Conditional Degree of Belief

Jan Sprenger*

October 5, 2017

Abstract

This paper articulates and defends a suppositional interpretation of conditional degree of belief. First, I focus on a type of probability that has a crucial role in Bayesian inference: conditional degrees of belief in an observation, given a statistical hypothesis. The suppositional analysis explains, unlike other accounts, why these degrees of belief track the corresponding probability density functions. Then, I extend the suppositional analysis and argue that *all* probabilities in Bayesian inference should be understood suppositionally and model-relative. This sheds a new and illuminating light on chance-credence coordination principles, the relationship between Bayesian models and their target system, and the epistemic significance of Bayes' Theorem.

*Contact information: Tilburg Center for Logic, Ethics and Philosophy of Science (TiLPS), Tilburg University, P.O. Box 90153, 5000 LE Tilburg, The Netherlands. Email: j.sprenger@uvt.nl. Webpage: www.laeuferpaar.de

1 Introduction

Bayesian inference is a well-established theory of inductive reasoning that represents an agent's epistemic attitudes—their degrees of belief—by the laws of probability (e.g., Jeffrey, 1971; Savage, 1972; de Finetti, 1972; Earman, 1992; Bovens and Hartmann, 2003; Howson and Urbach, 2006). For a scientific hypothesis H, $p(\mathrm{H})$ represents a rational agent's degree of belief that H is true. Upon learning a proposition E, the agent changes her degrees of belief in H according to Conditionalization: $p^{\mathrm{E}}(\mathrm{H}) = p(\mathrm{H}|\mathrm{E})$. That is, the agent's posterior degree of belief in H is set equal to the conditional degree of belief in H given E. This posterior probability $p(\mathrm{H}|\mathrm{E})$ is used as a basis for judgments of acceptance and decision-making and can be calculated from Bayes' Theorem:

$$p(\mathrm{H}|\mathrm{E}) = p(\mathrm{H})\,\frac{p(\mathrm{E}|\mathrm{H})}{p(\mathrm{E})} \qquad \text{(Bayes' Theorem)}$$

In other words, the degree of belief that a rational agent has in hypothesis H after learning E depends on three factors: her prior degree of belief in H, her prior degree of belief in E, and her conditional degree of belief in E given H.

In this paper, I investigate the semantics and epistemology of probabilities of the type $p(\mathrm{E}|\mathrm{H})$: conditional degrees of belief in an observation, given a hypothesis that makes statistical predictions (henceforth, "statistical hypothesis"). Notwithstanding the salience of these probabilities in Bayesian inference, this question is rarely investigated systematically.

The negative part of the paper argues why standard views of Bayesian inference cannot fully describe the epistemic role of these degrees of belief (Section 2). In particular, chance-credence coordination principles fail to explain why these conditional degrees of belief are rationally constrained by the corresponding probability densities. Then I discuss (and reject) the proposal to reduce conditional to unconditional degrees of belief (Section 3).

In the constructive part of the paper, I develop a suppositional analysis of conditional degrees of belief: conditional degrees of the type $p(\mathrm{E}|\mathrm{H})$ are the degrees of belief that we have in the occurrence of E upon supposing that the target system is fully described by H (Section 4). Then, I explore the wider implications of the suppositional analysis. I argue that *all* probabilities in Bayesian inference should be understood suppositionally and model-relative. Finally, I develop a new interpretation of Bayes' Theorem as a coherence constraint between different probability functions and I explain how suppositional and possibly counterfactual degrees of belief can contribute to reliable predictions about the real world (Section 5). I conclude by wrapping up the results of the paper (Section 6).

Overall, this research elucidates how an adequate semantics and epistemology of conditional degree of belief leads to a better understanding of Bayesian inference as a whole. The proposed analysis also clarifies the role of chance-credence coordination in scientific inference. Moreover, it solves some pertinent problems of Bayesian reasoning, such as the interpretation of degrees of belief in hypotheses that are part of highly idealized models. Finally, it provides the conceptual groundwork for applications of Bayesian inference in statistics and other domains of science.

The suppositional approach to conditional degrees of belief is not entirely novel. Ramsey (1926) has famously argued that conditional degrees of belief are determined by supposing the antecedent (i.e., the proposition that is conditioned on) and reasoning on that basis about the consequent (i.e., the proposition whose probability is evaluated). Also, it is related to philosophical discussions about the status of conditional probability (e.g., Hájek, 2003; Fitelson and Hájek, 2017). However, my paper is, to the best of my knowledge, the first to explore the implications of Ramsey's proposal for Bayesian inference and inductive reasoning.

2 The Equality and Chance-Credence Coordination

What constrains conditional degrees of belief in an observation E, given a hypothesis H? Suppose that we toss a coin in order to learn about its bias—a simple and transparent case of inductive inference. The hypotheses (H_θ, $\theta \in [0,1]$) describe the chance of the coin to come up heads on any individual toss. When we toss the coin N times, our sample space is $\mathcal{S} = \{0,1\}^N$. Under the assumption that the tosses are independent and identically distributed (henceforth, i.i.d.), we can describe the probability of observation E_k (=k heads and $N-k$ tails) by the Binomial probability distribution and the corresponding **probability density function** $\rho_{\mathrm{H}_\theta}(\mathrm{E}_k) = \binom{N}{k}\,\theta^k\,(1-\theta)^{N-k}$. More generally, the sample space $\mathcal{S}$ together with the set of probability distributions over $\mathcal{S}$ induced by the ρ_{H_θ} constitutes a **statistical model** (e.g. Cox and Hinkley, 1974; Bernardo and Smith, 1994; McCullagh, 2002).

Suppose $N = 3$ and let H be the hypothesis is that the coin is fair, H: $\theta = 1/2$. Then, the probability of observing two heads in three tosses (=E) is described by $\rho_{\mathrm{H}}(\mathrm{E}) = \binom{3}{2}\,(1/2)^2\,(1-1/2)^{3-2} = 3/8$. Bayesian reasoners in the sciences take these probabilities as plug-ins for their conditional degrees of belief $p(\mathrm{E}|\mathrm{H}_\theta)$ (e.g., Bernardo and Smith, 1994; Howson and Urbach, 2006). And since the former are uniquely determined, so are the latter. Agents with divergent priors will approach each other's views in the posterior distribution. Uniquely rational values for $p(\mathrm{E}|\mathrm{H}_\theta)$ have an important epistemic function: they work toward establishing (approximate) consensus on posterior distributions, and they lead to unanimous assessments of the strength of observed evidence. For instance, the Bayes factor between two hypothesis H_1 and H_0, a popular measure of evidence in Bayesian statistics (Jeffreys, 1961; Lee and Wagenmakers, 2013), is just a ratio of conditional probabilities: $\mathrm{BF}_{10}(\mathrm{E}) = p(\mathrm{E}|\mathrm{H}_1)/p(\mathrm{E}|\mathrm{H}_0)$. All this presupposes, however, a satisfactory answer to the the main question of the paper:

Main Question What justifies the equality between conditional degrees of belief and the corresponding probability densities?

$$p(\mathrm{E}|\mathrm{H}) = \rho_H(E) \qquad \text{(The Equality)}$$

A tempting way to justify The Equality consists in deriving it from a general epistemic norm: a **chance-credence coordination principle** that calibrates subjective degree of belief with objective chance (e.g., Lewis, 1980; Strevens, 1999; Williamson, 2007, 2010). For instance, according to the **Principle of Direct Inference (PDI)** (e.g., Reichenbach, 1949; Kyburg, 1974; Levi, 1977), if I know that a coin is fair, I should assign degree of belief 1/2 that heads will come up. David Lewis (1980) formalized a related intuition in his **Principal Principle (PP)**: the initial credence function of a rational agent, conditional on the proposition that the physical chance of E takes value x, should also be equal to x.

Do these principles apply to The Equality? Not straightforwardly so. Presumably, the Principle of Direct Inference (or its close cousin, the Principal Principle) would recommend to carry over the value of $\rho_{\mathrm{H}}(E)$ to the conditional degree of belief $p(\mathrm{E}|\mathrm{H})$.

After all, ρ_{H} does not depend on subjective epistemic attitudes, just on the properties of the statistical model. In this sense, it seems to qualify as an objective chance.

Note, however, that objective chances make empirical statements: their values depend on "facts entailed by the overall pattern of events and processes in the *actual* world" (Hoefer, 2007, 549, original emphasis), such as the setup of an experiment or the physical composition of the coin we toss. This is true of frequencies, propensities, and best-system chances alike. Our case is different: the truth value of sentences such as $\rho_H(\text{E}) = 3/8$ is *entirely internal to the statistical model*: if H denotes the hypothesis that the coin is fair and E denotes the observation of two tosses, then it is part of the *meaning* of H that the probability of E given H is 3/8. That is, the sentence

> When we perform three i.i.d. tosses of a fair coin, the chance of observing two heads is 3/8. (+)

has no empirical content—it has a distinctly analytical flavor. It does not refer to real-world properties or events; in fact, no fair coins need to exist for this sentence to be true. In other words, $\rho_H(\text{E})$ describes an objective, but not an ontic probability; its value is subject-independent, but not a physical chance (Rosenthal, 2004; Sprenger, 2010).

The Principle of Direct Inference or the Principal Principle do not apply to this case, at least not without further qualification. They coordinate our degrees of belief with properties of the actual world, such as observed relative frequencies or known propensities. But the objective probabilities in question, $\rho_H(\text{E})$, do not make any statements about the actual world. Therefore, the usual chance-credence coordination principles cannot justify The Equality.

The rest of the paper will be devoted to finding a semantics for conditional degree of belief that explains The Equality and elucidates the role of chance-credence coordination in statistical inference. We first discuss a radical proposal: to reduce conditional to unconditional degree of belief.

3 Ratio Analysis

According to most textbooks on probability theory, statistics and (formal) epistemology, the conditional probability of E given H is defined as the ratio of the probability of the conjunction of both propositions, divided by the probability of H (Jackson, 1991; Earman, 1992; Skyrms, 2000; Howson and Urbach, 2006).

$$p(\mathrm{E}|\mathrm{H}) := \frac{p(\mathrm{E} \wedge \mathrm{H})}{p(\mathrm{H})} \qquad \text{(Ratio Analysis)}$$

By applying a Dutch Book argument, we can transfer this reductive analysis of conditional probability to conditional degree of belief. It can be shown that the conditional degree of belief in E given H, operationalized by a bet on E that is called off if $\neg$H, must be equal to the ratio of the unconditional degrees of belief in E $\wedge$ H and H. Otherwise, the agent will be susceptible to a Dutch book (e.g., Easwaran, 2011a). Ratio Analysis is therefore uncontroversial as a mathematical constraint on rational conditional degrees of belief. What is at stake is a semantic question: can Ratio Analysis figure as a *definition* of conditional degree of belief? That is, do conditional degrees of belief $p(\mathrm{E}|\mathrm{H})$ just impose constraints on the ratio of $p(\mathrm{E} \wedge \mathrm{H})$ and $p(\mathrm{H})$ and have no independent significance? If this were the case, then our search for a semantics of conditional degrees of belief would pursue a *Scheinproblem*.

The classical line of attack against Ratio Analysis contends that we can often make sense of $p(\mathrm{E}|\mathrm{H})$ although $p(\mathrm{H}) = 0$ (e.g., de Finetti, 1972; Hájek, 2003; Fitelson and Hájek, 2017). Intuitively, such conditional degrees of belief seem to be well-defined, e.g., in the question "What is the probability that a point on Earth is in the Western hemisphere (E), given that it lies on the equator (H)?". To make sense of such questions, Ratio Analysis needs to be amended substantially, and then we are in the middle of a highly technical debate (e.g., Howson and Urbach, 2006; Easwaran, 2011b; Myrvold, 2015). Actually, I find the probability zero objection appealing, but it would be rash to conclude that it deals a lethal blow to Ratio Analysis. Going into the details of this debate would go beyond the scope of this paper. It is also not necessary: the thesis that conditional degrees can be reduced to unconditional degrees of belief has other shortcomings.

The first objection is mainly descriptive: Ratio Analysis fails to do justice to the cognitive role of conditional degrees of belief. Intuitively, Ratio Analysis seems to miss something essential: we usually assess conditional degrees of belief $p(\mathrm{E}|\mathrm{H})$ in a direct way, rather than by reasoning about $p(\mathrm{E} \wedge \mathrm{H})$ and $p(\mathrm{H})$, and calculating their ratio. The latter route carries a high cognitive load, and in many cases, such as the coin toss example, we can easily determine $p(\mathrm{E}|\mathrm{H})$ by reasoning about the likely consequences of H. Indeed, recent psychological evidence demonstrates that Ratio Analysis is a poor description of how people reason with conditional probabilities (Zhao et al., 2009).

Second, Ratio Analysis fails to grasp the normative role of conditional degree of belief in statistical inference. In the previous section, we have seen that it is part of the *meaning* of H to constrain $p(\mathrm{E}|\mathrm{H})$ in a unique way. Recall the example. For determining our rational degree of belief that a fair coin yields a particular sequence of heads and tails, it does not matter whether the coin in question is actually fair. Regardless of our degree of belief in that proposition, we all agree that *given that the coin is fair*, the probability of two heads in three tosses is 3/8. On Ratio Analysis, this feature of conditional degree of belief drops out of the picture. $p(\mathrm{E}|\mathrm{H})$ is constrained only via constraints on $p(\mathrm{E} \wedge \mathrm{H})$ and $p(\mathrm{H})$. But suppose we suspend judgment on $p(\mathrm{E} \wedge \mathrm{H})$ and $p(\mathrm{H})$, or we have imprecise degrees of belief about these probabilities. Even then, the (unique) constraints on $p(\mathrm{E}|\mathrm{H})$ would still be in place. Ratio Analysis therefore misses an important normative aspect of reasoning with conditional degrees of belief.

A consequence of these arguments is that conditional degree of belief is an equally fundamental a concept as unconditional degree of belief. The former cannot be reduced to the latter. The reverse direction—defining unconditional degrees in terms of conditional degrees of belief (e.g., Popper, 1959/2002; Rényi, 1970)—is more promising: in fact, on the view of Bayesian inference that we will end up defending, all relevant probabilities happen to be conditional degrees of belief.

4 The Suppositional Analysis

Between the lines, the previous sections have already anticipated an alternative analysis of conditional degree of belief: a **suppositional interpretation**. That is, we determine our degrees of belief in E given H by *supposing* that H is true.[1]

There are two great figures in philosophy of probability associated with this view.

[1] I call my account "suppositional" instead of "counterfactual" since in some cases, no counterfactual assumptions are necessary. On the other hand, many applications of statistical models in scientific inference imply straightforward counterfactual assumptions; and whenever I want to emphasize this point, I use the adjective "counterfactual" (analysis, interpretation, ...) instead of the more general "suppositional".

One is the British philosopher Frank P. Ramsey, the other is the Italian statistician Bruno de Finetti (1972, 2008). I will focus on Ramsey's view since de Finetti also requires that H be a verifiable event if $p(\mathrm{E}|\mathrm{H})$ is to be meaningful (de Finetti, 1972, 193). This verificiationism is unnecessarily restrictive.

Here is Ramsey's famous analysis of conditional degrees of belief:

> If two people are arguing 'if H will E?' and both are in doubt as to H, they are adding H hypothetically to their stock of knowledge and arguing on that basis about E. (Ramsey, 1926)

The above quote is ambiguous: it is about conditional degree of belief or the truth conditions of conditionals? Many philosophers, most famously Stalnaker (1968, 1975), were inspired by the latter reading and developed an epistemology of conditionals based on the idea that assessing the conditional "If H, then E" involves adding H to one's background knowledge.

I would like to stay neutral on all issues concerning conditionals and interpret Ramsey's quote as an analysis of conditional degrees of belief. Indeed, in the sentence that follows the above quote, Ramsey describes the reasoning procedure in question as

> We can say that they are *fixing their degrees of belief in E given H*. (ibid., my emphasis)

This makes clear that regardless of the possible link to the epistemology of conditionals, Ramsey intended that hypothetically assuming H would determine one's conditional degrees of belief in E, given H. That is, $p(\mathrm{E}|\mathrm{H})$ is the rational degree of belief in E if we suppose that H is true and draw conclusions about E on the basis of H.

When applied to Bayesian inference, this analysis of conditional degree of belief asks us to imagine a possible world ω_H where H is the probability law that governs the frequency of possible observations. The term "possible world" is intended in a quite restrictive sense: not as a full description of the actual world, but as specifying a way the target system (e.g., the coin toss) could be. Supposing a statistical hypothesis H means conceiving of the target system as governed by physical chance, with the probability of an observation E given by the probability density $\rho_H(\mathrm{E})$.

Let us return to the coin-tossing example. Let H denote the hypothesis that the coin is fair. We are now supposing H and reason, within ω_H, about the degrees of belief that we should have in the occurrence of certain sequences of heads and tails. Denote the function that describes our degrees of belief in these observations by $p_{\omega_H}(\cdot)$. From the above it follows that in ω_H, the (physical) chance of two heads in two tosses is 1/4, the chance of two heads in three tosses is 3/8, and so on. Within ω_H, we can now apply a chance-credence coordination principle such as the Principle of Direct Inference (PDI) or the Principal Principle (PP). After all, ρ_{H} qualifies as a physical chance within ω_H and we have no reason to challenge the rationality of PDI or PP in this case.[2] Thus, our (unconditional!) degrees of belief within ω_H should satisfy $p_{\omega_H}(\mathrm{E}) = \rho_H(\mathrm{E})$. By definition of the suppositional analysis, we also have $p(\mathrm{E}|\mathrm{H}) = p_{\omega_H}(\mathrm{E})$. Thus we obtain $p(\mathrm{E}|\mathrm{H}) = \rho_H(\mathrm{E})$ and we conclude that conditional degrees of belief track probability density functions. Hence, the suppositional analysis of conditional degree of belief directly yields The Equality, and it explains why sentences such as (+) appear analytic.

Supposing H may be in conflict with available background knowledge about the target system. For this reason, my interpretation of conditional degree of belief differs from Ramsey's in a crucial nuance: where Ramsey suggested that H is *added* to existing background knowledge, H may also *overrule* conflicting knowledge about the target system on my account. In such cases, we obtain a genuinely counterfactual interpretation of conditional degree of belief. This is often necessary: a given coin may not be perfectly fair, the tosses are not i.i.d., and so on. How to use counterfactual degrees of belief for drawing conclusions about the real world is a question that we will tackle in the next section.

It is important to understand the role of the Principle of Direct Inference and the Principal Principle in the suppositional analysis. Both principles apply to real-world, ontic chances, e.g., "the chance of this atom decaying in the next hour is 1/3" or "the chance of a zero in the spin of this roulette wheel is 1/37". The principles simply claim that degrees of belief should mirror such chances. Compare this to the picture that we sketch for conditional degree of belief. We do not deal with real-world chances; rather we observe that *in the world ω_H described by H*, the objective chance of E is given by

[2]Note that p_{ω_H} is really an initial credence function, as PP demands: information about the actual world $\omega_@$ that may conflict with H is overridden by supposing H and moving to ω_H. Complications induced by inadmissible information (e.g., Lewis, 1980; Hoefer, 2007) do not occur in most cases because the worlds ω_H are so simple and well-behaved. The events in ω_H are elements of the sample space $\mathcal{S}$ and H assigns a definite and unambiguous probability to them.

$\rho_H(E)$. In other words, we do not apply PDI/PP in the actual world $\omega_@$, but in the counterfactual world ω_H described by H, and within that world, we adapt our degree of belief in E to $\rho_H(E)$. By supposing a world where the occurrence of E is genuinely chancy, the suppositional analysis explains why our conditional degree of belief in E given H is uniquely determined and obeys The Equality.

Thus, we can explain why chance-credence coordination is so important for probabilistic reasoning, without committing ourselves to the existence or nature of physical chance in the actual world. This is a distinct strength of the proposed account of conditional degree of belief. Scientists need to calibrate their degrees of belief with objective probabilities gained from statistical models, but the soundness of their inference should not depend on establishing that these probabilities are limiting frequencies, propensities, or any other sort of objective chance to which PDI and PP apply. Given the prevalence of statistical inference in all areas of science, and the ongoing discussion about whether genuine chances exist outside some subfields of physics and biology, one may even conclude that the use of chance-credence coordination principles in hypothetical, idealized worlds is more important for science than their use for relating actual physical chances to our degrees of belief.

Let us wrap up the essence of my proposal. We interpret $p(\mathrm{E}|\mathrm{H})$, the conditional degree of belief in observation E given statistical hypothesis H, by supposing that H is the true model of the target system. In this hypothetical or counterfactual scenario, we coordinate our degree of belief in E with the known objective chance of E, given by $\rho_H(\mathrm{E})$ (Lewis, 1980). To repeat, we are talking about chance-credence coordination in hypothetical worlds where the space of possible events (=the sampling space) is determined by the statistical model, not about chance-credence coordination in the actual world $\omega_@$. Probability density functions inform our *hypothetical* degrees of belief, not our actual degrees of belief.

This interpretation of conditional degree of belief also matches the practice of non-Bayesian inference. Here are the thoughts of the great frequentist statistician Ronald A. Fisher on conditional probability in hypothesis testing:

> In general tests of significance are based on *hypothetical* probabilities calculated from their null hypotheses. They *do not lead to any probability statements about the real world*. (Fisher, 1956, 44, original emphasis)

That is, Fisher is emphatic that the probabilities of evidence given some hypothesis have hypothetical character and are not physically realized objective chances. Probabilities are useful instruments of inference, not components of the actual world. According to Fisher, statistical reasoning and hypothesis testing is essentially counterfactual—it is about the probability of a certain dataset if we suppose the tested "null" hypothesis. The null hypothesis usually denotes the absence of any effect, the additivity of two factors, the causal independence of two variables in a model, etc. In most cases, it is strictly speaking false: there will be *some* minimal effect in the treatment, some slight deviation from additivity, some negligible causal interaction between the variables (e.g., Gallistel, 2009). Our statistical inferences are based on the probabilities of events under a hypothesis which we know to be false—although it may be a good idealization of reality. The suppositional interpretation of conditional degree of belief naturally fits into the practice of statistical inference with its emphasis on testing idealized point hypotheses, e.g., in null hypothesis significance testing.

So far we have applied the suppositional analysis to exactly one variety of (conditional) degree of belief. We will now demonstrate that it is also a fruitful framework for interpreting the degree of belief in a hypothesis, given a statistical model.

5 Implications: Model-Relative Bayesian Inference and Bayes' Theorem

My proposal has been silent on the demarcation of the statistical hypothesis H vis-à-vis the background assumptions that are part of the statistical model $\mathcal{M}$. Consider the coin-tossing case. When we evaluate $p(\mathrm{E}|\mathrm{H})$ with H = "the coin is fair", we assume that the individual tosses of the coin are independent and identically distributed. However, this assumption is not part of H itself: H just describes the tendency of the coin on any particular toss. If we contrast H to some alternative H', we notice that differences between them are typically expressed in terms of parameter values, such as H: $\theta = 1/2$ versus H': $\theta = 2/3$, H'': $\theta > 1/2$, etc. Crucial assumptions on the experimental setup, such as independence and identical distribution of the coin tosses, do not enter the particular hypothesis we are testing. In a sense, $p(\mathrm{E}|\mathrm{H})$ is underdetermined because so much supplementary information about the experiment is missing.

Thus, we have to make the role of our background assumptions explicit. There are two layers of statistical modeling. First, there is the general statistical model $\mathcal{M} = (\mathcal{S}; \mathcal{P})$ where $\mathcal{S}$ denotes the sample space, and $\mathcal{P}$ denotes the set of probability distributions over $\mathcal{S}$. What counts as an element of $\mathcal{P}$ depends on several factors: if we describe the repeated coin toss by a set of Binomial distributions $B(N, \theta)$, then we also assume that the tosses are i.i.d. and that we toss the coin exactly N times. The choice of specific hypotheses such as H: $\theta = 1/2$, H': $\theta = 2/3$, and so on, is only possible after choosing a family of distributions for $\mathcal{P}$.

This implies that the conditional degree of belief $p(\mathrm{E}|\mathrm{H})$ is not only conditional on H, but also conditional on the general model $\mathcal{M}$. Indeed, a Bayesian inference about the probability of heads in the coin-tossing example takes $\mathcal{M}$ as given from the very start. For the analysis in the previous section, whenever we wrote $p(\mathrm{E}|\mathrm{H})$, we should have written $p(\mathrm{E}|\mathrm{H}, \mathcal{M})$, strictly speaking.

The same diagnosis applies to the assignment of prior probabilities $p(\mathrm{H})$: Bayesian inference regarding particular parameter values is relative to a model $\mathcal{M}$ into which all hypotheses are embedded (e.g., $\mathcal{M} = (\mathcal{S} = \{0,1\}^N; \mathcal{P} = B(N, \theta), \theta \in [0,1])$). Degrees of belief are distributed only over elements of $\mathcal{P}$. In particular, also the prior and posterior degrees of belief, $p(\mathrm{H})$ and $p(\mathrm{H}|\mathrm{E})$, are relative to a model $\mathcal{M}$. For instance, in a Binomial model, we distribute prior probabilities over the different values of θ, and over no other statistical hypothesis. Of course, there are also hierarchical models which involve sub-models of different types (e.g., the Binomial model vs. a model with dependent outcomes), but this does not change the general diagnosis: all probabilities in Bayesian inference are relative to a statistical model (see also Wenmackers and Romeijn, 2016).

This observation connects to a classical problem of Bayesian inference. Standardly, $p(\mathrm{H})$ is interpreted as the prior degree of belief that H is true. However, many statistical models are strong idealizations of reality. In those cases, it would be silly to have a strictly positive degree of belief in the truth of a hypothesis (e.g., a particular global climate model), even if it happens to make accurate predictions. It would be even more silly to bet on the truth of any particular hypothesis, as operational interpretations of subjective probability demand. Also the sample space that we use in statistical trials is typically idealized: measurement instruments may have limited precision and reliability, categorizing outcomes can be controversial (e.g., in the case of clinical disorders), and so on. For these reasons, Bayesian inference seems to be based on false and unrealistic premises: the interpretation of $p(\mathrm{H})$ as the degree of belief that H is *true* fails to make sense for many statistical models. In statistics, this is an issue shaping recent debates about applying Bayesian inference to complex statistical models (Gelman and Shalizi, 2013; Morey et al., 2013).

In a nutshell, we face the trilemma of having to reject one of the following three, jointly inconsistent propositions:

(1) $p(\mathrm{H})$ denotes an agent's degree of belief that H is true.

(2) Such a hypothesis H is part of a general probabilistic model $\mathcal{M}$ with a partition of hypotheses $\mathcal{H} = \{H, H_1, H_2, \ldots, \}$.

(3) Many of these models $\mathcal{M}$ are strong idealizations of reality—and likely or known to be false.

Since (2) and (3) are rather uncontentious and also part and parcel of scientific reasoning, we have to reject (1). There are various ways of doing so. For example, Vassend (2017) defends the view that the subjective probability of H expresses the degree of belief that H is most similar to the truth among all (false) hypotheses in $\mathcal{M}$ (see also Niiniluoto, 2011; Cevolani and Tambolo, 2013). Instead, I propose a suppositional analysis: $p(\mathrm{H}|\mathcal{M})$ expresses the degree of belief in H that we *would have* if we supposed that the target system is fully and correctly described by one of the hypotheses in $\mathcal{M}$. On this reading, the entire Bayesian inference is relative to $\mathcal{M}$.

However, supposing $\mathcal{M}$ does not yield a uniquely rational value for $p(\mathrm{H}|\mathcal{M})$. There is no objective chance of H in the hypothetical world $\omega_{\mathcal{M}}$. Does this mean that the suppositional analysis of conditional degree of belief is not applicable to prior degrees of belief?

The answer is no. It is the very essence of subjective Bayesian inference that different rational choices of the prior distribution are possible. If $p(\mathrm{H}|\mathcal{M})$ were determined in the same way as $p(\mathrm{E}|\mathrm{H}, \mathcal{M})$, this feature of subjective Bayesian inference would get lost. Moreover, the lack of unique determination by a probability density function does not imply that "anything goes", that all prior distributions are equally rational. Given what we know about the target system, some hypotheses are more plausible than others. For instance, almost nobody would assign a high prior to the hypothesis that an ordinary coin is highly biased towards heads ($\theta > .9$), regardless of whether she accepts or rejects the Binomial model as an adequate description of the coin toss. A possible explication of the suppositional interpretation for prior probabilities consists in adopting a broadly Lewisian perspective (Lewis, 1973): to determine the prior probability of H as a function of the estimated similarity of ω_H to the actual world $\omega_@$. Of course, different agents can come up with different similarity rankings, dependent on their goals, interests and the features of the target system they pick out as relevant.

This view of Bayesian inference squares well with the famous saying from statistics that all models are wrong, but some are illuminating and useful (Box, 1976). Having a prior (or posterior) over θ does not commit us to any degree of belief about the "true" value of θ, or to betting on some propositions about θ with specific odds. It just makes a statement about how our degrees of belief and betting odds *would be* if the coin toss were fully described by $\mathcal{M}$.

We now proceed to the remaining probabilities in Bayes' Theorem. Also the marginal likelihood $p(\mathrm{E}) = p(\mathrm{E}|\mathcal{M})$ can be interpreted in a model-relative, suppositional way. In fact, this follows from the Law of Total Probability. $p(\mathrm{E}|\mathcal{M})$ is the weighted average of the conditional probabilities of E, and thus our subjective expectation that E occurs if $\mathcal{M}$ is the case.

$$p(\mathrm{E}|\mathcal{M}) = \sum_{H_i \in \mathcal{H}} p(\mathrm{E}|\mathrm{H_i}, \mathcal{M}) \cdot p(\mathrm{H_i}|\mathcal{M}) \tag{1}$$

Suppose that $\mathcal{M}$ is an ensemble of hypotheses (e.g., global climate models) that make statistical predictions for specific events such as E: "On emission scheme S, global mean temperature in 2100 will be at least 2°C higher than in 1990." Suppose that $p(\mathrm{E}|\mathcal{M}) \approx 1$ because most hypotheses in $\mathcal{M}$ assign a high probability to E. Do these predictions affect our actual epistemic attitudes regarding the occurrence of E? After all, on the suppositional interpretation, all probabilities are relative to a statistical model $\mathcal{M}$. But we care for whether E is *actually* likely to occur, not whether it is likely to occur given a certain model. In other words: how does the suppositional analysis link model predictions to the real world? This question is especially urgent in contexts such climate science, where doubts about the adequacy of large-scale models for making probabilistic predictions are widespread:

> It is [...] inappropriate to apply any of the currently available generic techniques which utilize observations to calibrate or weight models to produce forecast probabilities for the real world. (Stainforth et al., 2007, 2145)

In general, confidence in the predictions of a statistical model $\mathcal{M}$ depends on whether its constituents capture the relevant aspects of the target system. When the target system is highly complex and hard to predict, we will have to preserve a healthy dose of skepticism toward the predictions of $\mathcal{M}$. On the other hand, the better the constituents of $\mathcal{M}$ describe relevant aspects of the target system, the more justification do we have for an inference by analogy, and for transferring their predictions to our actual epistemic attitudes. In this respect, Bayesian inference is just another form of model-based reasoning in science (Weisberg, 2007; Frigg and Hartmann, 2012; Suárez, 2016). According to the suppositional analysis, converting knowledge about the model into knowledge about the target system is equally difficult for Bayesian and non-Bayesian model builders.[3]

The interpretation of the posterior probability $p(\mathrm{H}|\mathrm{E}, \mathcal{M})$ opens up new questions. Why, and to which extent, are posterior probabilities constrained by the evidence? In the model-relative version of Bayes' Theorem, $p(\mathrm{H}|\mathrm{E}, \mathcal{M})$ depends on $p(\mathrm{E}|\mathrm{H}, \mathcal{M})$, $p(\mathrm{H}|\mathcal{M})$ and $p(\mathrm{E}|\mathcal{M})$:

$$p(\mathrm{H}|\mathrm{E}, \mathcal{M}) = p(\mathrm{H}|\mathcal{M}) \, \frac{p(\mathrm{E}|\mathrm{H}, \mathcal{M})}{p(\mathrm{E}|\mathcal{M})} = \left(\sum_{H_i \in \mathcal{H}} \frac{p(\mathrm{E}|\mathrm{H_i}, \mathcal{M})}{p(\mathrm{E}|\mathrm{H}, \mathcal{M})} \cdot \frac{p(\mathrm{H_i}|\mathcal{M})}{p(\mathrm{H}|\mathcal{M})} \right)^{-1} \tag{2}$$

We see that the impact of $p(\mathrm{H}|\mathcal{M})$ on the posterior probability diminishes when there is sufficiently strong evidence in favor of or against H, as expressed by the ratio of conditional probabilities $p(\mathrm{E}|\mathrm{H_i}, \mathcal{M}) / p(\mathrm{E}|\mathrm{H}, \mathcal{M})$. The value of the posterior probability $p(\mathrm{H}|\mathrm{E}, \mathcal{M})$ is rationally constrained because (i) Bayes' Theorem binds it to the value of other conditional degrees of belief; (ii) the conditional degrees of belief in E given the competing statistical hypotheses $\mathrm{H_i}$, $p(\mathrm{E}|\mathrm{H_i}, \mathcal{M})$, obey The Equality and track the corresponding probability densities.

The model-relative version of Bayes Theorem' deserves some further comments. It

[3]Moreover, Bayesian inference has some features that are useful in problematic contexts such as climate science: one can hedge model error by embedding different types of models into a hierarchical Bayesian model, and by averaging their predictions (e.g., Tebaldi and Knutti, 2007). This approach is, as the above quote indicates, not without problems, but it is hard to come up with better alternatives.

can be derived easily from applying Ratio Analysis to both $p(\mathrm{E}|\mathrm{H}, \mathcal{M})$ and $p(\mathrm{H}|\mathrm{E}, \mathcal{M})$. Operationalized in terms of bets, this means the following: if we understand the above probabilities as indicating fair odds for conditional bets that are called off if the proposition behind the dash is false, then the only way to avoid Dutch books is to make sure that one's betting odds (and the implied probabilities) respect Equation (2)—the model-relative version of Bayes' Theorem. It is not about the *definition* of $p(\mathrm{H}|\mathrm{E}, \mathcal{M})$, but about its coherence with other probabilities: it states how the different functions that represent our conditional degrees of belief ($p(\cdot|\mathcal{M})$, $p(\cdot|\mathrm{H}, \mathcal{M})$ and $p(\cdot|\mathrm{E}, \mathcal{M})$) need to coordinate in order to avoid a Dutch book. This is in fact very similar to the standard justification of Bayes' Theorem and Ratio Analysis (see Section 3). But now, rather than coordinating degrees of belief described by a single probability function, Bayes' Theorem coordinates different probability functions.

Most primers and encyclopedia entries on Bayesian inference describe the theorem as "a simple mathematical formula used for calculating conditional probabilities" (Joyce, 2008), that is, as a practical tool for being able to perform Bayesian inference and to compute posterior probabilities. On the suppositional analysis, it is more than that; it emerges as an epistemic coordination principle with genuine philosophical significance.

Note that Equation (2) is a synchronic, not a diachronic constraint: it describes the conditional degree of belief in H, given E and $\mathcal{M}$, not the degree of belief in H, given $\mathcal{M}$ and after learning E. Only when it is conjoined with Conditionalization $p^{\mathrm{E}}(\mathrm{H}|\mathcal{M}) = p(\mathrm{H}|\mathrm{E}, \mathcal{M})$, we obtain a rule for describing how our degrees of belief should change after a learning experience. Conditionalization emerges as an inference which connects two different modes of reasoning: learning E and supposing E. Exploring this connection is a topic for future research.[4]

In sum, Bayesian inference should be construed as a model-relative activity where we reason with counterfactual, and not with actual degrees of belief. As I have argued,

[4]For instance, one may ask how learning the observation E^* affects the conditional degree of belief in E, $p^{\mathrm{E}^*}(\mathrm{E}|\mathrm{H})$? In a nutshell, the answer is that we evaluate this probability relative to the counterfactual world ω_H, taking into account the occurrence of E^* within that world. Sometimes $p^{\mathrm{E}^*}(\mathrm{E}|\mathrm{H}, \mathcal{M})$ will just be equal to $p(\mathrm{E}|\mathrm{H}, \mathcal{M})$. For instance, if E and E^* denote two i.i.d. tosses of a coin, then H screens off both observations from each other and equality holds: $p^{\mathrm{E}^*}(\mathrm{E}|\mathrm{H}, \mathcal{M}) = p(\mathrm{E}|\mathrm{H}, \mathcal{M})$. In other cases, E^* may be informative and determine the conditional degree of belief in E. Time series are a good example. Consider $\mathcal{M} : Y_t = Y_{t-1} + X_t$ with $X_t \in \{-1, 1\}$ and the hypothesis H: $p(X_t = 1) = 1/2$. If $\mathrm{E}^* : Y_{t-1} = 3$ and E: $Y_t = 4$, then we can directly calculate that $p^{\mathrm{E}^*}(\mathrm{E}|\mathrm{H}, \mathcal{M}) = 1/2$. (The other possible observation is $Y_t = 2$.)

this neither impairs the functionality of Bayesian reasoning nor its normative pull. By contrast, the suppositional analysis leads to a more unified picture that resolves some persistent puzzles such as the problem of chance-credence coordination, and the interpretation of prior probabilities in statistical inference.

All this suggests that conditional degree of belief might be a more fruitful primitive notion in Bayesian reasoning than unconditional degree of belief. This diagnosis resonates well with Hájek's (2003) analysis of conditional probability. If we aim at a mathematical basis that is congenial to the epistemic applications, we need to replace Kolmogorov's three standard axioms ($p(\perp) = 0$; $p(A) + p(\neg A) = 1$; $p(\vee_{n=1}^{\infty} A_n) = \sum_{n=1}^{\infty} p(A_n)$ for mutually exclusive A_n) by an axiom system that takes conditional probability as primitive, such as the Popper-Rényi axioms (Popper, 1959/2002; Rényi, 1970; Fitelson and Hájek, 2017). Unconditional probability can then be obtained as a limiting case of conditional probability. Another way is to define conditional probability in terms of an expectation conditional on a random variable (Gyenis et al., 2017). It is up to future work to determine which road is the most promising one.

6 Conclusion

This paper set out to justify The Equality: why do conditional degrees of belief in an observation, given a statistical hypothesis, track the corresponding probability densities? In other words, why is $p(\mathrm{E}|\mathrm{H}) = \rho_H(E)$? I have argued that the suppositional analysis of conditional degree of belief provides the key to answering this question, and that it aligns well with our intuitive handling of conditional probabilities. On top of this, it explains the role of chance-credence coordination in scientific reasoning, the role of priors in highly idealized models, and the epistemic function of Bayes' Theorem. Moreover, it applies naturally to central debates in confirmation theory and the philosophy of probability, such as the Problem of Old Evidence. For example, Howson (1984, 1985) and Sprenger (2015) propose to solve that problem by means of counterfactually interpreted conditional probabilities. Other promising applications are the epistemology of conditionals and the role of conditional probability in theories of objective chance (for recent contributions, see Douven, 2016; Suárez, 2017). All in all, the suppositional analysis leads to a novel and innovative perspective on foundations and practice of Bayesian inference.

Finally, I would like to state the main results of the paper. Some of them may, taken individually, also be articulated elsewhere, but I am aware of no place where a coherent and unified picture is offered.

1. Ratio Analysis is a mathematical constraint on conditional degree of belief, but not a satisfactory philosophical analysis.

2. Conditional degrees of belief of the type $p(\mathrm{E}|\mathrm{H})$, with observation E and statistical hypothesis H, should be embedded into the context of a general statistical model $\mathcal{M}$ and be interpreted in the suppositional (and possibly counterfactual) way anticipated by Ramsey: we suppose that $\mathcal{M}$ and H are true and reason on this basis about the probability of E.

3. The Equality follows directly from the suppositional interpretation: conditional degrees of belief in an observation given a statistical hypothesis track the corresponding probability densities.

4. The suppositional analysis explains the seemingly analytical nature of many probability statements in science, and it agrees with how scientists view probabilities in inference: as objective, but hypothetical entities. Statistical inference need not be committed to particular interpretations of objective chance.

5. The suppositional analysis clarifies the role of chance-credence coordination principles in scientific inference, such as the Principal Principle or the Principle of Direct Inference.

6. The suppositional analysis explains the normative pull of Bayesian inference, by contributing to agreement on the value of Bayesian measures of evidential support such as the Bayes factor.

7. All probabilities in Bayesian inference are conditional degrees of belief: they are conditional on assuming a general (statistical) model. This position successfully deals with the objection that we should never assign positive degrees of belief to a hypothesis when we know it to be wrong.

8. Bayes' Theorem expresses an epistemic coordination principle for various probability functions that describe conditional degrees of belief.

9. Bayesian Conditionalization relates the posterior degrees of belief (learning E) to conditional degrees of belief (supposing E).

10. When it comes to transferring model predictions to the real world, the suppositional analysis makes Bayesian inference analogous to other model-based reasoning strategies in science.

References

Bernardo, J. M. and Smith, A. F. M. (1994). *Bayesian Theory*. Wiley, New York.

Bovens, L. and Hartmann, S. (2003). *Bayesian Epistemology*. Oxford University Press, New York.

Box, G. E. P. (1976). Science and Statistics. *Journal of the American Statistical Association*, 71:791–799.

Cevolani, G. and Tambolo, L. (2013). Progress as approximation to the truth: a defence of the verisimilitudinarian approach. *Erkenntnis*, 78:921–935.

Cox, D. and Hinkley, D. (1974). *Theoretical Statistics*. Chapman & Hall, London.

de Finetti, B. (1972). *Probability, Induction and Statistics: the Art of Guessing*. John Wiley & Sons, New York.

de Finetti, B. (2008). *Philosophical Lectures on Probability*. Springer, Berlin.

Douven, I. (2016). *The Epistemology of Indicative Conditionals*. Cambridge University Press, Cambridge.

Earman, J. (1992). *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory*. MIT Press, Cambridge/MA.

Easwaran, K. (2011a). Bayesianism I: Introduction and Arguments in Favor. *Philosophy Compass*, 6:312–320.

Easwaran, K. (2011b). The Varieties of Conditional Probability. In Bandyopadhyay, P. S. and Forster, M. R., editors, *Handbook of the Philosophy of Statistics*, pages 137–148. Elsevier, Amsterdam.

Fisher, R. A. (1956). *Statistical methods and scientific inference.* Hafner, New York.

Fitelson, B. and Hájek, A. (2017). Declarations of independence. *Synthese*.

Frigg, R. and Hartmann, S. (2012). Models in science. In Zalta, E., editor, *The Stanford Encyclopedia of Philosophy*.

Gallistel, C. R. (2009). The importance of proving the null. *Psychological Review*, 116:439–453.

Gelman, A. and Shalizi, C. R. (2013). Philosophy and the practice of Bayesian statistics. *British Journal of Mathematical and Statistical Psychology*, 66:8–38.

Gyenis, Z., Hofer-Szabó, G., and Rédei, M. (2017). Conditioning using conditional expectations: The Borel-Kolmogorov Paradox. *Synthese*, 194:2595–2630.

Hájek, A. (2003). What Conditional Probability Could Not Be. *Synthese*, 137:273–323.

Hoefer, C. (2007). The Third Way on Objective Probability: A Sceptic's Guide to Objective Chance. *Mind*, 116:549–596.

Howson, C. (1984). Bayesianism and support by novel facts. *British Journal for the Philosophy of Science*, 35:245–251.

Howson, C. (1985). Some Recent Objections to the Bayesian Theory of Support. *British Journal for the Philosophy of Science*, 36:305–309.

Howson, C. and Urbach, P. (2006). *Scientific Reasoning: The Bayesian Approach*. Open Court, La Salle, IL, 3rd edition.

Jackson, F. (1991). *Conditionals*. Oxford University Press, Oxford.

Jeffrey, R. (1971). *The Logic of Decision*. University of Chicago Press, Chicago.

Jeffreys, H. (1961). *Theory of Probability*. Oxford University Press, Oxford, 3rd edition.

Joyce, J. (2008). Bayes' Theorem. In Zalta, E., editor, *The Stanford Encyclopedia of Philosophy*.

Kyburg, H. E. (1974). *The Logical Foundations of Statistical Inference*. Reidel, Dordrecht.

Lee, M. D. and Wagenmakers, E.-J. (2013). *Bayesian Cognitive Modeling: A Practical Course*. Cambridge University Press, Cambridge.

Levi, I. (1977). Direct Inference. *Journal of Philosophy*, 74:5–29.

Lewis, D. (1973). *Counterfactuals*. Blackwell, London.

Lewis, D. (1980). A Subjectivist's Guide to Objective Chance. In Jeffrey, R. C., editor, *Studies in Inductive Logic and Probability*, pages 263–293. University of California Press, Berkeley.

McCullagh, P. (2002). What is a statistical model? *Annals of Statistics*, 30:1225–1310.

Morey, R. D., Romeijn, J. W., and Rouder, J. N. (2013). The humble Bayesian: Model checking from a fully Bayesian perspective. *British Journal of Mathematical and Statistical Psychology*, 66:68–75.

Myrvold, W. C. (2015). You Can't Always Get What You Want: Some Considerations Regarding Conditional Probabilities. *Erkenntnis*, 80:573–603.

Niiniluoto, I. (2011). Revising Beliefs Towards the Truth. *Erkenntnis*, 75:165–181.

Popper, K. R. (1959/2002). *The Logic of Scientific Discovery*. Routledge, London. Reprint of the revised English 1959 edition. Originally published in German in 1934 as "Logik der Forschung".

Ramsey, F. P. (1926). Truth and Probability. In Mellor, D. H., editor, *Philosophical Papers*, pages 52–94. Cambridge University Press, Cambridge.

Reichenbach, H. (1949). *The Theory of Probability*. University of California Press, Berkeley/CA.

Rényi, A. (1970). *Foundations of Probability*. Holden-Day, San Francisco.

Rosenthal, J. (2004). *Wahrscheinlichkeiten als Tendenzen*. mentis, Paderborn.

Savage, L. J. (1972). *The Foundations of Statistics*. Wiley, New York, 2nd edition. Originally published in 1954.

Skyrms, B. (2000). *Choice & Chance*. Wadsworth, Belmont, CA, 4th edition.

Sprenger, J. (2010). Statistical inference without frequentist justifications. In *EPSA Epistemology and Methodology of Science: Launch of the European Philosophy of Science Association*, pages 289–297. Springer, Berlin.

Sprenger, J. (2015). A Novel Solution of the Problem of Old Evidence. *Philosophy of Science*, 82:383–401.

Stainforth, D., Allen, M., Tredger, E., and Smith, L. (2007). Confidence, uncertainty and decision-support relevance in climate predictions. *Philosophical Transactions of the Royal Society (Series A)*, 365:2145–2161.

Stalnaker, R. (1968). A Theory of Conditionals. In *Studies in Logical Theory: American Philosophical Quarterly Monograph Series, No. 2*. Blackwell, Oxford.

Stalnaker, R. (1975). Indicative Conditionals. *Philosophia*, 5(3):269–286.

Strevens, M. (1999). Objective Probability as a Guide to the World. *Philosophical Studies*, 95:243–275.

Suárez, M. (2016). Representations in Science. In Humphreys, P., editor, *The Oxford Handbook of Philosophy of Science*, pages 440–459. Oxford University Press, New York.

Suárez, M. (2017). The Chances of Propensities. *British Journal for the Philosophy of Science*.

Tebaldi, C. and Knutti, R. (2007). The use of the multi-model ensemble in probabilistic climate projections. *Philosophical Transactions of the Royal Society (Series A)*, 365:2053–2075.

Vassend, O. B. (2017). Non-Standard Bayesianism: How Verisimilitude and Counterfactual Degrees of Belief Solve the Interpretive Problem in Bayesian Inference. Unpublished manuscript, avaliable at `https://sites.google.com/site/olavbvassend/research`.

Weisberg, M. (2007). Who is a Modeler? *The British Journal for the Philosophy of Science*, 58:207–233.

Wenmackers, S. and Romeijn, J. W. (2016). A New Theory About Old Evidence. *Synthese*, 193:1225–1250.

Williamson, J. (2007). Motivating Objective Bayesianism: From Empirical Constraints to Objective Probabilities. In Harper, W. and Wheeler, G., editors, *Probability and Inference: Essays in Honour of Henry E. Kyburg Jr.*, pages 151–179. College Publications, London.

Williamson, J. (2010). *In Defence of Objective Bayesianism*. Oxford University Press, Oxford.

Zhao, J., Shah, A., and Osherson, D. (2009). On the provenance of judgments of conditional probability. *Cognition*, 113:26–36.