

Chapter 2

All About Actual Accuracy-Dominance

2.1 Abstract

We have lots of good arguments for a variety of epistemic norms on how you should plan to change your credences or beliefs upon coming to possess new evidence. We don't have many good arguments for how you should actually change your credences or beliefs in response to receiving new evidence. Sure, we do have some arguments for actual epistemic norms, but none of them are the gold standard in the field, that is, none of them are accuracy-dominance arguments. Here we fill this gap. Doing so requires some conceptual development about good and bad ways to evaluate your epistemic performance. In short: your evidence, while not directly placing constraints on your rational attitudes, places a constraint on how you should evaluate your epistemic performance. If you possess evidence E , it seems, from your point of view, bad to take non- E worlds as relevant to the assessment of your epistemic performance. Using this idea, we develop an accuracy-dominance argument for Actual Conditionalization and a variety of other actual updating norms.

2.2 Introduction

Suppose that we consider an agent who has finitely-many credences. Perhaps they're 0.6 confident that it's going to rain tomorrow, perhaps not. The point is that we are going to represent our agent's degrees of belief or confidences using a credence function c_t , which is formally a function from finitely-many propositions to real-valued confidences that an agent has at time t . In this paper, we will be primarily concerned with how such agents are rationally required to respond to receiving new evidence. Clearly, not just anything goes. For example, if I know that I'm going to a magic show, and I see some trickery, I'm probably not going to change my fundamental beliefs about how the world works. Rabbits don't tend to spontaneously materialize in hats. But, if I were walking down the street on the day to day, and I saw such spontaneous rabbit materializations quite often, that might change the way I view the world. Examples like this suggest that rationally responding to learning new evidence by changing your credences depends upon your current credences. We begin by stating some well-known and widely adopted norms on how you are

required to plan to change your credences in response to receiving new evidence E from some evidential partition ϵ . Formally, your plan P is a function from ϵ to the set of credences, with the interpretation being that if you receive exactly evidence E between times t and $t + 1$, then you plan to adopt credences P_E .

Def¹: Plan Conditionalization: your prior and planning pair (c, P) is such that if $c(E) > 0$, then $P_E(.) = c(.|E)$.

Def: Plan Blackwell Condition: $P_E(E) = 1$.

Now, while these norms have a lot going for them in terms of delivering plausible verdicts in a wide variety of cases, it seems desirable to actually argue for them. After all, why think them a requirement of epistemic rationality? In order to answer this, we begin by reviewing how we are going to legitimately measure the accuracy of our credences.

2.3 Measuring Credal Accuracy

The following legitimacy conditions on credal accuracy-measures are fairly standard and well-known (Oddie, 1997) (Joyce, 1998) (Schervish, Seidenfeld, & Kadane, 2009) (Pettigrew, 2016). A is a function from credence-world pairs with the interpretation that $A(c, w)$ is the numerical accuracy of having credences c at possible world w . Further motivation and discussion of these legitimacy conditions can be found in (Pettigrew, 2016).

Def: Strict Propriety: accuracy-measure A is said to be strictly proper iff for every probability function p , only p maximizes the p -expected-accuracy over all credence functions, that is, p maximizes $EA(.|p) = \sum_w p(w)A(., w)$.

In slogan form, strict propriety says that, for A to be legitimate, p must think itself the best. It must expect itself to be the most accurate. While this intuitive slogan used to be the primary reason for accepting strict propriety, recent results in (Levinstein & Campbell-Moore, 2021) and (Williams & Pettigrew, 2023) have provided better arguments.

Def²: Continuity: $A(., w)$ is continuous for every $w \in W$.

Def: Additivity: $A(c, w) = \sum_p a(c(p), w)$

Additivity says that the accuracy of a credence function can be broken up into the sum of the accuracy of each individual credence. While there has been some recent and welcome advances in developing accuracy-dominance arguments without additivity, we are going to only consider additive accuracy-measures for now³ (Pet-

¹We adopt the notation-saving convention that pairs of doxastic attitudes that are not time-indexed are considered synchronic.

²See (Rudin, 1964) for the technical definition of this condition.

³Later on, we will find that we don't need it at all. Furthermore, I suspect that Additivity can be dropped even where it is assumed, though I haven't worked out the details.

tigrew, 2022a) (Nielsen, 2022).

Now that we know how to legitimately measure the accuracy of your credences at such-and-such possible world. But, how are we going to measure the accuracy of your credal plannings?

Def: $A(P, w) = A(P(E_w), w)$.

This condition just says that the accuracy of your credal plannings at world w is the accuracy of your credences that you plan to adopt at world w upon coming to possess evidence E . This definition works (at all worlds) because the possible evidence that you plan to receive comes from an evidential partition, so every world w is in exactly one E .

Finally, we need a way of legitimately measuring the accuracy of your credence/credal-planning pairs (c, P) . (Briggs & Pettigrew, 2018) suggest the following way:

Def⁴: Temporal Separability: $A[(c, P), w] = A(c, w) + A(P, w)$.

Basically, the idea here is that to legitimately measure the total accuracy of your credence/credal-planning pairs, just measure the accuracy of your credences and their plannings individually and then sum them up.

That's it; that is all the legitimacy conditions that we need. Let's put them to work. The idea here is that accuracy seems to be epistemically valuable, epistemic in the sense of considering our credences as representations, and only as representations, of the world. Even further, accuracy seems to be the only thing of intrinsic epistemic (representational) value. This position is called Veritism (Goldman, 1986) (Pettigrew, 2016). We adopt this position throughout the paper. Now, given Veritism, (Pettigrew, 2016) proposes the following norm connecting accuracy with epistemic rationality.

Def: Strong (Weak) Accuracy-Dominance: c is strongly (weakly) accuracy-dominated with respect to accuracy-measure A iff there exists a c' such that $A(c', w) > A(c, w)$ for all possible worlds w ($A(c', w) \geq A(c, w)$ for all w and $A(c', w) > A(c, w)$ for some w).

Def: **Strong (Weak) Undominated Dominance**: Doxastic state D is irrational if there exists another doxastic state D' such that D' strongly (weakly) accuracy-dominates D (with respect to legitimate accuracy-measure A) and, further, D' is not itself strongly (weakly) accuracy-dominated.⁵

Now, with this norm in hand, we have the following technical results, the undomi-

⁴It is perhaps a bit misleading to call this condition *Temporal* Separability because both the c and P are synchronic, so there is no adding of accuracies across different times, but there is adding of accuracies across different kinds of attitudes.

⁵I am here ignoring some subtleties about how we are supposed to understand/quantify over A . See (Pettigrew, 2016) and (Rooyakkers, ms) for a discussion of these details.

nated strong/weak accuracy-dominance theorems for Plan Conditionalization (and the Plan Blackwell Condition) (Briggs & Pettigrew, 2018) (Nielsen, 2021).

Briggs-Pettigrew Theorem: credence/credal-planning pair (c, P) satisfies Probabilism+Plan Conditionalization iff (c, P) avoids strong (undominated) accuracy-dominance for every legitimate accuracy-measure.

Briggs-Pettigrew-Nielsen Theorem: credence/credal-planning pair (c, P) satisfies Probabilism + Plan Conditionalization + the Plan Blackwell Condition iff (c, P) avoids weak (undominated) accuracy-dominance for every legitimate accuracy-measure.

These are substantive results; they can be used to develop an accuracy-dominance argument for Plan Conditionalization (and the Plan Blackwell Condition). Thus, we cannot rationally plan to change our credences all willy nilly: rational credal planning must follow Plan Conditionalization. But what about actual rational credal change and not just plannings to change. That is, are there any rational constraints on diachronic credal pairings (c_t, c_{t+1}) upon coming to possess new evidence? Must your actual future credences be your past credences conditional on your newly received evidence? More formally, are the following constraints rationally required?

Def: Actual Conditionalization (for evidence E): if between times t and $t + 1$ you receive exactly evidence E , your credal-pair (c_t, c_{t+1}) is such that if $c_t(E) > 0$, then $c_{t+1}(\cdot) = c_t(\cdot|E)$.

Def: Blackwell Condition: If, at time t , you possess evidence E , then $c_t(E) = 1$.

The above questions are good questions because it is not immediately obvious how to argue for Actual Conditionalization. This was first pointed out by (van Fraassen, 1989) in the context of developing a different kind of argument (namely, a dutch strategy argument) for Plan Conditionalization. He carefully observed that Plan Conditionalization is a purely synchronic norm, and thus that it couldn't be used just by itself to argue for a genuinely diachronic norm like Actual Conditionalization. This subtlety matters because Actual Conditionalization seems to be a very compelling norm; it is widely adopted and appealed to throughout the sciences (such as in Bayesian Confirmation Theory and Bayesian Statistics). Thus, we begin by surveying the best attempts at arguing for it available in the literature with the ultimate goal of developing a better argument, namely, an accuracy-dominance argument.

2.4 Some Arguments for Actual Conditionalization

2.4.1 The Argument from Diachronic Continenence

On pain of irrationality, you best follow your plans, or so says the principle of Diachronic Continenence (Paul, 2014) (Pettigrew, 2020). After all, why even make plans if there is no normative reason to follow them? Thus, with this new proposed

requirement of rationality in hand, we can directly give an argument for Actual Conditionalization from Plan Conditionalization. This is all well and good, but why think that the normative reason for following your plans amounts to rational requirement and not just rational permissibly? Further, it is not clear that the notion of rationality appealed to in Diachronic Contenance is of the epistemic kind, as opposed to the pragmatic kind. What is the epistemic reason for being required to follow your plans? Maybe there are pragmatic difficulties when you treat your plans finkishly; after all, it probably took a fair amount of effort, time, and cognitive resources to develop your plans, but the relevant question for us is whether or not there are also epistemic difficulties in treating your plans finkishly. It's not immediate that there such epistemic difficulties. After all, Diachronic Contenance makes no mention of our fundamental concern with accuracy, so it is unclear how it is connected with the rationality of our epistemic lives. Even further, there's a strong case to be made that sometimes we aren't rationally required to follow our plans. Taking inspiration from (Schultheis, forthcoming), if you believe that adopting such-and-such plan P_t will cause you to actually adopt credences c_{t+1} (upon actually receiving evidence E) and $P_{t,E} \neq c_{t+1}$, then, if you care about the actual accuracy of your future credences, there are situations in which it is rationally required to adopt plans that you wouldn't want to follow.

2.4.2 The Argument from Interpretation

Sometimes, clever interpretive moves can work wonders, or so says the argument from interpretation. The idea is that if we interpret our credal planning function in just the right way, we can get an argument for Actual Conditionalization. In order to do this, we interpret P as a disposition (or plan): upon receiving new evidence E (our evidential stimulus), you will adopt P_E at time $t + 1$ (or you will follow your plan). This gets us an argument for Actual Conditionalization because if you fail to actually conditionalize, then your plan fails to satisfy Plan Conditionalization and that is irrational. Thus, using our new interpretation, we can argue for Actual Conditionalization from Plan Conditionalization, just like the argument from Diachronic Contenance, but without appealing to any new rational requirements. How elegant!

The problem: this argument only works for a very restricted kind of agent: the kind of agent who has *deterministic* dispositions or plannings (Pettigrew, 2020). I don't know about you, but my plans go awry all the time. There's lots of stuff that I intend to do, but end up failing to do, like my laundry tonight. Nor do I think my credal dispositions so precise as to be fully determined by my confidences and the possible evidential stimulus. The point is that while the argument from interpretation can sometimes get us Actual Conditionalization, it only does so in a spotty way. Much of the time it doesn't work for agents like us, and that's pretty disappointing. The question is: can we do better?

2.4.3 The Evidentialist Expectation Argument

Evidentialism, as it is traditionally conceived, is the view that your evidence directly and primitively constrains your rational doxastic attitudes, in our case, your

credences. Here is an example: if you possess evidence E , then you rationally must be certain, $c(E) = 1$, of E . So, here's an idea: maybe we could give an expected-accuracy argument for Actual Conditionalization that is restricted to evidentially permissible future credences; in other words, it is rationally required that c_{t+1} maximize $EA(.|c_t)$ over the collection of credence functions that are certain of E . But why the restriction of only maximizing over evidentially certain credences? Well, according to Evidentialism, any credences that are not certain of their evidence are irrational anyways, so it seems reasonable to exclude them from consideration when maximizing expected-accuracy. So, does this approach succeed in giving an argument for Actual Conditionalization. Well, it turns out that sometimes it does and sometimes it doesn't, that is, it works for some legitimate accuracy-measures and it doesn't work for other legitimate accuracy-measures (Leitgeb & Pettigrew, 2010) (Pettigrew, 2016)⁶. The exact upshot of this formal result depends on how you conceive of the relevant "maximize (evidentially-constrained) expected-accuracy principle" along lines similar to those discussed by (Pettigrew, 2020) and (Rooyakkers, ms) in the context of accuracy-dominance principles. At best, like the previous argument, it only gets us Actual Conditionalization in a spotty way, and that is fairly disappointing. The only upside of this argument is that it is not parasitic on Plan Conditionalization. This is an advantage because Actual Conditionalization says nothing about credal plannings, so it allegedly applies to agents who don't even have a credal plan at time t . Again, the question is: can we do better?

2.4.4 The Expectation Argument from Value-Change

There are actually two expected-ish accuracy arguments for Actual Conditionalization. The first one is due to joint work by Hannes Leitgeb and Richard Pettigrew in (Leitgeb & Pettigrew, 2010). The idea is that when you possess evidence E , your future credences should maximize an evidentially-truncated expectation with respect to your current credences. This is in contrast to maximizing your expected-accuracy simpliciter. The idea is that your future credences should maximize $\sum_{w \in E} A(., w)c_t(w)$ in contrast to $\sum_{w \in W} A(., w)c_t(w)$. The reason for this is that we don't want to take non- E worlds into account because they are incompatible with your evidence. Now, with this proposed rational norm in hand, (Leitgeb & Pettigrew, 2010) go on to prove the following result:

Leitgeb-Pettigrew Evidentially-Truncated Expectation Theorem: if c_t is a probability function with $c_t(E) > 0$ and you receive total evidence E between times t and $t + 1$, then c_{t+1} maximizes $\sum_{w \in E} A(., w)c_t(w)$ iff $c_{t+1}(.) = c_t(.|E)$.

Thus, we have the following argument for Actual Conditionalization:

- (1): Veritism: accuracy is the only intrinsic epistemic value (for credences).
- (2): Legitimate accuracy-measures are strictly proper and continuous.
- (3): It is a requirement of rationality that (c_t, c_{t+1}) is such that c_{t+1} maximizes

⁶For example, (Leitgeb & Pettigrew, 2010) have explicitly shown that if we run the Evidentialist Expectation Argument with the Brier-score, then we end up with an updating rule that is not Actual Conditionalization. In fact, it can be shown that the only continuous, strictly proper, and additive accuracy-measure that gets us Actual Conditionalization in the Evidentialist Expectation Argument is a measure called the enhanced log-score (Pettigrew, 2020, July 3)

$$\sum_{w \in E} A(., w) c_t(w).$$

(4): **Leitgeb-Pettigrew Evidentially-Truncated Expectation Theorem.**

(5): Therefore, Actual Conditionalization.

In response to this argument, (Gallow, 2019) criticized premise (3). He points out that the thing being maximized in (3) is not a genuine expectation, so it seems to run against the usual advice of traditional decision theory to maximize expected value. And, further, there is no plausible story given about why that quantity should be maximized as opposed to a normal expectation, just leaving out non- E -worlds because they are incompatible with your evidence doesn't tell us why we can epistemically ignore them. While a strong case has recently been made in favor of premise (3) by (Pettigrew, 2023), it is still worth looking at Gallow's story behind why we should epistemically ignore non- E -worlds. (Gallow, 2019, pg. 17-18, original italics) proposes the following answer:

In general, new experiences can rationalize shifts in value...the reason for not valuing accuracy at the non- E possibilities, in spite of the fact that you think the non- E possibilities are likely, is that you have *learned* that those possibilities are not actual. And *that you've learned E* is a sufficient reason to not value accuracy at non- E possibilities, whatever your prior degrees of belief in E happened to be...This highlights an important feature of the present proposal: though it claims that the rationality of your degrees of belief is entirely a matter of whether you are rationally pursuing accuracy—though it denies that there are any evidential norms directly governing credence—it is consistent with there being substantive evidential norms governing the rational *evaluation* of credences.

What Gallow is getting at here is that (evaluative) “legitimacy” is not synonymous with “adequacy conditions on measuring accuracy”. It's one thing to explicate the concept of accuracy. It's another thing to specify permissible ways of evaluating your epistemic performance, even if the only thing of intrinsic epistemic value is accuracy. What Gallow is saying is that the latter should be sensitive to the evidence that you possess. Possessing evidence should change the ways in which you can legitimately evaluate your epistemic performance. If you possess evidence E , it seems, from your point of view, bad to consider non- E worlds relevant to the assessment of your epistemic performance. It is bad to let your evaluation of your epistemic performance be influenced by non-actual worlds *when such an influence can be avoided*, that is, when you possess evidence that some world is non-actual. After all, wouldn't it be nice if we could remove the influence of all the non-actual worlds in evaluating ourselves. With this in mind, and A being a credal accuracy-measure, Gallow vertistically suggests the following way of legitimately evaluating your epistemic performance when you possess total evidence E ⁷. We say that this way of evaluating your epistemic performance is E -respecting.

⁷We note that it is also possible to, instead of taking E to be your total evidence, take E to be your total processed evidence in the sense of (Dallmann, 2017). Under this reading, the following arguments show that even Dallmann's resource-bounded agents must, on pain of irrationality, actually update by conditionalization, albeit on their total processed evidence.

$$A_E(c, w) = \begin{cases} A(c, w) & \text{if } w \in E. \\ k_w & \text{if } w \notin E. \end{cases} \quad (2.1)$$

So, if w is an E -world, we measure the accuracy of any credence function at w in the usual way. But, if w is a non- E -world, then we assign some constant epistemic value k_w to every credence function. Thus, the epistemic value assigned at non- E -worlds does not vary with what credences you have. Now, given this E -respecting way of measuring epistemic value, Gallow proves the following:

Gallow Change-of-Value Theorem: if c_t is a probability function with $c_t(E) > 0$ and you receive total evidence E between times t and $t + 1$, then c_{t+1} maximizes $EA_E(\cdot|c_t)$ iff $c_{t+1}(\cdot) = c_t(\cdot|E)$.

Thus, with this result in hand, we can develop an E -respecting expected-accuracy argument for Actual Conditionalization:

- (1): Veritism: accuracy is the only intrinsic epistemic value (for credences).
- (2): Legitimate accuracy measures are E -respecting, strictly proper, and continuous.
- (3): It is a requirement of rationality that (c_t, c_{t+1}) is such that c_{t+1} maximizes $EA_E(\cdot|c_t)$.
- (4): **Gallow Change-of-Value Theorem.**
- (5): Therefore, Actual Conditionalization.

Both the Leitgeb-Pettigrew and Gallow arguments come with many advantages. They don't rely upon Plan Conditionalization at all because plans are nowhere to be found in their arguments for Actual Conditionalization. So, any worries about introducing some normative or deterministic connection between planning and actuality are moot. They also get Actual Conditionalization all the time, and thus not in a spotty way.

This is all well and good, but it is important to keep in mind that both the Leitgeb-Pettigrew and Gallow arguments are still expectation-style arguments, so any complaints against expectation arguments in general also apply to these arguments. In particular, (Pettigrew, 2016, pg. 200, 208), in the context of discussing the Greaves-Wallace Expected-Accuracy Argument for Plan Conditionalization found in (Greaves & Wallace, 2006), develops such a general complaint against any expectation-style argument:

Now, there is a tension here: I must adopt a new posterior credence function because my current one—namely, c —doesn't respect my evidence. It assigns credence less than 1 to E . So I know that c is defective. If it weren't I wouldn't need to adopt a replacement. Yet we assess the rationality of my possible posterior credence functions by appealing to that very credence function, the one we know to be defective...Indeed, the foregoing considerations seem to suggest that there can be no [expectation] epistemic argument for Diachronic [Actual] Conditionalization at

all. After all, [our argument for] Diachronic [Actual] Conditionalization requires us to set our posterior credences on the basis of our prior credences—it requires us to base our judgements at t' on our judgements at t . But, as we saw above, there is no epistemic reason that compels us to retain at t' any faith in the judgements we made at t .

The question is: can we do better?⁸ It turns out that we can. Accuracy-dominance arguments do not suffer from any of the objections raised against previous attempts to justify Actual Conditionalization. It gets us Actual Conditionalization all the time, and thus not in a spotty way. It doesn't introduce new non-epistemic rational norms. It doesn't require our agent to evaluate their current plans with their current credences. In fact, it doesn't even require our agent to have had credal plannings, let alone credal plannings over an evidential partition. It doesn't require our agent to use their past credences to assess their current credences. With this motivation, we now develop our accuracy-dominance argument for Actual Conditionalization.

2.5 Our Accuracy-Dominance Argument for Actual Conditionalization

Before getting to our new accuracy-dominance argument for Actual Conditionalization, we have to say how we are going to legitimately measure the epistemic value of diachronic credal-pairs (c_t, c_{t+1}) where at time $t + 1$ you possess evidence E :

Def: E -respecting Diachronic Additivity: $A[(c_t, c_{t+1}), w] = A(c_t, w) + A_E(c_{t+1}, w)$.

Given Veritism, this says is that the epistemic value of diachronic credal-pair (c_t, c_{t+1}) is the sum of the epistemic value of c_t and the E -respecting epistemic value of c_{t+1} . The idea here is that once you pick a continuous and strictly proper accuracy-measure A , that is, some evaluative standard, just additively use that standard at every time in an E -respecting way in order to evaluate your overall diachronic epistemic performance. With this in hand, here is our main technical result.

Actual Strong Accuracy-Dominance Theorem: credal-pair (c_t, c_{t+1}) satisfies Actual Conditionalization for evidence E iff (c_t, c_{t+1}) avoids strong accuracy-dominance for every legitimate accuracy-measure on (c_t, c_{t+1}) .

Proof:

“ \Rightarrow ”: We prove the contrapositive. Let A be any legitimate accuracy-measure. Suppose that (c_t, c_{t+1}) is strongly accuracy-dominated with respect to A . Now, observe that (c_t, c_{t+1}) can be extended to a conditionalizing credence/planning pair on an evidential partition including E iff (c_t, c_{t+1}) is a conditionalizing pair on evidence E . Furthermore, if (c_t, c_{t+1}) can be extended to such a conditionalizing credal-planning pair, then extending the dominating pair with those same plans contradicts the **Briggs-Pettigrew Theorem**. Thus, given the observation made above, (c_t, c_{t+1}) is not a conditionalizing pair on evidence E . (We call this method of proof the “extension-method”. Another proof can be given by applying the method found in

⁸Even if you are not convinced by the previous objection to expectation-style arguments, it is still better to have more arguments for a position than less.

(Nielsen, 2021) and using the **Gallow Change-of-Value Theorem**.)

“ \Leftarrow ”: We prove the contrapositive under the assumption that legitimate accuracy-measures are additive. Suppose that credal-pair (c_t, c_{t+1}) does not satisfy Actual Conditionalization for evidence E . Thus, (c_t, c_{t+1}) is strongly dutchbookable (via unconditional and E -conditional bets respectively) by results in (Lewis, 1999). But being strongly dutchbookable in this way is provably equivalent to being strongly accuracy-dominated for every additive legitimate accuracy-measure by results in (Schervish, Seidenfeld, & Kadane, 2009).

◇.⁹

Actual Weak Accuracy-Dominance Theorem: credal-pair (c_t, c_{t+1}) satisfies Actual Conditionalization for evidence E and $c_{t+1}(E) = 1$ (c_{t+1} is Blackwell) iff (c_t, c_{t+1}) avoids weak accuracy-dominance for every legitimate accuracy-measure on (c_t, c_{t+1}) .

Proof:

“ \Rightarrow ”: parody the above proof using the extension-method on the **Briggs-Pettigrew-Nielsen Theorem**.

“ \Leftarrow ”: trivial proof can be given directly.

◇.

With these results in-hand we can finally detail our accuracy-dominance argument for Actual Conditionalization.

- (1): Veritism: accuracy is the only intrinsic epistemic value (for credences).
- (2): Legitimate accuracy-measures are E -respecting Diachronically Additive, strictly proper, and continuous.
- (3): **Undominated Strong (Weak) Dominance**.
- (4): **Actual Strong (Weak) Accuracy-Dominance Theorem**.
- (5): Therefore, Actual Conditionalization (and the Blackwell condition).

2.6 Actual Reverse Conditionalization

Sometimes we forget things. Sometimes we lose our evidence. I don’t remember what I had for dinner five years ago. I can’t recall what it looked like, smelled like, tasted like, or whether its texture was smooth or jagged. But, what I can recall, is the moral of (Titelbaum, 2014): forgetting, as a phenomena, is not inherently irrational. Losing evidence does not, by itself, make me irrational. It’s just something that happens, just like gaining evidence is just something that happens. Yes, perhaps losing evidence is epistemically unfortunate, but irrationality need not accompany misfortune. To be clear, I’m not saying that it’s rationally permissible to forget willy-nilly; I’m saying that we should treat the phenomena of forgetting just as we treat the phenomena of learning: as something subject to rational evaluation. Enough talk, let’s get started.

In the spirit of treating forgetting like learning, we are interested in the follow-

⁹We are going to ignore (Schoenfield, 2017)-type concerns here, and, throughout the rest of the paper. That said, the relevant Schoenfield-corrections are readily apparent.

ing question: are there any rational norms on how your credences should change in response to forgetting exactly evidence E between times t and $t + 1$ (where you gain no evidence between these times), and, can we actually argue for such norms? (Titelbaum, 2014), based off the work of (Levi, 1980), proposes the following candidate norm on purely forgetting scenarios¹⁰:

Def: Actual Reverse Conditionalization (for evidence E): if between times t and $t + 1$, you lose exactly evidence E (and don't gain any evidence), then (c_t, c_{t+1}) is such that if $c_{t+1}(E) > 0$, then $c_t(\cdot) = c_{t+1}(\cdot|E)$.

In words, your actual future credences, conditional on your forgotten evidence, should just be your current credences. Put another way, if you were to relearn E like a good Bayesian conditionalizer, your credences should be the same as when you originally learned E . You should forget in reverse of how you learn. This is all well and good, but why think it a requirement of epistemic rationality? (Titelbaum, 2014) shows that this norm accords well with our intuitive judgements about rationality in a variety of different cases. While there is definitely something to be said in favor of this normative modelling argument for Actual Reverse Conditionalization (Titelbaum, 2014) (Titelbaum, forthcoming), it seems desirable to see if we can give a more traditional argument in its favor. (As far as I know, the only attempts at this have been developed by (Levi, 1987) who gives a composition argument using Actual Conditionalization and (Bradley, 2024) who suggests a dutchbook argument for Actual Reverse Conditionalization.)

2.6.1 Parodying some Arguments

A natural place to start is to see how far we can get by parodying the existing arguments for Actual Conditionalization.

We run into immediate difficulties in trying to parody the arguments from Diachronic Contingence and Interpretation. While it makes sense to be disposed/plan to change your credences upon losing evidence E , we don't have a clear analogue to Plan Conditionalization from which to argue. This is because any such "forgetting function" will not be on a partition; you can only forget pieces of your total evidence and these pieces clearly overlap. Furthermore, even if this difficulty could be resolved, all the complaints against the original arguments apply to these caricatures as well, just as they apply to any parody to the Evidential Expectation Argument and the Expectation Argument from Value-Change. Given these considerations, it seems desirable to develop an accuracy-dominance argument for Actual Reverse Conditionalization. Let's get to it.

¹⁰Some situations are purely learning events; you only gain evidence. Some situations are purely forgetting events: you only lose evidence. And some situations are mixed; you both learn and forget. In such mixed situations, (Titelbaum, 2014) proposes a norm called Generalized Conditionalization.

2.6.2 Our Accuracy-Dominance Argument for Actual Reverse Conditionalization

Without further ado, here is the main technical result that we will need:

Actual Reverse Strong Accuracy-Dominance Theorem: credal-pair (c_t, c_{t+1}) satisfies Actual Reverse Conditionalization for evidence E iff (c_t, c_{t+1}) avoids strong accuracy-dominance for every legitimate accuracy-measure on (c_t, c_{t+1}) .

Proof:

Just permute the time indices and parody the proof in the Actual Conditionalization case.

◇.

Actual Reverse Weak Accuracy-Dominance Theorem: credal-pair (c_t, c_{t+1}) satisfies Actual Reverse Conditionalization for evidence E and $c_t(E) = 1$ (c_t is Blackwell) iff (c_t, c_{t+1}) avoids weak accuracy-dominance for every legitimate accuracy-measure on (c_t, c_{t+1}) .

Proof:

Just permute the time indices and parody the proof in the Actual Conditionalization case.

◇.

With these results in-hand we can finally detail our accuracy-dominance argument for Actual Reverse Conditionalization. It is worth noting that this argument has the exact same non-mathematical premises as our accuracy-dominance argument for Actual Conditionalization. Thus, if you endorse that argument, then you must endorse this one, and vice versa.

- (1): Veritism: accuracy is the only intrinsic epistemic value (for credences).
- (2): Legitimate accuracy-measures are E -respecting Diachronically Additive, strictly proper, and continuous.
- (3): **Undominated Strong (Weak) Dominance.**
- (4): **Actual Reverse Strong (Weak) Accuracy-Dominance Theorem.**
- (5): Therefore, Actual Reverse Conditionalization (and the Blackwell condition).

2.7 Actual Joint Almost Lockean Completeness

In this section, we switch gears a bit. Instead of considering an agent who has numerical degrees of belief, we will be interested in an agent who has all-or-nothing beliefs. They either believe or disbelieve p . They either believe that it's going to rain today or they disbelieve it. Such qualitative doxastic attitudes are more like a light switch; they're "on" or "off". Of course, if we have an agent with both credences and all-or-nothing beliefs, then we might be interested in investigating their metaphysical and normative relationship. See (Weisberg, 2020) for the former and (Foley, 1992) (Lin & Kelly, 2012) (Easwaran, 2015) (Dorst, 2019) (Leitgeb, 2017) (Rothschild, 2021) (Kelly & Lin, 2021) (Mierzewski, 2022) (Rooyakkers, ms) for the latter. But, in this section, we will be solely concerned with all-or-nothing beliefs and their plannings. In particular, we will be interested in whether or not there

are rational norms on how you are required to actually change your all-or-nothing beliefs in response to receiving new evidence as opposed to just how you ought to plan to change your beliefs (Rooyakkers, ms). We begin this inquiry by reviewing how we are going to legitimately measure the accuracy of one's all-or-nothing beliefs B and their plannings β (which is a function from some evidential partition to the collection of belief-sets).

2.7.1 Measuring Qualitative Accuracy

The following legitimacy conditions on measuring the accuracy of one's all-or-nothing beliefs are fairly standard and well-known (Easwaran, 2015) (Dorst, 2019) (Leitgeb, 2017) (Rothschild, 2021) (Rooyakkers, ms). So, a legitimate qualitative accuracy-measure A , which is again a function from triples consisting of a proposition, a belief-set, and a possible world to the set of real numbers, must satisfy:

Def: Extensionality:

$$A(\{p\}, B, w) = \begin{cases} T & \text{if } p \in B \text{ is true at } w. \\ F & \text{if } p \in B \text{ is false at } w. \\ 0 & \text{if } p \notin B. \end{cases} \quad (2.2)$$

Here, real number T represents the value of believing truths while F represents the disvalue of believing falsehoods.

Def: Qualitative Additivity: $A(B, w) = \sum_{p \in B} A(\{p\}, B, w)$.

This condition just says that the total accuracy of one's belief-set is the sum of the accuracy of your individual belief in each proposition.

Def: Variable Conservativeness: $T > 0 > F$ and $|F| > T$.

This condition says that believing truths is strictly good (of positive epistemic value) and believing falsehoods is strictly bad (of negative epistemic value). Furthermore, it also says that the disvalue of believing falsehoods is greater than the value of believing truths.

Def: Belief-plan accuracy: $A(\beta, w) = A(\beta_{E_w}, w)$.

This condition just says that the accuracy of your belief plan at world w is the accuracy of your planned beliefs at world w .

Def: Qualitative (Plan) Additivity: $A[(B, \beta), w] = A(B, w) + A(\beta, w)$.

This condition says that the total accuracy of a belief/belief-planning pair is the sum of the accuracy of your belief-set and the accuracy of your belief-plan. Great, now consider the following norms.

Def: B is said to be Almost Lockean Complete at threshold t iff there exists a

probabilistic credence function c st. for every proposition p [if $p \in B$, then $c(p) \geq t$ and if $p \notin B$, then $c(p) \leq t$].

Def: belief-plan β is said to be Almost Lockean Complete at threshold t iff there exists a probabilistic credence function c st. for every proposition p and for every piece of evidence E if $c(E) > 0$, then [if $p \in \beta_E$, then $c(p|E) \geq t$ and if $p \notin \beta_E$, then $c(p|E) \leq t$].

Def: the pair (B, β) is said to be jointly Almost Lockean Complete at threshold t iff B is Almost Lockean Complete at threshold t with respect to c and β is Almost Lockean Complete at threshold t with respect to the same c .

Once again, why think joint Almost Lockean Completeness a requirement of rationality? (Rooyakkers, ms) proves the following:

Qualitative SSK-BP Theorem: (B, β) is jointly Almost Lockean Complete at threshold t iff there is no legitimate qualitative accuracy-measure with threshold t such that (B, β) is strongly accuracy-dominated.

Now, given this technical result, (Rooyakkers, ms) proceeds to develop an accuracy-dominance argument for joint Almost Lockean Completeness via **Evaluationist Non-Vacuous Dominance** (which is a different dominance principle than the usual **Undominated Dominance**). The idea is that the “legitimacy” of an accuracy-measure is to be understood as “being permissible to evaluate yourself with”. So, **Evaluationist Non-Vacuous Dominance** just says that it is irrational to be able to permissibly evaluate your doxastic performance as bad, that is, as accuracy-dominated, when you can avoid such a bad evaluation by having different doxastic attitudes. Formally,

Evaluationist Non-Vacuous Dominance: doxastic attitude D is irrational if
 (1): There exists a legitimate accuracy-measure A and there exists a D' which strongly accuracy-dominates D according to A . And,
 (2): There exists a D' such that D' is not strongly accuracy-dominated according to any legitimate accuracy-measure A .

At this point, we are going to adopt this standard of dominance reasoning for the rest of the paper. It is also important to note that we can use **Evaluationist Non-Vacuous Dominance** in place of **Undominated Dominance** to develop our accuracy-dominance arguments for Actual Conditionalization and Actual Reverse Conditionalization, so, by adopting **Evaluationist Non-Vacuous Dominance**, we don’t lose any of our results. (In fact, if we do this, we can get the further advantage of dropping Additivity from our list of legitimacy conditions on credal accuracy-measures.)

Now that we have an accuracy-dominance argument for joint Almost Lockean Completeness on all-or-nothing belief/belief-planning pairs, we could, just as in the credal case, argue for its diachronic counterpart in the usual “from-planning-to-actuality” ways: via Diachronic Incontinence or Interpretation. We could also develop qual-

itative versions of the Leitgeb-Pettigrew and Gallow arguments (Rooyakkers, ms). But, once again, the same complaints against these arguments arise. Time for another E -respecting accuracy-dominance argument! In particular, we will be arguing for:

Def: (B_t, B_{t+1}) is said to be actually jointly Almost Lockean Complete with threshold t on evidence E iff there exists a probabilistic c st. B_t is Almost Lockean Complete with respect to c at threshold t and if $c(E) > 0$, then B_{t+1} is E -conditionally Almost Lockean Complete with respect to c at threshold t .

2.7.2 An Accuracy-Dominance Argument for Actual Joint Almost Lockean Completeness

Firstly, in order to develop our accuracy-dominance argument for Actual Joint Lockean Completeness, we have to say how we are going to legitimately measure the accuracy of diachronic belief-pairs. So, let A be a legitimate qualitative accuracy-measure and E the evidence that you possess at $t + 1$. We can then define an E -respecting qualitative accuracy-measure in analogy to the credal case.

$$A_E(B, w) = \begin{cases} A(B, w) & \text{if } w \in E. \\ k_w & \text{if } w \notin E. \end{cases} \quad (2.3)$$

Thus, we can also get the qualitative analogue to E -respecting Diachronic Additivity.

Def: E -respecting Qualitative Diachronic Additivity: $A[(B_t, B_{t+1}), w] = A(B_t, w) + A_E(B_{t+1}, w)$.

The motivation for accepting these conditions, that is, as taking them to be a legitimacy condition, is the same as in the credal case: it seems bad to take non- E worlds as relevant to the evaluation of your qualitative epistemic performance when their influence can be avoided, that is, when you possess evidence E . Now, with these conditions in-hand, here is our main technical result.

Qualitative Actual (Strong) Accuracy-Dominance Theorem: belief-pair (B_t, B_{t+1}) is actually jointly Almost Lockean Complete with threshold t on evidence E iff there is no legitimate accuracy-measure with threshold t such that (B_t, B_{t+1}) is strongly accuracy-dominated.

Proof:

“ \Rightarrow ”: We prove the contrapositive. Suppose that (B_t, B_{t+1}) is strongly accuracy-dominated with respect to some legitimate accuracy-measure. Now, assume for contradiction that belief-pair (B_t, B_{t+1}) is actually jointly Almost Lockean Complete on evidence E . Such a pair can be extended to a belief/belief-planning pair that is jointly Almost Lockean Complete on some evidential partition. So, let’s extend it to a jointly Almost Lockean Complete belief/belief-planning pairing while similarly extending the dominating belief-pair. Thus, the extended dominating pair strongly

accuracy-dominates the extended (B_t, B_{t+1}) . But, this contradicts the SSK-BP Theorem. Thus, (B_t, B_{t+1}) is not actually jointly Almost Lockean Complete.

“ \Leftarrow ”: We prove the contrapositive. Suppose that belief-pair (B_t, B_{t+1}) is not actually jointly Almost Lockean Complete with threshold t on evidence E . Then, we can apply the Farkas-Rothschild Lemma to get that (B_t, B_{t+1}) is jointly strongly dutchbookable at threshold t (in the sense described in (Rooyakkers, ms)). Finally, just apply an E -wise version of the Qualitative SSK Theorem (as found in (Rooyakkers, ms)) to get a legitimate qualitative accuracy-measure such that (B_t, B_{t+1}) is strongly accuracy-dominated with respect to that accuracy-measure.

◇.

Thus, we get the following accuracy-dominance argument for Actual Joint Almost Lockean Completeness with threshold t :

- (1): Veritism: accuracy is the only intrinsic epistemic value (for beliefs).
- (2): Legitimate qualitative accuracy-measures with threshold t are E -respecting Diachronically Fully Additive, Extensional, and Variable Conservative.
- (3): **Evaluationist Non-Vacuous Dominance.**
- (4): **Qualitative Actual (Strong) Accuracy-Dominance Theorem.**
- (5): Therefore, Actual Joint Almost Lockean Completeness with threshold t .

For the sake of brevity, we note, but don't spell out in detail, that similar results can be shown for avoiding actual weak accuracy-dominance via results in (Rooyakkers, ms). And, further still, that the case of evidence loss for all-or-nothing beliefs can be handled in the same way as the credal case, but just using the relevant qualitative theorems and results in this paper.

2.8 Actual Jeffrey Conditionalization

Great, now back to the credal side. In *The Logic of Decision*, (Jeffrey, 1965) pointedly observed that Actual Conditionalization requires certainty in the evidence that you come to possess. This seems to be a pretty strong constraint. What about situations in which you become more confident in something but without going all the way to certainty in that thing. (See Jeffery's famous example of looking at the color of a sweater in a dimly lit room.) With this motivation in mind, Jeffrey considered situations in which your credences change over some “evidential” partition ϵ , with the scare quotes over ‘evidential’ because this shift in credences need not be understood as a response to receiving uncertain evidence about elements in ϵ .

Def: A Jeffrey shift over an “evidential” partition ϵ is a collection of non-negative real numbers $\alpha_\epsilon = \{\alpha_E\}_{E \in \epsilon}$ st. $\sum_{E \in \epsilon} \alpha_E = 1$.

The idea here is that your new confidence in, say, proposition E is α_E , that is, $c_{t+1}(E) = \alpha_E$. For technical reasons, we will restrict our focus to only positive Jeffrey shifts, that is, when $\alpha_E > 0$ for every $E \in \epsilon$. Now, (Jeffrey, 1965) proposed that our prior and posterior should, rationally, be related in the following way in the context of Jeffrey shift α_ϵ .

Def¹¹: Jeffrey-Pair: if $c_t(E) > 0$ for all $E \in \epsilon$, then (c_t, c_{t+1}) is a Jeffrey-pair on Jeffrey shift α_ϵ iff $c_{t+1}(\cdot) = \sum_{E \in \epsilon} \alpha_E c_t(\cdot|E)$.

Now that we have formulated our candidate norm, why should we think it a requirement of epistemic rationality? In order to answer this, we begin by reviewing the unpublished work of Professor Dunn (Dunn, 2017).

2.8.1 Jeffrey-Accuracy and Dunn's Theorem

Let A be a continuous and strictly proper credal accuracy-measure. (Dunn, 2017) proposed that, within a Jeffery context, we should evaluate our credences at $t + 1$ using Jeffrey-accuracy in place of just using A . He defined Jeffrey-accuracy with respect to some Jeffrey shift α_ϵ and credence function c_t such that $c_t(E) > 0$ for every $E \in \epsilon$ as:

Def: Jeffrey-accuracy: $A_J(\cdot, w) = \frac{\alpha_{E_w}}{c_t(E_w)} A(\cdot, w)$.

The idea here is that the Jeffrey shift leads to a reweighing of the importance of accuracy at different elements in the evidential partition that depends upon your current c_t . (Dunn, 2017) then went on to develop an expectation argument for Actual Jeffrey Conditionalization using his newly defined Jeffrey-accuracy-measure.

Dunn Expectation Theorem: if $c_t(E) > 0$ for every $E \in \epsilon$ and c_t is a probability function, then c_{t+1} maximizes $EA_J(\cdot|c_t)$ with respect to Jeffrey shift α_ϵ and any legitimate A iff (c_t, c_{t+1}) is a Jeffrey-pair on α_ϵ .

Now, Dunn's expectation argument is particularly interesting because it is unclear if it is subject to Pettigrew-style concerns. Recall that Pettigrew was concerned that you might not be rationally required to use your past credences to assess (in expectation) your posterior credences because your past credences are known to be defective, defective because you have received evidence that says as such. It's unclear if this is a concern in Jeffrey contexts because such contexts don't involve receiving new evidence, so it's not clear that your prior is epistemically defective in such situations. Perhaps we could interpret a Jeffrey shift as receiving some kind of "uncertain evidence". This would raise Pettigrew-style concerns about Dunn's expectation argument. But, I admit, that I don't really understand what that means. Nevertheless, however this matter turns out, it seems desirable to develop an accuracy-dominance argument for Actual Jeffrey Conditionalization.

2.8.2 An Accuracy-Dominance Argument for Actual Jeffrey Conditionalization

Now, before getting to our accuracy-dominance argument, we have to say how we are going to legitimately measure the accuracy of the diachronic credal-pairs within our Jeffrey context.

¹¹In a Jeffrey-pair (c_t, c_{t+1}) , c_{t+1} is often called the Jeffery conditionalization of c_t . It is well-known that c_{t+1} is a Jeffrey conditionalization of c_t with Jeffrey shift α_ϵ iff both c_{t+1} and c_t are probability functions, $c_{t+1}(E) = \alpha_E$, and $c_{t+1}(\cdot|E) = c_t(\cdot|E)$.

Def: Diachronic Jeffrey Additivity: for (c_t, c_{t+1}) with $c_t(E) > 0$ for every $E \in \epsilon$, $A[(c_t, c_{t+1}), w] = A(c_t, w) + A_J(c_{t+1}, w)$.

The idea here is that in order to calculate your total accuracy in a Jeffrey context (with respect to some α_ϵ) just calculate your accuracy and Jeffrey-accuracy individually and sum them up. It is very important to understand here that what accuracy-measures are legitimate varies with what c_t we are considering. So, it is only bad when some credal-pair (c_t, c_{t+1}) is accuracy-dominated with respect to an accuracy-measure that is legitimate *with respect to that* c_t . Thus, unlike the case of actual conditionalization, where the legitimacy of an accuracy-measure only depends upon the evidence that you possess, the legitimacy of an accuracy-measure in a Jeffrey context depends upon the relevant c_t . At this point, it might seem that this observation undercuts the motivation for even developing an accuracy-dominance argument because, assuming that they are relevant, it raises Pettigrew-style concerns, concerns that we raised against Dunn's expected-Jeffrey-accuracy argument. That is, why is it even permissible to use your "defective" prior credences to assess your posterior credences, as demanded by Jeffrey-accuracy? This is where the importance of the Jeffrey shift comes in for evaluating our posterior credences; perhaps it gives a kind of reweighing of the importance of your credences in such a way that you can permissibly use your prior credences over the evidential partition to assess your epistemic standing in the way demanded by a Jeffrey-accuracy-measure. At the very least, our accuracy-dominance argument is *less* vulnerable to Pettigrew's concerns than Dunn's expectation argument because the latter relies on the *entire* prior credence function while our accuracy-dominance argument relies *only* on your prior credences over the relevant evidential partition. With this in mind, here is our key technical result:

Actual (Strong) Jeffrey-Accuracy-Dominance Theorem: If $c_t(E) > 0$ for all $E \in \epsilon$, then $[(c_t, c_{t+1})$ is a probabilistic Jeffrey-pair on Jeffrey shift α_ϵ iff there does not exist a legitimate accuracy-measure A such that (c_t, c_{t+1}) is strongly accuracy-dominated].

Proof:

" \Rightarrow ": Suppose that (c_t, c_{t+1}) is such that $c_t(E) > 0$ for all $E \in \epsilon$ and is a probabilistic Jeffrey-pair on Jeffrey shift α_ϵ . Now, suppose, for contradiction, that there exists a legitimate accuracy-measure A such that (c_t, c_{t+1}) is strongly accuracy-dominated by some (c'_t, c'_{t+1}) . That is, $A[(c'_t, c'_{t+1}), w] > A[(c_t, c_{t+1}), w]$ for all worlds. Thus, multiplying each inequality by $c_t(w)$ and summing over worlds gets us that $EA[(c'_t, c'_{t+1})|c_t] > EA[(c_t, c_{t+1})|c_t]$. But this contradicts the legitimacy of A and the **Dunn Expectation Theorem**.

" \Leftarrow ": We prove the contrapositive¹². Assume that (c_t, c_{t+1}) is a probabilistic (which can be trivially argued for here) non-Jeffrey pair such that $c_t(E) > 0$ for all $E \in \epsilon$. Such a pair is not Bayes for any legitimate accuracy-measure via the **Dunn Expectation Theorem**. Now, the Complete Class I Theorem implies that (c_t, c_{t+1}) is weakly accuracy-dominated by a randomization over credal-pairs with respect to

¹²This direction of the proof uses definitions and theorems that we state, but don't prove, in the appendix. The approach we take uses the fancy machinery of complete class theorems. See (Lehmann, 1983) for an overview of this approach.

any legitimate accuracy-measure. Now, just pick our legitimate accuracy-measure to be strictly convex. The Strict Convexity Lemma of (Lehmann, 1983, book page 48) then implies that (c_t, c_{t+1}) is strongly accuracy-dominated by a (non-randomized) credal-pair for such an accuracy-measure.

◇.

Thus, we get the following accuracy-dominance argument for Actual Jeffrey Conditionalization:

- (1): Veritism: accuracy is the only intrinsic epistemic value (for credences).
- (2): Legitimate accuracy-measures in a Jeffrey scenario with Jeffrey shift α_ϵ are Diachronically Jeffrey Additive, strictly proper, and continuous.
- (3): **Evaluationist Non-Vacuous Dominance.**
- (4): **Actual (Strong) Jeffrey-Accuracy-Dominance Theorem.**
- (5): Therefore, Actual Jeffrey Conditionalization.

Future work in this area might try to generalize our results to the case of avoiding weak-accuracy dominance in Jeffrey contexts.

2.9 Actual Gallow Conditionalization

Sometimes our credal plans go awry by misfiring. Sometimes we think we received evidence E but we actually received evidence F , and thus we mistakenly change our confidences as if we received E . After all, it seems that we can be rationally uncertain about which evidence we actually possess. This position is called Externalism, in contrast to Internalism (that you rationally must be certain of what evidence you possess). It is an interesting question, then, how externalists should rationally plan to change their confidences in response to new evidence. (Gallow, 2021) proposes the following answer. Let TE say that “evidence E is a part of my total evidence” and UF say that “I’ve updated my beliefs on evidence F ”.

Def: (Externalist) Plan Gallow Conditionalization: if $c(TE), c(UE) > 0$ for every $E \in \epsilon$, then (c, P) satisfies Plan Gallow Conditionalization iff $P_F(p) = \sum_{E \in \epsilon} c(TE|UF)c(p|TE)$. We call such pairs Gallow-pairs.

But why think it a requirement of epistemic rationality? We begin by introducing (Gallow, 2021)’s proposed accuracy-measure for externalist contexts and review his expectation argument for Plan Gallow Conditionalization.

2.9.1 Gallow-Accuracy and the Externalist Expectation Theorem

Let A be a strictly proper, additive, and continuous accuracy-measure. (Gallow, 2021) proposed the following way of measuring the accuracy of your plans (with respect to your credences c such that $c(TE) > 0$ for every $E \in \epsilon$) when you are possibly uncertain of what evidence you have received, and thus which part of your plan you should follow.

Def: (Plan) Gallow-accuracy: $A_G(P, w) = \sum_{F \in \epsilon} c(UF|TE_w)A(P_F, w)$.

The idea here is that the overall accuracy of your plan at a world is the expected accuracy of having that plan under the possibility that you might mistake your actual evidence for some other possible evidence. Now that we have a way to legitimately measure the accuracy of our plans within a Gallow context, (Gallow, 2021) proceeded to prove the following expectation theorem in favor of Plan Gallow Conditionalization.

Externalist Expectation Theorem: If $c(TE), c(UE) > 0$ for every $E \in \epsilon$ and c is a probability function, then $[(c, P)]$ satisfies Plan Gallow Conditionalization iff for every legitimate accuracy-measure, P maximizes the expected Gallow-accuracy with respect to c .]

This result could be used to develop an expectation argument for Plan Gallow Conditionalization in the usual way. While this is great, it would be better to argue for, rather than just assume, that your credences are rationally required to be probabilistic (just as in the Plan Conditionalization case in the (Greaves & Wallace, 2006) paper). Thus we are motivated, once again, to develop an accuracy-dominance argument for Plan Gallow Conditionalization.

2.9.2 An Accuracy-Dominance Argument for Plan Gallow Conditionalization

Before getting to our accuracy-dominance argument, we have to say how we are going to legitimately measure the accuracy of credence/credal-planning pairs (c, P) in a Gallow context.

Def: (Plan) Gallow Additivity: $A[(c, P), w] = A(c, w) + A_G(P, w)$.

Without further ado, here is our main technical result.

Plan (Strong) Gallow-Accuracy-Dominance Theorem: If $c(TE), c(UE) > 0$ for every $E \in \epsilon$ (and further that $c(TE \& UF) > 0$), then $[(c, P)]$ is a probabilistic Gallow-pair iff there is no legitimate accuracy-measure such that (c, P) is strongly accuracy-dominated].

Proof:

“ \Rightarrow ”: Suppose that $c(TE), c(UE) > 0$ for every $E \in \epsilon$. Further, suppose that (c, P) is a Gallow-pair. Now, suppose, for contradiction, that there exists a legitimate accuracy-measure A such that (c, P) is strongly accuracy-dominated by some (c', P') . That is, $A[(c', P'), w] > A[(c, P), w]$ for all worlds. Thus, multiplying each inequality by $c(w)$ and summing over worlds gets us that $EA[(c', P')|c] > EA[(c, P)|c]$. But this contradicts the legitimacy of A and the **Externalist Expectation Theorem**.

“ \Leftarrow ”: Suppose that (c, P) is not a Gallow-pair (but that (c, P) is probabilistic). Such a pair is not Bayes. Thus, via the Complete Class I Theorem, (c, P) is weakly accuracy-dominated by some randomization over credence/credal-planning pairs for any legitimate accuracy-measure. In fact, such a pair is strongly accuracy-dominated by a (non-randomized) credence/credal-planning pair via the Strict Con-

vexity Lemma if we choose our accuracy-measure to be strictly convex.

◇.

Thus, we get the following accuracy-dominance argument for Plan Gallow Conditionalization:

- (1): Veritism: accuracy is the only intrinsic epistemic value (for credences).
- (2): Legitimate accuracy-measures are (Plan) Gallow Additive, additive, strictly proper, and continuous.
- (3): **Evaluationist Non-Vacuous Dominance.**
- (4): **Plan Gallow-Accuracy-Dominance Theorem.**
- (5): Therefore, Plan Gallow Conditionalization.

Future work in this area might try to generalize our results to the case of avoiding weak-accuracy dominance in Gallow contexts. This is all well and good, but how are we rationally required to actually change our confidences in Gallow contexts? We suggest the following.

Def: Actual Gallow Conditionalization: if $c_t(TE), c_t(UE) > 0$ for every $E \in \epsilon$ and you actually update on “evidence” F , then $[(c_t, c_{t+1})]$ satisfies Actual Gallow Conditionalization iff $c_{t+1}(p) = \sum_{E \in \epsilon} c_t(TE|UF)c_t(p|TE)$. We call such pairs actual Gallow-pairs.

We now develop our an actual accuracy-dominance argument for Actual Gallow Conditionalization.

2.9.3 An Accuracy-Dominance Argument for Actual Gallow Conditionalization

Before getting to our actual accuracy-dominance argument, we have to say how we are going to legitimately measure the accuracy of one’s actual posterior credences and their diachronic credal-pairs in a Gallow context. So, given credence function c_t such that $c_t(TE) > 0$ for every $E \in \epsilon$, we define:

Def: (Actual) Gallow-Accuracy: $A_G(c_{t+1}, w) = c_t(UF|TE_w)A(c_{t+1}, w)$.

Def: Diachronic Gallow Additivity: $A[(c_t, c_{t+1}), w] = A(c_t, w) + A_G(c_{t+1}, w)$.

The latter condition just says that we are going to measure the total accuracy of your credal-pairs as the sum of the accuracy/Gallow-accuracy of your respective credences. Once again, just as in the Jeffery case, we see that the legitimacy of a Gallow-accuracy-measure depends upon the relevant prior credences (and the “evidence” F that you actually update on). Without further ado, here is our key technical result.

Actual Gallow-Accuracy-Dominance Theorem: If $c_t(TE), c_t(UE) > 0$ for every $E \in \epsilon$ (and further that $c(UF \& TE) > 0$ for every $F, E \in \epsilon$), then $[(c_t, c_{t+1})]$ is a probabilistic actual Gallow-pair updating on F iff there is no legitimate accuracy-measure such that (c_t, c_{t+1}) is strongly accuracy-dominated].

Proof:

“ \Rightarrow ”: Use extension method with the above theorem.

“ \Leftarrow ”: Suppose that (c_t, c_{t+1}) is not an actual Gallow-pair (but that (c_t, c_{t+1}) is probabilistic). Such a pair is not Bayes. (This follows from results in (Gallow, 2021).) Thus, via the Complete Class I Theorem (as stated precisely in the Appendix), (c_t, c_{t+1}) is weakly accuracy-dominated by some randomization over credal-pairs for any legitimate accuracy-measure. In fact, such a pair is strongly accuracy-dominated by a (non-randomized) credal-pair via the Strict Convexity Lemma if we choose our accuracy-measure to be strictly convex.

◇.

Thus, we get the following accuracy-dominance argument for Actual Gallow Conditionalization:

- (1): Veritism: accuracy is the only intrinsic epistemic value (for credences).
- (2): Legitimate accuracy-measures are Diachronically Gallow Additive, additive, strictly proper, and continuous.
- (3): **Evaluationist Non-Vacuous Dominance.**
- (4): **Actual Gallow-Accuracy-Dominance Theorem.**
- (5): Therefore, Actual Gallow Conditionalization.

Future work in this area might try to develop the case of avoiding weak-accuracy dominance and investigate if we can drop the additivity condition on legitimate accuracy-measures. We also note here that our results have applications to Gallow-style forgetting scenarios in which you actually lose some evidence E but you nevertheless change your credences as if you lost evidence F .

2.10 Conclusion

This paper was all about actual accuracy-dominance. We developed accuracy-dominance arguments for: Actual Conditionalization, Actual Reverse Conditionalization, Actual Joint Almost Lockean Completeness, Actual Jeffrey Conditionalization, and Actual Gallow Conditionalization. Future work in this area might continue to pursue this line of inquiry for other possible actual norms (such as, for example, Adams Conditionalization (Dziurosz-Serafinowicz, 2023)). I expect them to find actual accuracy-dominance all around, as we did. Furthermore, it might be desirable to see to what extent we could drop the various Diachronic Additivity conditions (or argue for them) and still end up with the relevant accuracy-dominance theorems.

2.11 Appendix

2.11.1 Some Useful Theorems

Let A be an accuracy-measure, that is, a function from world-credence pairs to $(-\infty, 0]$.

Def: Bayes: credences c are said to be Bayes iff there exists a probabilistic p such

that c maximizes $EA(.|p)$.

Complete Class I Theorem: if c is not weakly-dominated (even by possible randomizations over credal space) and there are only finitely-many relevant possible worlds, then c is Bayes. (Hoff, 2013).

Strict Convexity Lemma: if A is strictly convex and $r(c)$ is a randomization over credences, then $r(c)$ is strictly accuracy-dominated by some (non-randomized) c' . (Lehmann, 1983).

These theorems proved useful in the proof of our accuracy-dominance theorems for Actual Jeffrey Conditionalization and Actual Gallow Conditionalization. Now, it turns out that these theorems can be strengthened from "there exists an accuracy-measure with a dominating pair" to "all legitimate accuracy-measures have such a dominating pair" via results in (Pettigrew, Campbell-Moore, & Konek, ms), specifically their Theorem D7. While this is an interesting technical observation, I'm unsure of its philosophical import. Our accuracy-dominance arguments for Actual Jeffrey/Gallow Conditionalization appealed to Evaluationism, and, if Evaluationism is good enough in the qualitative case, why is it not good enough in the quantitative case as well?

2.11.2 A Remark on Actual Dutchbooks

In this little section, we assume some familiarity with dutchbooks and the usual credal betting norms.¹³ Now, (van Fraassen, 1989) observed that the dutch strategy argument of (Lewis, 1999) only gets us an argument for Plan Conditionalization and not Actual Conditionalization. In fact, if we apply the usual credal betting norms, we end up with the result that it's irrational to actually change any of your credences, even in response to receiving new evidence. Surely, something has gone wrong with this kind of diachronic dutchbook argument. I propose that what is wrong is the credal betting norm. What I'm suggesting here is an evidence-sensitive credal betting norm. Basically, you should only take permissible E -conditional bets when you actually possess evidence E . Thus, not only are your confidences relevant to which bets are acceptable, your evidence is also relevant. It can be shown that this new evidence-sensitive betting norm can get us a dutchbook argument for Actual Conditionalization. (Just use the extension-method on (Lewis, 1999)'s Dutch Strategy Theorem and a result in (Pettigrew, 2020).)

2.11.3 A Remark on Actual Awareness Growth

In this little section, we assume some familiarity with the phenomena of awareness growth. Now, (Pettigrew, 2022) systematically went through how a bunch of the standard accuracy-dominance/dutchbook arguments apply to the case of awareness growth, and, a common difficulty that he ran into was that there isn't something analogous to a plan defined on an evidential partition. Furthermore, he is also concerned about the soundness of the relevant expectation arguments because he thinks that awareness growth might disrupt the normative requirement of using your prior

¹³If unfamiliar, check out (Pettigrew, 2020) for a pedagogical overview.

to assess (in expectation) your posterior credences in the same way that receiving new evidence disrupts this requirement in Gallow's expectation argument for Actual Conditionalization. Luckily, our accuracy-dominance approach for arguing for Actual Conditionalization doesn't run up against either of these concerns, so it is worth seeing what a similar accuracy-dominance approach says in the case of awareness growth. Answer: In cases of just actual awareness growth, credal-pair (c_t, c_{t+1}) is probabilistic and c_{t+1} is an extension of c_t . (Note: the accuracy of c_t at an extended-world is taken to be the accuracy of c_t at the non-extended part of that world.) That's it.

Now, it might be complained that something has gone wrong here. This norm is just too rigid. But, I'm not so sure. Most, if not almost all, cases of awareness growth *also* come with receiving new evidence. For example, if I see a tornado for the first time, I learn that tornadoes are a thing. I learn that tornadoes exist, say. Cases like this, in which I learn something that I didn't even have an attitude about before and which aren't entailed by any propositions that I had an attitude towards before, allow me to change my credences any-which-way. More technically, if we consider only additive accuracy-measures, no credal-pair (c_t, c_{t+1}) that is probabilistic and Blackwell is even weakly accuracy-dominated. (This follows, via a proof by contradiction, by extending c_t to a probability function on the bigger opinion set and setting the extended confidence in E to 0 and using the additivity of A combined with the **Actual Weak Accuracy-Dominance Theorem**.) So, moral of the story, in cases such as these, avoiding accuracy-dominance places no additional or special rational constraints on actual awareness growth with learning.

2.11.4 A Remark on Actual Guidance-Value

Guidance-value arguments improve on the usual dutchbook arguments by showing how failing to satisfy the relevant norm results in your credences guiding your actions comparatively poorly (Schervish, 1989) (Levinstein, 2017) (Pettigrew, 2020) (Konek, 2022). We remark here that we can develop an actual guidance-value argument for Actual Conditionalization using the same proof strategy in our **Actual Strong Accuracy-Dominance Theorem** and results in (Pettigrew, 2020).

2.12 Works Cited

- Bradley, Darren. (2024). How to lose your memory without losing your money: Shifty epistemology and dutch strategies. *Synthese*, 203(4). <https://doi.org/10.1007/s11229-024-04516-z>.
- Briggs, Ray. & Pettigrew, Richard. (2018). An accuracy-dominance argument for conditionalization. *Noûs*, 54(1), 162–181. <https://doi.org/10.1111/nous.12258>.
- Campbell-Moore, Catrin. & Levinstein, Ben. (2021). Strict propriety is weak. *Analysis*, 81(1), 8–13. <https://doi.org/10.1093/analys/anaa001>.
- Dallmann, Justin. (2017). When obstinacy is a better (cognitive) policy. *Philosophers' Imprint*, 17. <https://doi.org/http://hdl.handle.net/2027/spo.3521354.0017.024>.

Dorst, Kevin. (2019). Lockeans maximize expected accuracy. *Mind*, 128(509), 175–211. <https://doi.org/10.1093/mind/fzx028>.

Dunn, Jeffery. (2017). Jeffery conditionalization and scoring rules. ms.

Dziurosz-Serafinowicz, Patryk. (2023). Expert deference and Adams conditionalization. *Synthese*, 201(5). <https://doi.org/10.1007/s11229-023-04175-6>.

Easwaran, Kenny. (2015). Dr. Truthlove or: How I learned to stop worrying and love bayesian probabilities*. *Noûs*, 50(4), 816–853. <https://doi.org/10.1111/nous.12099>.

Foley, Richard. (1992). The epistemology of beliefs and the epistemology of degrees of belief. *American Philosophical Quarterly*. <https://doi.org/https://joelvelasco.net/teaching/3865/foley%2092%20-%20epistemology%20of%20belief%20and%20degrees.pdf>.

Gallow, Dimitri. (2019). Learning and value change. *Philosophers' Imprint*, 19. <https://doi.org/http://hdl.handle.net/2027/spo.3521354.0019.029>.

Gallow, Dimitri. (2021). Updating for externalists. *Noûs*, 55(3), 487–516. <https://doi.org/10.1111/nous.12307>.

Goldman, Alvin. (1986). *Epistemology and cognition*. Harvard University Press.

Greaves, Hilary. & Wallace, David. (2006). Justifying conditionalization: Conditionalization maximizes expected epistemic utility. *Mind*, 115(459), 607–632. <https://doi.org/10.1093/mind/fzl607>.

Hoff, Peter. (2013). Admissibility and complete classes. Reading. Retrieved from <https://www2.stat.duke.edu/pdhl0/Teaching/581/LectureNotes/admiss.pdf>.

Jeffrey, Richard (1965). *The logic of decision*. University of Chicago Press.

Joyce, Jim. (1998). A nonpragmatic vindication of Probabilism. *Philosophy of Science*, 65(4), 575–603. <https://doi.org/10.1086/392661>.

Kelly, Kevin. & Lin, Hanti. (2021). Beliefs, probabilities, and their coherent correspondence. *Lotteries, Knowledge, and Rational Belief*, 185–222. <https://doi.org/10.1017/9781108379755.011>.

Konek, Jason. (2022). Degrees of incoherence, Dutch bookability, & guidance value. *Philosophical Studies*, 180(2), 395–428. <https://doi.org/10.1007/s11098-022-01860-z>.

Lehmann, Erich. (1983). *Theory of point estimation*. Springer New York.

Leitgeb, Hannes. & Pettigrew, Richard. (2010). An objective justification of Bayesianism II: The consequences of minimizing inaccuracy. *Philosophy of Science*,

77(2), 236–272. <https://doi.org/10.1086/651318>.

Leitgeb, Hannes. (2017). *The stability of belief: How rational belief coheres with probability*. Oxford University Press.

Levi, Isaac. (1980). *The Enterprise of Knowledge*. The MIT Press.

Levi, Isaac. (1987). The demons of decision. *Monist*, 70(2), 193–211. <https://doi.org/10.5840/monist198770215>.

Levinstein, Ben. (2017). A pragmatist’s guide to epistemic utility. *Philosophy of Science*, 84(4), 613–638. <https://doi.org/10.1086/693444>.

Lewis, David. (1999). Why conditionalize? In *Papers in metaphysics and epistemology*. essay, Cambridge University Press.

Lin, Hanti. & Kelly, Kevin. (2012). Propositional reasoning that tracks probabilistic reasoning. *Journal of Philosophical Logic*, 41(6), 957–981. <https://doi.org/10.1007/s10992-012-9237-3>.

Mierzewski, Krzysztof. (2022). Probabilistic stability, AGM revision operators and Maximum Entropy. *The Review of Symbolic Logic*, 15(3), 553–590. <https://doi.org/10.1017/s1755020320000404>.

Nielsen, Michael. (2021). Accuracy-dominance and conditionalization. *Philosophical Studies*, 178(10), 3217–3236. <https://doi.org/10.1007/s11098-020-01598-6>.

Nielsen, Michael. (2022). On the best accuracy arguments for probabilism. *Philosophy of Science*, 89(3), 621–630. <https://doi.org/10.1017/psa.2021.43>.

Oddie, Graham. (1997). Conditionalization, cogency, and cognitive value. *The British Journal for the Philosophy of Science*, 48(4), 533–541. <https://doi.org/10.1093/bjps/48.4.533>.

Paul, Sarah. (2014). Diachronic incontinence is a problem in moral philosophy. *Inquiry*, 57(3), 337–355. <https://doi.org/10.1080/0020174x.2014.894273>.

Pettigrew, Richard. (2016). *Accuracy and the laws of credence*. Oxford University Press.

Pettigrew, Richard. (2020). *Dutch book arguments*. Cambridge University Press.

Pettigrew, Richard. (2020, July 3). Updating by minimizing expected inaccuracy—or: my new favourite scoring rule. M-PHI. September 17, 2024, <https://m-phi.blogspot.com/2020/07/updating-by-minimizing-expected.html>.

Pettigrew, Richard. (2022a). Accuracy-first epistemology without additivity. *Philosophy of Science*, 89(1), 128–151. <https://doi.org/10.1017/psa.2021.17>.

- Pettigrew, Richard. (2022b). How should your beliefs change when your awareness grows? *Episteme*, 1–25. <https://doi.org/10.1017/epi.2022.33>.
- Pettigrew, Richard. (2023, August 3). Reviving and old argument for conditionalization. *M-PHI*. <https://m-phi.blogspot.com/2023/08/reviving-old-argument-for.html>.
- Pettigrew, Richard, Campbell-Moore, Catrin., & Konek, Jason. (2025). Choosing how to choose. ms.
- Rooyakkers, Patrick. (2024). Dutch strategies, accuracy-dominance, lockean beliefs, and their plannings. ms.
- Rothschild, Daniel. (2021). Lockean beliefs, Dutch books, and scoring systems. *Erkenntnis*, 88(5), 1979–1995. <https://doi.org/10.1007/s10670-021-00438-1>.
- Rudin, Walter. (1964). *Principles of Mathematical Analysis*. McGraw-Hill.
- Schervish, Mark (1989). A general method for comparing probability assessors. *The Annals of Statistics*, 17(4). <https://doi.org/10.1214/aos/1176347398>.
- Schervish, Mark, Seidenfeld, Teddy., & Kadane, Jospeh. (2009). Proper scoring rules, dominated forecasts, and coherence. *Decision Analysis*, 6(4), 202–221. <https://doi.org/10.1287/deca.1090.0153>.
- Schoenfield, Miriam. (2017). Conditionalization does not (in general) maximize expected accuracy. *Mind*, 126(504), 1155–1187. <https://doi.org/10.1093/mind/fzw027>.
- Schultheis, Ginger. (forthcoming). Accurate Updating. *Philosophy of Science*.
- Titelbaum, Michael. (2014). *Quitting certainties: A bayesian framework modeling degrees of belief*. Oxford University Press.
- Titelbaum, Michael. (forthcoming). Normative Modeling. In J. Horvath (Ed.), *Methods in analytic philosophy a contemporary reader*. essay, Bloomsbury Academic.
- van Fraassen, Bas. (1989). *Laws and symmetry*. Clarendon.
- Weisberg, Jonathan. (2020). Belief in psyontology. *Philosophers' Imprint*, 20. <https://doi.org/http://hdl.handle.net/2027/spo.3521354.0020.011>.
- Williams, John Robert Gareth. & Pettigrew, Richard. (2023). Consequences of calibration. *The British Journal for the Philosophy of Science*. <https://doi.org/10.1086/725097>.