Peter Wakker. Nonexpected utility as aversion of information. *Journal of Behavioral Decision Making*, 1:169–175, 1988.
- Peter P Wakker. *Prospect theory: For risk and ambiguity*. Cambridge University Press, 2010.

A Proof of main result

This appendix proves a general result encompassing both theorem 5.2 and theorem 6.2, but making much weaker assumptions about the update procedure. We represent the update procedure by a pair of functions (f_E, f_{-E}) that each take a prior credal state to a new credence in X . The theorem is then

Theorem A.1. *Suppose:*

- $r : [0, 1] \rightarrow [0, 1]$ is:
 - a risk profile, i.e. increasing, $r(0) = 0$ and $r(1) = 1$.
 - differentiable

– such that $r(x) < x$

- $f_E, f_{\neg E}$ are functions from probability functions c over the algebra generated by X and E to $[0, 1]$ (though f_E may be undefined when $c(E) = 0$, and $f_{\neg E}$ undefined when $c(\neg E) = 0$) such that:
 - f_E and $f_{\neg E}$ are continuous.
 - If $c(E) = 1$ then $f_E(c) = c(X)$; and if $c(\neg E) = 1$ then $f_{\neg E}(c) = c(X)$.
 - $f_E(c) = c(X)$ only if $c(X) = c(X|E) = c(X|\neg E)$; and similarly for $f_{\neg E}$.
 - There is some t such that if both $c(X|E)$ and $c(X|\neg E)$ are $> t$ then $f_E(c)$ and $f_{\neg E}(c)$ are both $> 1/2$.

Then there is some probability distribution c such that for each \mathcal{A} where:

- $\mathcal{A} : [0, 1] \times \{0, 1\} \rightarrow \mathbb{R}$ is:
 - weakly truth directed, i.e.
 - * If $x_1 < x_2 < 1$, then $\mathcal{A}(x_1, 1) \leq \mathcal{A}(x_2, 1)$,
 - * If $x_1 > x_2 > 0$, then $\mathcal{A}(x_1, 0) \leq \mathcal{A}(x_2, 0)$.
 - If $x > 1/2$, then $\mathcal{A}(x, 0) \leq \mathcal{A}(x, 1)$; which follows from 0/1-symmetry,
 - r -proper, i.e. For all $y \neq c(X)$, $\text{RExp}_c^r \mathcal{A}(y) </\leq \text{RExp}_c^r \mathcal{A}(c(X))$.
(With $<$ or \leq depending on whether it is strict or weak propriety.)

we have $\text{RExp}_c^r \mathcal{A}(\text{GATHER}) </\leq \text{RExp}_c^r \mathcal{A}(\text{AVOID})$.

(With $<$ or \leq depending on whether \mathcal{A} is strictly or weakly proper.)

Proof. Choose some $d_E > t$; since $d_E > r(d_E)$, we can also choose $d_{\neg E} > t$ with $d_E > d_{\neg E} > r(d_E)$. To construct c , we set $c(X|E) = d_E$ and $c(X|\neg E) = d_{\neg E}$; fixing $c(E)$ as any $e \in (0, 1)$ then determines a unique probability distribution. We consider different choices for e .

By the Intermediate Value Theorem, there are three cases:

- (a) For some choice of e , $f_{\neg E}(c) = f_E(c)$
- (b) For every choice of e , $f_{\neg E}(c) < f_E(c)$
- (c) For every choice of e , $f_{\neg E}(c) > f_E(c)$

Suppose (a) obtains. Then, for this choice of e , $\text{RExp}_c^r \mathcal{A}(\text{GATHER}) = \text{RExp}_c^r \mathcal{A}(f_E(c))$. Since $c(X|E) \neq c(X|\neg E)$, our conditions on f_E entail $f_E(c) \neq c(X)$. So $\text{RExp}_c^r \mathcal{A}(\text{GATHER}) </\leq \text{RExp}_c^r \mathcal{A}(\text{AVOID})$ is just an instance of propriety.

Suppose (b) obtains. Since $d_E, d_{\neg E} > t$, we then have $1/2 < f_{\neg E}(c) < f_E(c) \leq 1$, so the states are ordered as in the example from the main text (theorem 5.2). Slightly generalizing the argument there shows that $\text{RExp}_c^r \mathcal{A}(\text{GATHER}) </\leq \text{RExp}_c^r \mathcal{A}(\text{AVOID})$ follows from

$$\frac{r[c(E \wedge X) + c(\neg E \wedge X)] - r[c(E \wedge X)]}{r[c(E \wedge X) + c(\neg E)] - r[c(E \wedge X)]} > r(f_E(c)). \quad (1)$$

We can write the left hand side of eq. (1) in terms of d_E , $d_{\neg E}$, and e ; and then use L'Hôpital's rule and the chain rule to determine its limit as $e \rightarrow 1$

$$\begin{aligned}
& \lim_{e \rightarrow 1} \frac{r(d_E e + d_{\neg E}(1 - e)) - r(d_E e)}{r(d_E e + (1 - e)) - r(d_E e)} \\
&= \lim_{e \rightarrow 1} \frac{r'(d_E e + d_{\neg E}(1 - e)) \cdot (d_E - d_{\neg E}) - r'(d_E e) \cdot d_E}{r'(d_E e + (1 - e)) \cdot (d_E - 1) - r'(d_E e) \cdot d_E} \\
&= \frac{r'(d_E) \cdot (d_E - d_{\neg E}) - r'(d_E) \cdot d_E}{r'(d_E) \cdot (d_E - 1) - r'(d_E) \cdot d_E} \\
&= d_{\neg E}
\end{aligned}$$

By our conditions on f_E , $f_E(c^*) = c^*(X) = c^*(X|E)$ if $c^*(E) = 1$. So, since f_E is continuous, $f_E(c) \rightarrow d_E$ as $e \rightarrow 1$; and thus $r(f_E(c)) \rightarrow r(d_E)$. But we chose $d_{\neg E} > r(d_E)$. So by choosing e sufficiently close to 1 we ensure that eq. (1) holds, and hence $\text{RExp}_c^r \mathcal{A}(\text{GATHER}) </\leq \text{RExp}_c^r \mathcal{A}(\text{AVOID})$.

Suppose, finally, that (c) obtains. We then have $1/2 < f_E(c) < f_{\neg E}(c) < 1$. Analogous reasoning shows that it is sufficient that

$$\frac{r[c(\neg E \wedge X) + c(E \wedge X)] - r[c(\neg E \wedge X)]}{r[c(\neg E \wedge X) + c(E)] - r[c(\neg E \wedge X)]} > r(f_{\neg E}(c)). \quad (2)$$

And an analogous argument shows that the left hand side of this converges to d_E as $e \rightarrow 0$. Since $f_{\neg E}(c) \rightarrow d_{\neg E}$ as $e \rightarrow 0$, and $d_E > d_{\neg E} > r(d_{\neg E})$, it follows that choosing e sufficiently close to 0 yields the desired result. \square

Suppose that the update rule is conditionalization. Then we get a sufficient condition:

Corollary A.2. *For all $d_E \in [0, 1] \setminus \{1/2\}$,
for all $d_{\neg E}$ between $1/2$ and d_E and close enough $_{d_E}$ to d_E ,
for all x between $d_{\neg E}$ and d_E and close enough $_{d_E, d_{\neg E}}$ to d_E :
If c has $c(X|E) = d_E$, $c(X|\neg E) = d_{\neg E}$, $c(X) = x$,
then provided $\mathcal{A}(x, 0) \geq \mathcal{A}(x, 1)$ for $x < 1/2$, which also follows from 0/1 symmetry, we have:*

$$\text{RExp}_c^r \mathcal{A}(\text{GATHER}) </\leq \text{RExp}_c^r \mathcal{A}(\text{AVOID})$$

Where what counts as close enough depends on the indicated variables.

Proof. For $d_E > d_{\neg E} \geq 1/2$, conditionalization ensures that we are in case (b); so we just check the choices made in the above proof, noting that choosing e close to 1 amounts to choosing x close to d_E . For $d_E < d_{\neg E} \leq 1/2$ the accuracy orderings are switched. The above argument can be modified to deal with this case: we find the sufficient condition

$$\frac{r[c(E \wedge \neg X) + c(\neg E \wedge \neg X)] - r[c(E \wedge \neg X)]}{r[c(E \wedge \neg X) + c(\neg E)] - r[c(E \wedge \neg X)]} > r(1 - f_E(c)). \quad (3)$$

and show this is met when $1 - d_{\neg E} > r(1 - d_E)$, and e is sufficiently close to 1. This establishes the result. \square

A.2 tells us that any interval will contain values for $c(X)$, $c(X|E)$ and $c(X|\neg E)$ leading to $\text{RExp}_c^r(\text{GATHER}) < \text{RExp}_c^r(\text{AVOID})$. It follows that we can weaken the assumptions on \mathcal{A} and r , so that they only hold on an interval, and still get A.1 for conditionalization. Similar reasoning shows that such weakened assumptions are sufficient for A.1 even if we don't assume conditionalization, provided the interval contains some values $> t$.