Philosophy of Science (2023), **90**, 744–749
doi:10.1017/psa.2023.17

PSA PHILOSOPHY OF SCIENCE ASSOCIATION

PHILOSOPHY OF SCIENCE

DISCUSSION NOTE

Supposition and (Statistical) Models

Corey Dethier

Leibniz Universität Hannover, Hannover, Germany
Email: coreydethier@gmail.com

(Received 14 February 2022; revised 04 January 2023; accepted 19 January 2023; first published online 17 February 2023)

Abstract

In a recent paper, Sprenger advances what he calls a "suppositional" answer to the question of why a Bayesian agent's degrees of belief should align with the probabilities found in statistical models. I show that Sprenger's account trades on an ambiguity between hypothetical and subjunctive suppositions and cannot succeed once we distinguish between the two.

1. Introduction

In a recent paper, Sprenger (2019) advances what he calls a "suppositional" answer to the question of why a Bayesian agent's degrees of belief should align with the probabilities found in statistical models. More precisely, he holds that we should interpret the probability density functions found in statistical models as providing us with insight into how the world would be on the supposition of the hypothesis and that this suppositional reading reveals why a Bayesian should assign degrees of belief in accordance with the probabilities found in the model. Sprenger's account trades on an ambiguity between hypothetical and subjunctive suppositions, however, and thus cannot succeed once we distinguish between the two.

I begin by briefly outlining Sprenger's suppositional account and the problem it is designed to resolve. The second part of the article distinguishes between the two different kinds of supposition, whereas the third argues that Sprenger's account fails once it's recognized that the two kinds of supposition come apart. Section 4 considers a pair of possible rejoinders, and I end by discussing what hangs on the success of the view.

2. Sprenger's suppositional account

Consider a coin and the problem of determining whether it is "fair." To address this problem, the statistician begins by setting up a "statistical model," essentially a specification of the probabilistic relationship between various hypotheses and (possible) sets of experimental results. Formally, in the (Bayesian) coin-flipping case, we can think of this statistical model as consisting of the following: (1) a set, $\mathcal{M}$, consisting of x hypotheses of the form $H_\mu : p(heads) = \mu$ and a σ-algebra, Σ, over that set; (2) a prior probability distribution on the hypothesis space, $p : \Sigma \to [0, 1]$; (3) a sample space, $S = \{heads, tails\}^n$, where n is the number of flips (the evidence, E, then consists of a particular realization of this sample space, that is, $E_k = k$ *heads* and $n - k$ *tails*); and (4) a probability density or mass function relating each hypothesis to (at least) the evidence that is actually observed, $\rho_{H_\mu}(E_k) = \binom{n}{k}(\mu)^k(1-\mu)^{(n-k)}$.

Much has been written about the first three elements of the statistical model and the various challenges that Bayesians face in (properly) specifying these elements. Sprenger (2019), by contrast, focuses on (4). As he notes, the use of Bayesian inference (at least as traditionally presented) requires treating the realization of the probability density function $\rho_{H_\mu}(E_k)$ as equivalent to the conditional degree of belief $p(E_k|H_\mu)$. The "main question" (Sprenger, 2019, 323) that Sprenger aims to answer is what justifies treating these two quantities as equivalent.

To give some teeth to this question, consider the fact that our statistical model is heavily idealized. We've (implicitly) assumed that the flips are independent and identically distributed, and although this assumption may be a good approximation, it's unlikely to hold exactly in real life. As Sprenger argues, more complex examples of statistical models—such as those found in climate science—contain even more unrealistic idealizations. In these cases, therefore, we can't justify the equivalence of $\rho_{H_\mu}(E_k)$ and $p(E_k|H_\mu)$ on the grounds that the statistical model accurately represents the world—after all, we know it doesn't.

Sprenger's solution is to follow Ramsey and hold that the justification of the equivalence rests on suppositional reasoning: "we evaluate the conditional degree of belief $p(E|H)$ by supposing the truth of the conditioning proposition H and by assessing the plausibility of E given this supposition" (Sprenger, 2019, 325). Connecting $p(E|H)$ to suppositional reasoning in this way allows Sprenger to run the following argument. Consider the set of worlds $\mathcal{W}$, "where the behavior of S is governed by the probability law H" (Sprenger, 2019, 325). In each of these worlds, the objective chance of E is given by $\rho_H(E)$. By the principal principle, we should assign degrees of belief that accord with the objective chances, meaning that if we've supposed that we're in $\mathcal{W}$, our degrees of belief should align with the probabilities given by the model: $p_{\mathcal{W}}(E) = \rho_H(E)$. Finally, given the suppositional analysis, the conditional degree of belief $p(E|H)$ should be equivalent to the probability that we assign on the supposition that we are in $\mathcal{W}$, or $p(E|H) = p_{\mathcal{W}}(E)$. The result is the desired equivalence between $p(E|H)$ and $\rho_H(E)$.

3. Two kinds of supposition

Since Adams (1975), the literature on supposition has recognized two distinct kinds of supposition: hypothetical and subjunctive. The difference is neatly brought out by Adams's own example:

Hypothetical: If Oswald didn't kill Kennedy, someone else did.
Subjunctive: If Oswald hadn't killed Kennedy, someone else would have.

In the former case, we suppose the proposition [[Oswald didn't kill Kennedy]] by adding it to the stock of things that we already know. Given that we know that Kennedy was in fact shot, it follows that on this hypothetical, someone else must have shot him. In the latter case, by contrast, we're supposing that [[Oswald didn't kill Kennedy]] in a manner that doesn't hold fixed the stock of things that we already know. In particular, though we know that [[Someone killed Kennedy]], this proposition is dropped from our stock of knowledge when we suppose in a subjunctive

https://doi.org/10.1017/psa.2023.17 Published online by Cambridge University Press

manner. There's much more to say here; the lesson is simply that subjunctive supposition allows for counterfactual possibilities such as those in which Kennedy was not in fact shot.

Just as the antecedent of a conditional can be supposed either hypothetically or subjunctively, so, too, can the proposition on which we condition in a conditional probability. On a standard interpretation, $p(A|B)$ is the probability of A on the hypothetical supposition of B, that is, in the situation where we add B to our stock of knowledge. By contrast, we'll follow Schwarz (2018) in using $p(A||B)$ to indicate the probability of A on the subjunctive supposition of B, which allows for B to overwrite our previous knowledge.[1] As we would expect, these two kinds of conditioning can and will come apart. Plausibly, the probability that someone else killed Kennedy given the hypothetical assumption that Oswald didn't is quite high, whereas the probability that someone else would have on the subjunctive supposition that Oswald didn't is at least substantially lower.

4. The central problem

Sprenger is explicit that his account allows the supposition of H to "overrule conflicting information" (Sprenger, 2019, 325), meaning that H is (at least sometimes) supposed in a subjunctive manner.[2] The problem for Sprenger's analysis is that once we account for this fact, his argument no longer shows what he intends it to. If the type of supposition employed is subjunctive, what his argument shows is that $\rho_H(E)$ is equivalent to $p(E||H)$; because $p(E||H)$ is not generally equivalent to $p(E|H)$, the argument no longer warrants the equivalence between the probability density or mass function and our degrees of belief.

To see how $p(E|H)$ and $p(E||H)$ can come apart in a manner that is relevant to Sprenger's analysis, consider again the famous example from Adams. Let H be [[Oswald didn't kill Kennedy]], and let E be the proposition [[The Warren Commission concluded that Oswald didn't kill Kennedy]]. Suppose that our agent does not yet know the content of the Warren Report but nevertheless has relatively high priors that Oswald killed Kennedy. Plausibly, they should assign a relatively high confidence to $p(E|H)$: if Oswald didn't kill Kennedy, the report is relatively likely to say so—certainly much more likely to say so than if he did. When we suppose H subjunctively, by contrast, we get the opposite result. Assuming that Kennedy was not assassinated in the (nearest) worlds where Oswald didn't kill him (or that the alternative assassination did not prompt the Warren Commission's investigation), the probability of E in these worlds—that is, $p(E||H)$—is quite low. The upshot is that we're not justified in equating the probability that some evidence has on the subjunctive supposition of H with its probability given H understood in the traditional (hypothetical) manner.

This is not an idle problem and does not depend on the qualitative character of this example. So consider studies in which climate scientists investigate whether and to what degree humans are responsible for climate change. The hypotheses in these

[1] The most well-developed use of subjunctive conditional probabilities is in "interventionist" accounts of causation, where such probabilities are used to evaluate the probability of an outcome if "we" were to intervene to change the world in a specific way. See Pearl (2009); Spirtes, et al. (2000).

[2] Indeed, Sprenger (2019, 325) suggests that this is the key innovation of his account.

https://doi.org/10.1017/psa.2023.17 Published online by Cambridge University Press

studies are statistical; many of these studies are essentially regressions, and Bayesian attribution studies typically involve determining which value for anthropogenic temperature change has the highest posterior probability (see, e.g., Ribes, et al. 2021).

On the hypothetical assumption that humans are not responsible for climate change, we should expect that some other factor is. So, for example, we would expect the data to exhibit patterns consistent with a massive increase in solar energy output. Call these patterns E and the hypothesis that humans are not responsible for climate change H_0; the claim is that $p(E|H_0)$ should be relatively high. But again, if we have relatively high priors that humans are responsible for climate change, the subjunctive supposition that we're not yields the opposite conclusion. On the subjunctive supposition that humans are not responsible for climate change, there would be no climate change, and so we would not expect to discover patterns that indicate a massive increase in solar energy output; $p(E||H_0)$ should be quite low.

As these examples illustrate, hypothetical and subjunctive suppositions differ in systematic ways. These systematic differences block Sprenger's proposed justification for equating $p(E|H)$ and $\rho_H(E)$: his argument succeeds only if $p(E|H)$ and $p(E||H)$ are generally equivalent, and we've just seen that they aren't.[3]

5. Potential rejoinders

A defender of Sprenger has two potential rejoinders that are worth addressing here.

First, the defender might respond by retreating to a purely hypothetical account: although Sprenger has indicated that the supposition should be subjunctive by allowing it to overrule known information, perhaps the account would be better without this commitment. Indeed, dropping the subjunctive element of the account would allow the original argument to go through. Unfortunately, however, the resulting account no longer has the generality for which Sprenger aims. As noted, many—some would say all (e.g., Box, 1976)—of the models used in statistical reasoning are idealized—that is, we know that some of their assumptions are, strictly speaking, false. So we cannot hypothetically suppose that these assumptions describe the objective probabilities of the world without endorsing a contradiction from which everything will follow. Because Sprenger explicitly aims to give an account of how statistical modeling functions in "highly idealized" cases (Sprenger, 2019, 321), this hypothetical version of the account cannot be said to succeed either.

Second, the defender could respond that $p(E|H)$ should actually be interpreted subjunctively in Sprenger's discussion—that is, as $p(E||H)$.[4] Indeed, there are a number of places in his paper where Sprenger suggests something along these lines by saying (for example) that his interpretation of degrees of belief is "genuinely counterfactual" (Sprenger, 2019, 325). In the latter half of the paper in particular, Sprenger retreats from equating $\rho_H(E)$ with $p(E|H)$ to equating $\rho_H(E)$ with $p_{\mathcal{M}}(E|H)$, "the degree

[3] Notice that this problem doesn't (obviously) affect other uses of subjunctive supposition in confirmation theory. Garber (1983), for instance, appeals to the change in the probability of H upon learning E in the counterfactual scenario where E isn't known to solve the "ahistorical" problem of old evidence. Garber's position doesn't require equating $p(A|B)$ with $p(A||B)$—in fact, it only makes sense because these two quantities come apart, as Garber (1983, 103) explicitly notes—and so isn't undermined by the argument just given.

[4] My thanks to an anonymous reviewer for pressing me on this point.

https://doi.org/10.1017/psa.2023.17 Published online by Cambridge University Press

of belief in the truth of H that we would have if we had supposed that the target system is fully and correctly described by one of the hypotheses in $\mathcal{M}$" (Sprenger, 2019, 330)—a retreat that suggests that his goal is ultimately the equation of $\rho_H(E)$ with a subjunctively supposed quantity.

It seems unlikely that this reinterpretation is accurate: Sprenger is explicit that his goal is to connect $\rho_H(E)$ with the conditional probabilities that we actually have, that is, with $p(E|H)$. And it's easy to see why: what we want is an account that tells us whether, why, or when we're justified in taking the probabilities generated using idealized statistical models to constrain our own beliefs. This is the intuitive appeal of Sprenger's "main question." Although the imagined retreat would avoid the problem raised in the previous section, therefore, the cost is that the suppositional account now tells us only that we should equate the probabilities generated using idealized statistical models with the degrees of belief that we would have in certain nonactual worlds. It's not clear why that's relevant to the confirmation of hypotheses in this world.

6. What now?

What are the stakes for the success of the suppositional account? That depends on what we take the goal of the suppositional account to be. On one hand, if the goal is to show why we should allow the probability density functions delivered by statistical models to constrain our degrees of belief—and if Sprenger is right that other attempts to answer this question also fail—the failure of the suppositional account is a serious problem for (Bayesian) philosophy of statistics. If Bayesian statistics doesn't provide us with at least rough constraints on our degrees of belief, it's not clear what it's doing.

On the other hand, there are reasons to think that Sprenger's goal is in fact more modest. As noted, Sprenger ultimately suggests that even though the suppositional account forges a connection between $p(E|H)$ and $\rho_H(E)$, it's rarely a good idea to "naively" calibrate our degrees of belief to the probabilities generated by the statistics (Sprenger, 2019, 332). Instead, how these probabilities should constrain our degrees of belief depends on facts about the relationship between the particular model and its target. If this nuanced reading is right, the stakes appear much lower. I've shown that the argument that Sprenger gives for the equation of $p(E|H)$ and $\rho_H(E)$ cannot succeed, but it's not clear why he needs this argument—after all, what's really doing the work are local facts, not the appeal to suppositional reasoning that I've undermined in this article.

Acknowledgments. Funding for this article was provided by the Deutsche Forschungsgemeinschaft (German Research Foundation; DFG), project 254954344/GRK2073. My thanks to Mack Sullivan and two anonymous reviewers for comments on an earlier version of the article.

References

Adams, Ernest. 1975. *The Logic of Conditionals: An Application of Probability to Deductive Logic*. Dordrecht: D. Reidel.

Box, George E. P. 1976. "Science and Statistics." *Journal of the American Statistical Association* 71 (356): 791–99.

Garber, Daniel. 1983. "Old Evidence and Logical Omniscience in Bayesian Confirmation Theory." *Minnesota Studies in the Philosophy of Science* 10: 99–132.

Pearl, Judea. 2009. *Causality: Models, Reasoning, and Inference*. Cambridge: Cambridge University Press.

Ribes, Aurélien, Saïd Qasmi, and Nathan P. Gillett. 2021. "Making Climate Projections Conditional on Historical Observations." *Science Advances* 7 (4): 1–9.

Schwarz, Wolfgang. 2018. "Subjunctive Conditional Probability." *Journal of Philosophical Logic* 47: 47–66.

Spirtes, Peter, Clark Glymour, and Richard Scheines. 2000. *Causation, Prediction, and Search*. 2nd ed. Cambridge, MA: MIT Press.

Sprenger, Jan. 2019. "Conditional Degree of Belief and Bayesian Inference." *Philosophy of Science* 87 (2): 319–35.

Cite this article: Dethier, Corey. 2023. "Supposition and (Statistical) Models." *Philosophy of Science* 90 (3): 744–749. https://doi.org/10.1017/psa.2023.17