CONJUNCTIVE EXPLANATIONS

Jonah N Schupbach and David H Glass

Reviewed by
Naftali Weinberger

Conjunctive Explanations: The Nature, Epistemology, and Psychology of Explanatory Multiplicity
Jonah Schupbach and David Glass
New York: Routledge, 2023, £38.99
ISBN 9781003184324

Cite as:
Weinberger, N. [2024]: 'Jonah Schupbach and David Glass's *Conjunctive Explanations*', *BJPS Review of Books*, **2024**

Explanation is a strange beast. While theories of explanation have been central to the last century of philosophy of science, explanation itself is not a first-order concept in most scientific theories. Analysing explanation is not a matter of unpacking an antecedently determinate scientific term, but rather part of a philosophical project of systematizing the methods, outputs, and aims of science. Unsurprisingly, there are many different approaches to

pursuing this, including those rooted in psychology, linguistics, history, probability theory, formal epistemology, and causal inference, and these approaches have been pursued largely independently. It thus seems fair to ask: is there any epistemic or cognitive benefit to pursuing this heterogeneous set of projects under the label of 'explanation'?

After reading Jonah Schupbach and David Glass's edited volume, *Conjunctive Explanations: The Nature, Epistemology, and Psychology of Explanatory Multiplicity*, I feel much more confident in answering this question with a yes. It certainly helps that every essay in the volume is well written and makes a substantive contribution. But the real test for an edited volume is whether it is greater than the sum of its parts. What makes this book so rewarding is that it allows one to repeatedly revisit a core set of issues from different angles and with distinct tools. In this review, I present what I take to be the unifying themes of the book, with the aim of conveying its broad interest to a wide range of philosophers of science.

Conjunctive explanations are those that combine multiple explanatory resources rather than appealing to a single one. Such explanations include: the appeal to multiple causes to explain an effect, the use of multiple types of explanatory norms to understand a phenomenon, and the combination of independent explanatory frameworks for navigating the world. Concrete examples from the volume include: attributing the extinction of the dinosaurs to the combination of the asteroid impact and volcanic eruptions, the combination of individual and aggregate reasoning in economics, the combination of folk and scientific theories in reasoning about COVID, the combination of religious and scientific beliefs in explaining the world, and the combination of functional and mechanistic reasoning in thinking about living kinds.

The volume is divided into three sections, corresponding to the nature, epistemology, and psychology of conjunctive explanations. The section on epistemology is the most technical, as it focuses mainly on probabilistic approaches to explanation, with an emphasis on inference to the best explanation. While the chapters in this section would be a challenge for someone with no background in probability theory, they nevertheless serve as paradigm examples of clear writing and pedagogy. All of the crucial technical details are also explained informally in a way that clarifies the key motivating intuitions.

Taking a step back from the details of the ten contributions, there are three sets of unifying questions that struck me as salient throughout the volume:

- Combination (chapters 5 and 6): How do the distinct components of a conjunctive explanation combine with one another, and what justifies the costs related to the increased complexity of the conjunctive explanation over its conjuncts?
- Abstraction (chapters 2, 3, and 7): What is the nature of explanations that represent a system in a coarse-grained way, and why might one be justified in appealing to multiple levels of explanation?
- Integration (chapters 1, 8, 9, and 10): When multiple explanatory schemas are invoked in an explanation, how does this interaction function and is it justified?

These issues are not fully independent of one another. Take abstraction and combination. A popular answer to the question of why a combined hypothesis might be preferable to its conjuncts is that the combined conjuncts are unified in terms of a more abstract property (for example, 'mass' instead of inertial and gravitational mass). And combination and integration are likewise clearly similar. Here I use 'combination' to involve hypotheses defined within the same framework (invariably denoted as H_1, H_2,...) and 'integration' for combining elements from different frameworks (for example, those appealing to religious and scientific evidential norms).

Tomoji Shogenji's contribution in chapter 4 is an outlier in being the least focused on conjunctive explanations, but provides the clearest motivation for discussing them at all. His focus is not on the comparison or combination of explanations, but on giving an account of when an observation calls for an explanation in the first place. He argues that the prior probabilities and likelihoods at the core of Bayesian epistemology can't distinguish between observations that intuitively do and do not call for an explanation, but his point here generalizes beyond Bayesian analyses. The issue is that because the probabilities are intended to reflect everything relevant to the truth of a hypothesis, it is unclear how explanatory norms not captured by these probabilities can play a 'substantive role in epistemic evaluation' (p. 88). This question of how explanatory virtues can serve as something more than a frill when combined with unambiguously epistemic ones such as truth and accuracy is always in the background when discussing conjunctive explanations. We'll see that both combination and abstraction invariably involve a trade-off between such virtues.

Shogenji highlights Schupbach and Sprenger's ([2011]) concession that their probabilistic measure of explanatory power cannot by itself measure explanatory strength. This is because their purely probabilistic account cannot capture explanatory asymmetries such as those between cause and effect. One might imagine that the solution is therefore to simply replace their conditional probabilities with causal ones (cf. Eva and Stern [2019]). That Shogenji goes a different route, but is still able to propose an account of the asymmetry of explanation, suggests a novel way of explicating the latter asymmetry as being more general than the causal one.

Full disclosure: Most of the people publishing on explanatory strength are friends and former colleagues of mine. I must nevertheless confess that prior to reading the chapter by (current colleagues) Stephan Hartmann and Borut Trpn, I was unsure why this project was anything more than a purely intellectual exercise in probabilifying yet another domain. This worry was assuaged by the clarity with which they motivate their own account as measuring the degree to which the individual hypotheses in a conjunctive explanation cohere with one another, and their contrasting this with the standard explication of explanatory strength as the amount by which the explanation reduces how surprising one's evidence is. The focus on the coherence of conjunctive explanations not only links explanatory strength to the issue of when complex explanations are improvements as opposed to merely *ad hoc*, but also builds a satisfying bridge between quantitative analyses of explanatory strength and the old-school problem of irrelevant conjuncts. Accounts of explanation have long struggled to yield the result that combining arbitrary hypotheses does not improve them. It is a virtue of Hartmann and Trpn's account that it has the resources to address this traditional problem and uses them to then develop a more general quantitative framework.

Hartmann and Trpn's account in the fifth chapter is fruitfully read alongside Glass and Schupbach's account in the sixth, which has many similarities, but which diverges in some key details. Glass and Schupbach present the most precise account of the trade-off between the explanatory gain from combining hypotheses and the explanatory loss resulting from the increased complexity. This trade-off can be seen as the core dilemma when theorizing about the set of questions related to combination, and arguably of analyses of conjunctive explanation more generally.

I now turn to abstraction. Leah Henderson (chap. 7) considers whether competing statistical models should be treated as mutually exclusive. Her discussion is situated in a debate going back to Popper ([2005]) regarding the justification for accepting statistical models with fewer degrees of freedom (for example, a linear model over a parabolic one), given that models with more degrees will, logically, always fit the data at least as well. This discussion has led philosophers to engage with information-theoretic model selection criteria such as the Akaike information criterion. Henderson argues that, contrary to appearances, models with strictly more degrees of freedom are not compatible with those with fewer. Her main point is one about abstraction, and relies on an analogy between statistical and causal models. Causal models with more causal arrows (and thus more adjustable parameters) are not mere extensions of those with fewer; rather, they diverge in terms of the proposed (more abstract) data generating mechanism. Similarly, even though the data sets compatible with a linear model are a

subset of those compatible with a parametric one, Henderson claims they correspond to different data generating mechanisms. I hope that philosophers will go on to further discuss this analogy, which raises foundational questions about the connection between causal models and model selection criteria.

Steve Petersen's contribution in chapter 2 similarly considers issues of abstraction within an information-theoretic context. His discussion is motivated by a unificationist view on which explanation is about finding patterns, and the key philosophical challenge is to understand the relationship between partial explanations and complete explanations (that is, Railton's ([1981]) 'ideal explanatory text'). In addition to compellingly rebutting an argument by James McAllister ([2007]) that algorithmic information theory cannot allow for multiple non-competing partial explanations, the chapter is enhanced by a clear discussion of Kolmogorov complexity and the focus on a time series example.

The final chapter focusing on issues of abstraction is Samuel G. B. Johnson and Michiru Nagatsu's 'Individual and Structural Explanations in Scientific and Folk Economics' (chap. 3). This chapter is also relevant to scholars working on structural explanation in other contexts (for example, discrimination) and, given the influence of economics in public policy and public opinion, has broad implications for teaching and scientific communication. Although this contribution resembles Andrew Shtulman's on COVID (chap. 10), insofar as both compare scientific and folk theories of a domain, the contributions engage differently with the topic of conjunctive explanation. In Johnson and Nagatsu's chapter, the focus is on the combination of both individual and aggregate explanations within economics, alongside evidence that non-scientists appreciate the individual much more than the aggregate elements. In Shtulman's chapter, the focus is on how both scientific and folk conceptions of disease are integrated within folk conceptions of disease transmission.

Continuing on the theme of integration, Telli Davoodi and Tania Lombrozo (chap. 9) empirically explore a form of conjunctive explanation in which people employ both religious and scientific explanations in accounting for phenomena. They provide the most detailed discussion in the book of the nuances involved in testing for conjunctive explanations. To show that people integrate religious and scientific reasoning, it is not enough to claim that some target phenomena are explained one way and others the other way. Rather, one must find evidence that the patterns of reasoning interact with one another in a combined explanation. Their findings are of interest to those interested in Steven Jay Gould's notion that science and religion have 'non-overlapping magisteria', as well as to those interested in the interaction of different types of values in scientific explanations.

Frank C. Keil (chap. 1) similarly explores the psychology of a type of integrated explanation, specifically, the relationship between functional and mechanistic explanations. While Davoodi and Lombrozo consider the psychology of reasoning, Keil's project fits within the psychology of concepts. Both contributions reveal the fruitfulness of studying conjunctive explanation through the lens of psychology. Because of concerns that functional explanations are teleological and thus, some might suggest, only available to theists, Keil's contribution is also relevant to understanding the relationship between science and religion. These two contributions are thus the most clearly relevant to the topic of the Templeton Foundation grant that funded the book, 'Conjunctive Explanations: How Science and Religion Can Work Together'.

I noted earlier that questions of integration differ from combination insofar as they involve combining explanations from distinct frameworks. Igor Douven's contribution in chapter 8 gets to the question of what a framework is. Building on work with Peter Gärdenfors ([2004]), he argues that conceptual frameworks can be understood using the mathematical theory of conceptual spaces. Using the colour space as a guiding example, Douven argues that there is a way to understand the partition yielded by a conceptual space as optimal without being committed to a robust theory of natural kinds. 'Optimal' here is determined by a trade-off between contrast (colours must be maximally easy to distinguish) and representativeness (each instance of a colour should be maximally similar to

other instances). The notion of optimality helps address the concern that conceptual framework-relativity implies that the external world places no constraints on our ontology, while still allowing that the optimal trade-off might differ across distinct frameworks. The account thus provides an interesting test case for theories of realism and objectivity.

Looking at the contributors, no one should be surprised that this book is required reading for any philosopher of science interested in model selection, Bayesianism, algorithmic complexity theory, inference to the best explanation, or measures of explanatory power. What perhaps is surprising is the range of core philosophy of science topics this volume touches on, including realism, values in science, science and religion, parsimony, scientific understanding and communication, and explanatory unification. This volume is thus of broad interest, and I could easily envision it being used to organize a graduate seminar in which each contribution was taught as part of a broader section on a standard topic from the philosophy of science. This volume shows that there is still a lot of exciting work to be done on the topic of explanation, and what makes it especially satisfying is the way that the contributions are situated within the broader tradition.

Acknowledgements

Thanks to Ina Jäntgen for feedback on an earlier draft.

Naftali Weinberger
Munich Centre for Mathematical Philosophy
naftali.weinberger@gmail.com

References

Eva, B. and Stern, R. [2019]: 'Causal Explanatory Power', *British Journal for the Philosophy of Science*, **70**, pp. 1029–50.

Gärdenfors, P. [2004]: *Conceptual Spaces: The Geometry of Thought*, Cambridge, MA: MIT Press.

McAllister, J. W. [2007]: 'Model Selection and the Multiplicity of Patterns in Empirical Data', *Philosophy of Science*, **74**, pp. 884–94.

Popper, K. [2005]: *The Logic of Scientific Discovery*, Abingdon: Routledge.

Railton, P. [1981]: 'Probability, Explanation, and Information', *Synthese*, **48**, pp. 233–56.

Schupbach, J. N. and Sprenger, J. [2011]: 'The Logic of Explanatory Power', *Philosophy of Science*, **78**, pp. 105–27.