Book of Abstracts

"New Perspectives on Causation in the Life Sciences" Conference

27 – 28 June 2022

Department of Philosophy and Centre for Reasoning

University of Kent, UK

Sponsored by the British Academy and Leverhulme Trust

Contents

Biological function, causal explanation, and natural selection 2

Proximate epigenetic effects and the role of causal graph theory 2

Mechanistic explanation and evolutionary causal chains 3

Cascade versus mechanism: The diversity of causal structure in science 4

Dynamical causes & explanations in systems biology 4

Between nature and culture: developing a third way for causal cognition 5

Causation, observation, and experiment: reflections on practical inquiry in the German lands around 1800 6

Causation in a world of process 6

Defending the good cause: phenotypes determine genotypes 7

Demystifying downward causation in biology 7

Natural selection and the metaphysics of causation 8

An emerging dilemma for reciprocal causation 8

Variable relativity of causation is good 9

Some philosophical and methodological issues raised by polygene risk scores 10

Biological function, causal explanation, and natural selection

Samir Okasha (University of Bristol, UK)

How to understand function-attributing statements (e.g. "the function of the cactus's spines is to deter herbivores) is a classic topic in the philosophy of biology. Proponents of the "selected effect" (SE) theory of function (also known as the etiological theory) claim that such statements make implicit reference to evolutionary history: a trait in a modern organism has a function in virtue of some effect the trait had in the organism's ancestors, which led to the trait being selected. The SE theory enjoys considerable support among philosophers and biologists alike. In a recent paper, J. Christie et al. (2022) argue that the theory rests on an overly simple model of evolution by natural selection and can't easily be generalized to more complex cases. In particular, they claim that the SE theory doesn't work under frequency-dependent selection (FDS), where the selective environment changes as the population evolves. Christie et al. claim that in such cases, citing a trait's SE function is not tantamount to giving an evolutionary explanation for why the trait exists, so what they see as the signature advantage of the SE theory fails.

In this paper I respond to Christie et al.'s critique. I argue that they are right that the SE theory, as usually formulated, rests on simplifying assumptions, however they can be fairly easily relaxed. In particular, the SE theory generalizes easily to FDS in cases where the population evolves to a stable equilibrium, so long as we are prepared to define adaptation in terms of "best response" rather than individual optimization. However, non-equilibrium dynamics pose a much more serious challenge. I end with some reflections on the teleosemantic project, which relies on the SE theory.

Proximate epigenetic effects and the role of causal graph theory

Tom Dickins (Middlesex University, UK)

By decomposing the Price equation into components that denote (a) the response to selection and (b) reproduction, Otsuka has shown where this statistical formulation can be rendered causal. This has been done by adopting causal graph theory. Otsuka has then extended the utility of the equation by incorporating proximate elements, including epigenetic processes, into a foundational causal graph and showing how this can be expressed in Pricean terms. For epigenetic processes Otsuka claims these additions to be to (b) and one of three demonstrations that proximate causes can affect ultimate causation. Whilst I agree with Otsuka that ultimate causation depends upon proximate causation, I disagree with some key assumptions made when constructing the graph. Otsuka's graph treats epigenetic marks (C) and genotypes (X) as causes that independently feed into the expressed phenotype (Z). However, epigenetic marks causally impact on the genotype directly, and it is via the genotype that they have their phenotypic effect. This means that for epigenetic heritability to work there must be consistent genetic heritability. Statistically we would note that Z(C|X), where Z takes a value when C, given that X. But causally we should note that Z(X|doC), where Z takes a value when X given that C is done. This presents a counterfactual argument supporting a view of epigenetic causes that preserves some fundamental aspects of Modern Synthetic thinking, which is consilient with Otsuka's expressed aims. But it also leads to some complexities when trying to re-render the causal diagram to incorporate C

appropriately. For example, C can be activated in response to an input from the environment (E) and can also be classed as a phenotypic response to E, something already within the foundational graph (E→Z). That phenotypic response then modifies the availability of DNA sequence for subsequent polypeptide chain formation, for example. Those modifications are constrained and conditional. C, E, X and Z can all carry slightly different interpretations under various schemes. For example, where X might be seen as the whole genotype, C, as an epigenetic mark, is affecting only one component of that. Are we modelling an allele with X at that point, or a genotypic state change? It is easy for the granular situation to proliferate and for clarity regarding evolutionary process to be lost. I will discuss this and claim that Otsuka's use of causal graph theory helps to expose this issue. But it also helps us to see that Mayr's rendition of the proximate-ultimate distinction was not an idealization, but rather an abstraction. I interpret abstraction as abstractness, in keeping with Levy's arguments about the relative richness of detail. Crucially, I see no incorporated falsehoods in Mayr's version of the distinction, which makes it less likely to be an idealization, in keeping with Potochnik. I will conclude with some broader comments about the tensions between those seeking to extend the Modern Synthesis, and those who resist this, noting that graph theory will not present a rapprochement but can help to clear the ground as long as we inspect our assumptions throughout.

Mechanistic explanation and evolutionary causal chains

Davide Vecchi (University of Lisbon, Portugal)

The thesis of the causal primacy of ultimate causation (Mayr 1961) is grounded on the adaptationist hypothesis that the evolution of complex traits is the result of cumulative, multi-generational natural selection. The corollary of this view is that the "reductionist" sciences studying proximate causes (e.g., biochemistry, molecular and developmental biology) can only, in principle, identify the small, often phenotypically insignificant steps of the long series of ontogenetic changes accumulated by "creative" selection. Rather than being untenable, this view is biased and speculative. It is biased because, while it is surely the case that the evolution of many complex non-molecular traits requires a causal explanation significantly appealing to cumulative, multi-generational selection, it cannot be assumed that this is the case for every trait. It is speculative because it is impossible to faithfully reconstruct the long series of events punctuating evolutionary causal chains, particularly those of special interest to Mayr and fellow "organismal biologists" (Dietrich 1998), mostly interested in "highly evolved" multicellular eukaryotes (Mayr 1988, p. 9723). Intriguingly, Francis (1990, p. 412) argued that, in order to "causally" understand evolutionary dynamics, the analysis should be "mechanistic". In this talk, I shall argue that, in some limiting cases, mechanistic analysis might be approximated. I shall focus on the evolutionary dynamics studied in, for instance, experimental evolution (Blount et al. 2012), ancestral protein reconstruction (Castro-Fernandez et al. 2017) and protein engineering (Alexander et al. 2017). I shall argue that, despite the abstraction and idealisation underlying the explanatory models proposed, the above research can be interpreted as decomposing evolutionary causal chains in terms of mutational, developmental and sorting events. Mutational and developmental events are at the basis of trait origin (a process more obviously mechanistically accessible by opening the variational black box), while sorting events (i.e.,

organismal selection and drift) cause the increase in frequency of the trait in the population or lineage. At least two aspects of the mechanistic account of evolutionary causation implicit in the molecular research on which this contribution focuses are philosophically relevant for understanding causation in evolutionary biology. First, this research suggests that, contrary to what Mayr argued, ultimate causation is in principle reducible to the proximate mechanistic causes at the basis of the mutational, developmental and sorting events affecting the ontogeny of the individual organisms constituting the population or lineage. Secondly, this research shows that it is sometimes clearly possible to distinguish between the sorting by selection and drift of the organisms possessing specific genomic and phenotypic properties (pace Walsh et al. 2002). Mayr would have probably rejected the significance of this kind of research, often focused on the molecular phenotypes of unicellular populations of organisms (Beattie 1994). Nevertheless, a general conceptualisation of evolutionary causal chains should also be able to account for molecular dynamics in unicellular lineages. If it cannot, then it is either incomplete or untenable. An alternative is that the project of seeking a unified understanding of evolution might be delusional (Stoltzfus 2017).

Cascade versus mechanism: the diversity of causal structure in science

Lauren Ross (University of California, Irvine, USA)

According to mainstream philosophical views causal explanation in biology and neuroscience is mechanistic. As the term "mechanism" gets regular use in these fields it is unsurprising that philosophers consider it important to scientific explanation. What is surprising is that they consider it the only causal term of importance. This talk provides an analysis of a new causal concept–it examines the cascade concept in science and the causal structure it refers to. I argue that this concept is importantly different from the notion of mechanism and that this difference matters for our understanding of causation and explanation in science.

Dynamical causes & explanations in systems biology

Russell Meyer (University of Wollongong, Australia)

In this talk, I attempt to resolve a debate within the philosophy of biology about the viability of non-mechanistic causal explanations in the face of criticisms from within the mechanistic approach to explanation.

Causal explanations in biology are often associated with Mechanism. Mechanism aims to explain phenomena by describing the causally active parts and activities that underlie and produce those phenomena. The transition from mere description to causal explanation, according to mechanists, requires a particular sort of description: one that is progressing towards an accurate and comprehensive accounting of the causal structure underlying the target phenomenon (Kaplan & Craver 2011).

One of the motivations for this mechanist approach is an emphasis on revealing this causal structure (Salmon 1984) – the web of causal relationships that are meant to underlie the

phenomena we observe and study. Describing the parts and activities that make up mechanisms is, on this view, a way to capture this causal structure in the life sciences. This approach to causal explanation, however, does potentially leave other kinds of model explanations out in the cold. Dynamical systems models, for instance, do not tell us anything (by themselves) about mechanisms. These models are instead a mathematical modelling tool used across a variety of research programs, capable of describing and predicting the behaviour of complex systems while abstracting away from the details of physical implementation (Stepp, Chemero & Turvey 2011).

Utility aside, mechanists dispute whether a model that does not describe mechanisms can adequately describe causes, capture causal structure, and consequently explain. A commonly held view within mechanism is that, while descriptively indispensable, dynamical models must be mapped onto an underlying mechanism – grafted onto a genuinely causally descriptive model – in order to say anything about causes (Kaplan 2015). This framing of the situation has become the received view around the viability of dynamical explanation. Here I argue against this position and show that dynamical models can indeed reveal the causal structure underlying biological phenomena and consequently produce non-mechanistic, dynamical explanations. Taking the example of cell fates from systems biology, I show how dynamical models, through modelling techniques like attractor landscape analysis, identify the causes of cell differentiation and explain why cells select particular fates. While dynamical models do not identify mechanistic parts or activities, they do identify causes in the interventionist (Woodward 2003) sense of adequately answering counterfactual what-if-things-had-been-different questions.

Moreover, the descriptions of cell fate differentiation provided by dynamical models better fit Woodward's (2018) criteria of specificity and proportionality, making them superior candidate causes of cell fates when compared to mechanistic models. I also show how these dynamical causes are not practically reducible to underlying mechanistic models, a process that proves cumbersome and potentially counterproductive.

Between nature and culture: developing a third way for causal cognition

Rachael Brown and Ross Pain (Australian National University, Australia)

What is the role played by causal cognition in the evolution of cumulative technological culture? Addressing this question leads to a further three questions: (i) do you need causal understanding of technology to pass it on to others?; (ii) do you need causal understanding of technology to gain it from others?; and (iii) do you need causal understanding of technology to innovate adaptively?

Theorising on the evolution of cumulative technological culture tends to fall into one of two camps: the nativist view and the cultural view. Nativists (e.g. Osiurak and Reynaud 2020) argue that causal cognition is a specialised, genetically inherited mechanism. In the case of complex technology, they respond "yes" to each of the above three questions. Proponents of the cultural view (e.g. Harris, Boyd, and Wood 2021) argue that cumulative cultural

technology is the product of high-level, 'blind' cultural evolutionary processes. They respond "no" to each of the above three questions, even in the case of relatively complex technologies.

We think there are important problems with each of these views. In the case of the cultural view: while causal information mightn't be necessary for cumulative technological culture, it's very plausible that causal cognition is relevant to understanding the actual evolution of cumulative technological culture. In the case of the nativist view: evolutionary theory tells us that the de-novo generation of an entirely new domain-specific complex cognitive mechanism via undirected means is at best highly unlikely, and at worst implausible.
In this paper, we build on previous work (Pain and Brown 2020) to develop a third way for the evolution of causal cognition and cumulative technological culture; one that commits to a gradualist evolutionary framework but also does justice to the significance of high-level cultural scaffolding. Central to our account is Starzak and Gray's (2021) recent proposal of a three-dimensional model for causal cognition, which splits causal cognition into three vectors: integration, explicitness, and sources of causal information. We argue that, in the case of cumulative technological culture, we need to focus on the evolution of integration. But we also make the case that causal cognition has a minimal precursor in an organism responding to its own interventions on the world.

Finally, following Woodward (2021), we suggest that thinking about the nature and evolution of causal cognition has implications for how we understand causation as a metaphysical concept.

Causation, observation, and experiment: reflections on practical inquiry in the German lands around 1800

Jutta Schickore (Indiana University Bloomington, USA)

This paper examines reflections on the nature of causation in German science in the decades around 1800, specifically in those sciences that concern human health and prosperity. Reflective practitioners, some academic philosophers, and educators discussed how to identify and exploit cause-effect relations in practical contexts such as making medicines, growing plants, extracting metals from ores, and warding off superstition. These inquirers developed heuristics as well as evaluative criteria for finding and assessing hypotheses about cause-effect relations and discussed how experiments and observations could assist in this endeavor. Their writings on "applied logic" [angewandte Denklehre] connect the practical problems of scientific inquiry with the lofty discourse on the metaphysics and epistemology of causation in Kantian philosophy.

Causation in a world of process

John Dupre (University of Exeter, UK)

In this talk I will briefly describe the world populated by persistent processes that I take to be ours. I shall focus especially on central examples of living processes, organisms and, to a

lesser extent, lineages, both of which I account as causally powerful processes. I claim that generally causation occurs when processes interact, and in such interactions they exercise capacities. Living systems are hierarchies of processes; as processes, they are sustained by activity; and that activity consists both of interactions among constituent processes, and the interactions of the whole with other processes. Articulation of this processual view of causation implies a pluralistic view of causation and suggests that there is no basis for the common scepticism about so-called downward causation.

Defending the good cause: phenotypes determine genotypes

David Haig (Harvard University, USA)

Cause is polysemous. Sometimes cause refers to the particular causes of particular events (token causation) and sometimes it refers to relations between kinds of events (type causation). Sometimes cause refers to how one thing impacts another thing (causation as mechanism) and sometimes it refers to differences that make a difference (causation as difference making). These differences of meaning can be sources of confusion even before considering the nuances of Aristotle's four causes or the disputed difference between proximate and ultimate causation. When one scratches beneath the surface, many disputes in evolutionary biology involve different commitments toward the status of Darwinian final causes or adaptive functions.

Darwinian final causes originate by natural selection among repeated differences. There can be no selection without a difference. This is not token causation but type causation. It involves difference making rather than mechanism. Some genetic differences are causes of phenotypic differences in the world. If a phenotypic difference is responsible, however indirectly, for preferential replication of the genetic difference, then the gene has been replicated, in part, because of this phenotypic effect which was a difference in the world that made a difference in replication. Phenotypes are causes of the replication and continued existence of genotypes. Thus, a gene's effects can be counted among its causes. A gene's purposes, or final causes, are the phenotypic effects 'for the sake of which' the gene exists. These causes are what worked in the past and are best predictions of what will work in the future. They are statistical summaries of historical difference-making.

Demystifying downward causation in biology

Yasmin Haddad (McGill University, Canada)

Despite the ambivalence of philosophers towards the notion of downward causation, it is a widespread concept in biology, where it is frequently used in an explanatory capacity to account for certain regularities and processes. Building on interventionist theories of causation, I propose a conceptual framework that demystifies what is meant by 'downward causation' using examples from the field of ecological evolutionary developmental biology. Downward causation in biology is thought to be problematic because it relies on the assumption that entities are connected by compositional hierarchies of levels of organization. I introduce examples of weak and strong compositional relations and argue that downward

causation becomes unproblematic if we use features of interventionist theories of causation. This requires a shift from entity-thinking to variable-thinking. I show that an interventionist account of downward causation successfully responds to the three central objections to downward causation in the philosophical literature and clarify the explanatory usefulness of the concept in biology by examining three empirical examples.

Natural selection and the metaphysics of causation

Charles H. Pence (Université catholique de Louvain, Belgium)

Contemporary arguments concerning the nature of causation in evolutionary theory, now often known as the debate between the "causalist" and "statisticalist" positions, have involved answers to a variety of independent questions – definitions of key evolutionary concepts like natural selection, fitness, and genetic drift; causation in multi-level systems; or the nature of evolutionary explanations, among others. In recent work, I have argued that a significant cluster of these questions, largely concerning causation, can be disconnected from the larger group and analyzed in isolation, allowing, I hope, for more fruitful engagement between the philosophy of biology and the metaphysics of science. Natural selection poses an under-studied but nonetheless not unique sort of causal structure that merits further exploration. In this talk, I will present some of these results, and highlight some of the exciting opportunities for future research that I believe this engagement could provide.

An emerging dilemma for reciprocal causation

Caleb Hazelwood (Duke University, USA)

In his recent monograph, The Causal Structure of Natural Selection, Charles Pence argues that the concept of "emergence" may be a helpful resource for untangling the causes of adaptive evolution. Within an emergence framework, natural selection is understood as a population-level process that arises from interactions among individuals and their environment. Pence additionally suggests that the emergent character of selection may help us understand concepts such as "reciprocal causation," i.e., the increasingly popular view that adaptive evolution is a bidirectional process, with selection and individual developmental processes such as niche construction co-causing one another.

In this talk, I follow Pence's suggestion and use the emergence framework to assess the concept of reciprocal causation. However, instead of helping us understand it, as Pence suggests, I will argue that the emergent character of natural selection puts reciprocal causation on the horns of a dilemma. Moreover, each horn appears to specifically threaten the claim that the concept captures a causal interdependency between population-level and individual-level processes.

The first horn embraces the emergent character of natural selection but denies its causal efficacy. This option resembles the "statisticalist" view of natural selection, whose proponents argue that the only genuine causes of evolution are those that operate on

individuals. Statisticalism easily handles the emergence of natural selection by rendering it a mathematical byproduct of individual births and deaths. But while this position can accommodate emergence, it is uninviting to the proponent of reciprocal causation, for whom niche construction and natural selection are interacting causal processes. Since statisticalism denies that natural selection is a cause, this view is incompatible with reciprocal causation.

The second horn embraces the causal efficacy of natural selection but struggles to accommodate its emergent character. This option resembles the "causalist" view, whose proponents argue that natural selection is a genuine causal process. But what kind of causal process? If selection is an emergent phenomenon, then it arises at the population level. On this view, reciprocal causation depicts individual-level processes (such as niche construction) as the causes and effects of a population-level process (natural selection). The problem, however, is that individual-level processes and population-level processes do not stand in a causal relation, but a compositional one. Because natural selection supervenes on the underlying individual processes, their relationship fails to meet an important criterion for establishing causation—namely, that causes are spatiotemporally distinct from their effects. Therefore, to say that natural selection and niche construction co-cause one another is to commit a category mistake.

To dodge the dilemma, the proponent of reciprocal causation may deny the emergent character of selection. If we can provide an individual-level account of selection, we avoid the category mistake and restore the causal feedback between selection and development. In response to this objection, I will argue that we have reason to be pessimistic about the possibility of an individual-level account of selection. At best, it is too costly to adopt, as it undermines the advantages of population thinking. At worst, it is conceptually incoherent.

Variable relativity of causation is good

Veli-Pekka Parkkinen (University of Bergen, Norway)

Interventionism is a theory with a pragmatic goal: to define causal concepts that are useful for reasoning about how things could be purposely manipulated. In (Woodward, 2003), interventionist causation is relativized to an analyzed variable set. In Woodward (2008) the definition of contributing cause or plain causal relevance is derelativized. This move has not gathered much attention, presumably because it is seen as an unproblematic way to honor the intuition that causal relations are objective. I argue that this move is problematic. Derelativization entails two concepts of unmediated causal relation that are not co-extensional, but which nonetheless do not entail different conclusions about manipulability within any variable set. This is in conflict with the pragmatic orientation of interventionism. I argue that no obvious solution for this problem exists that does not generate further problems, and conclude that interventionist causation should not be derelativized. I offer the following considerations to render that conclusion acceptable. When causation is defined as manipulability, causal concepts only become applicable when a distinction is drawn between a target system that one asks causal questions about, and its environment, from which it can be manipulated (Kuorikoski, 2014). This distinction is drawn by an agent that engages in causal reasoning, is influenced by the interests and knowledge of the agent, and amounts to a

decision to focus on one variable set rather than another. It is unclear what it means to say that the world as a whole has an objective and determinate causal structure qua a structure of manipulability relations, as the concept of causation as manipulability only applies to systems defined in some local context. Questions about manipulability relations only have determinate answers when the boundaries of the target system have been settled, and those answers may vary depending on how the boundaries are drawn. Hence, if causal concepts are to guide reasoning about manipulability, they should exhibit sensitivity to the choice of the analyzed variable set. This is not metaphysical relativism. Presumably the world has an objective structure at some fundamental level, studied by fundamental physics. As long as this structure gives rise to dependencies between some magnitudes representable as variables, causal concepts function as tools that causal reasoners use to parse those dependencies into ones that support manipulations of locally defined systems, and ones that do not. Once the boundaries of the target system are settled, questions about manipulability have objective answers determined by the underlying dependencies. This view is supported by the observation that causal concepts seem indispensable in disciplines like biology that study such locally bounded systems, but not similarly indispensable in physics. If one opposes interventionism because of variable relativity, one then opposes the very idea that causal concepts primarily function to guide reasoning about manipulability.

Some philosophical and methodological issues raised by polygene risk scores

James Woodward (University of Pittsburgh, USA)

This talk is based on joint work with Professor Kenneth Kendler (Virginia Institute for Psychiatric and Behavioral Genetics, Virginia Commonwealth University). It will explore some philosophical issues that arise in connection with the use of polygene risk scores (PRSs). These are increasingly widely used to predict various traits or disorders in populations from information about aggregate associations between those traits/disorders and many individual single nucleotide polymorphisms of very small individual effect (SNPs). I will discuss the methodological assumptions behind the use of such scores and whether they can be given a causal interpretation. I argue that one can think of the scores as functioning (roughly) as indicator variables-- they don't directly provide causal information but they track or indicate the presence of causal relationships between genetic variants (typically *different variants* than the SNPs incorporated into the PRS) and traits. I also argue that although PRSs themselves lack a direct causal interpretation they can aid in the discovery of other causal relationships, including those having to do with environmental influences. I also consider other relatively novel features of PRSs such as the comparatively "atheoretical" and inductive methodology involved in their construction and the fact that, although very useful, they do not provide "mechanistic" information.