

Fragmentation and Old Evidence

Will Fleisher

Penultimate version, forthcoming in *Episteme*
<http://dx.doi.org/10.1017/epi.2020.51>

Abstract

Bayesian confirmation theory is our best formal framework for describing inductive reasoning. The *problem of old evidence* is a particularly difficult one for confirmation theory, because it suggests that this framework fails to account for central and important cases of inductive reasoning and scientific inference. I show that we can appeal to the fragmentation of doxastic states to solve this problem for confirmation theory. This fragmentation solution is independently well-motivated because of the success of fragmentation in solving other problems. I also argue that the fragmentation solution is preferable to other solutions to the problem of old evidence. These other solutions are already committed to something like fragmentation, but suffer from difficulties due to their additional commitments. If these arguments are successful, Bayesian confirmation theory is saved from the problem of old evidence, and the argument for fragmentation is bolstered by its ability to solve yet another problem.

Bayesian confirmation theory (BCT) is our best formal framework for both describing and evaluating inductive reasoning.¹ One of the most pressing problems for BCT is the *problem of old evidence* (POE). This problem, at its base, is that the standard way of using Bayesian modeling entails that something one already knows cannot serve as evidence for any theory, new or old. But there are many instances in the history of science, and even in everyday inquiry, of thinkers considering how old evidence (things they already know) support, confirm, or provide evidence for a theory. Since BCT purports to be a comprehensive and unified model of inductive reasoning and confirmation, this is a significant problem. The problem is particularly unsettling because this is precisely what Bayesianism was supposed to do best.²

In this paper, I appeal to the growing literature on *fragmentation* to provide a solution to the problem of old evidence for Bayesianism. My purpose here is twofold. First, I want to show that fragmentation provides the best solution to the problem of old evidence for BCT. Second, I want to show that fragmentation has yet another application; one which is suggestive about a further set of

¹For more on BCT and its standing in contemporary philosophy of science, see Strevens (2006) and Earman (1992).

²For more on this thought, see Glymour (1980) and Garber (1983).

applications involving limited information access by design. This provides additional evidence for the value of the fragmentation framework for philosophy of mind and epistemology.

The first section of this paper provides a quick overview of Bayesian confirmation theory and the problem of old evidence. The second section explains fragmentation, and describes its independent support. Then, in section three, I provide the fragmentation solution to the problem of old evidence. I first suggest a general fragmentation strategy for solving the problem, before offering two ways this strategy can be concretely implemented.

Readers well-versed in the literature on the POE will immediately note the similarities between the fragmentation solution and several other proposed solutions to the problem, including Garber or GJN-style solutions (Garber 1983) and the counterfactual solution (Howson 1991; Howson & Urbach 2006). In the final section of this paper, I argue that these other candidate solutions to the POE are in fact already committed to some form of fragmentation (or something very much like it). This commitment has gone largely unnoticed in the literature. However, the additional features of each solution beyond fragmentation cause them unnecessary difficulties. The fragmentation solution does not suffer from these problems and is embedded in an independently successful research program. I will thus argue that fragmentation by itself is the best strategy for solving the problem.

1 Bayesian Confirmation and the Problem of Old Evidence

1.1 Bayesian Confirmation Theory

Bayesianism comes in many varieties, as there are many ways to add to the basic commitments of the Bayesian framework. Here, I will focus on a simple version of BCT which is characterized by three basic commitments.³ The first is *personalism*: confirmation, inductive reasoning, and evidence should be explained by appeal to the doxastic attitudes that inquirers have (and should have). What the formalism seeks to model is the rational update of a scientist's credences or degrees of belief as they learn new things. Credences are a kind of belief state, a mental attitude of the subject toward some proposition. They correspond to the subject's felt level of confidence in a proposition, and explain the subject's betting behavior. Thus, the norms of BCT are norms of rational belief update.

The second standard commitment of Bayesianism is known as *probabilism*. It requires that the credences of a rational subject be modeled (or represented) by a probability function. For the purposes of this paper, we can understand

³These basic commitments are adapted from textbook treatments in Earman (1992), Howson and Urbach (2006), Strevens (2006), and Titelbaum (2020). They are meant to reflect the traditional commitments of the view. For additional detailed introduction to Bayesian epistemology and BCT, see Talbott (2015), Crupi (2015).

the objects of credences as propositions. The probability function describing the credences is then defined over an algebra \mathcal{A}_Ω composed of propositions. Ω is the state space (or sample space) which is composed of a particular set of propositions, each of which describe fully determinate states of affairs in the world, i.e., possible worlds. The rest of the propositions in \mathcal{A}_Ω are obtained by taking the deductive closure of Ω using the standard five logical connectives. The probability function $Pr(\cdot)$ which represents the subject's credences assigns a real number to each proposition $p \in \mathcal{A}_\Omega$.⁴

In order for $Pr(\cdot)$ to count as a genuine probability function, it must satisfy Kolmogorov's axioms:

Non-negativity For any $p \in \mathcal{A}_\Omega$, $Pr(p) \geq 0$.

Normality For any logical truth $\top \in \mathcal{A}_\Omega$, $Pr(\top) = 1$.

Finite Additivity For any mutually exclusive $p, q \in \mathcal{A}_\Omega$, $Pr(p \vee q) = Pr(p) + Pr(q)$.

Bayesians also think that subjects have conditional credences: attitudes about how probable one proposition is, given that another proposition is true. These conditional credences are also modeled using a probability function, one connected to the unconditional probability function by the ratio definition:

Conditional Probability For any $p, q \in \mathcal{A}_\Omega$, if $Pr(q) > 0$, $Pr(p|q) = \frac{Pr(p \& q)}{Pr(q)}$.

The third standard Bayesian commitment is that learning should be modeled as *update by conditionalization*. Conditionalization requires that once the subject learns some proposition q , her new credence in p should be equal to her old credence in p , given q ($Pr(p|q)$). Importantly for our purposes, for any q with $Pr(q) > 0$, $Pr(q|q) = 1$. So, when a subject learns q her new credence in q should be $Pr(q) = 1$. It is this feature of BCT which led Glymour to recognize the Problem of Old Evidence (1980).⁵

This basic BCT framework provides a variety of benefits. Importantly, BCT is conducive to a simple and intuitive definition of the *evidential support* relation in terms of probabilistic relevance. A is evidence for B (for a subject) just when $Pr(B|A) > Pr(B)$. Evidential support is explained in terms of probability raising. This use of conditional probability combined with update by conditionalization as a way of modeling confirmation has led to the great successes of BCT.

BCT has both a normative and descriptive component. One reason to think BCT provides the right normative framework for understanding confirmation

⁴In setting things up this way I am following Joyce (1999). This is essentially equivalent to defining the probability function over a σ -algebra built from a sample space of outcomes or events (Hájek 2019; Strevens 2006). For some caveats in moving between talking about outcomes and propositions, see (Howson & Urbach 2006, pgs. 13–15).

⁵It is not essential to the Problem of Old evidence that evidence propositions receive credence 1 when learned. The POE arises even in frameworks that use Jeffrey conditioning and refuse to assign probability 1 to any contingent proposition. See Barnes 1999, sec. 6.2 for discussion of the "quantitative" POE. For Jeffrey Conditioning see Jeffrey (1990). However, I will largely focus on the strict conditionalization framework, and the credence 1 version of the problem, for ease of exposition. The fragmentation solution works just as well in both frameworks.

and evidence is that it seems to accurately describe cases of rational theory change in science. BCT justifies what are taken to be “sound methodological procedures” (Earman 1992, 63): it justifies the idea that a generalization is confirmed by its instances, that diverse evidence is more confirmatory than uniform evidence, and that novel prediction is more confirmatory than explanation after the fact. These are, plausibly, features of actual scientific practice (at least when things go well). Moreover, given suitable idealizations, BCT also does a good job describing historical episodes in which scientists changed their opinions about theories in light of confirmatory evidence, e.g., theory change in physics about electroweak interaction.⁶ Given that these historical instances involve excellent scientists who are (presumably) acting in epistemically appropriate ways, the fact that BCT describes them properly is reason to think that BCT is an appropriate account of (some) epistemic norms. Thus, BCT’s descriptive success in epistemically good cases provides support for the view as an account of epistemic norms.

1.2 The Problem of Old Evidence

Update by conditionalization is at the core of BCT, and while it is responsible for much of its success, it also leads directly to the problem of old evidence. The POE arises because the evidential import of some piece of evidence is fully accounted for by the update procedure. After the evidence is learned, and after the update occurs, the proposition learned will no longer satisfy the probabilistic definition of evidential support.

The canonical example used in the literature to describe the problem is the evidential support provided for the General Theory of Relativity (*GTR*) by the orbit of Mercury (*OM*) (Glymour 1980). Mercury’s orbit was long known to be different than Newtonian mechanics predicted, in particular due to the precession of its perihelion (Earman 1992, 119). *GTR* successfully explains this well-known anomaly regarding Mercury’s Orbit. Earman points out that, although general relativity has successfully made a number of further true predictions which confirmed it (e.g., gravitational lensing and gravitational waves), apparently most physicists came to accept it before this evidence was available (1992, 119). Also, historical surveys of scientists suggest they took *GTR*’s successful account of *OM* to be the most important factor in its confirmation (ibid).

Given the facts of this case, it is clear that the orbit of mercury is evidence for *GTR*. However, BCT does not deliver this result. Einstein already knew about the orbit of Mercury. So at the time he formulated *GTR*, his $Pr(OM) = 1$, and therefore $Pr(GTR|OM) = Pr(GTR)$. Thus, by the probabilistic relevance definition of evidence, BCT falsely suggests that *OM* is not evidence for *GTR*.

Following Eells (1985) and Sprenger (2015), we can distinguish between two different versions of POE that might arise for a theory *T* and evidence *E*:

Static Version The subject learned *E* previously, so their $Pr(E) = 1$, and so

⁶For support of the historical accuracy claim, the electroweak interaction case study, and other relevant examples, see Franklin (1990).

$Pr(T|E) = Pr(T)$. But it is now sometime later and they wish to evaluate E 's support for T .

Dynamic Version: E was known before the formulation of T , and it is now the time of the formulation of T (or barely after). Now the subject wishes to evaluate E 's support for T .⁷

Einstein and his contemporaries faced the dynamic version of the problem. Einstein knew about the orbit of mercury before formulating the general theory of relativity. An example of the static version can be generated for anyone who already knows about both GTR and OM , but subsequently wonders whether OM really provides good evidential support for GTR .

Thus, there are two central types of cases concerning evidential support which BCT gets wrong. In these cases, BCT implies that there is no evidential relation between two propositions when there clearly is. These two problems are significant for Bayesianism. Not only do we have the problem that something that once was evidence no longer is (static version), there are cases where things which intuitively should count as evidence never do (dynamic version). That is, according to BCT, if E is some proposition I already know, it cannot ever again be evidence for anything. This looks bad. Bayesianism needs an interpretation that allows it to account for evidential support relations that exist atemporally, while still being able to model learning over time. The simple, standard version of BCT cannot do both of these things at the same time.

What is at stake here is whether BCT is an appropriate model for describing and evaluating inductive reasoning. The problem is that we want to use this simple mathematical modeling framework to model actual subjects' evaluations of evidential support, as well as their rational doxastic evolutions, and to derive norms that govern these things.

What has gone wrong, I want to suggest, is not that we are using the wrong mathematical tools. Instead, it's that the idealizing assumptions being made in order to apply these tools are incorrect. If we are going to use BCT to model creatures like us, we need to loosen or change some of the assumptions so that BCT can offer the right evaluative verdicts in cases of old evidence. The solution is to stop thinking that the model should be applied to a subject globally. Instead, BCT should include fragmentation: the model should be applied to individual compartments of a subject's doxastic state. Once we do this, the model will apply without making false verdicts and bad prescriptions.

2 Fragmentation

Fragmentation is the idea that we should not model subjects as having one unified, global belief state. Subjects have access to different information at different times. Some information is accessible for one task, but not accessible

⁷I have followed Sprenger's formulation of this distinction. Eells makes a few other distinctions that are useful and important, but do not concern us here.

for another task. We can represent this by treating the subject as if they are compartmentalized, or fragmented. That is, distinct parts of the subject's mind contain different information, and are operative for different purposes. In this section, I will argue that fragmentation is independently motivated, and then detail how fragmentation accounts function.

2.1 Motivations for Fragmentation

Fragmentation is intuitively plausible for creatures like ourselves, who have limited computational abilities and imperfectly reliable learning methods. It also does a good job of explaining a variety of difficult cases. I will touch briefly on a few of the main motivations for fragmentation. My purpose here is to illustrate the independence of these motivations, as this is part of what makes fragmentation a compelling solution to the POE. It is not an ad hoc solution, but rather arises organically from a well-motivated and successful research program.

The notion of fragmentation first arose in the context of solving logical omniscience, consistency, and closure problems for possible worlds semantics. Lewis (1982) and Stalnaker (1984, 1999) both recognized these problems for a semantics where the meaning of a proposition is a set of possible worlds. Such a view has two features which lead to these problems: first, if a subject's information or doxastic state is represented by a set of possible worlds, this state must be both consistent and closed under deduction. Second, there is only one necessary proposition: the set of all possible worlds.⁸

As noted above, many Bayesian epistemologists define the probability function as I have here, in terms of an algebra of propositions built from a state space Ω which is comprised of possible worlds. This commits us to some of the same formalism as possible world semantics, and results in the same problems. However, because of the requirements of the probability axioms, any Bayesian epistemology has similar commitments. Probabilism requires that the subject have credence 1 ($Pr(\top) = 1$) in any logical truth, and that she have equal credence in any logically equivalent propositions. This leads to Bayesianism having the same logical omniscience, consistency, and closure requirements as possible worlds semantics.

There are a wide variety of extremely common cases which cause problems for these deductive requirements of omniscience, consistency and closure. I will just provide a few examples. Consider, for instance, Lewis's understanding of Princeton geography (1982):

I speak from experience as the repository of a mildly inconsistent corpus. I used to think that Nassau Street ran roughly east-west; that the railroad nearby ran roughly north-south; and that the two were roughly parallel... (436).

This, clearly, is an inconsistent triple. But Lewis's behavior, let's suppose, was not erratic or unpredictable. For some purposes, he behaved as if the street

⁸For more on possible worlds semantics, see Stalnaker (1984), Menzel (2016), King (2014).

and the railroad were parallel (perhaps while drawing maps, or while giving directions) while for others he believed them perpendicular (perhaps when navigating himself to Nassau street from the train station).

Lewis appealed to fragmentation to explain his own case: his inconsistent beliefs were kept quarantined, as parts of different psychological compartments or fragments which were active for different purposes. His behavior in different circumstances is rationalized by different fragments. This does a better job of explaining the systematic character of the mistake, rather than simply pointing out that it is a mistake.

There are also failures of closure that are well explained by fragmentation. Suppose Julie has just finished painting a square room, and now needs to lay carpet.⁹ She knows the room is square, and she knows how long the sides are: 9 feet. Moreover, she has a college education and knows that $9 \times 9 = 81$, and if you asked her how to calculate the area of a square she could tell you. However, when she is at the store buying carpet she realizes she doesn't know how much to buy.

Julie's case is not uncommon, and is certainly not impossible. But it is a clear failure of closure. The fragmentation framework has a simple way of explaining what is going on here: Julie has multiple doxastic fragments, and the fragment with the information about how to calculate the area of a square surface is not being accessed when she is in the store buying carpet. This explanation can be expanded to account for both failures of logical omniscience, and to explain mathematical achievement (Rayo 2013).

On top of the motivations for fragmentation that arise from these technical problems for possible worlds semantics and Bayesianism, there are a variety of other cases that are well explained by the framework. Some of the most compelling of these involve cases which seem endemic to the human condition: recall failures. Suppose I get asked the following question: (1) "What is a country in the world whose name contains three 'u's?" I might be stumped by this question, and not be able to answer. But suppose someone asks me: (2) "Is 'Uruguay' the name of a country that contains three 'u's?", to which I readily assent. How should we model my doxastic state? There is a sense in which I don't know that "'Uruguay' is the name of a country with three 'u's, because I couldn't answer question (1). But there is a clear sense in which I do know this proposition; the second phrasing of this question, in (2), elicits my knowledge.

Such examples are ubiquitous.¹⁰ We often find ourselves unable to recall information which we know. And it turns out this is hard to explain if we are using a global model of a subject's doxastic states. Do I know that "Uruguay is the name of a country with three 'u's"? Which proposition gets placed in a representation of my belief state? What credence do I have in this claim? There is no obvious way to model my doxastic state here using a global set of beliefs or global credence function.

Fragmentation offers an easy solution for this as well. We should model

⁹This example is like the farmer example from Rayo (2013).

¹⁰Cf. Egan (2008b), Elga and Rayo (2019), Rayo (2013) for more examples along these lines.

subjects as having different belief states under different “elicitation conditions” or tasks (Egan 2008b, Elga and Rayo 2019). When prompted with the first question, I don’t have access to the information about Uruguay. But when prompted with the second question, I do. This can be modeled by having two sets of beliefs, and/or two credence functions that govern my behavior under the two different elicitation conditions. When undertaking the task of answering the first question, I don’t have access to the information about Uruguay, but when answering the second question a different fragment is operative and I do have access.

There are many other kinds of examples which admit of the same fragmentation treatment. Fragmentation can explain how implicit bias functions: one fragment contains beliefs with the bias, while another explains the explicit commitments which contradict this (Elga and Rayo 2019). Fragmentation also provides an account of cases where “knowing that” and “knowing how” come apart, and fail to interact. For example, it’s well-known that one can be an excellent athlete, but make false claims about how one goes about exercising one’s athletic competences. The fragmentationist can account for this by proposing one fragment for the “know how” the subject displays, and another fragment which governs the subject’s (perhaps false) assertions regarding their own performances. Moreover, these examples fail to exhaust the framework’s usefulness.¹¹

In addition to providing an account of a variety of cases, fragmentation is simply intuitively plausible. It is a truism that we are limited creatures, with limited computational abilities. We don’t have the capability of holding all our information in our working memory at a single time.¹² Given these limitations, ideals of rationality that appeal to global doxastic states seem overly idealized, at least for many purposes. Moreover, models of our cognition which appeal to global doxastic states seem unlikely to be accurate. Fragmentation offers a way of providing more achievable normative goals, and of modeling subjects more realistically.

It’s important to note, however, that fragmentation is not motivated merely by the fact of computational limitations. There can be some cases in which even ideally rational agents will be better off being fragmented. Egan (2008a) provides examples of this kind of beneficial fragmentation. He suggests subjects have fallible “Spinozan” methods of belief-formation. A Spinozan method’s operation directly results in a belief, without that belief being subjected to evaluation before being accepted (Spinozan methods are contrasted with “Cartesian” methods, which do involve an evaluation stage prior to belief formation). Perceptual beliefs are plausibly formed by Spinozan methods: when I see a cup, I immediately form a belief that there is a cup. There is no intermediate evaluation stage between the deliverance of my perceptual system and my formation of a belief. In cases where Spinozan methods get things wrong, it

¹¹For more applications of the fragmentation framework, see Egan (2008a), Rayo (2013), Elga and Rayo (2019), Greco (2014), Yalcin (2018), and Johnson (2020).

¹²See Cherniak (1983) and Weisberg (2016) for support of this claim, though it hardly seems necessary to provide empirical support for such an obvious fact.

is better for an agent to be fragmented, so the errors don't seep into her entire doxastic state. Instead, they are quarantined within a few fragments, which can be corrected later.¹³

The fragmentation solution to POE will provide another instance in which even ideally rational agents will be better off if they are fragmented. The solution relies on fragmentation for the evaluation (and re-evaluation) of theories in the light of previously known evidence. This fragmentation is one of choice or design, rather than simply a way of dealing with computational limitations. And this will point the way to additional cases of fortunate and intentional fragmentation, as I will suggest in the concluding remarks.

2.2 How Fragmentation Accounts Work

Fragmentation is a framework in which we represent agents as being compartmentalized: their belief state is represented using more than one probability function.¹⁴

In order to apply the fragmentation framework to a Bayesian account, I will use indexed probability functions.¹⁵ Each subject has multiple fragments, each of which is represented by a probability function. Each probability function is indexed to indicate which fragment it is associated with. So instead of representing the subject's doxastic state with a single, global probability function, we instead represent the doxastic state with multiple probability functions, each of which represents an individual fragment's associated degrees of belief.

I will follow Rayo (2013) and index in terms of tasks.¹⁶ For our purposes, these will be tasks of evaluating individual theories. A fragment is represented by a triple containing a task, a probability function, and a set of available background knowledge: $\langle t, Pr_t(\cdot), K \rangle$. The t represents the task the fragment governs. The probability function will be indexed to this task: $Pr_t(\cdot)$. Thus, $Pr_t(P) = .5$ represents that the subject is .5 confident in P for the purpose of performing task t (where P is a proposition and an element of an algebra \mathcal{A}_Ω). Just as we saw above for a standard BCT model, each fragment's $Pr_t(\cdot)$ is defined over \mathcal{A}_Ω . For reasons that will become clear later, we will assume that this algebra, and its associated state space Ω , are the same for each fragment of an individual subject.

K is an element of \mathcal{A}_Ω that represents the accessible background knowledge of the subject, accessible from the relevant fragment.¹⁷ The use of K to

¹³For details, see Egan (2008a).

¹⁴The framework is compatible with all of the leading views about the nature of mental states, e.g., representationalism, functionalism, dispositionalism, and radical interpretation. See Schwitzgebel (2015) for a summary of these various accounts. The only view it is clearly incompatible with is eliminativism.

¹⁵Here, I am again following Egan (2008a); Rayo (2013) and Elga and Rayo (2019).

¹⁶The fragmentation strategy is compatible with a variety of ways to index. See Egan (2008a), Yalcin (2018), and ?.

¹⁷This is a slight departure from Rayo's (2013) use of this kind of notation, but the differences are unimportant here.

represent a subject's background evidence is standard in a Bayesian framework. A subject's probability function at a time is always conditional on their background evidence at that time: $Pr(P|K)$. The only difference here is that each of a subject's fragments can have a distinct set of background evidence K . In most Bayesian literature, this K is assumed and then suppressed, so that it is understood that a probability function written as $Pr(P)$ is equivalent to $Pr(P|K)$ (Earman 1992). The composition of K , however, is important for the fragmentation solution to POE, so I will continue to make this explicit. What counts as accessible background knowledge for a fragment, and so what will be in K , will depend on the fragment and the subject in question. I will return to the characterization of K , and how it should be constrained, in section 3.2 below.

This kind of framework can easily represent what is going on in the motivating cases from the previous section. For instance, we can account for the "Uruguay" recall case by appeal to two different indexed probability functions. When I am tasked with answering question (1), the version of the question that doesn't explicitly mention Uruguay, I am well represented by a probability function where I assign low credence to the claim (U) that Uruguay is a country whose name has three 'u's, so e.g. $Pr_1(U) = .01$. This is because I will behave, in answering the question, as though I have very low confidence that Uruguay is the right answer (specifically, I will not provide it as an answer). However, when performing the task of answering question (2), where Uruguay is explicitly named, I am well represented by a probability function that assigns very high probability to U , e.g. $Pr_2(U) > .99$. This explains the systematic way in which my recall failure occurs regarding Uruguay and the spelling of "Uruguay": different fragments are active for different tasks, and these fragments have access to different information, which is represented both by the prior probability function of the fragment, and the membership of K .¹⁸

With this fragmentation framework on the table, I will turn to showing how it can solve the problem of old evidence.

3 Fragmentation and the Problem of Old Evidence

3.1 The Fragmentation Strategy

I will first talk about a general *strategy* for using fragmentation to solve the problem of old evidence, before proposing particular concrete versions of this strategy which count as fully-fleshed out *solutions*. I proceed this way because of the variety of distinct versions of Bayesian confirmation theory that have been proposed. I think the fragmentation strategy should work to solve the

¹⁸Of course, I won't be in a position to report my low probability in U for task 1, which makes the view incompatible with a certain kind of accessibilist internalism about a subject's own credence function. But this actually makes sense, given the commitments of the framework: if I am asked to report my credence in U , that is a different task. And plausibly, if you mention U , I will no longer have a low credence in it, for the same reasons that I have a high credence in U when asked question 2.

problem of old evidence on any version of BCT. Which version of this strategy (i.e., which solution) one prefers will depend on one's background commitments. Thus, I think it is useful to discuss why the strategy works, before turning to some concrete implementations of the strategy.

The fragmentation strategy is to model subjects as engaged in the task of evaluating a theory by appeal to a fragment dedicated to this task. The basic idea is that we treat each evaluation (or re-evaluation) of a theory in light of available evidence as its own task. There will be a probability function associated with each such task, and a set of relevant background knowledge. Whenever a subject engages in such an evaluation, this involves a separate fragment. In this separate fragment, the previously known evidence need not be represented as having probability 1. I will first present this schematically, then apply it to the canonical example of general relativity.

Suppose a subject is deliberating in order to evaluate a theory T . She already knows E . However, she wants to consider E 's import for T . So, when turning to the task of evaluating T , she has (or constructs) a fragment which we represent with $\langle t, Pr_t, K \rangle$. Since she wants to evaluate E 's import for T , she does not include E in K . This is despite the fact that in other fragments (e.g., for the task of answering questions about scientific knowledge, call it "ask"), her credence is $Pr_{ask}(E) = 1$. In the new fragment, $Pr_t(E) < 1$. But there is still a sense in which the subject continues to possess the evidence: because she maintains credence 1 in the evidence in *ask*. The fragmentation framework has opened up the space to represent her evaluation of the new theory as a new fragment, one which does not assign credence 1 to the evidence.

This means that we can represent the subject as having E be evidence for T . We can do this in the traditional Bayesian way, using the standard definition of evidential support. So if E really is evidence for T (for this subject, given her priors and accessible background knowledge), then $Pr_t(T|E\&K) > Pr_t(T|K)$. This leaves BCT to function in its usual way. The standard Bayesian framework is preserved, but relativized to individual fragments in a subject's mental economy.

In any fragment that involves evaluating the import of old evidence on a theory, K will contain all of the appropriately accessible information, but will not contain the evidence being evaluated. This is what preserves the evidential relationships. In section 3.2, I will discuss how the subject's background knowledge K should be constrained in different fragments. The basic idea, however, is that K should contain all the background knowledge which is genuinely relevant for the task, except for the E to be evaluated.

In addition to determining whether some piece of old evidence counts as evidence for a theory, we may also be interested in modeling the subject's considered judgment about the theory after considering this old evidence. This can be modeled in the standard way, as the subject's posterior credence (but now limited to within a fragment). After the subject considers the evidence to be evaluated, E , it can be treated as though it was learned. That is, the fragment is updated by conditionalization. Thus, the posterior probability the subject has (within the fragment) will be given by $Pr_t(T|E\&K)$. This represents

the results of the subject's deliberative evaluation of E 's import for T .

Note that this K does not represent a set of background information that is limited only by the computational limitations of human subjects. Instead, it is being intentionally limited, for the purposes of evaluating a theory. This is another example (like Egan's 2008b) of fragmentation that is good even for ideally rational agents. This intentional limitation can be thought of as a necessary, rational part of giving a fair evaluation (or re-evaluation) of the theory. K is limited to what is appropriately accessible, not simply to what is accessible given the abilities of the subject. Contrast this with the imperfect recall cases from section 2.

This fragmentation strategy can explain both types of old evidence problem cases identified in section 1.2. As we will see, this gives it an immediate advantage over other proposed solutions. The strategy can be illustrated by variations on the canonical case involving Einstein and the general theory of relativity.

First, I will consider the dynamic version, the historical version that reflects Einstein's own situation. This version of the problem is where the evidence was known before the formulation of the theory, but it is now the time of the formulation of the theory (or barely thereafter). Recall that Einstein knew about the anomaly presented by the orbit of Mercury (OM) prior to Einstein's formulation of the General Theory of Relativity (GTR). Nonetheless, Einstein, his contemporaries, and subsequent physicists have taken GTR 's explanation of the orbit of Mercury to provide significant evidence for relativity.

The fragmentation framework explains this case as follows. For each individual who evaluates GTR in light of OM , and OM 's evidential import for GTR , we can represent this task of evaluation by a separate fragment of their doxastic state. Consider Einstein himself: when he first formulated and understood GTR , he already knew OM . However, when undertaking the task of evaluating GTR in the light of the evidence, we represent him as opening a new fragment of his doxastic state, indexed to this task. Let's call the task "evaluating relativity," or er .

The relevant fragment of Einstein's doxastic state is then represented by $\langle er, Pr_{er}, K \rangle$. K represents the background knowledge that is both relevant to the task and computationally accessible to Einstein on the occasion(s) of evaluation. Selecting the right K depends on substantive epistemic and psychological considerations. This selection is not easily operationalized, but is nonetheless of a kind which scientists and philosophers make routinely (more on this below).

Given what we know of the case, when K is appropriately selected this should provide us with a fragment in which Einstein's $Pr_{er}(OM|K) < 1$, and his $Pr_{er}(GTR|OM\&K) > Pr_{er}(GTR|K)$. Thus according to BCT's standard definition, the orbit of Mercury was evidence for the General Theory of Relativity. Meanwhile, in other fragments, such as the one indexed to the task of teaching astrophysics, Einstein maintains high credence in OM , and thereby continues to possess the evidence.

Thus, the fragmentation strategy allows BCT to make the right predictions

in the general relativity case, and for the dynamic version of the POE more generally.

The fragmentation strategy is essentially the same for the static version of the problem of old evidence. In this case, we again simply represent the subject confronted by the problem as having a new fragment, one indexed to the task of evaluating the theory at that time. As an example, consider a contemporary physicist. She already knows about *OM* and *GTR*. If she stops to (re-)consider what kind of evidence *OM* provides for *GTR*, we can represent this consideration as a task with its own associated fragment of her doxastic state. This fragment represents the task of evaluating relativity for her in 2020, so we might call it *er2020*. This fragment will have an associated probability function, $Pr_{er2020}(\cdot)$. From there, the example can be handled precisely the way we handled the above case for the dynamic version.

So, the fragmentation version of BCT provides the right predictions in both versions of the problem of old evidence. To recap, the solution to the POE is to model the evaluation of a theory as a separate fragment in a subject's doxastic state. This is meant to represent the act of evaluating as a separate cognitive task. As in any Bayesian story, this doesn't require that the subject actually calculate her probabilities, or think explicitly about what her available background knowledge is. The model of a fragment is instead meant to represent an agent's act of deliberation or evaluation, and to make predictions about their future behavior based on this. This model provides us both with an adequate description of the subject, and an appropriate way to (normatively) evaluate whether their beliefs are rational.

3.2 Fragmentation Solutions

In order to obtain a concrete solution from the general fragmentation strategy, we need an account of how fragmentation is to be constrained. Specifically, we need a story about how a subject's evaluation-task fragment should be related to her other fragments. This final step is required for a solution to be descriptively adequate in cases of rational belief change in science, and for it to provide normative guidance for how subjects should think.

Generally, fragmentation accounts require constraints on how many fragments we use in representing a subject's behavior. This is important, because it might be tempting to simply add a fragment for each behavior the subject engages in. But this would leave the framework rationalizing almost any behavior, and it would offer little predictive power. The fragmentation strategy outlined in the last section provides resources to help with this kind of constraint. Specifically, it ties individual fragments to the subject's acts of evaluation: of freshly and fairly evaluating theories with an open mind, in light of accessible evidence. This seems like a familiar occurrence, and it is plausible that we can recognize such cases. So, indexing to these particular tasks seems a plausible enough way of constraining the number of fragments we need to describe and evaluate a subject.

In order to create a solution from the fragmentation strategy, however, we

also need constraints on how a subject's different fragments (should) relate to one another. There are conflicting considerations here, pulling us in different directions. On the one hand, the framework is designed to deal with inconsistency and recall failure. Thus, fragments must be able to differ significantly from one another. On the other hand, if there are no constraints on what the new fragments look like, a fragmentation version of BCT might still not give us much in the way of predictive and explanatory power, and may just appear too permissive. A theorist could vindicate any conclusion about whether something is evidence for a subject by cherry-picking a fragment with the right probability function or background knowledge.

Thus, a solution to the POE requires constraints on (1) the prior probability for the function of the new fragment, and on (2) the set K . I will explain each of these in turn.

Constraining ur-priors

The notion of *prior probability* that must be constrained here is sometimes called the *ur-prior*. In the traditional Bayesian picture, this is the hypothetical probability function describing the subject before they learned anything at all. The subject's degrees of belief at a time are given by conditionalizing the ur-prior on all evidence and background knowledge the subject possesses at that time. On the fragmentation picture, each fragment will, formally speaking, have its own ur-prior. This is the prior probability before conditionalizing on K , the relevant set of background evidence.¹⁹

Despite being hypothetical, ur-priors are taken to be interesting because they encode the most basic *epistemic standards* to which a subject is committed (Callahan 2020; Meacham 2016; Schoenfield 2014; Titelbaum 2020). A subject's epistemic standards determine how she will modify her beliefs given a piece of evidence. They are commitments the subject has regarding how to respond to evidence. For example, suppose a subject is committed to enumerative induction. That she saw a long sequence of black ravens (L) will increase her confidence that the next Raven she sees will be black (N). In other words, $Pr(N|L\&K) > Pr(N|K)$. Such responses will be influenced by what the subject has in her background knowledge (i.e., by her previously collected evidence). However, background knowledge cannot be the only thing that determines such responses. No sequence of background evidence alone is enough to require that someone respect enumerative induction. In fact, it is nonsensical on the Bayesian picture to talk about "background evidence alone." Responses to evidence will also always depend on the subject's starting place, her ur-prior. The shape of this ur-prior, and its contribution to a subject's epistemic responses, are determined by, and encode, the subject's epistemic commitments.

¹⁹Note that the fragmentationist, like the traditional Bayesian, is not committed to there ever having been a time that the subject's beliefs are accurately described by the ur-prior. The idea that one should conditionalize the ur-prior on all of one's evidence at a time is sometimes called *ur-prior conditionalization* (Meacham 2016). This is contrasted with conditionalization that occurs sequentially on new evidence as it arrives.

Epistemic standards can be thought of either as beliefs the subject has about how they ought to respond to evidence, or as plans or intentions regarding such responses (Meacham 2016). On either understanding, they are thought to have normative force regarding how a subject should respond to their evidence. If the subject has a commitment to a standard, they ought to follow the standard. Bayesians disagree about what epistemic standards a subject is required to have, however. *Objective Bayesians* think there is only one acceptable set of epistemic standards, and so only one way of constructing a rational ur-prior. *Subjective Bayesians* allow for more variety, and some suggest that there are no constraints whatsoever beyond probabilism.²⁰

The first question of how to constrain fragment similarity, then, is a question about how the ur-prior of a new fragment must relate to the ur-prior of the subject's other fragments. This question concerns what kind of coherence the subject must have in her epistemic standards across fragments. There are a variety of reasonably plausible ways of answering this question, each of which would entail a different version of the fragmentation solution. I will focus on two ways. The first we can call *subjective consistency*, and the second *subjective similarity*.

Subjective consistency requires that each individual subject has their own ur-prior, and that this ur-prior remains constant in all of their fragments. In keeping with the tradition of subjective Bayesianism, we can allow that this prior be any regular probability function, that is, any probability function that only assigns probability 0 to logical contradictions.²¹ Subjective consistency is attractive because it ensures that the subject has consistent epistemic standards, while respecting the differences in information access across contexts that motivates fragmentation. Differences between fragments are then entirely explained by differences in the background knowledge K available in the fragment. This fits with the usual way fragmentation is used: to model differences in information accessibility while engaged in different tasks.

The goal of the subjective consistency is to avoid ad-hoc fragmentation selection that would allow the subject (or the theorist describing them) to justify any response to evidence. If E is evidence for H for a subject given background evidence K , this should be the result of combining K with the epistemic commitments the subject generally follows. These general commitments are reflected in the subject's single, general ur-prior.

A different kind of constraint on a fragment's ur-prior is *subjective similarity*: a subject is rationally permitted to have a variety of different ur-priors for different fragments, as long as each is close enough to the others. This requirement has the same goal as subjective consistency: to ensure that evaluations of evidential responses result from appeal to the subject's general epistemic commitments, and are not the result of an ad hoc selection of ur-priors. However,

²⁰For overviews regarding these distinctions, see Earman (1992), Strevens (2006), Talbott (2015). For a discussion of how they relate to uniqueness and permissivism, see Titelbaum (2020).

²¹This is what Weisberg calls "Initial Regularity" (2011). Alternatively, one could require that the ur-prior meet certain other requirements, e.g., respecting the principal principle or being inductivist.

this constraint is less demanding than subjective consistency. It still rules out ad hoc choices of standards which would rationalize any response to evidence whatsoever. At the same time, it avoids being excessively difficult to satisfy: it does not require the subject to adhere to precisely the same epistemic standards. It allows for some slippage between contexts. However, it still rules out radical differences in epistemic standards. A subject cannot have one fragment which generally relies on induction and another which encodes a general inductive skepticism (or counter-inductivism).²²

To determine what it means for ur-priors to be adequately similar, we can appeal to the same underlying distance measures used in the gradational accuracy literature (Joyce 1998; Pettigrew 2016). There, these measures are used to determine how accurate a subject's credences are (or would be) at any particular world, in order to argue for rational constraints such as probabilism and conditionalization. This is accomplished by measuring the distance between the subject's probability function and the perfectly accurate omniscient probability function at each world. The omniscient function assigns a probability of 1 to all the truths in that world, and 0 to all the falsehoods. A credence function is considered more accurate insofar as the values it assigns to each proposition are closer to what the omniscient function assigns. The most popular way of measuring accuracy is the Brier score, which is determined using a kind of divergence known as squared Euclidean distance measure (δ^2) (Pettigrew 2016, p. 8).

We can appeal to squared Euclidean distance to determine the similarity of ur-priors. This distance between probability functions Pr_{t_1} and Pr_{t_2} is determined by the equation:

$$\delta^2(Pr_{t_1}, Pr_{t_2}) = \sum_{w_i \in \Omega} (Pr_{t_1}(w_i) - Pr_{t_2}(w_i))^2$$

This works by first determining the difference between the probability assigned by each function to each world description in the state space Ω and squaring it. Then these squared differences are summed.²³

The subjective similarity requirement can then be summarized as follows. First, each of a subject's fragments must be representable using the same algebra \mathcal{A}_Ω . Ω must be rich enough, i.e., contain enough worlds, to represent all of the subject's doxastic possibilities in all of her fragments. These doxastic possibilities represent all of the differences in how the world could turn out to be by the subject's lights.²⁴ Second, the squared Euclidean distance between the

²²Support for this claim depends on the assumptions below and is discussed in more detail in the appendix.

²³I'm leaving aside a variety of technical details. For a more thorough explanation, see Pettigrew (2016).

²⁴There are a few technical issues to address with this first requirement of subjective similarity. Each w_i is a unique element of the maximally fine-grained partition of Ω , corresponding to a unique possible world. This partition (i.e., the relevant set of possible worlds) is determined by what possibilities the subject entertains as possible. Each possible world is maximally determinate: every event that is possible either does or does not occur at that world Stalnaker (1984).

probability function in each fragment t_i and any other fragment t_j must be below some threshold θ . The size of θ determines just how permissive subjective similarity is.

Squared Euclidean distance has a variety of beneficial features that make it apt for measuring the closeness of probability functions, which is why the Brier score is commonly used in the accuracy literature (Pettigrew 2016). For instance, δ_2 is continuous, so that small differences in probabilities assigned by each function result in only small changes in distance. Moreover, the overall distance between functions depends only on the differences between probabilities they assign to individual propositions. Furthermore, if the difference between $Pr_{t_1}(p)$ and $Pr_{t_2}(p)$ decreases, then the $\delta^2(Pr_{t_1}, Pr_{t_2})$ will decrease as well.²⁵

However, squared Euclidean distance is not the only option for measuring distance. In fact, many distance measures will arrive at essentially similar results for our purposes.²⁶ There is not space to fully defend a particular way of measuring distance between credence functions here, and such a choice may be unnecessary. My point here is conditional: if one is convinced that one (or more) of these divergences is appropriate for measuring accuracy, then one should be comfortable using it to define similarity for the purposes of constraining fragmentation. Hence, those who like the Brier score should like squared Euclidean distance for defining subjective similarity.

A distance-based definition of subjective similarity helps to rule out radical, ad hoc changes in epistemic standards. This is because an ur-prior represents the subject's epistemic standards in terms of conditional probability assignments. Ur-priors encoding different standards will assign different probability values to certain conditional probabilities. If the standards are radically different, as inductivism and counter-inductivism are, ur-priors which encode the two will display large differences in certain conditional probabilities. For inductivism/counter-inductivism, these will be conditional probabilities involving possible courses of experience (e.g., seeing fifty black ravens). Such ur-priors will therefore display significant distance between their probability functions. Since the encoded commitments are general, these differences will appear in many conditional probabilities across each function, leading to a large overall distance between them. Subjective similarity will therefore forbid

(The picture proposed here is compatible with thinking of the worlds as either concrete or abstract objects.) The w_i propositions in Ω thus pick out the finest-grained partition relevant to modeling a subject. I am appealing to substantive assumptions about subjects, and about the purpose of the model, to determine a privileged, unique finest-grained partition for Ω . However, if one is skeptical that this is enough to determine a plausible candidate for a unique partition of w_i propositions, there is still a way to use subjective similarity. One could instead make the first requirement existentially quantified: that there is *some* partition of w_i propositions in Ω which can be used to represent all of the subject's fragments.

²⁵Pettigrew (2016) provides a detailed discussion of each of these features and more.

²⁶In the appendix, I appeal to manhattan or taxicab distance for ease of calculation. Pettigrew identifies other examples from a class of similar measures to which δ^2 belongs, the additive Bregman divergences, which share many of its relevant features, before arguing that δ^2 is the best (Pettigrew 2016, pg. 65).

having fragments with significantly different general epistemic commitments. I provide a brief demonstration of how this applies in the case of inductivism and counter-inductivism in the appendix.²⁷

Both subjective consistency and similarity provide reasonable options for constraining ur-priors across fragments. Of course, there are other options for constraining ur-priors that are compatible with the fragmentation strategy. I chose subjective consistency and similarity as plausible ways of providing concrete solutions within the fragmentation strategy. However, it is primarily the strategy itself that I aim to support.

Constraining background knowledge

The second question regarding how the fragmentation of a subject should be rationally constrained concerns the membership of the set of background knowledge K in any fragment.²⁸ In order for fragmentation to solve the problem of old evidence, the fragment describing the subject's process of evaluation needs to exclude the evidence to be evaluated from their background knowledge. That is, K must exclude E . But there are many different potential sets of background beliefs that exclude E . Which ones are acceptable for the subject to have in the new fragment? How should this new set K be related to the subject's other fragments? Specifying an answer to this question is required for a concrete solution to the problem.

One potential source for an answer to this question is the literature on belief revision. The various belief revision frameworks, like AGM and ranking theory, have attempted to give an account of how belief states of a subject must rationally change when gaining or losing beliefs.²⁹ These accounts are meant to tackle precisely the kind of problem that concerns us: modeling the

²⁷Pettigrew's discussion about epistemic risk aversion and its relation to priors offers another helpful example (Pettigrew 2016, p. 166). It illustrates how different epistemic commitments result in prior probability functions that are significantly distant from one another. For Pettigrew, epistemic conservatives are those who are maximally averse to having inaccurate credences. These subjects will satisfy the principle of indifference when selecting their ur-prior, as a result of embracing a maximin strategy to avoid inaccuracy. An epistemic radical, on the other hand, is a (presumably hypothetical) subject who cares only about maximizing accuracy and is indifferent to risk of inaccuracy. Radicals will follow a maximax strategy to select an ur-prior with the highest *potential* accuracy score: a maximally opinionated probability function that assigns every proposition either 1 or 0. Notice that the principle of indifference, defined in Pettigrew's terms of minimizing worst case inaccuracy (Pettigrew 2016, p. 159), just is an attempt to be as far away from maximal opinionation as possible. Thus, an epistemically conservative ur-prior will have a large distance from an epistemically radical one defined over the same algebra. Subjective similarity (defined using δ^2 and an appropriate θ) will thus demand that subjects do not have one fragment with conservative commitments and another with radical commitments, as these will be too far apart.

²⁸Strictly speaking, K is a proposition variable on our setup, which would usually pick out a long conjunction. However, any such proposition will be equivalent to a set of propositions. In this section, I will ignore the distinction for ease of discussion, and talk of a set K .

²⁹See Hansson (2017) and Huber (2013a, 2013b) for an overview of the literature in belief revision, which was largely sparked by the proposal of the AGM theory of Alchourrón, Gärdenfors, and Makinson (1985).

removal of a proposition from a set of beliefs that is closed under logical entailment (i.e., removing the old evidence E from K). A formally defined operation which accomplishes this is known as *contraction*.

Unfortunately, there is no consensus on an acceptable contraction operator. This is because there are generally many ways to obtain a set of propositions closed under entailment that satisfy the constraint of not including E .³⁰ Belief revision theories attempt to find a particular contraction operator that acts as a function, taking in a belief set K and the proposition(s) to be removed (E), and providing a single set ($K - \{E\}$) as output. A contraction operator that works this way, and which does not have highly counter-intuitive consequences, has not been forthcoming. However, if one gives up commitment to the idea that such operators must act as functions which output a single new set, the various contraction operators seem more promising. If we are permissivist fragmentationists, we can treat the operators as providing constraints on the relationship between K and $K - \{E\}$, i.e., constraints on the relationship between the background knowledge in a subject's fragments.

So, one way of constraining K would be by appeal to one's favorite contraction operator. For instance, one could appeal to Gärdenfors' *entrenchment-based contraction* (1988). The basic idea of entrenchment contraction is that, when a subject is forced to give up some beliefs, they should choose to give up those with the least explanatory power and informational value. A rational subject would treat these more epistemically valuable propositions as more central, i.e., as more entrenched in their belief set. To represent this, Gärdenfors proposed that each agent be modeled as having an entrenchment ranking which orders all the propositions in K , with this ordering meeting certain constraints so as to reflect the subject's epistemic value judgments. So, when contracting K by E , the subject should choose to eliminate the least entrenched (least valuable) propositions which contribute to E -entailment, thereby retaining as many valuable propositions as possible in $K - \{E\}$.³¹

When applied to constraining fragmentation, entrenchment contraction would constrain the K_{t_j} for any new fragments by looking at the K_{t_i} from a previous theory-evaluation fragment. Then the K_{t_j} must be one obtained by an entrenchment contraction removing E from K_{t_i} . There may be more than one acceptable new K_{t_j} . Given that we accept permissivism, however, this is no problem for the solution. And this kind of constraint can be implemented using whichever contraction operator one prefers (e.g., the one proposed in Levi (2004b), or the probabilistic contraction operator found in Suzuki (2005),

³⁰To see why this is the case, consider $Cn(X)$ set obtained by taking all the logical consequences (entailments) of sentences in X , where $P \rightarrow R, P, R \in X$. Suppose we wish to remove R to obtain a new set closed under consequence (i.e, we want to contract by R). Because the new set will be closed under entailment, We must also remove any sentences which entail R . However, there are some sentences which do not individually entail R , but do together. For instance, $P \rightarrow R$ and P entail R . Which do we remove? Removing either will avoid the entailment of R . But it is unclear which we should prefer. Determining which to prefer is the main task of a theory of an appropriate contraction operator. There are many candidates, though none has achieved consensus acceptance. For more on this topic, see Gillies (2001) and Levi (2004a).

³¹For a helpful summary of entrenchment contraction, see (Hansson 2017, Sec. 2.2).

mentioned below in note 36).

A second way of constraining K may appeal to those with more objectivist sympathies, and who are unsatisfied with a permissive account. This option also appeals to an ordering of the epistemic value of propositions. However, this ranking would be applied to propositions generally, not merely to those in the subject's background evidence. Moreover, this ranking is objective, rather than being determined by the subject's own assessment of epistemic value. It ranks all propositions by their relevance to evaluating the theory at hand. On this account, what should be in K_{t_j} is all the *relevant* background evidence. When two potentially relevant pieces of information together entail E , and so must not be included together, which one makes it into K_{t_j} is determined by which one has the higher relevance ranking. What counts as relevant background knowledge for a task will depend on substantive epistemic and psychological considerations about the subject and the task at hand. According to this version of the fragmentation solution, this will not be something which can be determined formally, in full generality, without appeal to the particular features of the case. In Williamson's terms, it is not "operationalizable" (Williamson 2008).³²

Putting constraints together

To recap, I suggested there are two kinds of constraints that we need to specify to make the fragmentation strategy into a concrete solution: (1) constraints on ur-prior and (2) constraints on background knowledge K . I then suggested two possibilities for constraints of type (1), and two more for constraints of type (2). This suggests four possible fragmentation solutions (though these are not exhaustive), but I will focus on two which I take to be plausible.

The first fragmentation solution combines subjective consistency about ur-priors with an objective relevance ranking. We can call this the *consistent objective-ranking* version of the fragmentation solution. This solution provides a specific requirement on what a subject's new fragment should look like when they consider the old evidence. However, it relies on appeal to an objective relevance ranking that some might find implausible.

The second solution combines subjective similarity with an AGM-style contraction operator. Call this the *similar contraction* version of the fragmentation solution. This provides a permissivist approach to the problem. It provides robust constraints, but also meets the spirit of the subjective Bayesian project, because it allows for significant (rational) variability. I think that this provides the best version of the solution. It ensures that fragmentation does not rationalize just anything, while avoiding overly demanding rational requirements.

My goal in this paper is to show that the fragmentation strategy works to solve the problem of old evidence. The goal of this section was to show that

³²This kind of constraint is inspired by Williamson's (2000) idea of objective evidential probabilities. One need not go in for Williamson's whole picture to take on this kind of constraint, however.

the strategy does suggest specific, concrete solutions to the problem. Such solutions can be constructed by adding one’s own background commitments into the fragmentation picture.

4 Advantages Over Other Solutions

The fragmentation solution to the problem of old evidence is intuitively plausible, independently well-motivated, and gives the right verdicts in each kind of case. Moreover, it does all of this without significant deformation of the formal machinery of Bayesianism, and thus preserves BCT’s advantages. What remains to be shown, however, is that it provides the best solution. In this section, I will argue that it has significant advantages over the most prominent alternative solutions on offer. I will argue that each of these solutions is in fact already committed to a the fragmentation strategy. However, they also make unnecessary additional commitments which cause trouble.

4.1 GJN

The first alternative solution is due to Garber (1983), Jeffrey (1983), and Niiniluoto (1983). This is sometimes called the GJN-style solution, or the Garber-style solution.

The key move in the GJN solution is to suggest that the subject learns something else when they consider how old evidence relates to a new theory. That is, when they apply the old evidence to the new theory, they are not updating directly on the evidence. Instead, they are learning a different proposition, one that is about the relation between the old evidence and the theory. On this view, the subject doesn’t update her credence in T on E . Rather, they update on the proposition that E entails T . This gets represented by “ $E \vdash T$ ”, though here the “ \vdash ”, can refer to any kind of logico-mathematical entailment, not just the consequence relation of propositional logic. According to GJN, then, when E is old evidence for T , what the subject is actually doing is learning that E entails T . Thus, it’s possible that $Pr(T|T \vdash E) > Pr(T)$, which meets the definition of evidence.

However, this solution requires a way to relax logical omniscience requirements. Otherwise, the subject will be required to have probability 1 in $T \vdash E$ (because if it is true it is a logical truth), which would make it unhelpful in solving the POE. The GJN solution achieves this by interpreting propositions of the form $T \vdash E$ as logically independent atomic propositions in \mathcal{A}_Ω . This works by making the fact that “ \vdash ” means entailment opaque to the subject’s probability function.³³ So, according to GJN, sentences like $T \vdash E$ must get

³³In their version of this kind of solution, both Hartmann and Fitelson (2015) and Sprenger (2015) dispense with treating \vdash as entailment, instead viewing it as “explanation.” Their versions of the solution divorces the solution to POE from the problem of logical omniscience, which I think is the right move (though note the objections to this move in Kinney (2019)). However, the remaining objections to the GJN solution still apply.

treated just like any atomic sentence P . This opacity allows for a probability function that assigns $Pr(T \vdash E) < 1$, yet is not rendered incoherent by the failure to respect entailment relations. As a result, any logical relationship between E , T , and $T \vdash E$ must be “extra-systematic,” that is, a constraint that is added over and above the standard logical relations of the language we use to describe \mathcal{A}_Ω .³⁴

In order to solve the logical omniscience problem, the GJN solution requires that the state space, Ω , cannot correspond to the set of possible worlds. That is, Ω cannot be composed of w_i propositions that each uniquely identify a possible world, as we have so far understood them here. This is because some of the propositions in \mathcal{A}_Ω are now *negated* logical truths in disguise (e.g., $\neg(T \vdash E)$), which means Ω must now include states where such logical truths are false. But such *impossible* states are not among the *possible* worlds. To put it another way, if Ω consist only of propositions describing possible worlds, entailment relations like $T \vdash E$ will be true in all of the worlds still in the subject’s space of possibilities, because all necessary truths are true in every possible world. Then it would be impossible to be uncertain about $T \vdash E$.

The GJN solution thus tries to provide a solution to logical omniscience as a way of solving the POE. However, this solution to logical omniscience carries a number of worries with it. The first problem, recognized immediately by Garber (1983), is that it means the subject is required to be certain of arbitrarily complex tautologies expressed in the language describing \mathcal{A}_Ω , but allows them to be ignorant of extremely simple, even obvious logical truths. Supposing $P, Q, R \in \mathcal{A}_\Omega$, the subject would have to be certain of (the non-obvious tautology) $((\neg P \& P) \& Q) \rightarrow R$, while being allowed to be uncertain of $P \vdash P$.

In order to distinguish these two cases, Garber appeals to a distinction between “local” and “global” Bayesianism. Garber suggests that the framework of BCT should not be something that applies globally to a subject, but is applied to a particular problem using a problem-relative language L . That is, for describing a particular scientist dealing with a particular problem, we only require “local” omniscience, relativized to the language best suited for the problem at hand. This allows the language to have uncertainties about non-empirical truths, like $T \vdash E$, by making them atoms of the language, while still using the Bayesian framework (which requires logical omniscience within the language).

This local/global distinction should sound familiar: though he doesn’t use the term, Garber is essentially suggesting fragmentation. He does so to deal with these issues arising from logical omniscience. However, to solve the two

³⁴Even all these changes are not yet enough to allow the GJN proposal to solve the problem of old evidence: we also need some further extra-systematic constraints that ensure that the extra-systematic entailment claims actually bear the proper evidential support relations needed to make the right predictions in old evidence cases. There have been a profusion of suggestions for how this might work (Earman 1992; Hartmann & Fitelson 2015; Jeffrey 1983; Kinney 2019; Sprenger 2015). Most of the literature on the GJN solution focuses on giving an account of these extra-systematic constraints that gets the right results and is well-motivated. No version of these constraints has won consensus acclaim. I won’t focus on that problem for GJN here, as I think the proposal’s problems run deeper than that.

issues above, he is also changing the language of each fragment, giving up on possible worlds, and describing the way scientists update on old evidence in a way that is, I will suggest below, counter-intuitive.

Thus, the GJN solution is actually already committed to the fragmentation strategy. However, what its proponents have not realized is that the fragmentation is all that is necessary to solve the problem of old evidence. The extra features of the view involving changes to the language and algebra, additions of “extra-systematic” items and constraints, and treating logically complex statements as atoms, cause it unnecessary trouble. Moreover, the view tries to solve the POE by solving the more difficult problem of logical omniscience first, which is unnecessary.

Perhaps the most difficult problem the GJN solution faces is that it involves interpreting all old evidence cases as updating on propositions like $T \vdash E$. But it is plausible that actual scientific episodes involve cases where this doesn't seem to have been the case. In fact, Earman suggests that that Einstein took OM to be his evidence for GTR , not some fact about what is explained (1992). He suggests that OM and $GTR \vdash OM$ are “neither semantically nor extensionally equivalent” (130). We have evidence that Einstein took OM to be his evidence, but no evidence that he took $GTR \vdash OM$ to be evidence. And regardless of how it worked for Einstein, it does seem like someone could possibly take the one thing as evidence and not the other. Moreover, Earman suggests that the fact that GTR entails OM must also be old evidence for contemporary students who want to consider OM 's support of GTR . This is because the students almost always learn that GTR entails or explains OM before they really learn any details about the theory itself (130). So the GJN account claims that all old evidence cases are to be explained by appeal to updating on these entailment claims. But this is implausible in many cases.

This last case, involving physics students, raises another significant issue for GJN: it is only a solution to the dynamic version of the POE. GJN does not give the right prediction in static version cases. Once a subject updates on a $T \vdash E$ claim, it will be assigned probability 1 and will no longer meet the probabilistic relevance definition of evidence. The GJN solution does not offer any way to make it that something remains evidence for a theory at any time after it has been learned.

The fragmentation solution does not suffer from these problems, and has significant advantages. It is simpler, because all it requires is the fragmentation. It involves fewer commitments, and does not make us give up on possible world semantics. Meanwhile, it derives independent support because it is embedded in a larger fragmentation framework. Moreover, it provides a solution to both versions of the POE. Finally, the fragmentation solution allows OM itself to be evidence for GTR even when it is old evidence, and has an explanation for how to make this very same old evidence “new” and relevant again. Thus, the fragmentation solution is preferable to GJN.

4.2 Counterfactual Solution

The second alternative solution I will address is the counterfactual solution, from Howson (1991) and Howson and Urbach (2006). The basic idea behind this solution is that we can explain old evidence by appealing to a different probability function than the subject's current one. Specifically, this is the probability function the subject *would* have had, if they had not already learned the evidence. Since this counterfactual probability function would not treat the subject as having probability 1 in the old evidence, it would be able to have $Pr(T|E) > Pr(T)$, as is required for the definition of evidence. The trick is just to find the right counterfactual probability function. This would be a function that starts with the subject's own ur-prior, but is then conditional on a set of background knowledge containing all the subject's knowledge, but with the evidence to be evaluated, E , left out: $K - \{E\}$.

The counterfactual solution appeals to an additional probability function, distinct from the one that represents the subject's current belief state. Thus, there are two probability functions that describe the subject's doxastic state, rather than a single global one. Moreover, this function must be conditional on a set of background evidence K that does not contain E . So, the solution is already committed to something very much like the fragmentation strategy. Formally speaking, they are quite similar. The biggest difference is in interpretation of the formalism, which is where the counterfactual solution runs into a number of problems.

The first problem is that there might not be any probability function the subject would have had if she hadn't learned E . To take an extreme but illustrative example, it might be that learning E saved the subject's life, so that there are no nearby worlds in which the subject didn't learn E and has any doxastic state whatsoever.³⁵

Second, even if there are probability functions the subject would have counterfactually, there is no guarantee that there is only one. That is, it is very unclear that there is a single probability function that satisfies the description "the probability function the subject would have had if she had not learned E ." This is for the same reasons discussed in section 3.2 above concerning belief revision, as a result of the logical relations between E and other members of K . Without a contraction operator that always provides a single output set $K - \{E\}$, there will not be a fact of the matter about what a subject would have believed.

The problem for the counterfactual solution is that it is committed to there being one belief state the subject *would* have had, if they did not have the evidence E . The fragmentation solution makes no such commitment, and can be satisfied with a contraction operator that does not provide a single new set $K - \{E\}$. The account can instead be permissive.³⁶

³⁵See Earman (1992) and Strevens (2006) for more on this line.

³⁶In later work, Howson and Urbach (2006, Ch. 9c) address some of these issues for the counterfactual solution. First, they suggest appealing to a probabilistic contraction function found in (Suzuki 2005). Then, they argue that the solution is compatible with permissivism, and that any

But even if these problems for the counterfactual solution were solved, there would still be the question of why what the subject *would have* believed is relevant. BCT is supposed to tell us whether something *is* evidence for a theory, not whether it would be. Appealing to the counterfactual seems a bit like changing the subject.³⁷ So the counterfactual solution, like GJN, threatens to involve a kind of change of subject: it tells us whether some proposition would be evidence for another, if things had been different than they are. But we want our Bayesian framework to get the right answer about whether some subject has evidence for a theory, given the way the world actually is.

In sum: The counterfactual solution is committed to something very similar to the fragmentation strategy. However, the fragmentation strategy is all that is needed to solve the problem of old evidence. The additional commitments the counterfactual solution makes are about counterfactual relations between the subject and certain probability functions, relations which might not obtain. The counterfactual nature of these relations is what causes the problems, and this counterfactual relation is not needed to solve the problem of old evidence.

From a formal standpoint, the counterfactual solution in (Howson & Urbach 2006) is very similar to a fragmentation solution. In particular, it is similar to a fragmentation solution committed to subjective consistency about the *ur*-prior, and which appeals to AGM for constraints on *K*. However, where the solutions differ substantially is in the interpretation of the formalism, and in what kind of constraint there is between the *K* sets of different fragments. Fragmentation allows a permissive constraint, tied to different ways of understanding relevance, which is not committed to any extraneous counterfactual claims.

Fragmentation thus presents a different answer about the nature of the subject's different probability functions and their relations to one another. The appeal to a separate probability function is justified not because it is what the subject would have believed. Instead, it is justified by its use in describing something everyone admits is going on in the cases in question: the subject's deliberation about whether some evidence supports some theory. This avoids the worry that there is no relevant probability function. Moreover, it avoids any worry about changing the subject. The fragmentation solution uses a probability function that describes the actual state of the subject to explain why *E* is evidence for *T*. This function is just limited to a particular compartment of the subject's doxastic state.

Therefore, despite their similarity, the fragmentation solution has significant advantages over the counterfactual solution.³⁸

Bayesianism is committed to counterfactual considerations because of the importance of conditional probabilities. Insofar as this is correct, this makes their solution even more clearly a version of the fragmentation strategy. In fact, it would be formally identical to a fragmentation solution which is committed to subjective consistency and AGM-style constraints on *K*, i.e., to a consistent-contraction version.

³⁷See Earman (1992) and Strevens (2006) for similar worries.

³⁸There are two other proposed solutions to the POE that also rely on a kind of fragmentation. First, Sprenger (2015) offers a solution which combines the counterfactual solution with the GJN solution. It similarly relies on a kind of fragmentation, and ultimately suffers from both the prob-

Conclusion

The problem of old evidence threatens Bayesian confirmation theory by suggesting that it fails to deliver on one of its purported best attributes: describing and explaining the evidential support relation. The fragmentation solution provides a response to this worry that is independently motivated, intuitively plausible, and explanatorily fruitful. It is also compatible with possible world semantics. Although the fragmentation framework has been used in a solution to logical omniscience (Rayo 2013), the fragmentation solution to POE does not rely on solving this much harder problem first. Nor does it require significant alterations to the Bayesian framework. Instead, it simply relaxes an assumption often left in the background, and a poorly justified one at that: the assumption that a single probability function should be applied globally to a subject in all circumstances. The fragmentation solution preserves BCT by suggesting probability functions describe (and rationally govern) fragments of the subject's doxastic state.

The fragmentation solution also enjoys significant benefits over its main rivals as a solution to the POE. In fact, these rivals are best understood as imperfect versions of the fragmentation solution, which make unnecessary, problematic commitments. Thus, I suggest that the fragmentation solution does solve the Problem of Old Evidence. Moreover, it points us in the direction of solutions to other problems. For instance, I conjecture that the problem of novel theories ((Earman 1992; Strevens 2006) admits to a similar treatment.

Appendix

Here, I offer support for the claims made in section 3.2 regarding subjective similarity defined in terms of distance. Specifically, I will illustrate how placing a distance constraint on ur-priors for different fragments will proscribe radical changes in epistemic standards. I appeal to Fitelson's PrSAT decision procedure implemented in *Mathematica* (2008).³⁹ PrSAT is a *Mathematica* function that tests whether (arbitrary) sets of statements in the probability calculus are satisfiable. In other words, it will test for whether certain claims are true. In particular, I use it here to illustrate that a probability function defined over

lems I have raised for GJN and for the counterfactual solution. Second, Eells (1985), Eells and Fitelson (2000), and James Hawthorne (2005) propose a solution that posits two distinct probability functions used to describe an agent: one representing the agent's evidential support judgments, and one representing their current belief state as it is updated over time. Whether E is evidence for T depends on the first function, while whether E confirmed T depends on the history of the second function. This solution is also clearly committed to a form of fragmentation, but suffers from its own difficulties. For one thing, it severs the relationship between confirmation and evidence.

³⁹The *Mathematica* notebook containing this demonstration is available as a supplement to this article, or at <https://www.willfleisher.net/research>. For more information about PrSAT, and to download it, see <http://fitelson.org/PrSAT/>. For information about *Mathematica*, see <https://www.wolfram.com/mathematica/>. Thanks to Branden Fitelson for extensive help and discussion on the material in this appendix.

Table 1: Inductive ur-prior

	H	E	$Pr_m(\cdot)$
w_1	T	T	$1/4$
w_2	T	F	$1/4$
w_3	F	T	$1/32$
w_4	F	F	$15/32$

a small set of sentences that satisfies an inductivist constraint must be a certain significant distance from one that satisfies a counter-inductivist constraint.

It is beyond the scope of this paper to prove, in full generality, that there must be a certain distance between any inductivist prior and any counter-inductivist prior. This is primarily for philosophical reasons. The nature and norms of inductive inference are extremely controversial and I cannot settle the debate here. Instead, I will focus on a particular case which appeals to paradigmatic, uncontroversial instances of induction and counter-induction to illustrate why we should expect many (perhaps all) inductive and counter-inductive ur-priors to have significant divergence.

In this simplified case, the subject has probability functions defined over an algebra built from two atomic propositions, a hypothesis H and evidence E . While this is extremely simple compared to probability functions which model actual scientists, it helpfully illustrates how differences in epistemic standards result in significant distances between ur-priors. It will be easy to see how the point generalizes to probability functions defined over larger algebras.

Table 1 displays a probability function for a fragment m , using a stochastic truth table (Fitelson 2008). We assume that H is some hypothesis, and E is some relevant course of experience that inductively supports H . For instance, we can suppose that E is the proposition *I saw 100 ravens and all were black*, while H is the proposition *ravens are black*. This function, Pr_m (or just m), taken as an ur-prior, displays an inductivist commitment. This is because m assigns E a high degree of confirmation (or evidential support) for H . We can see this by appeal to the *Likelihood Ratio* confirmation measure (Fitelson 1999, 2007):

$$LR_m(H, E) = \frac{Pr_m(E|H)}{Pr_m(E|\neg H)}$$

Using PrSAT, we can calculate that the likelihood ratio for H and E in m is $LR_m(H, E) = 8$. This is a significant positive degree of confirmation, which is precisely what an inductivist standard would require in this case.⁴⁰ This gives us a paradigmatic example of an inductivist prior.

A paradigmatic counter-inductivist, on the other hand, would assign a high degree of *disconfirmation* in this case. That is, they would take E to strongly disconfirm H . An ur-prior Pr_{m^*} (or m^*) displaying this epistemic commitment

⁴⁰Other measures of confirmation will have similar results in this case. LR is chosen for illustration, but nothing hinges on the choice.


Table 2: Counter-inductive ur-prior

	H	E	$Pr_{m^*}(\cdot)$
w_1	T	T	$1/512 \times (191 - 15\sqrt{161})$
w_2	T	F	$1/4$
w_3	F	T	$1/32$
w_4	F	F	$3/512 \times (59 + 5\sqrt{161})$

will thus assign a high value for the likelihood ratio for $\neg H$ and E . Since we are already treating 8 as a significantly confirmatory likelihood ratio, we can then ask: what is the closest (ur-prior) probability function m^* to m such that $LR_{m^*}(\neg H, E) = 8$? PrSAT calculates the closest such m^* as displayed in table 2 (Note that $Pr_{m^*}(w_1) \approx 0.0013$ and $Pr_{m^*}(w_4) \approx 0.7174$). We can then calculate the distance between m and m^* . I will illustrate this point using Manhattan (taxicab) distance, since it is easier to calculate in PrSAT than δ^2 , but the basic thrust of the results will hold for both. Using PrSAT, we can calculate that the Manhattan distance between m and m^* is ≈ 0.25 . Thus, as required, there is a significant distance between m and m^* . It would be easy to set a threshold θ below this distance.

Moreover, for any m' , there is a monotonic relationship between the difference in likelihood ratio of m and m' and the distance between m and m' . To see this, consider a function f which calculates the closest m' to m for any value of $LR_{m'}(H, E)$. Figure 1 shows f plotted as $LR_{m'}(H, E)$ goes from $1/8$ to 8. It shows that distance decreases monotonically as difference in $LR_{m'}(H, E)$ does (the non-smoothness at $LR_{m'}(H, E) = 1$ is a result of the shift from disconfirmation to confirmation). In other words, the closer they are to one another in inductive commitments, the closer they are in distance. This shows that there is general relationship between strength of commitment to inductivism (as represented by the likelihood ratio) and distance. This is one reason that suggests we can probably generalize from cases like this one.

The upshot for subjective similarity is this: one could easily pick a threshold θ that would forbid having a fragment with an ur-prior Pr_m while having another fragment with ur-prior Pr_{m^*} . This allows for small differences between fragments' ur-priors, without allowing one to go from an epistemic commitment to induction in one fragment to commitment to counter-induction in another fragment. Although I focused here on a simple case, it is easy to see how this would generalize. Suppose a fragment's ur-prior is paradigmatically inductivist in this way across a wide range of propositions in a large algebra. This can only increase the distance between an inductivist ur-prior and another that is paradigmatically counter-inductivist. Each place they disagree in likelihood ratio would increase the distance between them (and adding more propositions to the algebra can never decrease the distance). Thus, it should be possible to choose a threshold θ for such an agent that would rule out having any (paradigmatically) counter-inductivist fragments. Of course, for large algebras, this kind of distance measure might not forbid having an inductivist ur-prior and

Figure 1: Plot of f 

having another ur-prior that is generally inductivist but very rarely prescribes counter-induction. But this is in keeping with the mildly permissivist spirit of subjective similarity, which is meant as a less stringent requirement. Its goal is to forbid agents from having radically divergent ur-priors.

Acknowledgements

For helpful discussion and comments on this paper, I want to thank Austin Baker, Laura Callahan, Adam Elga, Georgi Gardiner, Agustin Rayo, Susanna Schellenberg, Ernest Sosa, and Peter van Elswyk. Special thanks to D Black, Andy Egan, Branden Fitelson, Megan Feeney, and Anton Johnson for repeated readings and discussion. In addition, I want to thank several anonymous referees, especially one at this journal.

References

- Alchourrón, C. E., Gärdenfors, P., & Makinson, D. (1985). On the Logic of Theory Change: Partial Meet Contraction and Revision Functions. *Journal of Symbolic Logic*, 50(2), 510–530.
- Barnes, E. C. (1999). The Quantitative Problem of Old Evidence. *British Journal for the Philosophy of Science*, 50(2), 249–264.
- Callahan, L. F. (2020). Epistemic Existentialism. *Episteme*, forthcoming, 1–16. doi: 10.1017/epi.2019.25
- Cherniak, C. (1983). Rationality and the structure of human memory. *Synthese*, 57(2), 163–186.

- Crupi, V. (2015). Confirmation. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2015 ed.).
- Earman, J. (1992). *Bayes or bust?: A critical examination of Bayesian confirmation theory*. Cambridge, Mass: MIT Press.
- Eells, E. (1985). Problems of Old Evidence. *Pacific Philosophical Quarterly*, 66(3), 283.
- Eells, E., & Fitelson, B. (2000, December). Measuring Confirmation and Evidence. *The Journal of Philosophy*, 97(12), 663. doi: 10.2307/2678462
- Egan, A. (2008a). Seeing and believing: Perception, belief formation and the divided mind. *Philosophical Studies*, 140(1), 47–63.
- Egan, A. (2008b, July). Seeing and believing: Perception, belief formation and the divided mind. *Philosophical Studies*, 140(1), 47–63. doi: 10.1007/s11098-008-9225-1
- Elga, A., & Rayo, A. (2019). Fragmentation and information access. *Manuscript*.
- Fitelson, B. (1999). The Plurality of Bayesian Measures of Confirmation and the Problem of Measure Sensitivity. *Philosophy of Science*, 66(3), 378.
- Fitelson, B. (2007). Likelihoodism, bayesianism, and relational confirmation. *Synthese*, 156(3), 473–489. doi: 10.1007/s11229-006-9134-9
- Fitelson, B. (2008). A Decision Procedure for Probability Calculus with Applications. *Review of Symbolic Logic*, 1(1), 111–125. doi: 10.1017/s1755020308080039
- Franklin, A. (1990). *Experiment, Right Or Wrong*. Cambridge University Press.
- Garber, D. (1983). Old Evidence and Logical Omniscience in Bayesian Confirmation Theory. In J. Earman (Ed.), *Testing Scientific Theories* (Vol. 10). Minneapolis: University of Minnesota Press.
- Gärdenfors, P. (1988). *Knowledge in Flux: Modeling the Dynamics of the Epistemic States*. Bradford Book. MIT Press, Cambridge, MA.
- Gillies, A. S. (2001). *Rational Belief Change* (Unpublished doctoral dissertation). Dissertation.
- Glymour, C. (1980). *Theory and Evidence*. Princeton University Press.
- Greco, D. (2014). Iteration and Fragmentation. *Philosophy and Phenomenological Research*, 88(1), 656–673.
- Hájek, A. (2019). Interpretations of Probability. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2019 ed.). Metaphysics Research Lab, Stanford University.
- Hansson, S. O. (2017). Logic of Belief Revision. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2017 ed.). Metaphysics Research Lab, Stanford University.
- Hartmann, S., & Fitelson, B. (2015). A New Garber-Style Solution to the Problem of Old Evidence. *Philosophy of Science*, 82(4), 712–717.
- Hawthorne, J. (2005). Degree-of-Belief and Degree-of-Support: Why Bayesians Need Both Notions. *Mind*, 114(454), 277–320.
- Howson, C. (1991). The 'Old Evidence' Problem. *British Journal for the Philosophy of Science*, 42(4), 547–555.
- Howson, C., & Urbach, P. (2006). *Scientific Reasoning: The Bayesian Approach*. Open Court Publishing.
- Huber, F. (2013a, July). Belief Revision II: Ranking Theory. *Philosophy Compass*, 8(7), 613–621. doi: 10.1111/phc3.12047
- Huber, F. (2013b, July). Belief Revision I: The AGM Theory. *Philosophy Compass*, 8(7), 604–612. doi: 10.1111/phc3.12048
- Jeffrey, R. C. (1983). Bayesianism With A Human Face. In J. Earman (Ed.), *Testing Scientific Theories* (pp. 133–156). University of Minnesota Press.

- Jeffrey, R. C. (1990). *The logic of decision*. Chicago: University of Chicago Press.
- Johnson, D. A. (2020). *Deliberations on deliberation* (Unpublished doctoral dissertation). Rutgers University Department of Philosophy, New Brunswick, NJ.
- Joyce, J. M. (1998). A Nonpragmatic Vindication of Probabilism. *Philosophy of Science*, 65(4), 575–603.
- Joyce, J. M. (1999). *The Foundations of Causal Decision Theory*. Cambridge University Press.
- King, J. C. (2014). Structured Propositions. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2014 ed.).
- Kinney, D. (2019). Inductive Explanation and Garber–Style Solutions to the Problem of Old Evidence. *Synthese*, 196(10), 3995–4009.
- Levi, I. (2004a). *Mild Contraction: Evaluating Loss of Information Due to Loss of Belief*. Oxford University Press.
- Levi, I. (2004b). *Mild contraction: Evaluating loss of information due to loss of belief*. Oxford; Oxford; New York: Clarendon ; Oxford University Press.
- Lewis, D. (1982). Logic for Equivocators. *Noûs*, 16(3), 431–441.
- Meacham, C. J. G. (2016). Ur-Priors, Conditionalization, and Ur-Prior Conditionalization. *Ergo: An Open Access Journal of Philosophy*, 3. doi: 10.3998/ergo.12405314.0003.017
- Menzel, C. (2016). Possible Worlds. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2016 ed.).
- Niiniluoto, I. (1983). Novel Facts and Bayesianism. *British Journal for the Philosophy of Science*, 34(4), 375–379.
- Pettigrew, R. (2016). *Accuracy and the Laws of Credence*. Oxford University Press Uk.
- Rayo, A. (2013). *The Construction of Logical Space*. Oxford University Press.
- Schoenfield, M. (2014). Permission to Believe: Why Permissivism Is True and What It Tells Us About Irrelevant Influences on Belief. *Noûs*, 48(2), 193–218. doi: 10.1111/nous.12006
- Schwitzgebel, E. (2015). Belief. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Summer 2015 ed.).
- Sprenger, J. (2015). A Novel Solution to the Problem of Old Evidence. *Philosophy of Science*, 82(3), 383–401.
- Stalnaker, R. (1984). *Inquiry*. Cambridge: Cambridge University Press.
- Stalnaker, R. (1999). *Context and content: Essays on intentionality in speech and thought*. Oxford ; New York: Oxford University Press.
- Strevens, M. (2006). Notes on Bayesian confirmation theory. *manuscript*.
- Suzuki, S. (2005). The Old Evidence Problem and Agm Theory. *Annals of the Japan Association for Philosophy of Science*, 13(2), 105–126.
- Talbott, W. (2015). Bayesian Epistemology. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Summer 2015 ed.).
- Titelbaum, M. (2020). *Fundamentals of Bayesian Epistemology* (Vol. forthcoming). Oxford University Press.
- Weisberg, J. (2011). Varieties of Bayesianism. In D. Gabbay, S. Hartmann, & J. Woods (Eds.), *Handbook of the History of Logic* (Vol. 10). North Holland.
- Weisberg, J. (2016). Belief in Psychontology. *The Philosopher's Imprint*, *Forthcoming*.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.
- Williamson, T. (2008). Why Epistemology Cannot Be Operationalized. In Q. Smith (Ed.), *Epistemology: New Essays*. Oxford University Press.
- Yalcin, S. (2018). Belief as Question-Sensitive. *Philosophy and Phenomenological Research*, 97(1), 23–47. doi: 10.1111/phpr.12330