

New Theory about Old Evidence

A framework for open-minded Bayesianism

Received: date / Accepted: date

Abstract We present a conservative extension of a Bayesian account of confirmation that can deal with the problem of old evidence and new theories. So-called open-minded Bayesianism challenges the assumption—implicit in standard Bayesianism—that the correct empirical hypothesis is among the ones currently under consideration. It requires the inclusion of a catch-all hypothesis, which is characterized by means of sets of probability assignments. Upon the introduction of a new theory, the former catch-all is decomposed into a new empirical hypothesis and a new catch-all. As will be seen, this motivates a second update rule, besides Bayes’ rule, for updating probabilities in light of a new theory. This rule conserves probability ratios among the old hypotheses. This framework allows for old evidence to confirm a new hypothesis due to a shift in the theoretical context. The result is a version of Bayesianism that, in the words of Earman, “keep[s] an open mind, but not so open that your brain falls out.”

Keywords Bayesianism · Confirmation · Old Evidence Problem · Theory Change · Catch-all Hypothesis · Formal Epistemology

1 Introduction

Bayesianism offers a way to revise our degrees of belief in light of new evidence. However, it does not capture all the relevant belief dynamics: in the process of evaluating our evidence, we may want to consider a new theory, and so reconsider some of the assumptions on which all of our former degrees of belief co-depend. Standard forms of Bayesianism do not foresee the option of adopting a new theory in its formalism, so it seems that in such a case we have to start from scratch: assigning priors to the empirical hypotheses belonging to the new theories, and revising the degrees of belief in the face

of further evidence. In the current paper, we propose a conservative extension of Bayesianism that is able to encompass theory change, while retaining comparative aspects of probabilities that have been computed prior to this change.

1.1 Example: food inspector raising a new hypothesis

Throughout the paper, especially the more technical section 3, it may be helpful to keep in mind a simple example. For this purpose, we offer the following scenario (inspired by an example from Romeijn, 2005):

A food safety inspector wants to determine whether or not a restaurant is taking the legally required precautions against food poisoning. She enters the restaurant anonymously and orders a number of dishes. She uses food testing strips to determine for each of the dishes whether or not they are infected by a particularly harmful strain of Salmonella. She assumes that these tests work perfectly, interpreting a positive test result as a Salmonella-infected dish and a negative result as an uninfected one. She also assumes that in kitchens that implement the precautionary practices each dish has a probability of 1% of being infected, whereas this probability rises to 20% in kitchens that do not implement the practices. She orders five dishes from the kitchen and they all turn out to be infected. This prompts her to consider a third hypothesis: the test strips may have been contaminated, rendering all test results positive, irrespective of whether the dish is infected or not.

After considering this third option, the inspector will not order any additional dishes. Instead, she will take the old evidence (that five dishes out of five appeared to be infected) to confirm the new theory (that the test strips were infected) and it seems reasonable enough for her to do so. Our challenge is how to represent this positive confirmation of the old evidence for the new theory in (an extension of) the Bayesian framework.

1.2 Old evidence and new theories

The confirmation-theoretic model of this paper sheds new light on the problem of old evidence and new theories. This problem for Bayesianism was first identified by Clark Glymour (1980). The problem arises from the discrepancy between descriptive, historical examples, in which old evidence does seem to lend positive confirmation to new theories, and the normative, Bayesian position, in which old evidence cannot confirm new theories. In particular, by updating via Bayes' rule, taking into account evidence that has already been conditioned upon cannot change the probabilities. And since all expressions of confirmation hinge on differences in probabilities, it seems that old evidence cannot lead to confirmation of new theories. Many later authors have called Glymour's problem simply "the problem of old evidence". A minority of philosophers has stressed the importance of the other side of the problem: "the

problem of new theories” (for example Earman, 1992). In what follows, we will clarify that both problems can be resolved in open-minded Bayesianism.

Observe that new theories pose a problem for Bayesianism in itself, even outside the context of confirmation.¹ In fact, without a way of introducing a new theory into the domain of an agent’s degrees of belief, its prior and posterior degrees of belief simply do not show up in the model. In effect, as we will explain in section 3.3.2, those probabilities are set to zero. Either way, for want of a way to express non-zero probability assignments to a new theory, the problem of old evidence does not even occur—or it is worse than the usual presentations suggest. Therefore, we analyze the problem of new theories first and offer a conservative extension of Bayesianism to deal with this problem: a framework for open-minded Bayesianism. In the course of doing so, it will become clear what is missing to deal adequately with old evidence and to determine the confirmation it may give to a new theory. In particular, our model is compatible with Glymour’s observation that in important historical examples old evidence does offer positive confirmation to new theories.

Some proposals for addressing the problem of old evidence (in particular that of Garber, 1983)² observe that the crucial content that is being learned and that lends positive confirmation to a new theory, is not the old evidence itself, but rather the fact that this new theory implies or explains the old evidence. Recently, Sprenger (2014) has proposed a new solution along these lines. We are sympathetic to this approach.³ However, Sprenger’s results presuppose that the old evidence, the new theory and the relevant relation between the two are all elements of some algebra (see his Theorems 1 and 2). As such, this approach does not address a more fundamental question: how can a new theory (or a new relation between a theory and a piece of evidence) be incorporated in the algebra? This is the problem of new theories, which is especially pressing in the presence of old evidence, that we tackle here.

1.3 Bayesian confirmation theory

Since the problem of old evidence and new theories is ultimately a problem concerning Bayesian confirmation, we should first be clear on how we intend to measure confirmation of a hypothesis by a body of evidence. This is in itself an interesting problem in formal epistemology, and some reactions to the problem of old evidence are in fact proposals for a new measure of confirmation (*e.g.*, Christensen, 1999, Joyce, 1999). In qualitative terms, a piece of evidence E lends positive confirmation to a theory T if the posterior $P(T \mid E)$ exceeds the prior $P(T)$. To turn this into a quantitative notion, different measures of

¹ This has often been noted in the literature. See for example Gillies (2001).

² See for instance Easwaran (2011) for a recent overview of approaches to the problem of old evidence.

³ We also agree with Sprenger (2014) that, if we intend to capture objective confirmation in a scientific context, the relevant credence function belongs to an abstract agent representing any unbiased scientist in the relevant context, rather than a particular historical person.

confirmation have been proposed: for instance, the difference or (the log of) the ratio of posterior and prior.

However, our current investigation focuses on how to deal with new theories, which is a problem that besets Bayesianism more broadly, and quite independently of the chosen confirmation measure. Therefore, we will not opt for any such measure, and focus our attention on what they supervene on: the probability assignment over hypotheses themselves. Nothing in our exposition hinges on the precise measure of confirmation that may be grafted onto the probabilistic models.

1.4 The catch-all hypothesis

Our proposal of open-minded Bayesianism relies on the use of a catch-all hypothesis: given a set of explicit hypotheses, we introduce an additional hypothesis that is the negation of the union of the previous hypotheses. The possibility of currently unexplored theoretical alternatives is relevant not only for the formal framework of Bayesian confirmation theory, but also for philosophy of science in general. See for instance the discussion on the pessimistic meta-induction by Sklar (1981), who speaks of “unborn hypotheses”, and by Stanford (2006), who uses the term “unconceived alternatives”.

Shimony (1970) discussed the idea of a catch-all hypothesis in the context of his “tempered personalist” account of probability: he suggested it as a way to represent open-mindedness, which he regarded as a tempering condition to obtain a weakened form of Bayesianism adequate for scientific inference. Shimony (1970, p. 96) suggested not to assign numerical weights (priors) to the catch-all (in contrast to the other hypotheses).

Also Earman (1992) discussed the use of a catch-all to make room for later theory change. According to Earman (1992, p. 196), new theories are “shaven off” from the catch-all hypothesis, which thus “serves as a well for initial probabilities for as yet unborn theories, and the actual introduction of new theories results only in drawing upon this well without disturbing the probabilities of previously formulated theories.” However, he is not satisfied by the proposal of shaving off from a catch-all; according to Earman (1992, pp. 195–196) it leads to the assignment of successively smaller probabilities to later theories (*cf.* Romeijn, 2004), and shaving off does not give an adequate description of scientific revolutions (in the Kuhnian sense) that involve radically new theories. These reservations do not apply to the way in which we formalize the notion of a catch-all, as we will explain in section 4.

1.5 Assigning open-minded probabilities

When Bayesian ideas are applied within the sciences, the domain of the probability function tends to have a small scope: it is used to compare parametric models that apply to a single, well-delineated target system. In philosophy,

however, we often speak as if the domain of the probability function captures every thinkable thought. In particular, in philosophy of science and Bayesian confirmation theory, the probability function assigns values to scientific theories.

If all later changes to the probability assignment are to be due to conditioning, as standard forms of Bayesianism prescribe, we have to be able to specify the domain in such a way as to include all possible scientific theories, including those that are yet to be developed. Nevertheless, it may happen that genuinely new scientific theories do emerge. It is unclear how those can be incorporated in a domain that has to be defined upfront.

In section 2, we will make explicit what the domain of the probability function is on the standard account. Since probability functions assign values to scientific theories as well as to particular pieces of evidence, we have to define a domain that can represent all these objects, even though they are of very different kinds. Specifying this domain provides us with a good opportunity to formalize the notion of the catch-all hypothesis, and how it is used to change the domain of the probability function.

In section 3, we will introduce two forms of open-minded Bayesianism, called vocal and silent, which both employ a catch-all hypothesis. Both are based on the idea that we can remain open-minded about our probabilities by employing sets of probability functions rather than single functions. But both approach the (incomplete) assignment of probabilities in slightly different ways. Also the rule for updating on new theories takes a different form in both contexts.

In section 4, we evaluate and compare the proposals.

2 Bayesianism and the catch-all hypothesis

Upon the introduction of a new theory, the domain of the probability function may change. Before we decide how we will capture this change, let us first specify the domain for the standard form of Bayesianism. We will start this investigation by considering Bayes' theorem. This set-up will also prove fruitful to formalize the notions of hypothesis, evidence, and the catch-all, which prepares us for the subsequent treatment of domain changes and associated changes in probability.

2.1 Domain of the probability function

Bayes' theorem is often presented as follows:

$$P(H | E) = \frac{P(E | H) P(H)}{P(E)},$$

where H is a hypothesis and E is a piece of evidence. But what is P ? This function symbol appears four times in the equation, but can it be interpreted in the same way in all four appearances?⁴

We maintain that the four occurrences of P in Bayes' theorem do refer to the same probability function with the same domain. We take probability to be a one-place function, and we employ the standard definition of conditional probability to make sense of posterior and likelihood. As the common domain, we consider an algebra spanned by the Cartesian product of a set of elementary hypotheses, Θ , and a sample space, Ω (more on these in the following subsections):

$$\mathcal{A}(\Theta \times \Omega).$$

To be precise, we interpret the argument H of the prior and the posterior as shorthand for $H \times \Omega$ and the argument E of the marginal likelihood as shorthand for $\Theta \times E$. The interpretation of these elements of the algebra \mathcal{A} remains as before.

Dynamics: time stamps In Bayesian confirmation theory, we model the rational degrees of belief of an agent by a probability function. To capture the dynamics of the agent's degrees of belief, we consider a succession of probability functions, indexed by a time stamp: P_t is the probability function that represents the rational degrees of belief of the agent at time t . In standard Bayesianism, these belief states are linked by Bayes' rule, as detailed below.

2.2 Evidence and updating

Before we can start applying probability theory, we have to fix a particular sample space (or set of atomic events), Ω , which is chosen such that any result of a measurement can be represented as a subset of Ω . The sample space can be a Cartesian product of sets, which allows us to represent very different types of empirical data.⁵ We represent (actual and hypothetical) pieces of evidence as elements of an algebra on the sample space, $\mathcal{A}(\Omega)$. This set is usually called the event space, but in the current context it is better to call it the 'evidence space'.

⁴ Some introductory texts, such as Hogg (2012, p. 4) and Bertsekas and Tsitsiklis (2008) even argue that we are dealing with four different functions and suggest the use of subscripts to distinguish between them. But we follow a different approach.

⁵ In the discrete case, we may think of the sample space as the set of infinitely long sequences (ranging over temporal instants or individuals) of the values of a property (from a discrete set S) or a vector of properties (each from a discrete set S_i): $\Omega = S^I$, with S the possible values of a certain property or a Cartesian product set of such value sets $S = \prod_i S_i$ and I the infinite index set (*e.g.*, \mathbb{N}); see for instance Romeijn (2011, section 2). Considering the algebra spanned by the cylindrical subsets of this sample space allows us to represent measurements as initial segments of infinitely long streams of data.

Dynamics: Bayes' rule If an agent with probability function $P_{t=n}$ receives evidence E between $t = n$ and $t = n + 1$, then Bayes' rule prescribes that the agent has to adopt a new probability function $P_{t=n+1}$ that is equal to the posterior of the agent's previous probability function: $P_{t=n+1}(\cdot) = P_{t=n}(\cdot | E)$, which can be computed via Bayes' theorem.

2.3 Explicit hypotheses and the catch-all

In the Bayesian framework, probability functions range over evidence and hypotheses. Hence, in addition to specifying Ω , we need to define a set of hypotheses, \mathcal{H} , and an algebra over this set, $\mathcal{A}(\mathcal{H})$. The hypotheses are only specified up to their empirical content. The scientific theories that motivate them are not brought into view. The way to characterize an empirical hypothesis, H , is by specifying a likelihood function $P(\cdot | H)$ ranging over the evidence space, $\mathcal{A}(\Omega)$. Because the empirical content of hypotheses is spelled out in terms of probability functions over the data, the hypotheses are called statistical.⁶

Under a hypothesis we may also subsume an entire family (*i.e.*, a set) of likelihood functions, which have the same form except for a different value of a parameter (or vector of parameters).⁷ Henceforth, we will treat all hypotheses as sets of probability functions on the domain $\mathcal{A}(\Omega)$. Hypotheses that correspond with singleton sets will be called elementary hypotheses, others will be called composite. Observe that the hypotheses in \mathcal{H} need not be elementary in this sense.

Like the elementary events in Ω , the hypotheses in \mathcal{H} need to be mutually exclusive and jointly exhaustive. However, merely exhausting the union of the hypotheses in \mathcal{H} , which is the set of hypotheses that are being considered at a given point in time, may not suffice. In particular, it does not suffice once a new hypothesis emerges, because in that case we want to involve a hypothesis outside $\bigcup_{H \in \mathcal{H}} H$. As indicated before, if we do not offer a domain in which possibilities outside \mathcal{H} can be denoted, we cannot begin to formulate the problem of old evidence and new theories.

Our first and important deviation from what we call 'standard Bayesianism' is that we give the probability function a domain that includes hypotheses outside the set that is currently under consideration. We propose that the hypotheses ought to be mutually exclusive and jointly exhaustive of the vast set of all probability functions on the evidence space $\mathcal{A}(\Omega)$:⁸

$$\Theta = \{P : \mathcal{A}(\Omega) \rightarrow [0, 1] \mid P \text{ is a probability function}\}.$$

⁶ Using this terminology, this article deals with the problem of new hypotheses, rather than the problem of new theories.

⁷ See for instance Romeijn (2011, section 7). In such a case, it is more common to speak of a statistical model or a theory, but we stick to the term 'hypothesis', to avoid confusion with scientific theories.

⁸ It would be more accurate to label the set as $\Theta_{\mathcal{A}(\Omega)}$, but we omit the subscript to keep the notation light.

Then, we can represent an empirical, or statistical, hypothesis as a non-empty set of probability functions on $\mathcal{A}(\Omega)$; hypotheses are thus elements of an algebra on Θ .

Let us consider a collection of $N + 1$ hypotheses (with N a positive integer) that are mutually exclusive and jointly exhaustive: this partition of Θ contains N explicitly formulated hypotheses, H_0, \dots, H_{N-1} , and one catch-all, $\overline{\Theta_N}$. By an ‘explicitly formulated’ hypothesis, H_i , we mean an empirical hypothesis that is produced by a scientific theory. We do not discuss in detail the scientific theories themselves, or even how they lead to statistical hypotheses.⁹

We will denote the set of explicitly formulated hypotheses (previously indicated by \mathcal{H}) by

$$T_N = \{H_i \mid i \in \{0, \dots, N - 1\}\}.$$

T_N represents the ‘theoretical context’ against which hypotheses are being considered. We will denote the union of the hypotheses in T_N by:

$$\Theta_N = \bigcup_{i=0}^{N-1} H_i.$$

Hence, T_N is a partition of Θ_N . Θ_N is the subset of Θ that is currently being covered by some scientific theory. The catch-all, $\overline{\Theta_N}$, is the complement of Θ_N within Θ (so, $T_N \cup \{\overline{\Theta_N}\}$ is a partition of Θ): this hypothesis is the set of all the probability functions that are not produced by any known scientific theory. Whereas the other hypotheses come with a—possibly very intricate—theoretical background story, the catch-all $\overline{\Theta_N}$ has no content other than “none of the explicitly formulated hypotheses”. So, $\overline{\Theta_N}$ is the set $\Theta \setminus \bigcup_{i=0}^{N-1} H_i$ and that is all that can be said about it. In the same vein, we cannot say anything about the probabilities that the catch-all hypothesis assigns to the evidence.

Dynamics: shaving off In the previous subsection, we have seen that the incorporation of evidence leads to an update of the probability function governed by Bayes’ rule. Standard Bayesianism lacks an analogous procedure for revising the probability function in light of a new hypothesis. We will now discuss how the presence of the catch-all allows us to represent the dynamics of the set of hypotheses. This prepares us for the proposal of open-minded Bayesianism in the next section.

After a new scientific theory has been developed, the statistical hypothesis it produces may be added to the partition of Θ by “shaving off” from the catch-all (by the terminology of Earman, 1992, p. 196). At this point in time, the former catch-all may be decomposed into an additional explicitly

⁹ It is clear that an indeterministic theory can generate statistical predictions about measurable quantities. In the case of deterministic theories, such as Newtonian mechanics, it may be less clear how they lead to hypotheses that are expressed in terms of a probability assignment. However, when we combine such a theory with measured values for masses, velocities, *etc.* the associated measurement uncertainty can be represented in terms of probability distributions, which in turn leads to statistical predictions concerning other measurable quantities.

formulated hypothesis H_N (disjoint from the earlier hypotheses) and a new (smaller) catch-all, $\overline{\Theta}_{N+1}$. So, the algebra on $\Theta \times \Omega$ changes. Typically, this will happen because the evidence obtains a very low probability according to the hypotheses currently under consideration, *i.e.*, $P(E|H_i)$ is very small for every i . A principled decision to introduce a new theory may be based on the computation of a model score, or on the application of a model selection tool. But such scores and tools fall outside the scope of the present paper. The procedure for deciding to introduce a new theory is not intended to be a part of our model.

2.4 Summary of key ideas

We briefly recapitulate our approach so far and our use of the following terms: scientific theory, statistical hypothesis, sample space, evidence, and catch-all.

A scientific theory together with background assumptions produces an empirical, or statistical, hypothesis. (How this happens requires engaging with the details of a scientific theory, which falls outside the scope of our current framework.) Such an empirical or statistical hypothesis is a set, possibly a singleton, of probability functions. In order to compare hypotheses produced by different theories in the light of a common body of empirical data (and thus to compare their measures of confirmation or evidential support), their probability functions need to have a common domain. This domain is called the evidence space: it is an algebra on a sample space (which may be a Cartesian product set to allow for the representation of mutually independent measurable quantities).

The union of all statistical hypotheses produced by the currently available scientific theories (Θ_N) does not exhaust the set of *all* probability functions on the evidence space (Θ). The complement of the former set relative to the latter set is called the catch-all hypothesis ($\overline{\Theta}_N$): unlike the other hypotheses, it is not produced by a scientific theory, but rather it results from a meta-theory. The catch-all hypothesis is included to express that many other hypotheses are conceivable, each associated with a probability assignment or a set of such assignments over the evidence.

With the idea of a catch-all hypothesis in place, we can now turn to a full specification of open-minded Bayesianism. The inclusion of a catch-all hypothesis makes room for modeling the introduction of new hypotheses, namely by shaving them off from the catch-all. But this in itself is not sufficient: we still need to specify how shaving off influences probability assignments over the hypotheses. This is the task undertaken in the next section.

3 Open-minded probability assignments

In the previous section, we have introduced the formal framework of open-minded Bayesianism. It is a form of Bayesianism that requires the set of hypotheses to include a catch-all hypothesis. In the current section, we develop

the probability kinematics for open-minded Bayesianism. Two versions will be considered: vocal and silent. The two approaches suggest slightly different rules for revising probability functions upon theory change.¹⁰

3.1 Vocal and silent open-mindedness

In open-minded Bayesianism, hypotheses are represented as sets of probability functions. If prior probabilities are assigned to the functions within a set, then a single marginal probability function can be associated with the set. But without such a prior probability assignment within the set, the set specifies so-called imprecise probabilities (see, for instance, Walley, 2000).

We first clarify probability assignments over explicitly formulated hypotheses. In standard Bayesianism, prior probabilities are assigned to the hypotheses, which are all explicitly formulated. We can furthermore assign priors over the individual probability functions contained within composite hypotheses, if there are any. We call such a prior within a composite hypothesis a sub-prior. The use of sub-priors leads to a marginal likelihood function for the composite hypothesis.¹¹ Upon the receipt of evidence we can update all these priors, *i.e.*, those over elementary and composite hypotheses as well as those within composite hypotheses.

Now recall that in open-minded Bayesianism, the space of hypotheses also contains a catch-all, which is a composite hypothesis encompassing all statistical hypotheses that are not explicitly specified. In standard Bayesianism, this catch-all hypothesis is usually not mentioned, and all probability mass is concentrated on the hypotheses that are formulated explicitly. Within the framework of open-minded Bayesianism, we will represent this standard form of Bayesianism by setting the prior of the catch-all hypothesis to zero.¹²

Let us turn to open-minded Bayesianism itself. To express that we are prepared to revise our theoretical background, we assign a strictly positive prior, and perhaps a sub-prior, to the catch-all. However, we agree with Shimony (1970) that it is not sensible to assign *any* definite value to the prior of the catch-all. Since the catch-all is not based on a scientific theory, the usual “arational” considerations (to employ the terminology of Earman, 1992, p. 197) for assigning it a prior, namely by comparing it to hypotheses produced by other theories, do not come into play here. Moreover, it seems clear that the catch-all should give rise to imprecise marginal likelihoods as well, which suggests that we should refrain from assigning sub-priors to its constituents,

¹⁰ See Morey et al (2013) for a less rigorous exposition of open-minded Bayesianism, which they term humble Bayesianism, in a statistics context. The idea of open-mindedness is already present in what Lindley (1991, p. 104) called Cromwell’s Rule.

¹¹ In statistics this is known as hierarchical modeling (*cf.* Gelman et al, 2004). A useful philosophical angle on this is provided in Henderson et al (2010).

¹² If the option of a catch-all simply hasn’t been considered, one might intuitively expect its probability to be undefined rather than zero. However, if we represent Bayesianism without a catch-all within an open-minded framework, a probability has to be assigned to the catch-all and its value has to be zero: see section 3.3.2.

too. (Recall that the algebra on $\Theta \times \Omega$ cannot pick out any strict subset of the catch-all.) These considerations lead us to consider two closely connected forms of open-minded Bayesianism, which both avoid assigning a definite prior to the catch-all:

Vocal open-minded Bayesianism assigns an **indefinite prior and likelihood** to the catch-all hypothesis, $\overline{\Theta_N}$. We represent its prior by $\tau_N \in]0, 1[$ and its likelihood by $x_N(\cdot | E)$. To ensure normalization over all hypotheses (including the catch-all), the priors assigned to the explicitly formulated hypotheses are set equal to the value they would have in a model without a catch-all now multiplied by $(1 - \tau_N)$.

Silent open-minded Bayesianism assigns **no prior or likelihood** to the catch-all hypothesis, not even symbolically. To achieve this, all probabilistic statements are conditionalized on the algebra on Θ_N (shorthand for $\Theta_N \times \Omega$). Θ_N represents the union of the hypotheses in the current theoretical context. From the viewpoint of the algebra on $\Theta \times \Omega$, the probability assignments will be incomplete.

In both cases, we deviate from the standard Bayesian account in that we give a strictly positive prior to the catch-all, and then allow opinions to be **partially unspecified**: vocal open-minded Bayesianism retains the entire algebra but uses symbols without numerical evaluation as placeholders, whereas silent open-minded Bayesianism restricts the algebra to which probabilities are assigned (leaving out the catch-all). Formally, the partial specification of a probability function comes down to specifying the epistemic state of the agent by means of a set of probability assignments (*cf.* Halpern, 2003, Haenni et al, 2003).

3.2 A conservative extension of standard Bayesianism

As detailed in the foregoing, we aim to represent probability assignments of an agent that change over time. An agent's probability function therefore receives a time stamp t . We often talk about this as if the probability function changes over time, but it is more accurate to say that the entire probability function gets replaced by a different probability function at certain points in time. Accordingly, subsequent functions need not even have the same domain.

Standard Bayesianism has one way to replace an agent's probability function once the agent learns a new piece of evidence: Bayes' rule. It amounts to *restricting* the algebra to those sets that intersect with the evidence just obtained. Equivalently, it amounts to setting all the probability assignments outside this domain to zero. If at time t an agent learns evidence E with certainty, Bayes' rule amounts to setting $P_{t=n}$ equal to $P_{t=n-1}(\cdot | E)$. If E is the first piece of evidence that the agent learns, this amounts to restricting the domain from an algebra on $\Theta \times \Omega$ to an algebra on $\Theta \times E$ and redistributing the probability over the remaining parts of the algebra according to Bayes' theorem.

In addition to this, open-minded Bayesianism requires a rule for replacing an agent's probability function once the agent learns information of a different kind: the introduction of a new hypothesis. This amounts to *expanding* the algebra to which explicit probability values are assigned (from an algebra on $\Theta_N \times E$ to an algebra on $\Theta_{N+1} \times E$). Or in other words, it amounts to *refining* the algebra on $\Theta \times E$. On both views, the new algebra is larger (*i.e.*, it contains more sets). What is still missing from our framework is a principle for determining the probability over the larger algebra. In analogy with Bayes' rule, one natural conservativity constraint is that the new probability distribution must respect the old distribution on the preexisting parts of the algebra.

Viewed in this way, our proposal does not introduce any radical departure from standard Bayesianism. Open-minded Bayesianism respects Bayes' rule, but this rule already concerns changes in the algebra, namely reductions. The only new part is that we require a separate rule for enlarging the algebra (extending Θ_N or refining the partition of Θ) rather than for reducing it (restricting Ω). The principle that governs this change of the algebra again satisfies conservativity constraints akin to Bayes' rule. As detailed below, silent and vocal open-minded Bayesianism will give a slightly different rendering of this rule.

3.3 Updating due to a new hypothesis

In this section, we consider how the probability function ought to change upon the introduction of a new hypothesis after some evidence has been gathered. We first consider an abstract formulation of a reduction and extension of the domain, as well as an example of such an episode in the life of our epistemic agent. After that we consider both versions of open-minded Bayesianism as developments of the standard Bayesian account.

3.3.1 Reducing and enlarging: setting the stage

The epistemic episode that we aim to model has three stages:

- ($t = 0$) *N* explicit hypotheses At time $t = 0$, the theoretical context of the agent consists of N explicit hypotheses: $T_N = \{H_0, \dots, H_{N-1}\}$. The union of the hypotheses in T_N is Θ_N . The catch-all is the complement of the latter (within Θ): $\overline{\Theta_N}$.
- ($t = 1$) Evidence E At time $t = 1$, the agent receives evidence, E . The initial likelihood of obtaining this evidence given any one of the hypotheses H_i ($i \in \{0, \dots, N - 1\}$) is a particular value

$$P_{t=0}(E \mid H_i).$$

- ($t = 2$) New hypothesis H_N At time $t = 2$, a new scientific theory is introduced, which produces a statistical hypothesis that is a subset of $\overline{\Theta_N}$: call this additional hypothesis H_N . The new set of explicit hypotheses is thus

$T_{N+1} = \{H_0, \dots, H_{N-1}, H_N\}$. The union of the hypotheses in T_{N+1} is $\Theta_{N+1} \supset \Theta_N$. The new catch-all is the complement of Θ_{N+1} : $\overline{\Theta_{N+1}} \subset \overline{\Theta_N}$. In other words: in the algebra on Θ , the old catch-all $\overline{\Theta_N}$ is replaced by two disjoint parts, H_N and $\overline{\Theta_{N+1}}$. The new explicit hypothesis H_N is shaven-off from the old catch-all, $\overline{\Theta_N}$, leaving us with a smaller new catch-all, $\overline{\Theta_{N+1}}$.

Our first question is how the agent ought to revise her probability assignments at $t = 2$. The second question is whether the old evidence (E obtained at $t = 1$) can lend positive confirmation to the new hypothesis (H_N formulated at $t = 2$). We will consider these questions in the context of standard Bayesianism and both forms of open-minded Bayesianism. As will be seen, the probability assignments that result from open-minded Bayesianism will show the relevant similarities with those of standard Bayesianism: within Θ_N both have the same proportions among the probabilities for the hypotheses H_i .

Food inspection example While reading our general treatment of the three stages, it may be helpful to keep in mind the example of section 1.1. In this example, the number of explicit hypotheses is $N = 2$. The hypotheses H_0 (meaning, informally, “the kitchen is clean”) and H_1 (“this kitchen is not clean”) can be made formal in the following way: the distribution of infections follows a binomial distribution with bias parameter $p_0 = 0.01$ (H_0) or with bias parameter $p_1 = 0.2$ (H_1). The sample space is the same for both hypotheses: $\Omega = \{0, 1\}^{\mathbb{N}}$, where 0 means that a dish tested as negative and 1 means that a dish tested as positive. In this case, the evidence takes the form of initial segments of the sequences in the sample space (cylindrical sets of $\{0, 1\}^{\mathbb{N}}$).¹³ At $t = 1$, the inspector tests five dishes and receives as evidence an initial segment of five times ‘1’. The initial likelihood of obtaining this evidence E given hypothesis H_0 is

$$P_{t=0}(E | H_0) = p_0^5 = 10^{-10},$$

and given hypothesis H_1 the initial likelihood of the evidence is

$$P_{t=0}(E | H_1) = p_1^5 = 3.2 \times 10^{-4}.$$

At $t = 2$, the inspector considers a new hypothesis, H_2 , which can be modeled as a binomial distribution with $p_2 = 1$.

3.3.2 No update rule for standard Bayesianism

Standard Bayesianism works on a fixed algebra on a fixed set $\Theta_N \times \Omega$. On this view, none of the probabilities can change due to hypotheses that are external to Θ_N .

¹³ Since the inspector assumes that the test is perfect, instead of representing the test results, she may just as well represent these data in terms of dishes being infected or not (such that 0 means that a dish is not infected and 1 that a dish is infected.) This illustrates how data and evidence may come apart: we regard evidence as interpreted data, where the interpretation depends on the sample space that is used in a hypothesis. We do not pursue this complication here.

($t = 0$) *N explicit hypotheses* Each explicit hypothesis receives a prior probability, $P_{t=0}(H_i)$. If we assume that, initially, the agent is completely undecided with regard to the N hypotheses, she will assign equal priors to them: $P_{t=0}(H_i) = 1/N$ (for all $i \in \{0, \dots, N - 1\}$).¹⁴

($t = 1$) *Evidence E* The marginal likelihood of the evidence can be obtained via the law of total probability:

$$P_{t=0}(E) = \sum_{j=0}^{N-1} P_{t=0}(H_j) P_{t=0}(E | H_j),$$

which is about 1.6×10^{-4} for the example. The posterior probability of each hypothesis given the evidence can be obtained by Bayes' theorem:

$$P_{t=0}(H_i | E) = \frac{P_{t=0}(H_i) P_{t=0}(E | H_i)}{P_{t=0}(E)} \text{ (for all } i \in \{0, \dots, N - 1\}\text{)}.$$

In the example, this is about 3.1×10^{-7} for H_0 and $1 - 3.1 \times 10^{-7}$ for H_1 . According to Bayes' rule, upon receiving the evidence E , the agent should replace her probability function by $P_{t=1} = P_{t=0}(\cdot | E)$. In the example, the inspector should now assign a probability to H_1 that is more than three million times higher than the probability she assigns to H_0 .

In the example, the confirmation is positive for H_1 and negative for H_0 .

($t = 2$) *New hypothesis H_N* Suppose a new hypothesis is formulated: some $H_N \in \Theta_N$. In terms of the example: the inspector was in a situation in which she could have received evidence with a much higher initial probability than that of the evidence she actually received, and we might imagine that this makes her decide to take the hypothesis H_2 on infected test strips into consideration. Now since, in general, the new hypothesis H_N is not a part of the theoretical context, T_N , *the intersection of H_N with Θ_N is empty*. Hence, the probability assigned to H_N is zero, simply because $P(T_N) = 0$. And because the prior of this hypothesis is zero, the confirmation of this hypothesis is zero as well. In other words, standard Bayesianism simply does not allow us to represent new hypotheses (other than by the empty set). In this sense, the problem of old evidence does not even occur: new theories cannot be taken into account.

3.3.3 Update rule for vocal open-minded Bayesianism

This version of Bayesianism employs a refinable algebra on a fixed set $\Theta \times \Omega$. In this view, none of the previous probability assignments change upon theory change, but new probabilities can be expressed and earlier expressions can be rewritten accordingly.

¹⁴ The assumption of equal priors is not essential for the framework. The agent may assign different priors, based on considerations that are external to the Bayesian framework, such as relevant base rates (where the usual reference class problem emerges; cf. Hájek, 2007).

($t = 0$) *N explicit hypotheses* Each explicit hypothesis receives a prior, $P_{t=0}(H_i)$ (and, where appropriate, sub-priors). The proposal of vocal open-mindedness is to assign an undefined prior, $\tau_N \in (0, 1)$, to the catch-all hypothesis, $\overline{\Theta}_N$:

$$P_{t=0}(\overline{\Theta}_N) = \tau_N.$$

No subsets of the catch-all receive (sub-)priors at $t = 0$, but certain subsets of the catch-all will receive a prior later on. To ensure normalization over all hypotheses (including the catch-all), the priors assigned to the explicitly formulated hypotheses are set equal to the value they had in the model without a catch-all now multiplied by $(1 - \tau_N)$; for each $i \in \{0, \dots, N - 1\}$:

$$P_{t=0}(H_i) = (1 - \tau_N) P_{t=0}(H_i | \Theta_N).$$

Although the value of τ_N is unknown, the $N + 1$ priors sum to unity. In the example, we have as prior of the catch-all $P_{t=0}(\overline{\Theta}_2) = \tau_2$ and as prior for the two explicit hypotheses $P_{t=0}(H_0) = 1/2 \times (1 - \tau_2) = P_{t=0}(H_1)$.

The likelihood functions of the explicit hypotheses H_i are the same as in the usual model. Regarding the likelihood of the catch-all, the proposal is to represent it by an undefined weighted average of functions in $\Theta \setminus \Theta_N$: $P_{t=0}(\cdot | \overline{\Theta}_N) = x_N(\cdot)$.

($t = 1$) *Evidence E* The marginal likelihood of the evidence has an additional term as compared to the standard model:

$$P_{t=0}(E) = \sum_{j=0}^{N-1} P_{t=0}(H_j) P_{t=0}(E | H_j) + \tau_N x_N(E).$$

Due to the presence of undetermined factors associated with the catch-all, $P_{t=0}(E)$ cannot be evaluated numerically. As a result, also the updated probability function, $P_{t=1}(\cdot) = P_{t=0}(\cdot | E)$, contains unknown factors. These are the posteriors for H_i (for all $i \in \{0, \dots, N - 1\}$):

$$\begin{aligned} P_{t=0}(H_i | E) &= \frac{P_{t=0}(H_i) P_{t=0}(E | H_i)}{P_{t=0}(E)} \\ &= \frac{(1 - \tau_N) P_{t=0}(H_i | \Theta_N) P_{t=0}(E | H_i)}{\sum_{j=0}^{N-1} (1 - \tau_N) P_{t=0}(H_j | \Theta_N) P_{t=0}(E | H_j) + \tau_N x_N(E)}. \end{aligned}$$

Although this expression cannot be evaluated numerically, some comparative probability evaluations can be computed since the unknown factors cancel. In particular, the ratio of two posterior probabilities assigned to explicit hypotheses can still be obtained; for $i, j \in \{0, \dots, N - 1\}$:

$$\frac{P_{t=1}(H_i)}{P_{t=1}(H_j)} = \frac{P_{t=0}(H_i | \Theta_N) P_{t=0}(E | H_i)}{P_{t=0}(H_j | \Theta_N) P_{t=0}(E | H_j)}.$$

In the example, it can still be established that after receiving evidence E the inspector should assign a probability to H_1 that is more than three million times higher than the probability she assigns to H_0 . Similarly, we can still establish that both hypotheses have a very small likelihood for the evidence

that is obtained. And this may be enough to motivate the introduction of a new hypothesis.

In the context of vocal open-mindedness, any expression of the belief change will contain unknown factors, and the implications are worse than for the posteriors: if the change is measured as the difference between posterior and prior, both terms have different unknown factors ($\frac{1-\tau_N}{P_{t=0}(E)}$ and $1 - \tau_N$, respectively).

($t = 2$) *New hypothesis H_N* Recall that the old catch-all $\overline{\Theta_N}$ is replaced by two disjoint parts: the hypothesis that is shaven off, H_N , and the remaining part of the catch-all Θ_{N+1} . Finite additivity suggests to decompose the prior that was assigned to $\overline{\Theta_N}$ into two corresponding terms:

$$\tau_N = P_{t=0}(H_N) + \tau_{N+1},$$

where $P_{t=0}(H_N)$ is the prior of the new hypothesis H_N and $\tau_{N+1} \in]0, \tau_N[$ is the (indefinite) prior of the remaining catch-all Θ_{N+1} , both of which are assigned retroactively. Although the value of τ_{N+1} is unknown, the $N + 2$ priors sum to unity.

The priors for the hypotheses in T_N can thence be written in three ways:

$$\begin{aligned} P_{t=0}(H_i) &= (1 - \tau_N) P_{t=0}(H_i | \Theta_N) \\ &= (1 - \tau_{N+1}) P_{t=0}(H_i | \Theta_{N+1}) \\ &= (1 - \tau_{N+1}) (1 - P_{t=0}(H_N | \Theta_{N+1})) P_{t=0}(H_i | \Theta_N), \end{aligned}$$

where $P_{t=0}(H_N | \Theta_{N+1})$ is some definite number $\in]0, \tau_N[$. For instance, if we had a uniform prior over T_N and we want to keep a uniform prior over T_{N+1} , we have to set $P_{t=0}(H_N | \Theta_{N+1}) = \frac{1}{N+1}$.

Now that H_N is an explicit hypothesis, its likelihood is a definite function $P_{t=0}(\cdot | H_N)$ (also specified retroactively). In the example, the likelihood for obtaining the evidence $P_{t=0}(E_{1(5)} | H_2)$ is 1 on the new hypothesis. We assign an undefined likelihood to the new catch-all: $P_{t=0}(\cdot | \overline{\Theta_{N+1}}) = x_{N+1}(\cdot)$. This allows us to rewrite the previous expression obtained for the marginal likelihood:

$$\begin{aligned} P(E) &= \sum_{j=0}^{N-1} (1 - \tau_{N+1}) (1 - P_{t=0}(H_N | \Theta_{N+1})) P_{t=0}(H_j | \Theta_N) P_{t=0}(E | H_j) \\ &\quad + P_{t=0}(H_N) P_{t=0}(E | H_N) + \tau_{N+1} x_{N+1}(E), \end{aligned}$$

where the last two terms equal $\tau_N x_N(E)$.

At this point, we can also rewrite the expressions for the posteriors (for all $i \in \{0, \dots, N - 1\}$):

$$P_{t=2}(H_i) = \frac{(1 - \tau_{N+1}) (1 - P_{t=0}(H_N | \Theta_{N+1})) P_{t=0}(H_i | \Theta_N) P_{t=0}(E | H_i)}{P(E)}.$$

Moreover, we can now assign a posterior to H_N :

$$P_{t=2}(H_N) = \frac{(1 - \tau_{N+1}) P_{t=0}(H_N | \Theta_{N+1}) P_{t=0}(E | H_N)}{P(E)}.$$

Although it is still not possible to evaluate these posteriors numerically, we can compute new comparative probability evaluations for ratios involving H_N . For all $i \in \{0, \dots, N-1\}$:

$$\frac{P_{t=1}(H_N)}{P_{t=1}(H_i)} = \frac{P_{t=0}(H_N | \Theta_{N+1}) P_{t=0}(E | H_N)}{(1 - P_{t=0}(H_N | \Theta_{N+1})) P_{t=0}(H_i | \Theta_N) P_{t=0}(E | H_j)}.$$

In the case of uniform priors, additional factors cancel:¹⁵

$$\begin{aligned} (1 - P_{t=0}(H_N | \Theta_{N+1})) P_{t=0}(H_i | \Theta_N) &= \left(1 - \frac{1}{N+1}\right) \frac{1}{N} \\ &= \frac{1}{N+1} \\ &= P_{t=0}(H_N | \Theta_{N+1}). \end{aligned}$$

For the example, we can compute $\frac{P_{t=2}(H_2)}{P_{t=2}(H_1)} = \frac{1}{p_0^5} = \frac{1}{3.2 \times 10^{-4}} = 3125$. So, in the new theoretical context (T_3) the posterior of the new hypothesis (H_2) given the old evidence E , namely the sequence of five positive tests, is more than three thousand times higher than that of the hypothesis that was best confirmed (H_1) within the old theoretical context (T_2).¹⁶

At $t = 1$, no belief can be expressed for H_N , but at $t = 2$ the beliefs regarding H_N at $t = 1$ can be expressed and the expressions for the old hypotheses H_i can be rewritten. We are still left with two terms that have different unknown factors, which do not simply cancel out.¹⁷ At any rate, degrees of confirmation can be evaluated if we first condition the probability assignments on the current theoretical context, Θ_N . We return to this point below.

3.3.4 Update rule for silent open-minded Bayesianism

This form of Bayesianism employs an algebra on a set $\Theta_N \times \Omega$, which may be extended to $\Theta_{N+1} \times \Omega$ (and beyond). On this view, when the theoretical context changes, new conditional probabilities become relevant to the agent.

Let us briefly motivate the silent version as an alternative to vocal open-mindedness. We have seen that the vocal version comes with a heavy notational load. Given that in the end, we can only compute comparative probabilities, it seems desirable to dispense with the symbolic assignment of a prior and a likelihood to the catch-all hypothesis. Silent open-mindedness achieves this by conditioning all evaluations on Θ_N , the union of the hypotheses in the theoretical context. This allows us to express the agent's opinions concerning the relative probability of H_i and H_j (for any $i, j \in \{0, \dots, N-1\}$) without

¹⁵ Since these factors are all known at $t = 2$, it is not a problem if they do not cancel.

¹⁶ Observe that the catch-all $\overline{\Theta}_2$ is strictly larger than the family of binomial distributions with $p \in [0, 1] \setminus \{0.01, 0.2\}$. The binomial distribution only applies to situations that can be thought of as having a fixed bias and producing independent outcomes. The catch-all should be large enough to allow the agent to also reconsider these assumptions at a later point in time.

¹⁷ This may be a reason to consider a particular measure of confirmation, such as $P(H | E) - P(H | \overline{E})$ (cf. Christensen, 1999, Joyce, 1999), for which the factors do cancel out.

saying anything, not even something in terms of free parameters, about the absolute probability that they have. Opinions about the theories in the current theoretical context T_N are thus comparative only.

(t = 0) N explicit hypotheses Instead of assigning absolute priors to $P_{t=0}(H_i) = P_{t=0}(H_i | \Theta)$, humble Bayesianism suggests to only assign priors that are conditionalized on the theoretical context, $P_{t=0}(H_i | \Theta_N)$.

(t = 1) Evidence E Since $H_i \subseteq \Theta_N$, the likelihoods of explicit hypotheses are statistically independent of the theoretical context:

$$P_{t=0}(E|H_i \cap \Theta_N) = P_{t=0}(E|H_i).$$

Silent open-mindedness suggests not to assign a likelihood to the catch-all. This “probability gap” is not problematic (by the terminology of Hájek, 2003), since all the other probability assignments are conditionalized on Θ_N . The agent can update her comparative opinion in the usual Bayesian way, as long as she conditionalizes everything on this context:¹⁸

$$P_{t=1}(H_i | \Theta_N) = P_{t=0}(H_i|E \cap \Theta_N) = P_{t=0}(H_i|\Theta_N) \frac{P_{t=0}(E|H_i)}{P_0(E|\Theta_N)}.$$

(t = 2) New hypothesis H_N After a new hypothesis has been introduced, the silently open-minded Bayesian has to start conditionalizing on the expanded (union of the) theoretical context Θ_{N+1} rather than on Θ_N . Once H_N gets formulated, its likelihood will be known too. We require that the probability evaluations conditional on the old context Θ_N do not change. In this way we cohere with standard Bayesianism and with the vocal open-minded variant.

We can treat $P_{t=2}(H_N | \Theta_{N+1})$ much like a ‘postponed prior’, and give it a value based on arational considerations that are not captured by constraints within the (extended) Bayesian framework. In particular, we can engage in the kind of reconstructive work as is done in vocal open-mindedness, but this is not mandatory here. We might also determine the posterior probability of H_N and so reverse engineer what the prior must have been to make this posterior come out after the occurrence of E . In any case, when moving to a new context the other posteriors need to be changed accordingly (such that the $N+1$ posteriors sum to unity): $P_{t=2}(H_i | \Theta_{N+1}) = (1 - P_{t=2}(H_N | \Theta_{N+1}))P_{t=1}(H_i | \Theta_N)$. So, the move from T_N to T_{N+1} essentially amounts to a kind of recalibration of the posteriors.

Importantly, we can compute all known confirmation measures using the priors and posteriors that are conditional on a particular theoretical context. Once the context changes, this clearly impacts on the confirmation allotted to the respective hypotheses. The price for this transparency is of course that we can only establish the confirmation of a hypothesis relative to a theoretical context Θ_N . The natural use of a degree of confirmation thus becomes

¹⁸ Recall from section 2 that we interpret E as shorthand for $\Theta \times E$, so $E \cap \Theta_N$ should be understood as $(\Theta \cap \Theta_N) \times E = \Theta_N \times E$.

comparative: it tells us which hypothesis among the currently available ones is best supported by the evidence, but there is no attempt to offer an absolute indication of this support.

4 Evaluation and conclusion

In this section we critically evaluate open-minded Bayesianism. We clarify our views on it, and conclude that it provides a handle on the problem of old evidence: it explains how old evidence can be used afresh without violating Bayesian coherence norms. Towards the end, we sketch a number of ideas and problems that deserve further exploration.

4.1 Evaluation of open-minded Bayesianism

It may be argued that open-minded Bayesianism fails to provide us with the required normative guidance. In the silent version, it only concerns suppositional reasoning and hence cannot inform our unconditional beliefs. In metaphorical terms, the worry is that the agent cannot keep hiding behind the conditionalization stroke. In the vocal form, the same worry arises in relation to the use of factors with indefinite numerical values, which are introduced to represent the prior and likelihood of the catch-all hypothesis, but which soon ‘infect’ all probability assignments and measures of confirmation. Either way, it may seem that the agent must come clean on her absolute commitments at some point.

We respond to this worry by biting the bullet. If we want to allow new theories to enter the conceptual scene, then we will have to provide room for this in our formal framework. There are attempts to accommodate (other forms of) theory change in a Bayesian framework that employ fully specified probability assignments (*e.g.*, Romeijn, 2004, 2005). In this paper, by contrast, we have offered a framework that creates room for new theories by leaving part of the probability assignment unspecified. We accept that this leads to a model that only concerns conditional belief.

We should emphasize that we are not alone in preaching an open-minded form of Bayesianism. We already mentioned the proposal for tempered Bayesianism by Shimony (1970), who suggested the use of a catch-all hypothesis in this context. This suggestion was also discussed by Earman (1992), who introduced the evocative terminology of shaving off new hypotheses from the catch-all. Furthermore, our proposal of humble Bayesianism is related to earlier work by Salmon (1990) and Lindley (1991). Morey et al (2013) recently introduced what they call humble Bayesianism in a debate over the nature of statistical model comparisons.

The latter paper lends further supports to open-minded Bayesianism. The point of Morey et al (2013) is that an agent may want to use Bayesian methods to evaluate statistical models, without buying into the conviction, implicit in

the standard Bayesian framework, that one of the theories under consideration is true. After all, a standard Bayesian will have the probabilities of the hypotheses under consideration add up to one, and so judges herself to be perfectly calibrated (*cf.* Dawid, 1982). The standard Bayesian is overly confident, hence a more open-minded form of Bayesianism seems called for.

The price to pay is that the epistemic attitudes for which the framework of the open-minded Bayesian provides the norms are different: they have a conditional nature. Whether we spell out the details using a vocal or a silent open-mindedness, the normative framework tells the agent what to believe only if she temporarily supposes, without committing to it, that the true theory is among those currently under consideration.

4.2 The old evidence problem

Now that we have bitten the bullet, we better make sure that we do so for good reasons. In this section, we argue that open-minded Bayesianism provides a new handle on the problem of old evidence, by explaining how old evidence can be re-used.

In his encyclopedia entry on Bayesianism, Talbott (2008) summarizes the Bayesian problem of new theories as follows: “Suppose that there is one theory H_1 that is generally regarded as highly confirmed by the available evidence E . It is possible that simply the introduction of an alternative theory H_2 can lead to an erosion of H_1 ’s support. [...] This sort of change cannot be explained by conditionalization.” It is precisely this “erosion” of support that can be captured by the update rule for open-minded Bayesianism, since both approaches make the agent reconsider the posteriors of the old hypotheses. The strong point of open-minded Bayesianism is that this reconsideration of the posteriors does not render the agent probabilistically incoherent.

When writing about the operation of shaving off new hypotheses, Earman (1992, p. 195) worried that a point may be reached “where the new theory has such a low initial probability as to stand not much of a fighting chance.” This worry, however, does not apply to our framework. Notice that we do not assign an explicit value to the prior of the current theoretical context. We may think of the prior associated with the catch-all hypothesis as a number extremely close to unity—and the humbler we are, the closer to unity we can imagine it to be. At any rate, no matter how large the discrepancy between the posteriors of the old hypotheses and the new one, the impact that the decomposition of the catch-all has on the catch-all’s posterior will remain unknown, or indefinite. Of course, once we pin down a value for the probability τ_N , the worry of Earman becomes a live one. But lacking such a definite value, the problem that the catch-all gets crowded out by explicit hypotheses does not arise.

There are, however, differences in how the vocal and silent approaches to open-minded Bayesianism deal with reassessing the posteriors, and in what role they give to old evidence. The vocal approach requires us to assign a prior to the new hypothesis H_N after the fact, and to compute its current

posterior on the basis of this assignment. The other posteriors are obtained via a renormalization.¹⁹ This approach requires us to evaluate probabilities retroactively: priors have to be set *post hoc*, for hypotheses that were not known at the time. To our mind this need not be a cause of concern though. One cannot unlearn the evidence that has been gathered, but it is still possible to use base rates or other sources of objective information to determine the priors retroactively.²⁰

The silent approach, by contrast, requires us to assign a posterior to the new hypothesis H_N without offering an explicit recourse to the prior probability assignments over the old hypotheses. The point here is rather subtle. It is in virtue of a prior probability assignment of τ_N to the old catch-all Θ_N that we can meaningfully claim, as part of the vocal approach, that the prior of the old catch-all is decomposed into the prior of a new hypothesis H_N and the prior of a new catch-all Θ_{N+1} . Since the silent approach remains silent precisely on this prior, it is hard to see how we can retroactively decompose it. So in this approach, it is not clear whether old evidence ever confirms new theories. Unless we have set the value of $P_{t=2}(H_N | \Theta_{N+1})$ by means of a reconstruction that ultimately depends on $P_{t=0}(H_N | \Theta_{N+1})$, its value is not obtained via conditionalization on E . In humble Bayesianism, the old evidence is therefore not given a new role.

Now that we have discussed the role of evidence in two forms of open-minded Bayesianism, it is time to take stock. Both approaches suffer some drawbacks. The vocal proposal comes with the complication of a heavy notational load that hampers the evaluation of the degree of confirmation. The silent proposal allows too much freedom in the assignment of a posterior to the new hypothesis—so much freedom, that it is not clear that the old evidence has any impact. For these reasons, we propose a hybrid approach to open-minded Bayesianism, that combines the best elements of both.

On our hybrid proposal, the open-minded Bayesian remains in the silent phase,²¹ except for the times at which her theoretical context changes. Unlike a standard Bayesian, the open-minded Bayesian is allowed to change the algebra to which probabilities are assigned and thus to assign non-zero probabilities to the new hypothesis, which is impossible without a catch-all. Then she enters the vocal phase: she engages in assigning a prior to the new hy-

¹⁹ More accurately, the decomposition into definite and indefinite factors changes in a way that is reminiscent of a renormalization.

²⁰ Vocal open-minded Bayesianism can be compared with the analysis of the problem of old evidence given by both Garber (1983) and Jeffrey (1983), who concluded that what is discovered is the fact that the new theory entails the old evidence. To model agents who discover a statement of this kind, they proposed weakening the Bayesian background assumption of logical omniscience. The vocal approach paints a similar, reconstructive picture, though it is not logical omniscience that fails the agent: what is discovered upon the change in the algebra at $t = 2$, is how to express the posterior (and hence the confirmation) of the new hypothesis given the old evidence, which was inexpressible at $t = 1$.

²¹ We might call the silently open-minded Bayesian a relativized standard Bayesian: the probabilities conditionalized on the theoretical context appearing in the humble approach equal the corresponding unconditional probabilities of the approach without a catch-all.

pothesis (retroactively) and computing its posterior given the evidence (also retroactively) and renormalizing the other priors.

Open-minded Bayesianism thus offers a particular perspective on the use of old evidence for confirming a new theory. On the conceptual level, it shows how our perception of evidence and confirmation changes if we move from one theoretical context to another. Relative to one set of hypotheses, the data were telling towards one particular candidate hypothesis, and so counted as evidence that confirms this candidate. But with the inclusion of a new hypothesis, the data may tell against the formerly best candidate, and so count as evidence that disconfirms it. We take it to be a virtue of our model that it brings out this context-sensitivity of evidence and confirmation.

4.3 Illustration of the hybrid approach

To make our proposal for a hybrid approach more vivid, we apply it to the food inspection example. Initially, when the food inspector implicitly assumes her equipment to be working properly, she can be described by the silent approach to open-minded Bayesianism. Within the initial context, she only needs to take into account two explicit hypotheses: the kitchen is clean or it is not. She assigns prior probabilities to these hypotheses and she computes posteriors, but these assignments are conditional on her implicit assumption that the testing strips are uncontaminated (as well as many other background assumptions). So far, she acts much like any Bayesian would; her open-mindedness will surface only when provoked.

The result, that five dishes out of five appear to be infected, was initially unlikely on both of her explicit hypotheses. (Recall that the initial likelihood was 10^{-10} in the case of a clean kitchen and 3.2×10^{-4} in the case of an unclean kitchen.) Computing the posterior probabilities, which implicitly requires us to assume that the correct hypothesis is among the two hypotheses being considered, leads to a value close to zero (3.1×10^{-7}) for a clean kitchen and a value near to unity ($1 - 3.1 \times 10^{-7}$) for an unclean kitchen. If the priors were equal (or at least of the same order of magnitude), then on any measure of confirmation, the evidence provides very strong confirmation for the hypothesis that the kitchen was unclean.

The observation that it is highly unlikely even for an unclean kitchen to produce five infected dishes may suggest that there is an even better hypothesis ‘out there’ that has not yet been taken into account. Indeed, seeing the result prompts the inspector to reconsider one of her implicit assumptions and she turns its negation into a new theory (and associated statistical hypothesis): the testing strips may not have been clean after all (bias = 1). (Of course, this is still but one out of many other alternative hypotheses.) Our framework for open-minded Bayesianism is able represent this formally.

In the vocal phase, the agent shaves off her third hypothesis from the catch-all and revises her probability assignments: she retroactively assigns a prior to the new hypothesis, adjusts the priors of the two old hypotheses by a suitable

factor, and computes the likelihood of the old evidence on the new hypothesis (as described in section 3.3.3). All this leads her to reassess the posteriors of the old hypotheses and to assign a posterior to the new hypothesis. Assuming equal priors, the final result is this: within the new theoretical context, the posterior of the new hypothesis given the old evidence is more than three thousand times higher than that of the hypothesis that was best confirmed within the old theoretical context. Irrespective of the details of the confirmation measure and assuming priors of at least equal orders of magnitude, this implies that the old evidence strongly confirms the new hypothesis and disconfirms the others. This illustrates that *it is the shift in theoretical context itself that may cause old evidence to confirm a new hypothesis*.

Once the agent is satisfied that, for the evidence currently at hand, the new theoretical context includes all the relevant hypotheses, she may start to conditionalize all her findings on this context and thereby enter a new silent phase. The remaining catch-all hypothesis need not be mentioned again until new doubts arise.

In Kuhnian terminology, the silent version of open-minded Bayesianism is sufficient for describing episodes of normal science (and if the conditionalization on the theoretical context remains implicit, it is indistinguishable from the usual Bayesian picture), but the vocal version of open-minded Bayesianism is required to model revolutionary changes in the theoretical context.

4.4 Further research

With the foregoing, we believe we have only scratched the surface of the matter at hand. Many avenues for further research lay open for exploration. In what follows, we briefly mention a number of these avenues. With this we showcase our ongoing research on this, we invite the reader to join in, but mostly we indicate where we ourselves feel that our account is lacking.

One important consideration that has not received sufficient attention in the foregoing concerns degrees of confirmation. Our goal with this paper was to show that we can accommodate the introduction of a new theory and hence a new empirical hypothesis in the Bayesian framework, and that old evidence can play a role in the determination of the posterior probability of this new hypothesis without violating probabilistic coherence. We have been mostly silent on how the posteriors may be used to compute a degree of confirmation, so that the impact of old evidence can be expressed more precisely. A complete account of open-minded Bayesianism will certainly involve more detail on degrees of confirmation.

Another aspect of the process of theory change targeted in this article also deserves a more detailed normative treatment: the decision to introduce a new theory. In the foregoing, we have treated this decision as completely external to the model. However, we also indicated that the search for new theories may be motivated by a so-called statistical model selection criterion, *e.g.*, by a measure of the predictive performance of the agent, or by some other score

that attaches to the data and the hypotheses currently under consideration. We think that our account, which may provide rationality constraints on the transition from one theoretical context to another, can be combined fruitfully with an account of how theoretical contexts are evaluated and selected.

Furthermore, we should stress that we have only considered one type of theory change—a change that may be captured by shaving off new hypotheses from a catch-all hypothesis.²² In general, theory change may lead to other types of change to the domain of the probability function, $\mathcal{A}(\Theta \times \Omega)$, in various ways. For one, we have not explicitly considered changes in the space Ω of empirical possibilities. Notice that such changes are generally more radical than changes in the theoretical realm: theories obtain their empirical content in terms of hypotheses that are formulated by means of Ω . One captivating question concerns the exact reach of our account of new theory and old evidence. Specifically, can we assume at the outset that Θ and Ω are rich enough to accommodate all conceivable theory changes? An answer to this question requires us to survey a rich landscape of theory changes.²³

We would like to mention one other aspect to theory change that is related to two issues discussed above, namely the decision to introduce a new theory and the type of theory change effected by that. It concerns the notion of *awareness*. Hill (2010) and Dietrich and List (2013) have recently argued that a decision problem obtains new dimensions when the agent is made aware of considerations that were previously not live to her. We think that roughly the same can be said about the epistemic problems an agent faces, and that the foregoing offers a natural model for an agent that becomes aware of a theory while performing a predictive, or more generally an epistemic task. It seems natural to combine the frameworks for modeling awareness.

Finally, we briefly mention two possibilities that open-minded Bayesianism offers, when it is combined with ideas on relative infinitesimals (in the sense of Wenmackers, 2013). One is that the framework allows us to model radically sceptical yet empiricist epistemic attitudes: all the priors and posteriors of explicit hypotheses, old and new ones, may be very small, indeed infinitesimally small, compared to the probabilities associated with the catch-all. That is, we may choose τ_N to be some number very close to one. Despite that, a

²² Earman (1992, p. 196) has introduced a distinction between two forms of theory change: “The mildest form occurs when the new theory articulates a possibility that lay within the boundaries of the space of theories to be taken seriously but that, because of the failure of logical omniscience [...], had previously been unrecognized as an explicit possibility. The more radical form occurs when the space of possibilities is itself significantly altered.” Although this is a helpful way of categorizing theory change, it is not absolute: the kind of theory change that we have discussed can be reconstructed as a radical one by humble Bayesianism (in which Θ_N is extended to Θ_{N+1}) or as a mild one by Bayesian Tauism (in which the partition on Θ is refined). Presumably, radical changes that can be reconstructed as mild changes are best considered as intermediate cases: there are both milder and more radical changes possible.

²³ Recall that we have defined Θ as the set of *all* probability functions on a common domain, $\mathcal{A}(\Omega)$. Arguably, it may suffice to choose a smaller set Θ , namely the set of all *computable* probability functions on the domain $\mathcal{A}(\Omega)$. This is the idea behind the celebrated theory of universal prediction by Solomonoff (1967).

particular theory may have a large prior or posterior relative to the other theories in the theoretical context. The framework thus allows us to model a radical sceptic who is nevertheless sensitive to differences in empirical support. On the other side of the spectrum, the framework of open-minded Bayesianism allows us to model practical certainty without spilling over into dogmatism. We may be aware of the existence of certain hypotheses, but we might choose not to include them in our considerations: they may seem irrelevant to the kinds of evidence under study (assuming statistical independence), they are deemed highly unlikely,²⁴ including them may require too high a number of computations, or because of other pragmatic reasons. However, upon receiving falsifying or strongly disconfirming evidence, we might want to reconsider some of these omissions.²⁵ The catch-all hypothesis with infinitesimal probability may then serve as a reservoir for the hypotheses that seemed dispensable at one point in time, but that later on turn out to be relevant. Falsifying or strongly disconfirming evidence may lead to a situation in which the probability of the catch-all is no longer regarded as a relative infinitesimal: the marginal likelihood becomes so small that it becomes comparable to the probability of the catch-all.

The above list of research topics indicates that our resolution to the problem of old evidence and new theories leaves much to be done. However, the list also suggests that the framework of open-minded Bayesianism provides access to several interesting aspects of belief dynamics that fall outside the scope of orthodox Bayesianism. We call to mind what Sue says in Earman (1992, p. 235): “By all means keep an open mind, but not so open that your brain falls out.” It seems to us that open-minded Bayesianism does precisely that.

References

- Bertsekas DP, Tsitsiklis JN (2008) Introduction to Probability. 2nd edition
Christensen D (1999) Measuring confirmation. *Journal of Philosophy* 96:437–461
Dawid AP (1982) The well-calibrated Bayesian. *Journal of the American Statistical Association* 77:605–610
Dietrich F, List C (2013) A reason-based theory of rational choice. *Noûs* 47:104–134
Earman J (1992) Bayes or bust? A critical examination of Bayesian confirmation theory. MIT Press, Cambridge, MA

²⁴ In a probabilistic framework, very few theories (or better: the associated statistical hypotheses) can ever be refuted completely, yet some theories—say, phlogiston theory—may become so unlikely that no scientist ever considers them again once a better alternative has been found.

²⁵ A formal model of the Lockean thesis in terms of context-dependent infinitesimals is given by Wenmackers (2013). Pacuit et al (2013) provide a different example of the use of infinitesimal probabilities for modeling the revision of practical certainties. See also Schwitzgebel (2014) on what he calls “1%-scepticism” for a less formal treatment of related issues.

- Easwaran K (2011) Bayesianism II: Applications and criticisms. *Philosophy Compass* 6:321–332
- Garber D (1983) Old evidence and logical omniscience in Bayesian confirmation theory. *Minnesota studies in the philosophy of science* 10:99–131
- Gelman A, Carlin JB, Stern HS, Rubin DB (2004) *Bayesian Data Analysis*, Second Edition. Chapman and Hall, Boca Raton
- Gillies D (2001) Bayesianism and the fixity of the theoretical framework. In: Corfield D, Williamson J (eds) *Foundations of Bayesianism*, Kluwer, pp. 363–379
- Glymour C (1980) Why I am not a Bayesian. In: Glymour C (ed) *Theory and Evidence*, Princeton University Press, Princeton, pp. 63–93
- Haenni R, Romeijn JW, Wheeler G, Williamson J (2003) *Probabilistic Logics and Probabilistic Networks*, Synthese Library: Studies in Epistemology, Logic, Methodology, and Philosophy of Science, vol 350. Springer, Dordrecht, The Netherlands
- Hájek A (2003) What conditional probability could not be. *Synthese* 137:273–323
- Hájek A (2007) The reference class problem is your problem too. *Synthese* 156:563–585
- Halpern JY (2003) *Reasoning about Uncertainty*. MIT Press, Cambridge, MA
- Henderson NL, Goodman JD, Tenenbaum JB, Woodward F (2010) The structure and dynamics of scientific theories: A hierarchical Bayesian perspective. *Philosophy of Science* 77(2)
- Hill B (2010) Awareness dynamics. *Journal of Philosophical Logic* 39:113–137
- Hogg DW (2012) Data analysis recipes: Probability calculus for inference, URL <http://arxiv.org/abs/1205.4446v1>
- Jeffrey R (1983) Bayesianism with a human face. *Minnesota studies in the philosophy of science* 10:133–156
- Joyce J (1999) *The Foundations of Causal Decision Theory*. Cambridge University Press, New York, NY
- Lindley D (1991) *Making Decisions* (2 ed.). Wiley.
- Morey R, Romeijn JW, Rouder JN (2013) The humble Bayesian: Model checking from a fully Bayesian perspective. *British Journal of Mathematical and Statistical Psychology* 66:68–75
- Pacuit E, Pedersen AP, Romeijn JW (2013) When is an example a counterexample? In: Schipper BC (ed) *TARK XIV proceedings*, ACM Digital Library, pp. 156–65
- Romeijn JW (2004) Hypotheses and inductive predictions. *Synthese* 143(3):333–364
- Romeijn JW (2005) Theory change and Bayesian statistical inference. *Philosophy of Science* 72:1174–1186
- Romeijn JW (2011) Statistics as inductive logic. In: Bandyopadhyay P, Forster M (eds) *Philosophy of Statistics, Handbook for the Philosophy of Science*, vol 7, North Holland (Elsevier), Oxford, UK, pp. 751–774
- Salmon WC (1990) Rationality and Objectivity in Science or Tom Kuhn meets Tom Bayes. In: Savage CW (ed.) *Scientific theories*, Minnesota studies in

- the philosophy of science, Vol. 14, pp. 175–204.
- Schwitzgebel E (2014) 1% Skepticism. Unpublished manuscript, <http://www.faculty.ucr.edu/~eschwitz/SchwitzPapers/1%25Skepticism-140512.pdf>
- Shimony A (1970) Scientific inference. In: Colony RG (ed) *The Nature and Function of Scientific Theories*, The University of Pittsburgh Press, Pittsburgh, pp. 79–172
- Sklar L (1981) Do Unborn Hypotheses Have Rights? *Pacific Philosophical Quarterly*, 62, 1729.
- Sprenger J (2014) A novel solution to the problem of old evidence. Unpublished manuscript, <http://philsci-archive.pitt.edu/10643/>
- Stanford K (2006) *Exceeding our grasp: Science, history, and the problem of unconceived alternatives*. Oxford: Oxford University Press.
- Talbott W (2008) Bayesian epistemology. In: Zalta N (ed) *Stanford Encyclopedia of Philosophy*, <http://plato.stanford.edu/entries/epistemology-bayesian/>
- Walley P (2000) Towards a unified theory of imprecise probability. *International Journal of Approximate Reasoning* 24:125–148
- Wenmackers S (2013) Ultralarge lotteries: Analyzing the lottery paradox using non-standard analysis. *Journal of Applied Logic* 11:452–467