

# Book of Abstracts

“Alternative Approaches to Causation: Beyond Difference-Making  
and Mechanism” Conference

28 – 29 June 2021

Department of Philosophy and Centre for Reasoning

University of Kent, UK

Sponsored by the British Academy and the Leverhulme Trust

## Contents

<b>Causal Inquiry: Hume’s Real Legacy (Simon Blackburn)</b> .....	2
<b>Factual Agent Causation (Amit Pundik)</b> .....	2
<b>A Dispositional Account of Causation: Application to Biological Sciences (Rani Lill Anjum and Elena Rocca)</b> .....	3
<b>Three Errors about Causal Efficacy (Rögnvaldur D. Ingthorsson)</b> .....	4
<b>What We Do to Each Other: Folk Psychology, Interventionism, and Special Causal Concepts (Julian Bacharach)</b> .....	5
<b>Novel Information-Theoretic Measures of Causation in Cognitive Science (Joe Dewhurst)</b> .....	6
<b>Causal Pluralism and Carnapian Explication (Matthias Rolffs)</b> .....	7
<b>Causal Pluralism: A New Perspective (Ned Hall)</b> .....	8
<b>Epistemic Causality in the Social Sciences (Jon Williamson)</b> .....	8
<b>An Epistemic Approach to Causation in Social Scientific Case Study Research (Rosa W. Runhardt)</b> .....	9
<b>Information Transmission and the Mosaic of Causal Theory (Federica Russo)</b> .....	9
<b>Of Powers and Causes: Groundwork for Causal Dispositionalism (Joaquim Giannotti)</b> .....	10
<b>Defending Non-Natural Intuitions About Causation: Why it’s Not Wrong to make Causal Judgements Based on Norms (Jon Bebb)</b> .....	11
<b>Resolving the Tension Between Different Approaches to Causation Within a Dynamical Systems Framework (Gergely Kertész, Balázs Gyenis, Gábor Hofer-Szabó, and Peter Fazekas)</b> .....	12
<b>An Inferentialist Definition of Causation (Christopher Clarke)</b> .....	13
<b>Causal Process Tracing Theory of Change (Nancy Cartwright and John Pemberton)</b> ..	13

## Causal Inquiry: Hume's Real Legacy (Simon Blackburn)

What is the best description of Hume's theory of causal commitments? Some have billed him as a sceptical realist, others as a regularity theorist, and yet others as inconsistent or unnecessarily hostile to natural science. In my talk I intend to take up a magisterial contribution to these debates by Peter Millican, in which he engages critically with Galen Strawson's recent revision of his book *The Secret Connexion*. Although I applaud Millican's offensive against Strawson, I am less enthusiastic about his final resting place, and in this paper I hope to explain why.

## Factual Agent Causation (Amit Pundik)

According to the theory of fact causation, the causal relata are neither events nor substances – but facts, namely true propositions (Bennett (1988) and Mellor (1995)). Under this view, the causal relata are abstract so they cannot be described in spatiotemporal terms, while spatiotemporal entities, such as events, are only truth-makers for the causal relata. Consequently, the causal relation is also not located in the space and time. These causal relata may well be real, as long as other abstract entities such as numbers may be real, and the causal relation may be fundamental, namely that it is irreducible to other non-causal phenomena or unanalysable in non-causal terms such as regularities of counterfactual dependencies.

Proponents of fact causation hold that its major advantage over event causation is its ability to account for causation by absences without denying that it is genuine causation (Bennett 1988, 140-41; Mellor 1995, 131-5; for responses, see, e.g., Dowe 2000, ch 6). In this paper, I suggest another advantage: when based on fact causation, an agent causal theory of free will can overcome some of the objections levelled against it. Agent causalist theories hold that actions may be free only if they are caused by the agent themselves and, consequently, they are often criticised for their reliance on substance causation, namely the view that some events, namely free actions or choices, are caused not by other events, but by certain substances, namely agents.

Indeed, however the causal relation between a substance as a cause and an event as its effect is spelled out, problems loom large. Clarke suggests an integrated account in which event and substance causation work in parallel, so every free action or choice is caused both by an event and an agent (Clarke 2003, 133-49). Yet critiques rightly point out that such parallel event causation renders agent-causation redundant (Ginet 2002, 397) and threatens freedom (O'Connor 2000, 78). O'Connor's own view seems to be that only substances are causes (O'Connor 2014, 33-34), but Clarke raises ten objections against the very possibility of substance causation (most of which also apply to his own integrated account) and cautiously concludes that 'there are, on balance, reasons to think that substance causation is impossible' (Clarke 2003, 221).

An agent causalist theory that is based on fact causation would not take the causal relation to hold between the agent itself as a concrete substance and the event of their freely choosing a certain course of action. Instead, the causal relation would hold between two relata of the same nature: the cause is the fact that (a) a certain agent exists and (b) the conditions that enable the agent to freely choose between certain courses of action are satisfied, and the

effect is the fact that they freely chose to act in a certain way. After briefly outlining what such an agent causal theory of free will would look like, I would turn to examine whether it can overcome the objections levelled against substance causation, and in particular, the three major problems that Clarke identifies ('the temporality of causation, the influence that causes have on the probability of their future effects, and the structured nature of entities that are directed in time', Clarke 2003, 209). I would seek to show that these objections stem from the dubious connection that substance causation makes between two entities of different ontological type (substance and event), and, consequently, an agent causal theory of free will that is based on fact causation is not vulnerable to these objections.

## A Dispositional Account of Causation: Application to Biological Sciences (Rani Lill Anjum and Elena Rocca)

There is a gap between the causal knowledge gained in the context of the lab and for what the knowledge is needed. What works in the lab doesn't always carry over to the context of application. This is a problem of low external validity, found in many areas of biology. We argue that this practical problem cannot be solved within the standard methodological and ontological framework. Instead, we need to start from a discussion of the nature of causation. We here propose that a dispositionalist theory of causation, as developed in Anjum and Mumford (2011, 2018), is better suited as a foundation for the real life contexts with which biologists are concerned.

We consider three examples from biology, where external validity of causal findings remains a methodological challenge.

The first challenge is the so-called 'translation crisis' in medicine, which refers to the problem of translating results from the lab to practical use in a population. For instance, less than 10 percent of the drug studies ever make it to the market, meaning that most of the molecules that work in pre-clinical research are clinically inefficient (Ioannidis 2005 PLoS Med).

The second challenge is the case of gene technology, where the application of results from tested environment to other environments represents a significant problem. Although genetically modified (GM) plants are assessed for molecular stability before it enters the market, there is no requirement of monitoring of commercial GM hybrids after they are cultivated in different environments. Adaptation to different environments might alter the phenotypic performance of the plant, which reduces the external validity of pre-commercial experiments.

The final challenge concerns the risk assessment of ecological impact of chemical substances. A chemical substance can be found safe in the lab, but numerous ongoing lawsuits from local communities against governments and industry suggests that these conclusions cannot always be trusted (Wickson & Wynne 2012 Ethics Pol Env). We here consider the shortcomings of standard risk assessment frameworks for predicting environmental effects of chemical mixtures.

We propose here that the problem of external validity, as presented here, has its origin in a scientific methodology influenced by Humean ontology and the regularity theory of

causation. Perfect regularities might exist in theoretical models or abstractions, including isolation of interferers and perfectly controlled experimentations. The problem is to assume that this framework could translate into the real life settings of biology. We here suggest that by changing the conception of causation to a dispositionalist one, one must also change scientific methodology and practice.

Causal dispositionalism is developed by Anjum and Mumford (2011 *Getting Causes from Powers* and 2018 *Causation in Science*), which offers a notion of causation that is very different from the Humean regularity theory. For instance, rather than being an extrinsic relation between two separated events, causation is a continuous process of change. This causal process involves the manifestation of intrinsic dispositions. On this view, causation is singular, complex, context-sensitive and irreducibly tendential. Dispositions, or causal powers, are seen as active and dynamic, and any form of stability is produced by counteracting powers keeping a situation in equilibrium. New dispositions can emerge (or demerge) as a result of the causal process (Anjum and Mumford 2017: *Emergence and Demergence*).

From a dispositionalist starting-point, scientific approaches should change focus away from static, changeless factors, to be studied in separation and isolation within some closed system, assuming additive composition of causes. Instead, the default expectation should be that causation is a complex and dynamic temporal process, involving nonlinear interactions within open systems. Causal dispositionalism also emphasises the importance of contextual and local knowledge. To focus only on a single causal factor when making predictions about the outcome, will be to disregard most of the causally relevant factors of the causal interaction.

Risk assessment approaches in pharmacovigilance, gene technology and environmental risk have traditionally followed a Humean model. However, researchers from various disciplines have urged that one moves beyond assessment of single stressors, outcomes, sources and pathways, instead acknowledging the crucial role of interactions between stressors and environmental factors for the insurgence of harmful effects in complex systems.

We here argue that the dispositionalist theory of causation offers support to a different risk methodology. For instance, rather than seeing complexity and contextual interference as obstacles for establishing the causal role of a stressor, a dispositionalist would urge that scientists use these features to their advantage. For a dispositionalist, complexity and the possibility of interference are essential features of causation. Stressor focused approaches to risk would then tell us more about the causal dispositions of a stressor from observing it in different types of interactions than from separation and isolation. But this can only be done in a realistic contextual setting, where we can observe real causal processes in their full complexity.

### Three Errors about Causal Efficacy (Rögnvaldur D. Ingthorsson)

Among the many extant accounts of causation, only four can be understood as attempts to make sense of efficient causation: (i) transmission/causal process accounts, (ii) mechanistic accounts, (iii) powers-based accounts, and (iv) my account of causation in terms of interaction. In this paper I want to sort out some mistakes that hamper the contemporary debate about causal influence. First, it is widely assumed that causal influence is unidirectional: either a conserved quantity transmitted between causal processes, or an

influence exerted by an active object on a passive receiver. This idea violates the established scientific fact that all interactions are reciprocal; nothing influences anything else without the latter simultaneously influencing the former in the same way and to the same degree. I will illustrate why the reciprocity of interactions guarantee that transmission accounts fail to generalize to all the cases; they fail to make sense of symmetrical interactions, say, collisions between two identical billiard balls moving with the same speed in opposite directions. The reciprocity of interactions falsifies the assumed dichotomy of interacting objects into active and passive components. It implies instead that in every interaction objects simultaneously affect each other in quantitatively proportionate and similar ways.

Second, philosophy has long confused two senses of ‘reciprocity’. There is the pre-Newtonian idea that cause and effect are reciprocal in the sense that the cause must lose as much as the effect gains, and the Newtonian idea that interacting things affect each other equally in opposite directions. As I will show, it is the pre-Newtonian idea that is the basis of the conviction that the concept of simultaneous causation is paradoxical. When thinking in the Newtonian mode, simultaneous causation is unproblematic.

Third, it is widely assumed that influence is transmitted between events rather than between continuants (there are exceptions). As I will show, this idea is relatively recent and is the result of a mistake; that of inadvertently fusing a Humean constant conjunction conception of causation, with the idea that cause transmits an influence to the effect. Before the rise of empiricism, all the major schools of thought took influence to be exerted between continuants; even Hume assumed that his contemporaries (mistakenly) believed that influence is passed between continuants. The idea that influence passes from one event to another only takes hold following the publication of Russell’s ‘On the Notion of Cause’ in 1912. In that paper, causal realism is presented as the claim that causation is a two-place relation between successive events of which the first influences the second. But, as far as I can tell, no causal realist has ever put that view forward, and it remains an anomaly in light of the Newtonian concept of reciprocal action and with the scientific understanding of influence in general.

## What We Do to Each Other: Folk Psychology, Interventionism, and Special Causal Concepts (Julian Bacharach)

In a series of papers, John Campbell has argued that the interventionist approach to causality, pioneered recently by Woodward and others, can be fruitfully applied to the psychological realm. One of the attractions of this approach is that it shows how causal claims can be liberated from the more heavy-duty constraints that have traditionally given rise to the problem of ‘mental causation’: for instance, the requirement that causes be backed by laws, or that there be an appropriate mechanism connecting cause and effect. Moreover, it promises to secure the causal relevance of psychological variables on their own terms, with relying on theoretical identifications with states or properties at a non-psychological level. To this extent, interventionism appears well-poised to vindicate the causal character of ordinary, personal-level folk-psychological explanation.

In this paper I shall argue that this promise remains unfulfilled. While it is plausible that our interest in psychological causal claims is in some contexts---for instance clinical psychiatric contexts---an interest in how a subject would respond to various hypothetical interventions,

this is not the case for ordinary personal and interpersonal contexts. This is because, in such contexts, our lives are too closely causally intertwined, and our interactions too much dependent on a shared background of mutual intelligibility, for normal actions towards one another qualify as interventions in the relevant, quasi-technical sense. Moreover, I shall argue that such contexts are the primary setting in which folk-psychological causal explanations are originally to be understood; the interventionist paradigm is applicable only by introducing idealisations which severely distort the original nature of our causal interest in one another.

This is not just a technical difficulty for the interventionist approach; rather, seeing why the paradigm is inapplicable here sheds light on the intuitive sense that interpersonal folk-psychological understanding is not detached or disengaged, but rather grounded in mutuality. One response would be to hold that interpersonal understanding, insofar as it is so grounded in mutuality, is not causal. In the latter part of this paper I sketch an alternative approach. I suggest that our natural sympathy and openness to one another may allow for a direct awareness of psychological causality, one which is to some extent analogous to the perceptual awareness of causality in mechanical transactions.

Finally, I close with some discussion of what commitments this suggestion incurs regarding the metaphysics of cause. The proposal might be fleshed out in terms of the idea that ordinary awareness of psychological causality involves what G. E. M. Anscombe once called ‘special causal concepts’: a concept of a type of causal process or interaction that is not reducible to a generic notion of cause plus a restriction to a specific domain. The general challenge for this approach is a unity challenge: to explain how diverse special causal concepts are nevertheless somehow guises of the same general concept of cause. Taking a cue again from Anscombe, I shall suggest that what unifies different special causal concepts is the possibility of their coming together in relations of interruption and prevention. Thus, viewing ordinary folk-psychological attributions as genuinely causal is a matter of seeing them as picking out processes which are subject to external, non-psychological distorting influences.

## Novel Information-Theoretic Measures of Causation in Cognitive Science (Joe Dewhurst)

The aim of this paper is to present and analyse three different information-theoretic measures of causation that have recently been proposed from within cognitive science, and to argue that they are best understood as offering a deflationary or instrumentalist account of causation. The formal measures of causal strength or influence that these accounts provide are useful tools for the analysis of complex systems, but such tools cannot necessarily do the philosophical work that more traditional, metaphysically robust accounts of causation are expected to. It would thus be a mistake to draw any strong conclusions about topics such as causal emergence or mental causation from these novel measures alone. Nonetheless, once this limitation has been recognised, the tools developed by these accounts can provide a valuable foundation for a formal philosophy of the special sciences, guiding our analysis of existing issues such as explanatory autonomy and levels of explanation.

The idea that there might be a connection between information and causation is not new. It was discussed, for instance, by Wiener (1956), developed formally by Granger (1969), and explored more recently in different contexts by Collier (1999, 2011), Griffiths et al (2015), and Andersen (2017). The three accounts that I will consider here, developed respectively by

Flack (2017), Hoel et al (2013, 2016; see also Hoel 2017, Albantakis et al 2019), and Rosas et al (2020), aim to use this connection between causation and information to resolve traditionally philosophical problem surrounding causation in cognitive science. I will argue that the formal measures offered by these accounts cannot by themselves do any substantial philosophical work, but must first be buttressed by additional metaphysical assumptions about the nature of causation and its relation to information. Nonetheless, these approaches do provide a useful formal toolbox for investigating existing philosophical problems, and I will demonstrate one such application to the question of explanatory autonomy and levels of explanation in cognitive science.

The kinds of measure offered by Flack, Hoel et al, and Rosas et al are better understood as measures of explanatory power than measures of causation per se. These could be compatible with a deflationary, pragmatic, or instrumentalist account of causation such as those offered by Price (2001) or Woodward (2007), but are unlikely to provide any compelling insights into what might be considered more traditionally metaphysical (rather than scientific or epistemic) issues to do with causation, such as strong emergence or mental causation. What they can do, however, is provide a formal framework for assessing the relative explanatory value of different levels or scales of explanation. In this capacity they can be understood as offering an account of epistemic (or weak) emergence. According to such an account, causation might itself be best understood as an instrumental construct for making sense of the world, and a formal demonstration of causal emergence could serve as a guide to the scale at which we might best study some phenomenon, even if it is not emergent in any strong or ontological sense. I will conclude by considering whether this approach should simply be understood as a pragmatic gloss on the concept of causation, or alternatively (and more boldly) as the foundation for a novel, scientifically informed scale-relative ontology in the tradition of Ross, Ladyman, & Collier 2007.

## Causal Pluralism and Carnapian Explication (Matthias Rolffs)

Conceptual causal pluralism can roughly be characterized as the thesis that there is more than one concept of causation. For example, Ned Hall (2004) maintains that there are two concepts of causation: production and dependence.

Three potential problems for this thesis are the following: First, if there are multiple concepts of causation “it is not clear anymore why and in virtue of what all these concepts are concepts of causation” (cf. Psillos 2010). Call this problem the Problem of Unity. Second, if there are multiple concepts of causation, it should be expected that “actual usage should contain certain kinds of evidence for this” (Godfrey-Smith 2009). However, such linguistic evidence of ambiguity is absent. Call this the Problem of Ambiguity. Third, Hall’s thesis that there are two concepts of causation is confronted with further counterexamples: There are intuitive cases of causation that are neither cases of production nor cases of dependence, for example cases of pre-empted double-prevention (cf. Hall 2004, Longworth 2006). Call this the Problem of Further Counterexamples.

The aim of my talk is to defend a version of conceptual causal pluralism against the three above-mentioned objections. My defense of conceptual causal pluralism rests on methodological considerations: I argue that Carnapian explication (cf. Carnap 1950, Brun 2016) is a plausible and well-motivated methodology for the philosophy of causation. Central to this methodology is a distinction between the explicandum ‘causation’ and explicata of this



explicandum. The explicandum is part of everyday language. It is but the starting point for an explicatory inquiry. The aim of an explicatory inquiry is the introduction of a new concept – an explicatum – that is similar to the explicandum, but more precise and more fruitful, that is, more useful for certain theoretical purposes.

Given this distinction, conceptual causal pluralism can be interpreted either as a thesis about the explicandum ‘causation’ (-> pluralism about explicanda) or as a thesis about explicata of this explicandum (-> pluralism about explicata) (cf. Olsson 2017). I argue for pluralism about explicata of the term ‘causation’. Production and Dependence can be seen as two different legitimate explicata for the explicandum ‘causation’: Both are concepts that are sufficiently similar to the explicandum and that are more precise and more fruitful for different theoretical purposes (cf. Hall 2006 for a somewhat similar approach).

Furthermore, I argue that this version of conceptual causal pluralism can be defended against the above-mentioned three problems: As to the Problem of Unity, Production and Dependence are both concepts of causation because they are explicata of the same explicandum. The Problem of Ambiguity can be avoided because pluralism about explicata does not imply that our everyday term ‘causation’ is ambiguous. Finally, the Problem of Further Counterexamples can be avoided because explicata need not be exactly similar to their explicandum: Intensional differences between an explicatum and its explicandum do not refute the legitimacy of the explicatum.

## Causal Pluralism: A New Perspective (Ned Hall)

One route into causal pluralism is via good old-fashioned conceptual analysis, some of which has appeared to uncover different and seemingly incompatible conceptual pressures on our “ordinary” concepts of cause and effect. But a very different and (I will try to argue) more illuminating route is by front-loading the following question: What is a concept of causation *good for*? I will argue that there are distinct (if related) epistemic uses to which we can put such a concept, and that carefully distinguishing them can reveal the need for similarly distinct precisified concepts, tailored to different styles and domains of inquiry.

## Epistemic Causality in the Social Sciences (Jon Williamson)

The epistemic theory views causality as a tool that helps us to predict, explain and control our world, rather than as a relation that exists independently of our epistemic practices. This paper first provides an introduction to the epistemic theory, showing that it is motivated both by the failure of standard theories of causality and by a careful consideration of the epistemology of causality. Then I argue that the epistemic theory provides a very natural account of causality in the social sciences, not least because other accounts struggle to accommodate the great variety of kinds of causal relations and causal relata in the social sciences. I consider a case study in sociology in some detail and go on to explore the connection between epistemic causality and formal methods for establishing causality in the social sciences, including structural equation modelling, the potential outcomes approach, Granger causality, and multilevel analysis.

## An Epistemic Approach to Causation in Social Scientific Case Study Research (Rosa W. Runhardt)

In this paper, I argue for an epistemic approach to causation in case study research in the social sciences, focusing in particular on political science. Social scientists have, for the last two decades, become more focused on causal mechanisms in case study research. However, I argue that the most promising methodological theories about such mechanisms see mechanisms not as an ontological property of the social world but rather as a theoretical construct that helps researchers make sense of the social world. To impose on mechanisms in social science the general philosophical theories of mechanism borrowed from e.g. the life sciences would ignore the fundamental differences between how mechanisms are used in, for example, political science versus biology.

I begin the paper by setting out these fundamental differences. I show that the mechanistic account of causation in the life sciences and natural sciences links cause-effect relations with physical processes, such as entities and activities (cf. Machamer, Darden and Craver 2000). In contrast, I argue for an ‘anti-realist’ position on social causation: social mechanisms are not features of the social world but rather parts of social theory. In the second part of the paper, I show that my theory is a special case of Williamson’s Epistemic Theory of Causality (cf. Williamson 2005), which generally argues that causation is an epistemic property rather than an ontological property. Following Williamson, I argue that social scientists “have (...) causal beliefs because of their utility as a basis for inference, not because there is any physical connection between cause and effect that [they] perceive or infer” (Williamson 2006: 272). In the third part of the paper, I briefly derive implications from Williamson’s Epistemic Theory of Causality for how case study researchers can strengthen their current methods of causal analysis.

Throughout the paper, I use a paradigmatic example of social mechanisms to illustrate my claims. In their seminal book *Dynamics of Contention* (2001), Doug McAdam, Sidney Tarrow, and Charles Tilly theorize that four of the same mechanisms (brokerage, category formation, object shift, and certification) produce a variety of events in contentious politics (including social movements, strikes, and revolutions). I show that their work demonstrates that social mechanisms are best thought of as epistemic constructs rather than ontological claims and ask what a strong epistemology of contentious politics ought to look like.

## Information Transmission and the Mosaic of Causal Theory (Federica Russo)

Causality is a central notion in the sciences. It is at the core of a number of epistemic practices such as explanation, prediction, or reasoning. The recognition of a plurality of practices calls, in turn, for a pluralistic approach to causality. In the ‘mosaic’ approach, as developed by Illari and Russo (2014), we need to select the causal account that best fits the practice at hand, and in the specific context. For instance, the concept of (causal) mechanism helps with explanatory practices in fields such as biology or neuroscience. Or, the concept of (causal) process helps with tracing ‘world-line’ trajectories in physics contexts or in social science. While no single notion of causality can simultaneously meet the requirements for a good explanation, prediction, or reasoning across different contexts and practices, a

pluralistic approach towards the epistemology of causality seems to be the most plausible and attractive solution.

But beyond having epistemic significance, causality is arguably the ‘cement of the universe’, to borrow the expression from the seminal work of John Mackie. What this means exactly is however made difficult in the light of the overspecialization of the sciences: is this cement the same in high energy physics and in molecular biology? Is it the same cement at the basis of social bonds and of disease onset? Two issues further complicate the picture. First, contemporary science, more often than not, crosses disciplinary boundaries, trying to establish causal relations across relata of different nature; for instance, we attempt to explain mental health conditions, such as depression, invoking biological and environmental factors, and how the two interact. Second, more often than not, contemporary science is *techno-science*, where instruments arguably allow for deeper and greater epistemic access to the (portion of the) world under investigation, but they do so by (partly) ‘constructing’ the object of study. For instance, the process of detection and measurement of biomarkers is not a simple and direct process giving access to a clearly identified entity, but is instead a complex process in which the ‘thing’ biomarker is much constructed via the technologies and theories employed.

How can we make sense of this complicated metaphysical picture? How can we make a metaphysics of causality compatible with an epistemology? In this chapter, I explore the prospects of an informational approach to causality, as one that can offer a thin metaphysics, derived from an epistemology of techno-scientific practices. Specifically, I shall try to support the view that the mosaic view of causality, as an epistemology of causality, needs to be accompanied by a (thin) metaphysics in which causality is cashed out as information transmission. This combination, I shall argue, helps make sense of causality across different techno-scientific contexts and domains.

## Of Powers and Causes: Groundwork for Causal Dispositionalism (Joaquim Giannotti)

Dispositionalism is a theory about the natural properties of our world. On this view, all or many natural properties are essentially dispositional: it is true in virtue of the nature of such properties that things that instantiate them have distinctive dispositions and causal roles. In a sense, essentially dispositional properties empower their bearers in characteristic ways. Let us call these properties powers.

Classic examples of powers are putative fundamental physical properties such as charge, mass, and spin. Dispositionalists would say, for instance, that an electron has the disposition to produce a force—as specified by Coulomb’s law—by virtue of having the essentially dispositional property of charge. Such a disposition is manifested in distinctive circumstances—for example, when the electron interacts with other charged particles—and is amenable to a counterfactual treatment.

It is claimed that dispositionalism brings us a major advantage: it offers a framework for a non-Humean theory of causation, one which admits genuine causal movers and shakers in nature. In broad strokes, the core idea is that causation is a matter of manifesting powers (e.g., Mumford and Anjum 2011, Marmodoro 2017, Corry 2019, Williams 2019, Bird 2020, Ingthorsson 2021). Label this framework causal dispositionalism.

Unfortunately, there is in-house disagreement among causal dispositionalists. A key issue concerns the question of how to understand the notion of a cause within the causal dispositionalist framework.

Some identify causes with powers. (e.g., Mumford and Anjum 2011, Williams 2019). These causal dispositionalists would say, for example, that the cause of the shattering of the window is its power to break in characteristic circumstances. An advantage of this view is that it vindicates the claim that powers are causal doers. However, it faces an important objection: since powers are properties, and properties are abstract, causes would be abstract too (e.g., Groff 2020).

Others identify causes with the triggering or exercising of powers (e.g., Marmodoro 2017, Bird 2020). To use the previous example, these causal dispositionalists would say that it is the hitting of the window that causes the shattering. However, also this view faces a serious problem: it seemingly misplaces the locus of causation into triggers and stimuli of powers (e.g., Williams 2019, 123–129).

My aim is to pave the way to a middle-ground position, one which pays heed to the explanatory benefits of causal dispositionalism. I will argue that we ought to distinguish between causal doers and causal explainers, and that powers are best regarded as the latter. Namely, entities that play a certain theoretical role in our causal explanations. This view, I will argue, permits us to maintain that causal doers are concrete propertied things, such as events or substances or processes or states of affairs, without jeopardizing the causal relevance of powers. After introducing a distinction between partial and full causal explainers, I will suggest that insofar powers are taken to be ineliminable partial causal explainers, we can enjoy the benefits of causal dispositionalism over the Humean alternatives.

## Defending Non-Natural Intuitions About Causation: Why it's Not Wrong to make Causal Judgements Based on Norms (Jon Bebb)

Through numerous studies and experiments, Joshua Knobe (alongside many collaborators) has provided a considerable amount of evidence that our attributions of (actual/token) causation are influenced by norms (Knobe & Fraser, 2008) (Hitchcock & Knobe, 2009) (Phillips, et al., 2015). Throughout these experiments it was found that, all things being equal, subjects were more likely to agree that an event was a cause of some effect if that event is considered to have violated a (moral/social/legal/etc.) norm, and more likely to disagree that an event was a cause of some effect if that event is not considered to have violated a norm. This phenomenon has been dubbed the 'Knobe Effect', and it has sparked a considerable amount of discussion regarding the role that norms may or may not have in an account of causation. For some philosophers, the causal judgements that Knobe et al. identify as being influenced by norms are considered to be instances of the concept of causation being applied competently. They take the fact that the subjects were more inclined to agree that an event caused some effect when that event was norm-violating as resulting from the subjects having a clear understanding of what causation is and then proceeding to apply that knowledge appropriately. Yet others disagree, arguing instead that the judgements made by these subjects are the result of some kind of error. Theories for what this error might be are varied, with the more common explanations being that the subjects have fallen victim to some cognitive bias (Blanchard & Schaffer, 2017) (Rose, 2017) or have simply confused the concept of causation for that of responsibility (Sytsma, et al., 2012). What all these theories

have in common is their commitment to the claim that the experimental subjects have somehow got it wrong when it comes to their judgements of causation.

In this paper I argue against this view that the subjects were mistaken in their judgements, and so will consequently be defending the view that these judgements arise from a competent use of the concept of causation. My argument will be that the non-competence view arises out of a particular intuition about the nature of causation: the intuition that it is a natural relation. The view that causation is a natural relation is the view that causal relations are ‘out there’ in the world: they obtain in virtue of facts about the natural world and are knowable a posteriori. This view is explicitly stated by Strawson (1992), but it has been largely assumed by many who have written on the topic of causation. However, I argue that when we consider the purpose of our concept of causation, we find this purpose is not served by holding that the causal relation is natural. Consequently, there is no positive reason for believing that causation is a natural relation, and so there is also no good reason to prefer our intuitions about this over the intuitive causal judgements made by the subjects of Knobe et al.’s experiments.

## Resolving the Tension Between Different Approaches to Causation Within a Dynamical Systems Framework (Gergely Kertész, Balázs Gyenis, Gábor Hofer-Szabó, and Peter Fazekas)

The tension between physical (Dowe 2000, 2009) and difference-making accounts (Schaffer 2004, 2016) of causation is one of the factors that have driven the recent proliferation of pluralistic attitudes towards the nature of causation (e.g.: Hall 2004). In this paper our aim is to show that such pluralism is not inevitable—that there exists a framework within which the central ideas behind physical and difference making approaches can be reconciled.

Our proposed idea could be interpreted as a realization of the project outlined by Ney (2009) that aimed to bring these two seemingly inconsistent approaches together. In her view, difference-making talk can be anchored at the physical level without reminder regardless of the causal status of fundamental physical processes. She highlighted negative causation as the main source of the tension and suggested some workable options that might overcome the problem. Our theory also offers a solution to this problem in a similar spirit but based on starting points that deviate from those utilized in well-known physical theories of causation. The framework we have built is based on dynamical systems broadly construed. In this model, the properties postulated by a causal claim correspond to sets of physical states (points in the phase space of the underlying system) that instantiate these properties. A causal claim is true if the cause and effect property descriptions postulated by the causal claim pick out sets of physical states that are systematically related by the dynamics, i.e. the time evolution of the underlying system. The features of this systematic relation determine the robustness of causal claims and ground successful causal discovery and inference methods (Fazekas et al. 2019). The fundamental characteristics of our framework make it possible to reconcile different descriptions of the same system formulated from the perspective of physical processes and difference-making, and to identify the truth-makers of everyday and special science causal claims.

This framework can anchor negative causal statements at the physical, dynamical level avoiding the problem created by the mismatch between the language of difference-making

causation and theories of physical causation. In our framework, a well-chosen, higher-level property-based causal description can pick out a region in the phase space of the system in question such that the time-evolution of the physical states that constitute this region evolve into a relevant effect region within a characteristic time frame. Such cause and effect regions are sometimes picked out by negative descriptions, i.e. descriptions that we usually call absences and omissions. When that happens, we have a case of negative causal talk grounded by the physical dynamics of the underlying system. This is what this paper aims to showcase.

## An Inferentialist Definition of Causation (Christopher Clarke)

One alternative to standard approaches to causation is inferentialism (Cartwright 2007, Reiss 2012), which can either be treated as form of primitivism about causation, or as a form of projectivism. Inferentialism about causation abandons the project of defining causation in terms of necessary and sufficient conditions. Instead, it defines causal hypotheses by (a) describing how evidence can warrant belief in causal hypotheses, and (b) describing how causal knowledge can warrant belief in other propositions (including propositions about what actions one ought to perform or what policies one ought to implement). Inferentialism is often motivated by a use-centred approach to philosophical theorizing: when deciding how to define a concept, one ought to define the concept in question in whatever way maximizes that concept's theoretical or practical usefulness (Dupre 1993).

But inferentialist accounts of causation can disagree about the extent to which (a) and (b) can be described via precise context-invariant principles, rather than by a host of imprecise context-variant rules of thumb. And Cartwright herself advocates the latter. For Cartwright, methods such as regression analysis or randomized experimentation, for example, are only warranted (that is, internally valid) if contextual conditions are just right.

I propose, pace Cartwright, that there are three (relatively) precise and context-invariant principles that describe (a) and (b). These principles are namely the (qualified) principle of faithfulness, the (qualified) principle of the common cause, and the (qualified) principle of the common effect. These qualified principles, I claim, avoid the objections that Cartwright and others have levelled against their unqualified versions (Sober 2001, Hoover 2001). And they make clear how causal knowledge can be gained and then used. One noteworthy feature of these principles (the qualified principle of faithfulness in particular) is that they are anthropocentric: in a sense, they make causation relative to our limitations as human inquirers.

In defending these principles (the qualified principle of the common cause in particular) from objections, I will note the following tension within Cartwright's position. On the one hand, Cartwrightian inferentialism seems to be motivated by the use-centred approach to philosophical theorizing. But on the other hand, by use-centric standards one of Cartwright's main objections to the principle of the common cause fails.

## Causal Process Tracing Theory of Change (Nancy Cartwright and John Pemberton)

This is not an account of what causation is ontologically nor of what people mean by the term—we are radical pluralists about that. Rather we aim to describe the mix of things that go on in the world when one event produces another and to provide a framework for laying this

out. We propose what we call a *pToc* for doing this—a *causal-process-tracing theory of change*. A *pToc* for a cause/effect pair weaves together the process causation through which the cause delivers the effect with the structural causation that ensures that the *causal principles by which it does so can all obtain*. Because of their level of detail *pToc*s are especially useful for causal prediction and post hoc evaluation. We illustrate with an example of a real negative causal feedback loop identified in the *Munro Review of Child Protection*.