The Value of Evidence in Decision-Making

Ru Ye, Wuhan University

Forthcoming in *The Journal of Philosophy*

Abstract

The Value of Evidence thesis (VE) tells us to gather evidence before deciding in any decision problem, if the evidence is free. This apparently plausible principle faces two problems. First, it fails on evidence externalism or nonclassical decision theories. Second, it's not general enough: it tells us to prefer gaining free evidence to gaining no evidence, but it doesn't tell us to prefer gaining more informative evidence to gaining less informative evidence when both are free.

This paper defends an alternative value-of-evidence principle that solves the above two problems. We call it 'the Comparative Value of Evidence for the Resolute (CVER).' CVER not only generalizes VE to the comparative case but is also compatible with evidence externalism and nonclassical decision theories. Our argument, if successful, broadens the scope of the value-of-evidence principle and shows that it's more robust than previously imagined.

Keywords: Value of Evidence; Dynamic Decision; Evidence Externalism, Nonclassical Decision Theory

1 Introduction

Should you always prefer gathering more evidence before making a choice in a decision problem? Presumably, no. Gathering and processing evidence can be costly, and you may doubt whether you will respond to the gathered evidence rationally (as too much evidence may cloud your judgment).[1] But if the evidence is free—if the cost of gathering and processing evidence is negligible—and if you are certain to respond to the evidence rationally, should you prefer gathering more evidence before making a decision? Presumably, yes. It seems that more evidence generally leads to better decisions. As the saying goes, it's often wise to 'look before you leap.' So, the following principle, which is famously proved by Good (1967), seems intuitive:

> *Value of Evidence (VE)*
>
> For any decision problem, if one can gather some evidence at a negligible cost before making a decision, and if one is certain that one will respond to the evidence rationally, one should prefer to gather the evidence first.[2]

As plausible as VE sounds, a number of writers have recently called it into question. They contend that Good's proof of VE rests on a bunch of nontrivial assumptions, including the assumption that one's evidence is internalist and that rational preferences obey the classical, expected utility theory. If we relax these assumptions—if evidence is externalist or rational preferences obey nonclassical

[1]Throughout this paper, 'respond to evidence rationally' means 'update beliefs epistemically rationally and act on those beliefs in an instrumentally rational way.'

[2]In this paper, 'prefer' refers to 'weakly prefer' unless otherwise noted. Besides, here and throughout this paper, the costs of an experiment should be understood broadly to include costs in money, time, and opportunity, but not so broadly to trivialize VE. More precisely, an experiment is cost-negligible if whether to perform the experiment and what the experimental outcome is don't affect (either positively or negatively) the utility of any outcome in any decision to be faced. (As we will clarify in §3.2, saying that the experimental outcomes don't affect utilities implies that one's utility can be defined as a function over $A \times \Omega$; there is no need to define it as a function over the enlarged space $A \times \Omega \times E$.) Thanks to an anonymous reviewer for urging us to clarify the notion of costs.

decision theories—Good's proof no longer works, and VE fails in some instances.[3] Some writers have interpreted the failure of VE on externalism of evidence or nonclassical decision theories as a serious cost of these theories (Briggs, 2015; Wakker, 1988). For some other writers, it's a cause to abandon the idea that free evidence is always valuable (Das, 2020).

In addition to the tension with nonclassical decision theories and evidence externalism, VE faces another problem, which has gone unnoticed so far: it's not general enough in the sense that it merely tells us to prefer 'gaining evidence' to 'not gaining evidence' when the evidence is free, but it doesn't tell us to prefer 'gaining more informative evidence' to 'gaining less informative evidence' when both are free.

In this paper, we propose an alternative value-of-evidence principle that solves the above two problems of VE. We call this alternative principle 'the Comparative Value of Evidence for the Resolute (CVER).' It says that more evidence is always preferable to less evidence in decision-making if one is certain that, among those action plans made available by gathering evidence, one will follow the plan that's ex ante optimal. As we will argue, this principle is not only more general than VE, but it's also compatible with evidence externalism and nonclassical decision theories. Our argument, if successful, broadens the scope of the value-of-evidence principle and shows that it's more robust than previously imagined. It also helps to answer a question important for zetetic epistemology, the question about what kind of evidence we should gather in our inquiries.[4]

[3]For VE failure on decision theories for imprecise credence, see Bradley & Steele 2016, Das 2023, and Kadane et al. 2008. For VE failure on risk-weighted expected utility theory, see Buchak 2010 and Campbell-Moore & Salow 2020. For VE failure on evidence externalism, see Ahmed & Salow 2019, Das 2020, and Dorst 2020. In addition, Neth (forthcoming) argues that VE can fail for modest agents who are uncertain that they will update beliefs rationally. Since we primarily focus on immodest agents, we will ignore this complication. We will also ignore situations involving world-act dependence (see Skyrm 1990 for VE failure in such cases.)

[4]A few attempts to save the value-of-evidence principle have been made. Ahmed and Salow (2019) replace VE with a principle called 'Conditionality.' Das (2023) proposes a specific decision rule for agents with imprecise credence and restricts VE to agents following that rule. Bradley and

Here is the plan. §2 reviews the recent arguments that VE fails on externalism of evidence or nonclassical decision theories; this section also explains why VE should be generalized to the comparative case. In §3, we argue that gaining evidence has a necessary effect: it's equivalent with increasing flexibility in action planning; that is, the more evidence one has, the more action plans are available, and vice versa. §4 argues that this connection between gaining evidence and increasing flexibility implies our value-of-evidence principle CVER, which generalizes VE to the comparative case. §5 argues that CVER is acceptable to externalists of evidence and nonclassical decision theorists. §6 argues that it's more powerful than the existing replacement of VE. §7 argues that our defense of CVER has an additional benefit: it generates a new argument for the accuracy value of evidence. We conclude the paper by briefly discussing CVER's implications for zetetic epistemology, specifically for the question of what kind of evidence we should gather in our inquiries.

2 Two Problems With VE

In this section, we explain VE's two problems. The first problem is that, as many scholars have pointed out, VE fails on externalism of evidence and nonclassical decision theories. By externalism of evidence, we mean those theories of evidence that allow evidence to be nontransparent: sometimes, one's evidence includes *p* but doesn't include 'my evidence includes *p*,' or one's evidence doesn't include *p* but doesn't include 'my evidence doesn't include *p*.' A well-known example is Williamson's (2000) theory that evidence is knowledge and that knowledge is not transparent. By nonclassical decision theories, we mean those decision theories that reject the independence axiom in the classical expected utility theory (which

Steele (2016) defend the 'Non-Dominated Set' decision rule for imprecise credence, arguing that it eases the tension with VE. Our approach to saving the value-of-evidence principle is most similar to Ahmed and Salow's. However, as we will argue in §6, our principle CVER is more powerful than Conditionality.

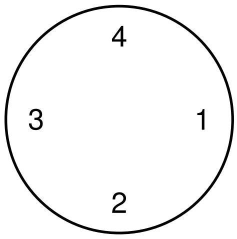

Figure 1: An Unmarked Clock.

says that for any actions a, b, c and any event E, the mixed action $a_E c$ is preferred to $b_E c$ if and only if a is preferred to b). A well-known example is Buchak's (2013) risk-weighted expected utility theory.

To see why VE can fail when evidence is nontransparent, consider the following example, which is based on Gallow's (2021) presentation of Williamson's (2014) unmarked clock example. You are facing an unmarked clock with a single digit pointing in one of four uniformly spaced directions labelled 1 to 4 (see Figure 1). Suppose that, for any digit position in {1, 2, 3, 4}, taking a look at the clock from a certain distance gives you limited evidence: it allows you to reliably judge that the digit is pointing at it or its two adjacent positions, but nothing more specific than that. So, if the digit is in fact pointing at 1, then by looking, you will get evidence revealing that it's either at 1, 2, or 4, but nothing more specific than that. So, your evidence doesn't tell you that this is your evidence: your evidence doesn't rule out the possibility that the digit is pointing at 2; in this alternative scenario, your evidence tells you that the digit is pointing at 1, 2, or 3, which differs from the evidence you get when the digit is pointing at 1. So, your evidence is nontransparent.

Now, consider this decision problem (Ahmed & Salow, 2019). Let E = {2, 4} and O = {1, 3} be the proposition that the digit is pointing at an even (odd) number respectively. Let EVEN be the bet that pays 0.4\$ if E but −0.6\$ otherwise, and ODD be the bet that pays 0.4\$ if O but −0.6\$ otherwise. Also, assume that your prior credence before looking is evenly distributed over the four positions. You are faced with a decision: accept EVEN, accept ODD, or neither. Suppose that taking a look at the clock before making a decision is free, you update beliefs by conditionalization, and your preferences conform with the classical, expected utility theory. Then you should strictly prefer not to take a look. You know that, for any i in $\{1, 2, 3, 4\}$, if i is even, looking will result in a 2/3 confidence in O and a 1/3 confidence in E, leading you to accept ODD and get −0.6$; if i is odd, looking will result in a 2/3 confidence in E and a 1/3 confidence in O, leading you to accept EVEN and get −0.6$. So, if you look and then decide, you will be guaranteed to lose 0.6$; if you don't take a look, you won't accept any bet and thus will not lose or win. So, VE fails in this case.

For an example of VE failure on nonclassical decision theories, consider the verdict of Buchak's (2013) risk-weighted expected utility theory in a case where the decision maker is risk-averse.[5] Suppose that you face a choice between the two lotteries described in Table 1. For the sake of concreteness, let's suppose that your

Table 1: Two Lotteries.

	$w_1(1/6)$	$w_2(1/3)$	$w_3(1/2)$
L_1	2	1	0
L_2	−1	4	0

risk function is $r(p) = p^2$. Let $\mathrm{E} = \{w_1, w_2\}$.

Suppose that you can perform a cost-free experiment to determine whether E is true before making a decision. Should you prefer to perform the experiment first? The risk-weighted expected utility theory says 'no.' If you don't perform the experiment, you will choose L_1 because L_1 has a greater risk-weighted expected utility

[5] If one's risk function is r, the risk-weighted expected utility of an act a is defined as:

$$u_1 + \sum_{i=2}^{n} r\left(\sum_{j=i}^{n} p_j\right) \cdot (u_i - u_{i-1})$$

where $u_1 < u_2 < \cdots < u_n$ are the utilities of the possible outcomes listed in ascending order, and $p_1, p_2, \ldots, p_n$ are the probabilities of the corresponding outcomes.

according to your present belief ($REU(L_1) = 1(1/2)^2 + (2-1)(1/6)^2 > REU(L_2) = (-1) + 1(5/6)^2 + 4(1/3)^2$, where 'REU' stands for 'risk-weighted expected utility'), and thus your payoff profile will be $\{w_1\colon 2; w_2\colon 1; w_3\colon 0\}$, which is the same as L_1's. But if you perform the experiment, two things can happen. If the experimental outcome says E, you will prefer L_2 over L_1, as the former has a greater risk-weighted expected utility relative to your updated belief (L_1's REU becomes $1 + (2-1)(1/3)^2$, whereas L_2's REU becomes $(-1) + 5(2/3)^2$). Assuming that you will choose according to your post-experiment preferences, you will choose L_2. If the experimental outcome says not-E, you will be indifferent between L_1 and L_2, and, no matter what you choose, you will get 0. So, if you perform the experiment, your payoff profile will be $\{w_1\colon -1; w_2\colon 4, w_3\colon 0\}$, which is the same with L_2's profile. But by the light of your present belief, this payoff profile has a lower risk-weighted expected utility than the one generated by not performing the experiment. Therefore, you should strictly prefer not to perform the experiment. Again, VE fails.

Similar arguments have been made concerning other nonclassical decision theories (see fn. 3 above). Wakker (1988, pp. 172-173) identifies a key reason why these theories often fail VE: if one violates the independence axiom by strictly preferring $a_E c$ to $b_E c$ but strictly preferring b to a for some actions a, b, c and some event E, then, investigating whether E before choosing between $a_E c$ and $b_E c$ has the same payoff profile with $b_E c$ (assuming that you will still strictly prefer b to a if you learn E), while not conducting an investigation has the same payoff profile with $a_E c$; since one strictly prefers $a_E c$ to $b_E c$, one strictly prefers not to conduct an investigation.[6]

So, we have seen that VE is in tension with externalism of evidence and non-

[6] Wakker's argument becomes more complicated if one of the strict preferences is replaced by indifference. If $a_E c$ is strictly preferred to $b_E c$ while a is indifferent to b, then VE still fails, assuming the continuity and transitivity of preference: there is some positive ϵ such that $(a-\epsilon)_E c$ is strictly preferred to $b_E c$ but b is strictly preferred to $a-\epsilon$. But if $a_E c$ is indifferent to $b_E c$ while b is strictly preferred to a, then without further assumptions (such as a strong dominance principle that requires one to strictly prefer $b_E c$ to $(b-\epsilon)_E c$ for any positive ϵ), Wakker's argument only gives us an indifference on whether to conduct an investigation, which is consistent with VE.

classical decision theories. The second problem with VE is its lack of generality: it only tells us to prefer a free experiment to deciding right away, but what if we can choose between two free experiments, one of which is more informative than the other? VE is silent on which one to prefer, whereas intuitively, we should prefer the more informative experiment. To put it another way, we should adopt the following principle:

> *Comparative Value of Evidence (CVE)*
>
> For any decision problem one faces, if one can choose at a negligible cost between performing a more informative experiment before deciding and performing a less informative experiment before deciding, and if one is certain that one will respond to the experimental outcome rationally, one should prefer the more informative experiment.

CVE is more general than VE: CVE implies VE since we can understand the choice 'deciding now, without gathering evidence' as performing a trivial, least informative experiment (e.g., observing the realization of a random variable that's a constant function). But VE doesn't imply CVE: the claim 'one should prefer a nontrivial experiment to a trivial one' doesn't imply 'one should prefer a more informative experiment to a less informative one.'

Although CVE is more general than VE, it's no less plausible. A more informative experiment generates more evidence about the world, and it seems that one should favor more evidence over less evidence before making a decision, if both experiments are costless and if one is certain to respond rationally to the experimental outcome. For example, if one can choose between checking a more reliable weather forecast before deciding whether to bring an umbrella and checking a less reliable forecast before deciding (assuming that both experiments are costless), it seems that one should prefer the former because it's intuitively more informative.

Of course, since CVE implies VE, it doesn't avoid VE's tension with evidence externalism and nonclassical decision theories. So, in what follows, we focus on how to revise CVE to avoid the tension. To motivate this revised principle, we first consider a necessary effect of gaining evidence.

3 More Evidence Implies Greater Flexibility

3.1 A Motivating Example

Gaining evidence before making a decision can have various effects: your beliefs about the world can change; you may feel overwhelmed by the additional information; you may become more risk-averse or less risk-averse; and so on. However, all these effects are contingent: even if they occur to you, they may not occur to other persons who face the same decision. Is there any effect of gaining evidence that's necessary, i.e., an effect that will occur to anyone who gains evidence before making a decision? The answer is positive: for any person and any decision problem the person faces, gaining evidence will make the person more flexible in decision-making, in the sense of giving the person more options to choose from.

Consider this everyday decision problem. You are deciding whether to bring an umbrella to work tomorrow. Suppose that you are given the opportunity to check the weather forecast for free before deciding. If you don't check the weather forecast, your available action plans are 'bring an umbrella' and 'don't bring an umbrella.'[7] But if you check the weather forecast first and then make a decision according to the result, you will have two more action plans available—you can now plan to tailor your actions to the weather forecast. So, you can make the plan 'if rain is forecasted, bring an umbrella; if not-rain is forecasted, don't bring an

[7]We only consider deterministic actions for now. We'll consider non-deterministic actions (like carrying an umbrella with an 80% chance) later in a formal proof that more evidence means greater flexibility.

umbrella.' You can also make the plan 'if rain is forecasted, don't bring an umbrella; if not-rain is forecasted, bring an umbrella.' And the original two action plans are still available—'(don't) bring an umbrella' can be understood as '(don't) bring an umbrella no matter whether rain or not-rain is forecasted.' So, in this case, it's intuitive that gaining evidence implies having more available action plans.

However, the intuition is not so clear in other cases. Consider the following variant of the umbrella case. Like before, you are deliberating whether to bring an umbrella to work tomorrow. But this time, you can choose between checking two different forecasts before deciding, one of which is more reliable than the other. Intuitively, checking the more reliable forecast will give you more evidence about the weather than checking the less reliable one; that is, it's a more informative experiment. But it's not immediately clear that checking the more reliable forecast leads to more available action plans compared to checking the less reliable one. It seems that, no matter which forecast you check, your available action plans are the following four: 'If rain is forecasted, bring an umbrella; otherwise don't'; 'if not-rain is forecasted, bring an umbrella; otherwise don't'; 'bring an umbrella no matter whether rain or not-rain is forecasted'; and 'don't bring an umbrella no matter whether rain or not-rain is forecasted.'

However, as we will see in §3.2, checking the more reliable weather forecast does give you more available action plans. This holds true generally: performing a more informative experiment give you more available action plans compared to a less informative one.

3.2 More Evidence Implies Greater Flexibility: The Proof

In this section, we explain why more evidence implies more available action plans. The key is to distinguish between two kinds of action plans. In the umbrella example, action plans take the form 'if rain is forecasted, bring an umbrella; if not-rain is forecasted, don't bring an umbrella.' This kind of plan makes actions depend on experimental outcomes. We call them 'evidence-contingent action plans,' since they map each evidence (i.e., experimental outcome) to some action.

But there is another kind of action plan, which we call 'state-contingent action plans.' Unlike evidence-contingent action plans, state-contingent action plans make actions directly depend on the underlying world states rather than the experimental outcomes. So, in the umbrella example, 'if rain is forecasted, bring an umbrella; if not-rain is forecasted, don't' is an evidence-contingent action plan, whereas 'if it rains, bring an umbrella; if it doesn't rain, don't' is a state-contingent action plan. It can be understood as a representational 'as if' plan. A person follows this plan if, from a third-person perspective, she acts as if she always brings an umbrella on rainy days and never brings an umbrella on dry days.

A state-contingent action plan is said to be 'available under an experiment' if it can be followed indirectly as a result of following some evidence-contingent action plan. This informal characterization, which will be made more precise later, is based on the idea that every evidence-contingent action plan, if followed, will realize following some (possibly non-deterministic) state-contingent action plan. For example, if a weather forecast is 100% reliable, then following the evidence-contingent plan 'if rain is forecasted, bring an umbrella; otherwise, don't' will realize following the perfect state-contingent plan 'if it rains, bring an umbrella; otherwise don't.' However, if the weather forecast is only 70% reliable, then this perfect state-contingent action plan becomes unavailable: there exists no evidence-contingent action plan of the form 'if rain (not-rain) is forecasted, bringing an umbrella with probability x (y)' such that, if I follow the plan, I will generally bring an umbrella if and only if it rains.

This highlights an important claim: which state-contingent action plans are available under an experiment depends on how informative the experiment is;

the more informative the experiment, the more state-contingent action plans are available. Even if two experiments generate the same set of evidence-contingent plans, one of them might make more state-contingent plans available. So, we may draw the following conclusion about the relationship between gaining evidence and increasing flexibility:

> *Gaining Evidence Implies Increasing Flexibility*
>
> For any two experiments A and B and any decision problem, if A is more informative than B, then the set of state-contingent action plans made available by performing A contains the set of state-contingent action plans made available by performing B.

It's worth noting that the converse of the above claim is also true. That is, if an experiment gives you greater flexibility in action-planning for any decision problem, then it's more informative. So, gaining evidence and increasing flexibility in action planning are equivalent. We think the equivalence itself is interesting and teaches us something important about the nature of evidence, even setting aside implications for the value of evidence. In what follows, we present a simple proof of this equivalence.[8] Along the way, we also provide some intuitive motivations for the formal criterion of informativeness we use.

Before we start, however, we need to address a worry about our definition of available state-contingent plans. In defining an available state-contingent plan as one realizable by *some* evidence-contingent plan, we have implicitly assumed that *all* evidence-contingent plans can be followed. Does this assumption commit us to the transparency of evidence?

The answer is negative. In this paper, evidence as the input of evidence-contingent plans should be understood as 'interpreted evidence' or 'what you take

[8]Blackwell (1951) first proved a similar principle. The proof in this paper follows de Oliveira's (2018) approach and is much simpler.

your evidence to be'; more exactly, it's whatever level of evidence-interpretation that actually guides your action. To illustrate, suppose that you interpret your evidence at multiple levels: your actual evidence is e, you take it to be f, you take it that you take your evidence to be g, etc., and all these levels might be nontransparent. However, suppose that it's always the first-order interpretation (that is, what you take your evidence to be) that guides your action.[9] Then an evidence-contingent plan will be of the form $\{T_e$: act a; T_f: act $b\}$, where 'T_e' means 'taking evidence to be e.' Such a plan is always followable, in the sense that if you adopt this plan, you won't fail to follow it—more exactly, you won't fail for the reason that you might misinterpret its input. For even if you misinterpret the input by taking your taking to be f when your taking is actually e (that is, even if TT_f but T_e), you won't end up following a different plan $\{T_e$: act b; T_f: act $b\}$: by supposition, it's your taking of evidence rather than your taking of taking that guides your action. (Of course, if you are of a different kind of person who is actually guided by the second-order taking rather than the first-order one, then the input of evidence-contingent plans will be revised into second-order taking.)

Now, we can start to prove the equivalence between gaining evidence and increasing flexibility. Let Ω be a state space. Let an experiment $\mathcal{E}$ be a pair (E, σ), where E is a set of possible experimental outcomes and $\sigma : \Omega \to \Delta(E)$ is the probability distribution the outcomes in E obey.[10] So, for any $w \in \Omega$ and any $e \in E$, $\sigma(e|w)$ is the probability that the experiment produces outcome e in state w. Let a decision problem $\mathcal{D}$ be a pair (A, u), where A is the agent's available actions and $u : A \times \Omega \to R$ is the agent's utility function that maps each pair (a, w) to a real number.[11] For simplicity, we assume that E, Ω, and A are all finite. But we don't

[9]So, here, 'taking your evidence to be e' is similar to Gallow's (2021) 'updating on evidence e' and Schultheis's (forthcoming) 'guessing the evidence to be e': all are supposed to pick out the level of interpretation that actually guides your action.

[10]So, σ is a stochastic function. A stochastic function $g : X \to \Delta(Y)$ maps each element in X to a probability distribution over Y.

[11]As an anonymous reviewer points out, defining utility as a function over $A \times \Omega$ assumes that

assume that E is a set of subsets of Ω, let alone a partition of Ω. Ω and E can be any non-empty (finite) sets. [12]

For a given experiment $\mathcal{E} = (E, \sigma)$, let 'an evidence-contingent action plan' be a function $\alpha : E \to \Delta(A)$; that is, α maps each experimental outcome in E to a probability distribution on the action set A. So, for each $e \in E$ and each $a \in A$, $\alpha(a|e)$ is the probability of choosing a when the experimental outcome is e.

Let 'a state-contingent action plan' be a function $\lambda : \Omega \to \Delta(A)$. To specify what kind of state-contingent plan is available under an experiment, we first define the composition of two stochastic functions. Suppose $\gamma : X \to \Delta(Y)$ and $\beta : Y \to \Delta(Z)$ are two stochastic functions. Define the composition $\beta \circ \gamma : X \to \Delta(Z)$ by $(\beta \circ \gamma)(z|x) = \sum_{y \in Y} \beta(z|y)\gamma(y|x)$. (In the degenerate case where both γ and β are deterministic functions, the composition reduces to the familiar form: $(\beta \circ \gamma)(x) = \beta(\gamma(x))$.) For any experiment $\mathcal{E} = (E, \sigma)$ and any decision problem $\mathcal{D} = (A, u)$, call a state-contingent action plan $\lambda : \Omega \to \Delta(A)$ 'available under $\mathcal{E}$' if there exists an evidence-contingent plan α such that $\lambda = \alpha \circ \sigma$. This captures the above-mentioned idea that a state-contingent plan is made available by performing an experiment if it can be indirectly followed by following some evidence-contingent plan.

To see this in action, consider the umbrella case involving two different weather forecasts again. Suppose one weather forecast has a 90% reliability and the other has an 80% reliability. If the agent checks the more reliable forecast before deciding, the state-contingent action plan {rain: bring an umbrella with probability 0.9; not-rain: bring an umbrella with probability 0.1} is available, since it's the composition

one doesn't care about what one learns. This assumption is included in our characterization of 'cost-negligible' experiments (see fn. 2).

[12]Note that, even when E is not a set of subsets of Ω, we can still use Kripke models to define externalist theories of evidence precisely, using an epistemic accessibility relation defined over the enlarged space $\Omega \times E$, where the accessibility relation captures what one is aware of or what one can be certain of in an enlarged state $< w, e >$. Internalist theories of evidence can be characterized by the requirement that the accessibility relation R satisfy the 'same-evidence' condition, that is, $< w, e > R < w^*, e^* >$ only if $e = e^*$. And externalist theories of evidence deny this requirement on R. Thanks to an anonymous reviewer for the clarification.

of the evidence-contingent action plan {rain is forecasted: bring an umbrella with probability 1; not-rain is forecasted: bring an umbrella with probability 0} with the experimental distribution given by σ(rain is forecasted | rain) = 0.9 and σ(rain is forecasted | not-rain) = 0.1. However, it's easy to verify that this state-contingent action plan is not available if the agent checks the less reliable forecast before deciding: it's not the composition of any evidence-contingent plan with the experimental distribution given by σ'(rain is forecasted | rain) = 0.8 and σ'(rain is forecasted | not-rain) = 0.2. But the opposite is true: any state-contingent plan that's available by checking the less reliable forecast is also available by checking the more reliable one. (As we will see, the reason is that you can always add some noise to the more reliable forecaster's prediction in order to get the less reliable forecaster's prediction.) For example, consider the state-contingent action plan {rain: bring an umbrella with probability 0.8; not-rain: bring an umbrella with probability 0.2} which is available with the less reliable forecaster. This state-contingent plan is also available with the more reliable forecaster: it's the composition of the evidence-contingent plan {rain is forecasted: bring an umbrella with probability 7/8; not-rain is forecasted: bring an umbrella with probability 1/8} with the more reliable forecaster's experimental distribution.

To show that more evidence generally leads to more available state-contingent action plans, we need to define the concept of 'more informative' experiments precisely. When is one experiment $\mathcal{E}_1 = (E_1, \sigma_1)$ more informative about the state space Ω than another experiment $\mathcal{E}_2 = (E_2, \sigma_2)$?

In the simple case where both experiments are deterministic,[13] there is an intuitive answer: $\mathcal{E}_1$ is more informative than $\mathcal{E}_2$ if it induces a finer partition of the state space, that is, if $\{\sigma_1^{-1}(e) \subseteq \Omega : e \in E_1\}$ is a finer partition of Ω than $\{\sigma_2^{-1}(e) \subseteq \Omega : e \in E_2\}$ (note that, for a deterministic experiment, we can understand the stochastic func-

[13] An experiment is deterministic when, for each state, it assigns probability 1 to some experimental outcome.

tion σ as a function $\sigma : \Omega \to E$). But many experiments are nondeterministic, such as checking an imperfectly reliable weather forecast or measuring the length of a table with a ruler.

Of course, formally, we can always reduce a nondeterministic experiment to a deterministic one: we just enlarge the underlying state space Ω into $\Omega \times E$, and when there are multiple nondeterministic experiments in question, we can just introduce a common state space by enlarging Ω into $\Omega \times E_1 \times \cdots \times E_n$. This, however, won't make it easier to compare the informativeness of two nondeterministic experiments. The above criterion of 'inducing a finer partition' won't do. For some pairs of experiments (E_1, σ_1) and (E_2, σ_2), one is intuitively more informative than the other and yet doesn't induce a finer partition of the common state space $\Omega \times E_1 \times E_2$. The umbrella case involving two different weather forecasting sources is an example. So, we need a condition more general than 'inducing a finer partition of the state space.'

There is another important reason why nondeterministic experiments are interesting: seemingly deterministic experiments whose outcomes are nontransparent can be modeled as nondeterministic experiments. Consider the unmarked clock case again. The experiment of taking a look appears to be deterministic: if the digit points at 1, taking a look will produce evidence {1, 2, 4} with probability 1 and other evidence with probability 0, and likewise for other digit positions. However, since the evidence is nontransparent, when a person receives evidence {1, 2, 4}, there is some chance that she will take her evidence to be {1, 2, 3} or {1, 3, 4}. So, even if actual evidence will be deterministically produced, what she will take her evidence to be might be nondeterministic. And if the person happens to be of the type whose actions are guided by her taking of the evidence rather than the actual evidence, then the experimental outcomes, which are understood as providing inputs of plans that guide actions, will be nondeterministic.

So, understanding the informativeness condition of nondeterministic experiments is important. In this paper, we will use a condition that can be called 'garbling,' a term inspired by Blackwell's (1951) discussion of comparative informativeness. Intuitively, one experiment garbles another if the distribution of the first experimental outcomes can be replicated by first performing the second experiment and then adding some 'noise'. As we will see, this condition matches our intuitive judgments on informativeness in a wide range of cases. Here is the definition: experiment (E', σ') is a garbling of experiment (E, σ) if there is a stochastic function $g : E \to \Delta(E')$ such that for any $w \in \Omega$ and any e' in E', $\sigma'(e'|w) = \sum_{e \in E} g(e'|e)\sigma(e|w)$. We say that (E, σ) is 'weakly more informative' than (E', σ') if the latter is a garbling of the former. We say that (E, σ) is 'more informative' than (E', σ') if the latter is a garbling of the former and not vice versa.

Garbling can be used as a condition of comparative informativeness because adding noise doesn't add information (of course, the noise must be state-independent, so that it doesn't carry information about the world). For example, imagine that the teacher writes the exam date on a piece of paper and gives it to Ben. Ben, who has poor eyesight, looks at the paper and writes the message on another piece of paper. Then Ben's writing cannot be more informative about the exam date than the teacher's writing: it's a result of the teacher's writing 'garbled' by his poor eyesight.

The garbling condition matches our intuitive verdicts on comparative informativeness in a wide range of cases. For example, it implies:

(a) Between two experiments whose outcome spaces both partition the underlying state space, the one that partitions the state space more finely is more informative.

(b) Testimonies from a more reliable person are more informative than those from a less reliable one.

(c) Suppose a trait occurs in the population with an unknown distribution. Then a larger-sized random sampling of the population is more informative about the true distribution of the trait than a smaller-sized random sampling.[14]

(d) Any experiment that doesn't have the same evidence distribution across all states is more informative than a trivial experiment (an experiment whose outcome is certain).

In each case, the intuitively less informative experiment is a garbling of the intuitively more informative one, and not vice versa. (It's worth noting that the ordering of experiments by garbling is not complete, which is also intuitive.)

However, there is a potential worry about the garbling condition: if we use this condition to define comparative informativeness, the individuation of state space will be important. For example, checking a 90% reliable forecast makes more state-contingent plans available than checking an 80% reliable weather forecast, but that's partly because we have decided that 'what the weather forecast says' is not part of our state-space. If we enlarge the state space from Ω to $\Omega \times E_1 \times E_2$, where E_1 and E_2 are the sets of experimental outcomes generated by checking the

[14]Here is an illustration. Suppose that there are two world states, w_1 and w_2, and the trait occurs with probability p in w_1 and with probability q in w_2. Consider two random samplings, one with size one and the other with size two. Let X and Y be two random variables representing the number of times the trait occurs in each of the two sampling results. The distributions of X and Y are shown in Table 2. We can see that X's distribution is a garbling of Y's distribution by the stochastic matrix depicted in Table 3.

Table 2: The Distributions of X and Y

	X = 1	X = 0
w_1	p	$1-p$
w_2	q	$1-q$

	Y = 2	Y = 1	Y = 0
w_1	p^2	$1-p^2-(1-p)^2$	$(1-p)^2$
w_2	q^2	$1-q^2-(1-q)^2$	$(1-q)^2$

Table 3: The Garbling Function

	X = 1	X = 0
Y = 2	1	0
Y = 1	0.5	0.5
Y = 0	0	1

two forecasts, respectively, then it's no longer true that checking the more reliable forecast makes more state-contingent plans available.[15]

In response, we agree that the individuation of state space is important; however, we don't think that this makes the garbling condition problematic. First, it's not problematic to make comparative informativeness relative to a state space. It's intuitively plausible that whether one experiment is more informative than another depends on how we answer the question, 'Informative about what?' For example, it's intuitive that, compared with checking an 80% reliable weather forecast, checking a 90% reliable forecast is more informative about *the weather* but less informative about the question 'what will the 80% reliable forecaster predict'?

Second, it's also not a problem that, with the garbling condition, intuitively comparable experiments often turn out to be incomparable (i.e., neither one turns out more informative than the other) if we enlarge the state space to include evidence. For our intuition about comparative informativeness is often with regard to 'informativeness about the underlying, small state space that generates the evidence' rather than informativeness about the enlarged state space. For example, we don't have the intuition that checking a more reliable weather forecast is more informative about 'will it rain, what will the more reliable forecaster predict, and what will the less reliable forecaster predict?'

Now, we are ready to state and prove the equivalence between gaining evidence and increasing flexibility.

> *Gaining Evidence Is Equivalent With Increasing Flexibility* (EF):
>
> For any experiment $\mathcal{E}$ and any decision problem $\mathcal{D}$, let $\mathcal{R}_{\mathcal{E}}^{\mathcal{D}}$ denote the set of state-contingent action plans available under experiment $\mathcal{E}$ for decision problem $\mathcal{D}$. Let $\mathcal{E}_1 = (E_1, \sigma_1)$ and $\mathcal{E}_2 = (E_2, \sigma_2)$ be two experiments. Then the following two claims are equivalent:

[15] Thanks to an anonymous reviewer for raising this worry.

(1) $\mathcal{E}_1$ is more informative than $\mathcal{E}_2$.

(2) For any decision problem $\mathcal{D}$, $\mathcal{R}^{\mathcal{D}}_{\mathcal{E}_2} \subseteq \mathcal{R}^{\mathcal{D}}_{\mathcal{E}_1}$; for some decision problem $\mathcal{D}$, $\mathcal{R}^{\mathcal{D}}_{\mathcal{E}_2} \subset \mathcal{R}^{\mathcal{D}}_{\mathcal{E}_1}$.

Proof. See Appendix.

4 Greater Flexibility Implies Greater Value

So far, we've argued for EF, the claim that gaining evidence is equivalent with increasing flexibility in action planning. This result is significant because it supports CVE, which is more general than VE as we have argued in §2. Recall

> *Comparative Value of Evidence (CVE)*
>
> For any decision problem one faces, if one can choose at a negligible cost between performing a more informative experiment before deciding and performing a less informative experiment before deciding, and if one is certain that one will respond to the experimental outcome rationally, one should prefer the more informative experiment.

Now, it may appear that EF implies CVE straightforwardly. The reasoning is simple. EF says that more evidence always leads to more, and sometimes strictly more, available action plans (hereafter, 'action-plans' refer to 'state-contingent action-plans' unless otherwise stated). Since an available action plan is an option you can realize, EF implies that, for any decision problem, more evidence gives you more options to choose from. But more options can never harm you in decision-making (if you are certain that you will be rational)—the option that's optimal in a given set cannot be worse than the option that's optimal in its subset.

To illustrate, let A and B be two experiments, where B is more informative than A. Suppose that the set of action plans available under A is $\{a, b\}$ and those available under B are {a, b, c}. Suppose b is optimal in {a, b} and c is optimal in {a, b, c}. Then, making a decision after performing experiment A means choosing b, whereas making a decision after performing experiment B means choosing c. But c is not worse than b—by supposition, c is optimal in the larger set. So, choosing from the larger set cannot be worse than choosing from the smaller set.

But this reasoning rests on an important assumption: the standard of optimality doesn't change as you gain evidence. For ease of explanation, let's suppose that your preference over action-plans is representable by some value function V that maps each plan to a real number, such that a plan is preferred to another just in case it has a greater V-value. Suppose that gaining evidence expands the available action plans from {a, b} to {a, b, c}, but also changes the value function from V to a new one, V*. Then this situation can happen: for plans a, b, and c, $V(c)>V(b)>V(a)$, but $V^*(a)>V^*(b)>V^*(c)$. Then, if one's choice of plans after gaining evidence is determined by the new value V*, one will follow plan a after gaining evidence. But this choice is suboptimal according to the old value function V. From the standpoint of V, one can do better by not gathering evidence and instead choosing the available plan b, which has a greater V-value than a.

Indeed, one's value function over action-plans can change quite often even in the classical decision theory. The value of an action plan depends on its payoff distribution, namely, the probability distribution of the payoffs that will result from following the plan. So, the value can change as one's probabilities change due to enriched evidence.

However, although your value function over action plans can change, the change doesn't always affect your actual choice of actions. When your value function changes from V to V* after gaining some evidence e, the action plan that's optimal according to V and the action plan optimal according to V* may agree on e-states, so it doesn't matter whether your choice of action after gaining evidence

e is dictated by V or V*. This is precisely what happens in the classical settings, where the value function is expected utility and the evidence is partitional. Here, the old optimal plan and the new optimal plan recommend the same action in states compatible with the evidence.[16] So, in the classical settings, if you are certain that you will be rational, you can be certain that your post-experimental choices conform with your pre-experiment value.

But outside of the classical settings, the change in value function can affect your choice of action. This is what happens in the risk-aversion case mentioned in §2. Before gaining evidence $\{w_1, w_2\}$, your value function, which computes the risk-weighted expected utility of action plans relative to the present credence, ranks $\{w_1$: L_1; w_2: L_1; w_3: $L_1\}$ as optimal, but after gaining evidence $\{w_1, w_2\}$, your new value function ranks $\{w_1$: L_2; w_2: L_2; w_3; $L_1\}$ and $\{w_1$: L_2; w_2: L_2; w_3: $L_2\}$ as optimal. The two sets of plans disagree on what to do when $\{w_1, w_2\}$ obtains: the former recommends L_1, while the latter recommends L_2.

So, outside of the classical settings, in order for you to benefit from having more available action plans resulted by gaining evidence (beneficial from the standpoint of your pre-experiment value function), you must ensure that, even if your value function changes, your post-experiment choice of action aligns with your pre-experiment value function. That is, no matter which experiment you perform,

[16] Assume the classical decision theory and conditionalization. Let $\mathcal{D}$ be a decision problem, $\mathcal{E}$ be an experiment whose outcome space E partitions the state space Ω, and $\mathcal{R}^{\mathcal{D}}$ be the set of state-contingent action plans available under experiment $\mathcal{E}$. Let V be one's value function before performing $\mathcal{E}$ and, for any $e \in E$, let V_e be one's value function after gaining e. The claim here is that for any $e \in E$, if $p^*_e = \arg\max_{p \in \mathcal{R}^{\mathcal{D}}} V_e(p)$, and $q^* = \arg\max_{p \in \mathcal{R}^{\mathcal{D}}} V(p)$, then q^* agrees with p^*_e on all e-states. For simplicity, we only prove the case for deterministic action plans here.

Proof: Let the agent's prior be c and utility function be u. Then for any $e \in E$, $p^*_e = \arg\max_{p \in \mathcal{R}^{\mathcal{D}}} V_e(p)$
$= \arg\max_{p \in \mathcal{R}^{\mathcal{D}}} \sum_{w \in \Omega} c(w/e) u(p(w), w)$
$= \arg\max_{p \in \mathcal{R}^{\mathcal{D}}} \sum_{w \in e} c(w/e) u(p(w), w)$
$= \arg\max_{p \in \mathcal{R}^{\mathcal{D}}} \sum_{w \in e} c(w) u(p(w), w)$.
But notice $q^* = \arg\max_{p \in \mathcal{R}^{\mathcal{D}}} V(p)$
$= \arg\max_{p \in \mathcal{R}^{\mathcal{D}}} \sum_{w \in \Omega} c(w) u(p(w), w)$
$= \arg\max_{p \in \mathcal{R}^{\mathcal{D}}} \sum_{e \in E} \sum_{w \in e} c(w) u(p(w), w)$.
Then for any $e \in E$, q^* must agree with p^*_e on all e-states: otherwise, there will be some $e \in E$ such that q^* can be improved by changing its recommendation on e-states into p^*_e's recommendation on e-states and thus won't be optimal in V.

you will follow the plan that's *ex ante* optimal among the plans available under the experiment ('ex ante' meaning the pre-experiment perspective). For simplicity, let's call this 'the resoluteness condition.' The resoluteness condition essentially means that you are committed to following the action plan you deemed optimal at an earlier time, even if your preferences change later.

Now, some scholars have argued that the resoluteness condition should hold for any rational person (Gauthier, 1997; Machina, 1989; McClennen, 1997). According to these scholars, in a dynamic choice problem, a rational person treats all future decisions as a single choice problem, makes a 'life plan' that's optimal from the present perspective, and commits to following it at all later times. So, if right now you think the optimal plan is the one that says 'if event C happens, choose *a*,' then you should choose *a* when gaining evidence C. You should do so even if you might have been inclined to act differently in that moment without the pre-existing plan.

Other scholars hold a different view on rational dynamic choices. They have argued that a rational person should be 'sophisticated' rather than resolute (Maher, 1992; Seidenfeld, 1988). In deciding what to choose at a later decision node, the issue of which plan has previously been deemed optimal is irrelevant. Instead, one should treat that decision node as the first in a decision tree, and then one should use backward induction, treating the choice at the latter time as given in determining what to choose at the earlier time.

It's an ongoing debate whether rational people should be resolute or sophisticated.[17] We can feel the force of arguments on both sides. On the one hand, there is some inconsistency in thinking of the plan 'if this thing happens, choose *a*' as

[17]The sophistication view seems to be the dominant position (at least in the value-of-evidence literature); the main argument is that resoluteness sometimes requires counter-preferential choice and thus is irrational (Maher, 1992; Steele, 2010). But there are also abundant defenders of the resoluteness view. Gauthier (1997), Machina (1989), and McClennen (1997) all argue that resoluteness avoids the dominance problem that sophisticated agents face. Thoma (2019) argues that resoluteness enables risk-averse agents to align single-case preferences with long-run ones; Buchak (2015) suggests that the resoluteness view fits better with a holistic picture of agency.

optimal and then not choosing a when you see it happens. It seems that a rational person should be consistent over time and follow whatever plan they have deemed the best. On the other hand, it seems reasonable that a person's updated value function should have some influence over what to choose at the later time. As some writers have insightfully commented, the debate over sophistication and resolution may hinge on a deeper issue about the nature of agency, specifically, the issue of whether we should understand an agent as a collection of different time-slices with separate interests or as a single and whole person extending over time (Buchak, 2013, pp. 180–2; Buchak, 2015, p. 14). The sophistication view fits better with the time-slice view of agency, whereas the resolution view fits better with the holistic view.

It will take us too far afield to delve into the debate over resolution and sophistication. To move forward, we will simply sidestep the issue and propose a variant of CVE that remains neutral on this debate:

> **The Comparative Value of Evidence for the Resolute (CVER)**
>
> For any decision problem one faces, if one can choose at a negligible cost between performing a more informative experiment before deciding and performing a less informative experiment before deciding, and if one is certain that one will follow the action-plan that's ex ante optimal among the plans available under the experiment to be performed, one should prefer the more informative experiment.

Again, 'ex ante optimal' means optimal according to the pre-experiment perspective. CVER is silent on whether rational persons should be resolute. It simply says that if one is resolute—or, more exactly, if one is certain that one is resolute—more evidence is always better.

CVER differs from CVE in replacing the rationality condition 'certain to respond to experimental outcome rationally' with the resoluteness condition 'certain to conform to ex ante optimal plans.' The resoluteness condition makes the rationality condition redundant: as long as you are certain to be resolute, more evidence will be more valuable, even if you are uncertain whether you will update beliefs rationally. Again, the reason is simple: more evidence brings more options, and from your pre-experiment perspective, more options can't harm you if you are certain to follow the ex ante optimal plan. So, even if you are uncertain whether you will update beliefs rationally (perhaps you simply don't know how), you can still benefit from gaining evidence as long as you are certain to be resolute in following action-plans. And certainty of resoluteness without certainty of rational belief-updating is possible: even if you are uncertain how you will update beliefs, you can still be certain that you will be resolute by, say, employing some commitment device or a device that ensures that you will not redeliberate after experiment and will just automatically follow the action plan that you now think is optimal.

Of course, updating beliefs rationally can sometimes help you be resolute. In the classical settings (where evidence is partitional and the decision theory is classical), rational belief updating typically means following conditionalization, and as we've shown in fn. 16, maximizing expected utility relative to the conditioned belief guarantees resoluteness—it results in conforming to the ex ante optimal plan since the plan agrees with the ex post optimal plan in states compatible with the evidence gained. However, outside of the classical settings, updating beliefs rationally doesn't guarantee resoluteness (we'll discuss this further shortly). When rational updating and resoluteness come apart, it's resoluteness that guarantees the value of free evidence.

5 CVER is Acceptable to Externalists and Nonclassical Decision Theorists

We've argued that EF implies CVER in the last section. In this section, we argue that externalists of evidence and nonclassical decision theorists can both embrace CVER. To begin, note that the reasoning from EF to CVER is fully general, not restricted to internalist evidence or classical decision theories. Again, the core idea of the reasoning is simple: more evidence gives you more options, and more options cannot be harmful from the pre-experiment perspective if you are certain to be resolute.

Of course, for convenience, the above discussion has assumed that one's preference is representable by a real-valued function that outputs a complete ranking for any set of options. But this assumption is not necessary. Even for decision theories that only select a subset of permissible options rather than producing a complete ranking (which is common among decision theories designed for imprecise credence), the inference from EF to CVER remains valid as long as the decision theory satisfies the Chernoff condition (also known as 'Sen's property alpha' or 'independence of irrelevant alternatives.') This condition states that if an option is permissible relative to a set of options, it must also be permissible relative to any of its subset. In other words, expanding the option set never makes an impermissible option permissible. This condition is fairly weak, as it's met by most nonclassical decision theories in the literature.[18] Therefore, it's fair to say that CVER is acceptable to most nonclassical decision theories in the literature.

Moreover, note that existing arguments claiming that externalist evidence or nonclassical decision theories lead to evidence-aversion all rely on the assumption

[18]These include Disappointment Theory (Loomes & Sugden, 1986), Rank-Dependent Utility Theory (Quiggin, 1982), Cumulative Prospect Theory (Tversky & Kahneman, 1992), Choquet Expected Utility Theory (Schmeidler, 1989), and many decision theories for imprecise credence. An exception is the regret theory proposed by Bell (1982) and generalized by Sugden (1993).

of sophistication. They don't show that a resolute person committed to following the ex ante optimal action plans can be evidence-averse. For example, in the risk-aversion case discussed in §2, the argument assumes that you will maximize risk-weighted expected utility later and take the later choice as given when reasoning backwards about whether to accept the free evidence at the earlier time. Similarly, in the unmarked clock case, the argument assumes that you will maximize expected utility later and take that choice as given when reasoning backwards about the early offer of free evidence.

In both cases, you can predict that you won't follow the ex ante optimal plan after gaining evidence. In the risk-aversion case, the sophisticated agent will choose L_2 after gaining evidence $\{w_1, w_2\}$, whereas your ex ante optimal plan recommends L_1. In the unmarked clock case (where the decision theory is classical but the evidence is nontransparent), you are assumed to update beliefs by conditionalization after looking and thus follow the evidence-contingent plan:

EP = {{1, 2, 4}: EVEN; {1, 2, 3}: ODD; {2, 3, 4}: EVEN; {1, 3, 4}: ODD}

When the experiment of taking a look at the clock is understood as a deterministic experiment (as it's often taken to be), following the evidence-contingent plan EP will result in following the state-contingent action plan

{1: EVEN; 2: ODD; 3: EVEN; 4: ODD}

This plan is clearly not ex ante optimal among all available plans since it leads to a guaranteed loss, whereas some other available plans don't. For instance, the plan of choosing neither bet in all states is clearly ex ante better, with an expected utility of 0, and this plan is available: you can follow it by following the evidence-contingent plan 'choosing neither bets no matter what your observation is.'

What if taking a look is understood as a nondeterministic experiment, as we have suggested in §3.2? Suppose that, when the digit points at 1 and thus taking a look will give one evidence {1, 2, 4}, one will take the evidence to be {1, 2, 4} with a probability of 0.8 and will take it to be {1, 2, 3} and {1, 3, 4} each with a probability of 0.1, and likewise for other digit positions. Then, following the evidence-contingent plan EP will result in following the state-contingent plan

SP = {1: (0.8 EVEN, 0.2 ODD); 2: (0.2 EVEN, 0.8 ODD); 3: (0.8 EVEN, 0.2 ODD); 4: (0.2 EVEN, 0.8 ODD)}

where the act (0.8 EVEN, 0.2 ODD) means choosing EVEN with a probability of 0.8 and choosing ODD with a probability of 0.2.

While SP avoids a guaranteed loss, it's still not ex ante optimal. It's expected utility is –0.4$, whereas the opposite plan

SP* = {1: (0.2 EVEN, 0.8 ODD); 2: (0.8 EVEN, 0.2 ODD); 3: (0.2 EVEN, 0.8 ODD); 4: (0.8 EVEN, 0.2 ODD)}

has an expected utility of 0.2$. SP* is available: it can be realized by following an evidence-contingent plan that's exactly the opposite of EP, that is, a plan that maps ODD instead of EVEN to {1, 2, 4} and likewise for other evidence. It's easy to show that, given your reliability of 0.8 in interpreting evidence, SP* rather than SP is ex ante optimal among all available state-contingent plans.

The fact that you are irresolute in the unmarked clock case—that is, you won't follow the ex ante optimal plan after looking—is not surprising. The culprit here is your belief-updating strategy. You are assumed to update beliefs through conditionalization (which leads you to follow the problematic evidence-contingent plan EP) in the case. But as mentioned earlier, conditionalization only guarantees resoluteness in classical settings, where evidence is partitional and the decision theory is expected utility, whereas the evidence in this case is not partitional. Furthermore, as several scholars have pointed out, when evidence is nonpartitional, conditionalization is not ex ante epistemically optimal: it doesn't maximize expected accuracy relative to your pre-experiment beliefs (Schoenfield, 2017). When your beliefs after gaining evidence are not ex ante optimal, it's not surprising that your actions after gaining evidence also won't be ex ante optimal.

This raises an interesting question: when evidence is nonpartitional but the decision theory remains classical, is there a belief-updating strategy that guarantees resoluteness? An extensive discussion is beyond the scope of this paper, but we'll note this: among existing externalist updating strategies that have been shown to maximize expected accuracy, the strategy called 'Update Conditionalization' by Gallow (2021, p. 500) guarantees resoluteness.[19] If you update beliefs by this strategy, then for each possible experimental outcome, the action plan optimal according to the post-experiment value function agrees with the action plan optimal according to the pre-experiment value function in all states in which you update beliefs on the experimental outcome.[20] This guarantees that, no matter whether you choose according to the pre-experiment value function or the post-experiment value function, the ex ante optimal plan will be conformed to. For proponents of resoluteness in dynamic decision-making, this can be seen as a strength of Update Conditionalization, but we'll leave further exploration for another time.

6 Comparing to Ahmed and Salow's 'Conditionality'

So far, we have argued that CVER offers a compelling alternative to VE. It's not only more general in telling us to prefer more informative experiments to less informative ones (not just preferring nontrivial experiments to trivial ones), but it's also acceptable to externalists and nonclassical decision theorists. In this section, we will compare CVER with another existing proposal that aims to replace VE. The comparison will further highlight the strengths of CVER.

Ahmed and Salow (2019) have proposed to replace VE with

[19]This strategy says that when the total evidence gained is e, c_{new} should be $c_{old}(/U_e)$, where 'U_e' is the proposition that one updates on evidence e.

[20]The proof is similar to the one provided in fn. 16 above: we can replicate the entire argument by substituting $\{U_e : e \in E\}$ for E and U_e for e throughout.

Conditionality

> A rational agent will make her action depend on the (true) answer to any question, whenever she can freely do so.

According to Ahmed and Salow (2019), Conditionality implies VE when and only when gathering evidence before deciding is a way of making one's action depend on the answer to a question. This condition, as they argue, is not met in the unmarked clock case or the risk-aversion case discussed in §2. In the unmarked clock case, for instance, taking a look at the clock is not a way of answering a question—Ahmed and Salow identify a question as a partition of the underlying state space, and the outcomes of taking a look doesn't partition the space.

Now, we think that replacing VE with Conditionality is on the right track—Ahmed and Salow's basic motivation for Conditionality is that making one's action depend on the answer to a question creates a larger option space, and our principle CVER also rests on the idea that more evidence means a larger option space.

However, we argue that CVER offers a better value-of-evidence principle compared to Conditionality for two reasons. First, due to Ahmed and Salow's restriction that a question must partition the underlying state space, Conditionality is simply silent on the value of nonpartitional evidence. But ideally, we would like a principle that tells us whether and when such evidence is valuable. Second, and more importantly, Conditionality is too weak: even within classical settings where evidence is partitional and the decision theory is classical, Conditionality doesn't tell us to gather free evidence before deciding.

To see this, note that it's very easy for you to make your actions depend on the answer to a question—waiting for evidence is one way to achieve this, but so is making a decision right away. Just imagine a trivial question, a question whose answer you already know for certain, such as 'would this double-headed coin land on heads or tails if flipped?' Deciding right away is essentially the same as making your action depend on the answer to this trivial question: when you are certain of heads, 'deciding now and choosing *a*' is equivalent to planning 'if heads, choose *a*; if tails, choose *b*.' So, even within the classical settings, Conditionality doesn't tell you to wait for the free evidence, given that both waiting and deciding right away make your actions depend on the answer to a question. The problem, as we can see, is that Conditionality only emphasizes making actions depend on the answer to a question, but it doesn't specify whether that question should be informative.

Compared to Conditionality, our principle CVER is more general: it tells us to prefer a more informative experiment to a less informative one and, unlike Ahmed and Salow's notion of a question, an experiment as we understand it doesn't have to partition the state space. This also makes CVER more powerful, since it tells us that even nonpartitional evidence can be valuable in certain conditions and the conditions are *the same* as the ones in which partitional evidence is valuable.

7 Conclusion

In this paper, we've proposed to replace VE with CVER, the principle that a more informative experiment should be preferred to a less informative one, if one is certain that one will be resolute. We have defended CVER by defending EF, a claim of equivalence between gaining evidence and increasing flexibility in action planning. We've argued that CVER has several virtues: it generalizes VE to the comparative case; it's acceptable to externalists and nonclassical decision theorists; and it's more powerful than the existing VE substitute in the literature. Our argument, if successful, shows that the value-of-evidence principle is more robust than previously thought and broadens the scope of such principles.

We will conclude this paper with two more implications of EF, the claim of equivalence that's key to our case for CVER. First, it generates a new and stronger argument for the accuracy value of evidence than those found in the literature. All existing arguments for evidence's accuracy value must assume that accuracy measures are strictly proper and that the decision theory is classical (Dorst, 2020; Dorst et al., 2021; Myrvold, 2012; Oddie, 1997). These assumptions will not be required in an EF-based argument for the accuracy value of evidence. By EF, gaining evidence will give one more available credal plans (which are maps from states to credence functions); having more available credal plans can never make one worse off accuracy-wise by the light of the present beliefs, if one is certain to follow the ex ante optimal credal plan.

Second, EF has an important implication for zetetic epistemology, in particular, for the question 'what evidence should we gather in our inquiry, that is, what kind of experiments should we set out to perform in our inquiry?'

A natural answer is this: setting aside considerations of costs, we should choose the most informative experiment among those we can perform. But what makes one experiment more informative than another? A good criterion should not only match our intuitive judgments of informativeness, but also explain why people should prefer more informative experiments to less informative ones before making decisions. In this paper, we have relied on the garbling condition and have used it to explain the value of evidence through establishing EF. This argument can be seen as supporting the use of garbling as a criterion of informativeness. Of course, we're open to exploring alternative ways to measure the informativeness of experiments, and it will be interesting to see what value-of-evidence principle might emerge from those alternative measures.[21] We hope, however, that our discussion has at

[21] For instance, here is a natural alternative approach: we can measure the informativeness of an experiment for an agent by the expected reduction in the entropy of the agent's priors. Unlike the garbling condition, this measure is agent-relative—the same experiment will be informative relative to some priors but not others; it also doesn't imply a general value-of-evidence principle: it's not true that, for any agent S, for any experiments $\mathcal{E}_1$ and $\mathcal{E}_2$, if performing $\mathcal{E}_1$ results in a greater expected reduction in the entropy of S's priors than performing $\mathcal{E}_2$ does, then for any decision problem $\mathcal{D}$, S should prefer performing $\mathcal{E}_1$ over performing $\mathcal{E}_2$ before making a decision about $\mathcal{D}$.

An anonymous reviewer suggests a second alternative: we can say that an experiment is more least shown that considering the value of evidence in decision-making can shed light on research on informativeness measures and the broader question of what kind of evidence we should gather in our inquiries.[22]

Bibliography

Ahmed, A. & Salow, B. (2019). Don't look now. *The British Journal for the Philosophy of Science, 70*(2), 327–350.

Bell, D. (1982). Regret in decision making under uncertainty. *Operations Research, 30*, 961–981.

Blackwell, D. (1951). Comparison of experiments. *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability, 2*, 93–102.

Bradley, S. & Steele, K. (2016). Can free evidence be bad? Value of information for the imprecise probabilist. *Philosophy of Science, 83*, 1–28.

Briggs, R. (2015). Costs of abandoning the sure-thing principle. *Canadian Journal of Philosophy, 45*, 827–840.

Buchak, L. (2010). Instrumental rationality, epistemic rationality, and evidence-gathering. *Philosophical Perspectives, 24*(1), 85–120.

Buchak, L. (2013). *Risk and Rationality*. Oxford: Oxford University Press.

Buchak, L. (2015). Revisiting risk and rationality: A reply to Pettigrew and Briggs. *Canadian Journal of Philosophy 45*(5), 841–862.

Campbell-Moore, C. & Salow, B. (2020). Avoiding risk and avoiding evidence. *Australasian Journal of Philosophy, 98*(3), 495–515.

informative just in case it has a greater expected value relative to any priors and utilities. This ordering turns out to be equivalent to the Blackwell ordering by garbling (Blackwell, 1951). However, we think that it's unsatisfying to directly define comparative informativeness in terms of greater value, because the definition makes it trivially true that a more informative experiment is more valuable.

[22]This paper has greatly benefited from discussions with Yifeng Ding, Kuno Zhao, two anonymous referees for *The Journal of Philosophy*, and audiences at NYU Shanghai, Shanghai Jiao Tong University, and Beijing Normal University.

Das, N. (2020). The value of biased information. *British Journal for the Philosophy of Science*, *74*, 25–55.

Das, N. (2023). Credal imprecision and the value of evidence. *Noûs, 57*(3), 684–721.

de Oliveira, H. (2018). Blackwell's informativeness theorem using diagrams. *Games and Economic Behavior, 109*(C), 126-131.

Dorst, K. (2020). Evidence: A guide for the uncertain. *Philosophy and Phenomenological Research, 100*(3), 586–632.

Dorst, K., Levinstein, B. A., Salow, B., Husic, B. E., & Fitelson, B. (2021). Deference done better. *Philosophical Perspectives, 35*(1), 99–150.

Flores, C. & Woodard, E. (2023). Epistemic norms on evidence-gathering. *Philosophical Studies, 180*(9), 1–25.

Friedman, J. (2020). The epistemic and the zetetic. *The Philosophical Review, 129*(4), 501–536.

Gallow, J. D. (2021). Updating for externalists. *Noûs, 55*(3), 487–516.

Gauthier, D. (1997). Resolute choice and rational deliberation: A critique and a defense. *Noûs, 31*(1), 1–25.

Good, I. J. (1967). On the principle of total evidence. *British Journal for the Philosophy of Science, 17*(4), 319–321.

Hughes, N. (2021). Epistemic feedback loops (or: how not to get evidence). *Philosophy and Phenomenological Research, 106*(2), 368–393.

Kadane, J. B., Schervish, M., & Seidenfield, T. (2008). Is ignorance bliss? *Journal of Philosophy, 105*(1), 5–36.

Loomes, G. & Sugden, R. (1986). Disappointment and dynamic consistency in choice under uncertainty. *Review of Economic Studies, 53*, 271–282.

Machina, M. J. (1989). Dynamic consistency and non-expected utility models of choice under uncertainty. *Journal of Economic Literature, 27*(4), 1622–1668.

Maher, P. (1992). Diachronic Rationality. *Philosophy of Science, 59*, 120–141.

McClennen, E. F. (1997). Pragmatic rationality and rules. *Philosophy and Public Affairs 26*(3), 210–258.

Myrvold, W. C. (2012). Epistemic values and the value of learning. *Synthese, 187*(2), 547–568.

Neth, S. (forthcoming). Rational aversion to information. *British Journal for the Philosophy of Science*.

Oddie, G. (1997). Conditionalization, cogency, and cognitive value. *British Journal for the Philosophy of Science, 48*, 533–541.

Quiggin, J. (1982). A theory of anticipated utility. *Journal of Economic Behavior and Organization, 3*, 323–343.

Schmeidler, D. (1989). Subjective probability and expected utility without additivity. *Econometrica, 57*, 571–587.

Schoenfield, M. (2017). Conditionalization does not (in general) maximize expected accuracy. *Mind, 126*, 1155–1187.

Schultheis, G. (forthcoming). Accurate updating. *Philosophy of Science*.

Seidenfeld, T. (2004). A contrast between two decision rules for use with (convex) sets of probabilities: Γ-maximin versus e-admissibilty. *Synthese, 140*, 69–88.

Seidenfeld, T. (1988). Decision theory without 'Independence' or without 'Ordering': What's the difference? *Economics and Philosophy, 4*, 267–290.

Steele, K. (2010). What are the minimal requirements of rational choice? Arguments from the sequential decision setting. *Theory and Decision, 68*, 463–487.

Sugden, R. (1993). An axiomatic foundation for regret theory. *Journal of Economic Theory, 60*, 159–180.

Thoma, J. (2019). Risk aversion and the long run. *Ethics, 129*, 230–253.

Tversky, A. & Kahneman, D. (1992). Advances in prospect theory: cumulative representation of uncertainty. *Journal of Risk and Uncertainty, 5*(4), 297–323.

Wakker, P. (1988). Nonexpected utility as aversion of information. *Journal of Behav-*

ioral Decision Making, 1(3), 169–175.

Williamson, T. (2014). Very improbable knowing. *Erkenntnis, 79*(5), 971–999.

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

Appendix

Gaining Evidence Is Equivalent With Increasing Flexibility (EF).

For any experiment $\mathcal{E}$ and any decision problem $\mathcal{D}$, let $\mathcal{R}_{\mathcal{E}}^{\mathcal{D}}$ denote the set of state-contingent action plans available under experiment $\mathcal{E}$ for decision problem $\mathcal{D}$. Let $\mathcal{E}_1 = (E_1, \sigma_1)$ and $\mathcal{E}_2 = (E_2, \sigma_2)$ be two experiments. The following two claims are equivalent:

(1) $\mathcal{E}_1$ is more informative than $\mathcal{E}_2$.

(2) For any decision problem $\mathcal{D}$, $\mathcal{R}_{\mathcal{E}_2}^{\mathcal{D}} \subseteq \mathcal{R}_{\mathcal{E}_1}^{\mathcal{D}}$; for some decision problem $\mathcal{D}$, $\mathcal{R}_{\mathcal{E}_2}^{\mathcal{D}} \subset \mathcal{R}_{\mathcal{E}_1}^{\mathcal{D}}$.

Proof. To prove the equivalence between (1) and (2), it's sufficient to prove the equivalence between the following two claims:

(1)* $\mathcal{E}_1$ is weakly more informative than $\mathcal{E}_2$.

(2)* For any decision problem $\mathcal{D}$, $\mathcal{R}_{\mathcal{E}_2}^{\mathcal{D}} \subseteq \mathcal{R}_{\mathcal{E}_1}^{\mathcal{D}}$.

(1)* implies (2)*:

Fix a decision problem $\mathcal{D} = (A, u)$. Suppose a state-contingent action plan λ is in $\mathcal{R}_{\mathcal{E}_2}^{\mathcal{D}}$. Then there is an evidence-contingent action plan $\alpha : E_2 \to \Delta(A)$ such that $\lambda = \alpha \circ \sigma_2$. Since (E_1, σ_1) is more informative than (E_2, σ_2), there is a garbling function $\gamma : E_1 \to \Delta(E_2)$ such that $\sigma_2 = \gamma \circ \sigma_1$. Let β be the evidence-contingent action plan: $E_1 \to \Delta(A)$ such that $\beta = \alpha \circ \gamma$. Then $\beta \circ \sigma_1 = (\alpha \circ \gamma) \circ \sigma_1 = \alpha \circ (\gamma \circ \sigma_1) = \alpha \circ \sigma_2 = \lambda$. The second equality holds because the composition between stochastic functions is associative.[23] So, there exists an evidence-contingent action plan $E_1 \to \Delta(A)$ such

[23]For any $a \in A$ and $w \in \Omega$, $((\alpha \circ \gamma) \circ \sigma_1)(a|w)$

that λ is the composition of σ_1 with that plan. So, $\lambda \in \mathcal{R}^{\mathcal{D}}_{\mathcal{E}_1}$. Since λ is an arbitrary plan in $\mathcal{R}^{\mathcal{D}}_{\mathcal{E}_2}$, we have $\mathcal{R}^{\mathcal{D}}_{\mathcal{E}_2} \subseteq \mathcal{R}^{\mathcal{D}}_{\mathcal{E}_1}$.

(2)* implies (1)*:

Let $\mathcal{D}$ be a decision problem in which the action set A has the same cardinality with E_2, and let α be a one-to-one map from E_2 to A. Then $\alpha \circ \sigma_2 \in \mathcal{R}^{\mathcal{D}}_{\mathcal{E}_2}$. Since $\mathcal{R}^{\mathcal{D}}_{\mathcal{E}_2} \subseteq \mathcal{R}^{\mathcal{D}}_{\mathcal{E}_1}$, $\alpha \circ \sigma_2 \in \mathcal{R}^{\mathcal{D}}_{\mathcal{E}_1}$. This means that there is some $\beta : E_1 \to \Delta(A)$ such that $\alpha \circ \sigma_2 = \beta \circ \sigma_1$. Then $\alpha^{-1} \circ (\alpha \circ \sigma_2) = \alpha^{-1} \circ (\beta \circ \sigma_1)$. Since $\alpha^{-1} \circ (\alpha \circ \sigma_2) = (\alpha^{-1} \circ \alpha) \circ \sigma_2 = \sigma_2$, we have $\sigma_2 = \alpha^{-1} \circ (\beta \circ \sigma_1) = (\alpha^{-1} \circ \beta) \circ \sigma_1$. Since $(\alpha^{-1} \circ \beta)$ is a stochastic function from E_1 to $\Delta(E_2)$ (for any $e_1 \in E_1$, we have $\sum_{e_2 \in E_2} (\alpha^{-1} \circ \beta)(e_2|e_1) = 1$), $\mathcal{E}_2$ is a garbling of $\mathcal{E}_1$. This completes the proof. ∎

$= \sum_{e_1 \in E_1} ((\alpha \circ \gamma)(a|e_1))\sigma_1(e_1|w)$
$= \sum_{e_1 \in E_1} (\sum_{e_2 \in E_2} \alpha(a|e_2)\gamma(e_2|e_1))\sigma_1(e_1|w)$
$= \sum_{e_2 \in E_2} \alpha(a|e_2)(\sum_{e_1 \in E_1} \gamma(e_2|e_1)\sigma_1(e_1|w))$
$= \sum_{e_2 \in E_2} \alpha(a|e_2)((\gamma \circ \sigma_1)(e_2|w))$
$=(\alpha \circ (\gamma \circ \sigma_1))(a|w)$.