Forthcoming in A. Wilson & K. Robertson (eds.) *Levels of Explanation*, OUP 2024.

From Multilevel Explanation to Downward Causation

The principle that everything that happens within the physical domain has a sufficient cause within that same domain—the *causal closure of the physical*—poses a familiar causal exclusion problem for the special sciences: special science properties are distinct from their physical realizers, but if the physical domain is causally closed, then what causal work is left for such properties to do?[1] Here I argue that causal closure in fact poses no exclusion problem for the special sciences. I focus on the simple case of vector composition and argue that it involves irreducibly multilevel causation. Discussion of such simple physics may not seem like the most promising way of defending the autonomy of chemistry, biology, or psychology. My aim, however, is to persuade you that this case has profound implications for how we should think about causal closure. The way in which multiple vector fields compose, I shall argue, falsifies any closure principle according to which the course of physical events is entirely determined by properties at the fundamental level. I shall argue that the strongest closure principle that is consistent with vector composition allows for a particular form of downward causation, and so makes room for an irreducible causal role for special science properties. Hence, there is in principle plenty of causal work left for special science properties to do.

I shall assume an ontology of fundamental physical particles interacting in spacetime by means of fundamental physical forces. Macroscopic objects, I assume, are fully composed of such particles and macro-causal interactions are fully grounded in fundamental particle-particle interactions.[2] This might seem like a lot to assume, especially given that true fundamental physical reality may turn out to be very different. Why then do I set things up this way from the outset? One simple reason is that these are the terms of the causal exclusion debate. If particular special sciences entities are fully constituted

[1] Kim (1992, 1998).

[2] For present purposes we can adopt a standard notion of grounding as a transitive, irreflexive and asymmetric relation of metaphysical explanation holding between entities such as properties, states, objects, or events.

by physical particles and their properties and relations, and everything that happens to those particles is fully determined by fundamental physical forces according to law, then what could special science entities be *doing*? And if special science entities have nothing to do, it seems just obvious that special sciences themselves are at most a simple and perspicuous way of grouping together fundamental physical causes. An ontology of fundamental particles, governed by a closed system of fundamental laws, is what gives rise to the exclusion problem in the first place. My aim here is to solve that problem on its merits: even assuming that everything is fully composed of physical particles, and that only fundamental physical forces are capable of accelerating such particles, there is still room for a kind of downward causation that would render special science autonomy unmysterious. Indeed, I shall argue, the idea that multiple fundamental forces compose to produce macroscopic effects *requires* downward causation of the kind in question.

What if it turns out—as many suppose it will—that the fundamental ontology of completed physics is radically different from that suggested by current theory? What if the fundamental ontology is not even spatiotemporal? Will the fundamental domain be in some sense causally closed, whatever causation turns out to be, and will that closure principle give rise to an exclusion problem? I confess that I don't know the answers to these questions, but since my intention here is to defend a solution to the exclusion problem as it is typically framed by those who think it is a serious problem worthy of attention, I don't see this as a major flaw. The reader may think of the proposed solution as follows: even if everything in our world were fully constituted by particles recognisably similar to those of current particle physics, and even if everything that happened to such particles were due to the exertion of fundamental forces similar to those we find in current physics, that would be no threat to the autonomy of the special sciences. In *that* kind of ontology—one, that is, in which the course of events is determined by the interaction of multiple vector fields in spacetime—downward causation is built in from the outset.[3] Now if it turns out that spacetime, particles and forces are emergent and that really, the fundamental dynamics have nothing to do with vector fields as we current conceive them, then the solution defended here is unlikely to solve exclusion problems that may arise between the fundamental ontology, whatever it turns out to be, and the emergent reality we seem to inhabit. But even if current physics is itself a special science, it's still interesting to consider whether the interlevel relations between physics and even higher level special sciences bring with them problems of causal exclusion. If physics and everything above ends up being in some way excluded by whatever lies beneath, perhaps we can cross that bridge when we come to it.

Before proceeding, let me clarify the overall strategy of the paper. I intend to offer a solution to the exclusion problem that works by showing that causal closure does not rule out downward causation, so I will work with the strongest closure principle that could plausibly be justified by the available evidence. That closure principle alone, I shall argue, does not rule out downward causation and poses no problem for the autonomy of the special sciences. There is an even stronger closure principle in the vicinity, which I will also discuss, and which does pose such problems. However, I will argue, that principle is *falsified* by the available evidence rather than supported by it. The overall aim of my discussion is to explore the kind of downward causation that is consistent with the strongest plausible closure principle on offer, and thereby to cast light on what it is that special science properties could possibly be doing in order to earn their autonomy from physics. There are bound to be alternative perspectives on causal closure to the framework adopted here and I do not claim to be solving any exclusion problems that may arise based on the resulting principles, however they may be defined. The proof, as always, is in the pudding: tell me what closure principle you are working with, why you believe it, and why you think it rules out downward causation as I shall conceive it here.

[3] The central arguments of this paper go through *mutatis mutandis* in field-theoretic ontologies according to which fields are the fundamental constituents of physical reality, with particles derivative. I lack the space to defend that claim here and assume an ontology of fundamental particles for simplicity.

The plan of the paper is as follows. In section 1, I first define and motivate the causal closure principle I shall be working with and then discuss the idea that causal closure entails causal-*explanatory* closure—the principle that everything that happens within the physical domain can in principle be fully *explained* without appealing to anything outside it. I do this because the case I consider involves irreducibly multilevel explanation and I want to draw conclusions from it concerning causation. I argue in section 2 that causal explanations in physics that involve multiple forces are irreducibly multilevel, involving both force-generating basic physical properties and higher-level properties that partially determine how forces compose. In section 3 I discuss the implications of multilevel explanations for the causal closure of the physical and suggest that the strongest empirically well-supported closure principle we can formulate is consistent with a certain kind of downward causation; I also discuss some options for understanding the metaphysics of this kind of downward causation in terms of causal powers. In section 4 I discuss two potential sources of downward causation so conceived in terms of the debate between kinematic and dynamic theories of the origins of spacetime symmetries. I conclude in section 5 with some reflections on special science autonomy.

1. From causal closure to causal-explanatory closure

Here is a typical way of stating the causal closure of the physical:

> CC: Every physical event E that has a cause at *t* has a sufficient physical cause C at *t*.[4]

According to CC, wherever in its causal history a given physical event has a cause at all, it has a sufficient physical cause. The content of 'physical' comes from physical theory, so that the closed domain is determined by our best fundamental science. There is undoubtedly more physics to be discovered and those yet-to-be-discovered things are holes in the causal structure of current physics.

[4] In this formulation and those that follow, the notion of a sufficient cause may be read in probabilistic terms, i.e. as a cause that suffices to determine the chances of the effect.

Nonetheless, I will assume that there is a complete theory in the vicinity and that it resembles current theory closely enough for CC to be non-vacuously true.[5] The intended scope of 'physical' in 'physical event' is not limited to the ontology of a completed physics. Rather, CC quantifies over the *broadly* physical, where an entity is broadly physical iff it is either part of the ontology of completed physics or appropriately related to such entities, where "appropriate relations" include *grounding* relations such as composition, constitution, and realization.[6] The deflection of a particle in a magnetic field counts as a physical event under this definition and so does the eruption of a volcano.

CC is already quite a strong closure principle, but I will work with a much stronger one.[7] Dialectically, the reason for this is as follows. My aim is to show that the strongest closure principle that is not falsified by evidence from simple physics still does not generate a causal exclusion problem for the special sciences. The reader may well suspect that the principle I eventually settle on is not actually supported by the available evidence, such is its strength. Fine with me—but if even that principle leaves room for downward causation, it should be clear that the weaker principles stated here—those which, the reader may suspect, have a chance of being true—do so as well. I shall thus formulate a sequence of closure principles of increasing strength, each one consisting in the preceding principle conjoined with an additional thesis. If a given principle in the sequence leaves room for downward causation, then so a fortiori do all those weaker than it. I note before proceeding that I do not claim that all possible closure principles can be ordered according to strength. As such, there may be principles that do not fit into the ordering developed here. The strongest true principle in my ordering

[5] According to Hempel's dilemma, causal closure principles are either false (because indexed to current physics), or vacuous (because indexed to a future physics whose content is obscure). I am betting on future physics being sufficiently close to current physics to avoid the threat of vacuity. See Crane & Mellor (1990) for the dilemma; see Papineau and Spurrett (1999) for arguments that we don't need to define 'physical' to formulate contentful closure principles, provided we are clear about what the fundamental ontology will *not* include; and see Wilson (2006), for arguments that betting on future physics is the way to go, but with the addition of a 'no fundamental mentality clause' to rule out sui generis mental properties counting as physical.

[6] See Crook & Gillett (2001) for this way of thinking about the physical. Note that the scope of 'physical' in 'physical event' and 'sufficient physical cause' can be varied independently. I will do just that presently, in order to formulate the strongest closure principle I can.

[7] For further details of the arguments that follow, see Yates (2009).

does not give rise to an exclusion problem, but that entails nothing about whether exclusion problems arise from such hypothetical principles. No matter, for I am trying to solve the exclusion problem that arises from typical causal closure principles, not hypothetical problems that may arise from others.

An immediate worry with CC is that causal sufficiency is too metaphysically coarse-grained. Suppose events to be property-instances. A physical event C might be causally sufficient for another such event E at *t* by means of an intermediary effect. Suppose C is synchronically sufficient for a *non*-physical event e*—an instance of a sui generis phenomenal property, say—and that it's only C and e* together that have the power to bring about E. This is an instance of traditional emergent downward causation, in which C synchronically causes an instance of a sui generis conscious property e* and e* then contributes a novel causal power, which (perhaps in combination with the powers of C), causally suffices at *t* for E. Given that CC does not rule this kind of situation out, we need a stronger closure principle. One way to secure such a principle is to add a clause stating that the physical cause in question has the power to bring about the effect in and of itself, in virtue of its physical properties alone. In the example sketched above, it's not the case that C's physical properties directly bestow upon C the power to cause E. Rather, they bestow upon C the power to cause e* and then the combined properties of C and e* bestow upon them the power to cause E. A causal power bestowal clause rules out this particular kind of emergent downward causation; I refer to the formulation below as *strong causal closure*:

> SCC: Every physical event E that has a cause at *t* has a sufficient physical cause C at *t and* C's physical properties bestow upon C all the powers needed to cause E.

According to SCC, every physical event that has a cause at any point in its history has a sufficient physical cause at that time, whose sufficiency is entirely in virtue of its physical features.[8] SCC is a *very* strong closure principle, but I propose to make it even stronger, by narrowing the scope of 'physical

[8] I shall say more presently about the idea that a certain set of causal powers are *needed* to cause a given effect.

properties' to *basic physical properties*. What is a *basic* physical property? Recall that we are taking "physical" to refer to entities that are part of the ontology of completed physics, or grounded by such entities. A basic physical property is a physical property whose instances are *not* grounded in their bearers' having any distinct physical properties. Examples may include electric charge, mass-energy, spin, entanglement, and spatiotemporal relations. In my usage, a *higher-level* property is any whose instances are grounded in the instantiation of further natural properties. This is a very weak condition on being a higher-level property, so weak that some such properties may well be within the ontology of completed physics: non-elementary values of quantitative properties such as electric charge will come out higher level, for instance, because they are instantiated in virtue of other properties, in this case, having a certain number of particles of unit charge. Not all cases of higher-level properties are higher-level in such an uninteresting way. Some higher-level properties, such as molecular geometry and neural firing patterns, are the proper subject matter of special sciences.

With this understanding of 'basic physical' properties in mind, this is the closure principle I shall be working with , which we may call *very strong causal closure*:

> VSCC: Every physical event E that has a cause at *t* has a sufficient physical cause C at *t* and C's basic physical properties bestow upon C all the powers needed to cause E.

Because C is a broadly physical event, the broadly physical particulars to which it happens are fully composed of fundamental physical particles. The idea behind VSCC is that the basic physical properties of these particles are responsible for C's power to cause E, in that they bestow upon C's basic physical proper parts all the powers that manifest as E's occurrence.[9] We can illustrate VSCC as follows. My raising my glass to take a drink of wine is fully grounded, on some occasion, by a certain plurality of

[9] For indeterministic causation, we can think of sufficient physical causes as causes that suffice to fully determine the chances of occurrence of their effects. Note that VSCC as formulated here does not depend on the claim that particles are basic physical and is entirely consistent with a field-theoretic ontology in which particles are metaphysically grounded. I lack the space here for a detailed treatment, but the central arguments of this paper will go through *mutatis mutandis* if (as seems to be the case) fields are more fundamental than particles.

basic physical events. Each of those events has a fully sufficient physical cause, in the sense that the power to cause each one is due entirely to the cause's basic physical properties. This then is the sense in which basic physical properties bestow all the powers that are *needed* for E to occur—given the combined manifestation of all those powers, *nothing else* needs to happen for me to sip my wine.

The central burden of this paper is to argue that *even VSCC* doesn't give rise to a problem of causal exclusion. Because my arguments depend on causal explanation, let's now turn to the issue of causal-*explanatory* closure. It seems obvious, reflecting on VSCC, that if we want to *fully explain* why some broadly physical event E happens, we do not have to appeal to anything outside the basic physical domain. Given that causal explanations cite causally relevant properties of causes, and all the causal powers C needs to bring about E are due to its basic physical properties, then it seems to follow that (whether or not we are able to formulate it) there is a complete causal explanation of E available in principle that cites only basic physical properties of its sufficient cause C. Thus, it seems that VSCC entails the following principle of causal-explanatory closure:

> CEC Every physical event E that has a cause at *t* has a sufficient physical cause C at *t* and E can be fully causally explained in terms of C's basic physical properties.[10]

Let's say that a *full* causal explanation is one that cites all the causal work that was done to bring about the effect—whatever contributed to an effect's occurrence, a full causal explanation thereof will refer to it. It's now fairly easy to see why CEC poses a problem for the autonomy of the special sciences. On the assumption of physicalism, special sciences such as chemistry, biology and psychology are in the business of explaining broadly physical events and processes—such as bonding, digestion and cognition. But according to CEC, all broadly physical events can in principle be fully explained without leaving the basic physical domain. Given CEC, it seems that if special sciences are autonomous, it is

[10] Things are somewhat complicated here if E's cause C only brings about E by sufficing to determine its probability. In cases where multiple outcomes are equally probable, the cause won't explain why the actual outcome occurred rather than some other. However, where there is truly indeterministic causation, explanations of this kind are the best we can do and full in the sense that there is nothing missing from them.

not in virtue of what they explain, but the manner in which they explain it. A popular strategy for defending special science autonomy is to accept this consequence of CEC and attempt to give an account of why special science explanations of certain phenomena are better than the corresponding basic physical explanations, which are assumed to be in principle available.[11]

Despite the appearances, CEC does not follow from VSCC. In the brief argument I gave above for CEC based on VSCC, I assumed that all causally relevant properties of C in relation to E—that is, all the properties a full causal explanation of E in terms of C must cite—are properties that bestow upon C the power to cause E. In other words, I assumed that the causal work that properties do in relation to some effect *consists in* bestowing the power to cause it. Given this assumption, VSCC entails that basic physical properties *do all the causal work* involved in C's causing E, which in turn implies that we can fully explain E in terms of C. We can state the assumption required to derive CEC from VSCC as follows:

> CW: All the causal work that properties do consists in causal power bestowal.

VSCC says that C's basic physical properties bestow all the causal powers required to cause E. If that's all the causal work there is to do in relation to E, it follows that there is a full causal explanation of E in purely basic physical terms. The conjunction of VSCC and CW entails CEC. We can think of this conjunction as defining a causal closure principle, which we may call *super-strong causal closure*:

> SSCC: Every physical event E that has a cause at *t* has a sufficient physical cause C at *t* and C's basic physical properties bestow upon C all the powers needed to cause E; and all the causal work that properties do consists in causal power bestowal.

SSCC entails CEC and for that reason, is just the kind of closure principle needed to undermine the autonomy of the special sciences. It is also plausibly the kind of principle Kim has in mind when he insists that closure entails that there is *no causal work left* for special science properties to do.

[11] Versions of this strategy can be found in Fodor (1974, 1997); LePore & Loewer (1987); Jackson & Pettit (1990); Yablo (1992); List & Menzies (2010); Wilson (2011); and Yates (2012). Kim (1998) regards all such accounts as a free lunch and I am now inclined to agree, which is why I am here defending downward causation.

However, as I shall now argue, CEC is falsified by some very simple physics, hence SSCC is false. It follows in addition that either (1) VSCC is false, or (2) CW is false. If we choose (1), then there are causal powers beyond those that basic physical properties bestow; if we choose (2), there is more to causation than causal power bestowal. Either way, I shall suggest, causal closure poses no problem for the autonomy of the special sciences. In the remainder of this paper, I will assume that VSCC is true and develop a position based on option (2), the idea that there is more genuine causal work to be done even after all the relevant causal powers have been bestowed.[12]

2. The role of geometric structure in vector composition

When philosophers rely on scientific cases to justify claims of emergence or downward causation, they typically pick sophisticated cases like ferromagnetic phase transitions in condensed matter physics, or self-organization in systems biology. By contrast, the case considered here is simple, but I think it tells us a lot about causation, causal powers, and causal closure. In fact, I think it gives us very good reason to believe in a kind of downward causation, which in turn calls for a metaphysical explanation. The benefit of simplicity is that if I'm wrong about all this, it should be easy for readers to see why.

In this section, I will try to convince you that downward causation is as ubiquitous as vector composition. I shall leave the task of explaining how this kind of downward causation is possible to sections 3 and 4, where I will discuss its metaphysical nature and ultimate source, respectively, in order to render it unmysterious and unproblematic. The result, I hope, will be a position that shows how something close to traditional emergent downward causation—close enough to defend the autonomy of the special sciences—is possible within a robustly reductionist metaphysical framework.

---

[12] I defended a version of (1) in relation to molecular geometry in Yates (2016), but now prefer (2), for the reasons given in section 3.

Figure 1[13] shows the calculation of the resultant field E due to two point charges, q_1 and q_1, at a distance r_1 from q_1 and r_2 from q_2, where the dotted lines from the two charges meet, marked 'X'. This is an illustration of the *parallelogram rule*.

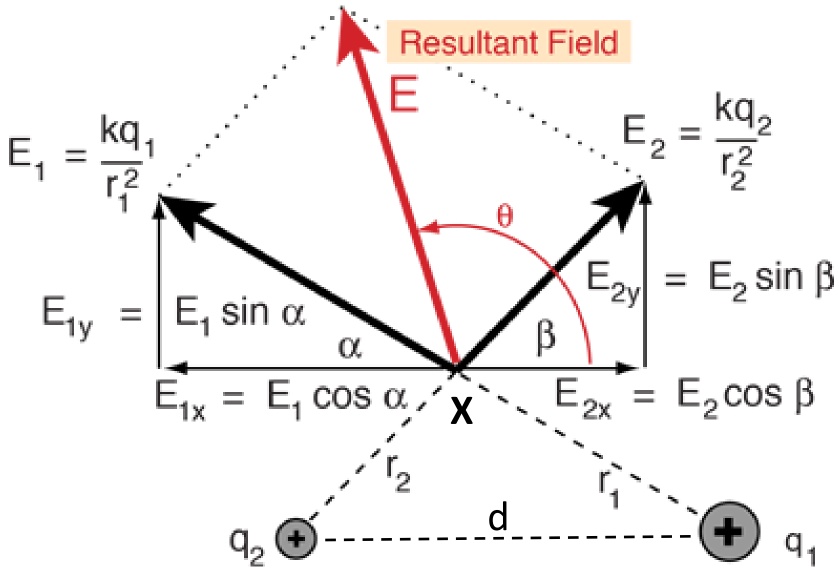

Figure 1: Calculation of the electric field E due to two point charges

The calculation proceeds by resolving E_1 and E_2 into their horizontal and vertical components. The x- and y-components of the resultant field E are given by:

- $E_x = E_{1x} + E_{2x}$
- $E_y = E_{1y} + E_{2y}$

The magnitude and direction of the resultant field E are then given, respectively, by:

1. $E = \sqrt{(E_x^2 + E_y^2)}$
2. $\tan\theta = (E_y/E_x)$

Suppose this calculation forms part of the explanation of the acceleration of a charged particle located at X. My central claims here are that (1) the explanation in question is *irreducibly multilevel*, which violates CEC; and that (2) this violation is due to a kind of downward causation, whose consistency

[13] Reproduced with permission from url: http://hyperphysics.phy-astr.gsu.edu/hbase/electric/mulpoi.html#c3.

with VSCC depends on how we think about causal powers. More on that presently. Figure 1 depicts both basic physical and higher-level properties of the two charges. It shows (A) the specific spatial relations that obtain between q_1 and q_2 and the point X at which we want to calculate the resultant (being separated by r_1 and r_2 respectively); and the distance d between the particles. But in addition to this, it shows (B) what I call the *geometric structure* of the particles in relation to X, in this case given by the angles α and β that we used to factor the component vectors into their horizontal and vertical components. Why isn't geometric structure basic physical? Simply put, because the geometric properties in (B) are instantiated *in virtue of* the basic physical properties given in (A) and hence grounded therein. The specific spatial relations in (A) are not the only way to achieve the geometric structure in (B). We can vary r_1 and r_2 independently while holding α and β fixed, resulting in the same geometric structure. Thus, geometric structure is multiply realizable.[14] I shall now argue that in the present case, geometric structure has a distinctive and irreducible causal-explanatory role compared to the basic physical properties that realize it.[15]

As noted above, we can vary basic physical properties without changing geometric structure. Suppose we move q_1 away from X without changing α. The x- and y-components of E_1 will decrease in magnitude as $1/r_1^2$, but they will remain in constant proportion to each other, since the direction of E_1 will remain the same. This change will alter both the direction and magnitude of the resultant field E, in a manner given by equations (1) and (2) above. This shows that the basic physical spatial relations of the particles make a difference to the resultant field independently of the geometric structure they realize. However, *the converse is also true*. Imagine now moving q_1 towards q_2 in a circle of radius r_1

[14] This kind of multiple realizability—in which a property P is instantiated in virtue of basic physical properties and there are many different configurations of basic physical properties that are sufficient but not necessary for P—won't count as genuine multiple realization for those, like Shapiro (2000), who see multiple realization in terms of structurally heterogeneous ways of implementing a given function. For others, such as Gillett (2003), multiple realization is easier to come by. I am with Gillett on this, but won't take a stand on the issue here, as my arguments only require that geometric structure is a higher-level property.
[15] What follows is a simplified and (I hope) improved version of an argument I gave in Yates (2016), to show that the molecular geometry of water plays an irreducible role in determining its dipole moment.

about X. This won't change the magnitude of E_1, but as α varies, its direction will change, resulting in different values for E_{1x} and E_{1y}, hence altering the magnitude and direction of E according to equations (1) and (2). This, then, is a way for geometric structure to make a difference independently of the values of r_1 and r_2, which remain constant. There is a complication, however. Because geometric structure is realized by, hence supervenient on, basic physical properties, it isn't possible to change the former without some change in the latter. It might be suspected, then, that it is this basic physical change that is the real difference maker in respect of E.

It's important to get clear about which basic physical changes occur. We start out with q_1 and q_2 located a distance r_1 and r_2 from X respectively and a distance d apart. As we move q_1 towards q_2, it remains r_1 from X and q_2 remains r_2 from X. The only basic physical parameter we change is the distance between q_1 and q_2. But there's nothing special about moving the particles from d to d' apart that explains *why* the magnitude or direction of E changes in the way that it does. Rather, what explains the changes in E is that we move q_1 towards q_2 *holding fixed r_1 and r_2*. In order to do that, we have to change the geometric structure of q_1, q_2 and X. It's only given r_1 and r_2 that a certain separation between q_1 and q_2 determines the magnitude and direction of E, because it's only *together* that these three parameters suffice to determine a geometric structure. The basic physical property we have to change in order to change the geometric structure makes a difference to the resultant field only *because* it makes a difference to the geometric structure. In and of itself, the basic physical change alone doesn't explain the change in the resultant.

What follows from this? I think the example shows that dynamical causal explanations that feature vector composition are irreducibly multilevel and hence that CEC is false, from which it follows that SSCC is also false. An explanation of the acceleration of a charged particle at point X must involve not only basic physical properties such as the charges of the particles q_1 and q_2 and the spatial relations r_1 and r_2, but also an ineliminable appeal to the geometric structure formed by q_1, q_2 and X. There's a simple reason for this: part of the explanation of the magnitude and direction of E is the degree to which the fields due to q_1 and q_2 point in the same direction at X. And *pointing in the same direction* is an irreducibly geometric property. If I am correct that geometric structure is grounded in basic physics and multiply realizable, then an entire family of familiar, simple dynamic explanations involve both basic physical and higher-level properties. Crucially, the higher-level properties are not merely a proxy for their basic physical realizers. As we've seen, in the case of the difference-making role of geometric structure, the basic physical properties that realize that structure are explanatory only insofar as they realize the geometric structure in question. This doesn't settle the issue of what it is that geometric structure actually does, or how it does this—thus far I have only argued that CEC and hence SSCC are false, leaving open the following options: (1) geometric structure is itself a powerful property, which would violate VSCC as well; or (2) geometric structure does novel causal work without bestowing novel powers and hence without violating VSCC. In the next section I defend option (2) and say more about the kind of downward causation involved.

3. What kind of downward causation?

To facilitate the discussion that follows, it will be useful to have the relevant principles of section 1 to hand. Recall that I am working with the following causal closure principle:

VSCC: Every physical event E that has a cause at *t* has a sufficient physical cause C at *t* and C's basic physical properties bestow upon C all the powers needed to cause E.

The conjunction of VSCC with the following principle about causal work:

CW: All the causal work that properties do consists in causal power bestowal.

entails a principle of causal-explanatory closure:

CEC Every physical event E that has a cause at *t* has a sufficient physical cause C at *t* and E can be fully causally explained in terms of C's basic physical properties.

If the arguments of section 2 are correct, then CEC is false, so we must reject either CW or VSCC. I shall discuss both strategies in what follows. If we say that the causal-explanatory role of geometric structure stems from its bestowing novel causal powers, then we must reject VSCC. This strategy seems plausible if we have an antecedent commitment to CW. If a property's causal-explanatory role stems solely from its bestowing causal powers on its bearers, then if a property P is one without which a full causal explanation of some token effect can't be given, P must bestow a novel causal power to bring about that effect on that occasion. One way to make sense of properties like geometric structure bestowing powers is to think in terms of Shoemaker's *conditional* powers.[16] For *x* to have the power simpliciter to ϕ is for *x* to be disposed to ϕ, under certain conditions C. For *x* to have a conditional power to ϕ is for *x* to be such that if it had certain other properties, it would have the power to ϕ simpliciter, where the other properties in question are not independently sufficient for this.

As Shoemaker notes, conditional powers enable us to isolate the causal contributions of individual properties to a power simpliciter when that power is jointly bestowed by several properties. It's plausible that basic physical properties such as electric charge bestow certain powers simpliciter, but that doesn't seem to be the case with geometric structure. There is no obvious power simpliciter that all n-tuples with the same geometric structure possess. However, the idea that spatiotemporal properties like geometric structure bestow *conditional* powers is more compelling. The difference-making role of geometric structure described in section 2 suggests that it determines the extent to which the fields of two charged particles are co-oriented at X, which in turn determines a range of possible resultant fields, of different magnitudes and directions. We might then say that geometric structure bestows upon q_1 and q_2 the power to accelerate a positively charged particle located at X at a certain rate in the direction of the resultant E, conditionally on the values of q_1, q_2, r_1 and r_2. Note that a corresponding conditional power is also bestowed on q_1 and q_2 by their charge, in this case conditionally on r_1, r_2 and geometric structure. On this approach, geometric structure bestows novel

[16] Shoemaker (2001), pp. 25-26.

conditional powers, in line with the novel difference it makes to the resultant field. This kind of novelty resembles strong emergence, traditionally conceived, for we have a dependent property with powers that are not inherited from its physical realizer. [17]

The alternative is to deny CW and keep hold of VSCC. On this approach, all the causal *powers* involved in causing a physical effect are bestowed by the basic physical properties of a sufficient physical cause, but other properties may also be involved in *causing* the effect in question. The causal powers that manifest when a particle accelerates in the resultant field of multiple charged particles are bestowed by electric charge, but geometric structure is involved in determining the direction of the resultant field and hence in determining *how* the powers in question manifest on some occasion. We are now faced with the question of how this kind of causal work relates to powers and their manifestations. For present purposes, the following simple account will suffice. Why not simply say that geometric structure is among the *manifestation conditions* of the powers (simpliciter or otherwise) of the two charged particles? On this account, all causal powers are due to basic physical properties, but at least some such powers have irreducibly geometric conditions on their manifestation.

It is widely held that causal powers have manifestation conditions—there are certain things that need to happen for a power to produce its characteristic effect. The power of a knife to cut butter, for instance, will manifest if the knife and butter are brought into contact in the right way. A simple way of thinking about the arguments of the present paper is then as follows: some powers are such that what you have to do to get them to manifest is arrange their bearers in a certain kind of geometric pattern. On this interpretation, we may say that q_1 and q_2 have the power to accelerate a positively charged particle located at X at a certain rate in the direction of the resultant E, that they have this power entirely in virtue of their basic physical properties, but that if you want to get the power in question to *manifest*, you have to fix r_1, r_2 and the relevant geometric structure. Given that vectors

[17] This was the strategy I employed in Yates (2016) to argue that the causal novelty traditionally attributed to strongly emergent properties is consistent with physical realization.

compose parallelogram-wise, the idea that geometric properties might be among the manifestation conditions of fundamental forces is, I think, rather compelling.[18]

We have seen two ways of making sense of the causal-explanatory novelty of geometric structure in vector composition. We can either: (1) say that geometric structure bestows novel conditional powers in line with its novel difference making role and because of this violates VSCC; or: (2) all causal powers are due to basic physical properties, but geometric structure is among the manifestation conditions of such powers, which in turn violates CW but leaves VSCC intact. There is at least one reason to prefer option (2). If we choose (1), there is an odd symmetry between what electric charge does and what geometric structure does, since both are regarded as bestowing conditional powers, which combine to yield a power simpliciter. But that is to gloss over a significant difference in the contributions of these properties—electric charge also bestows powers simpliciter, whereas geometric structure does not. There is no power that all bearers of a certain geometric structure have in common, regardless of their other properties. Since we are clearly already committed to causal powers due to properties such as electric charge, it would be extravagant to posit geometric conditional powers as well if we can understand the causal role of geometric structure in terms of just the former. In what follows, for brevity I will refer to the causal role of being a manifestation condition as *conditioning*.

It might be objected that conditioning isn't really a causal role, so downward conditioning—a higher-level property being a condition on the conditional powers of a basic physical property—isn't really downward *causation*. The conditioning role in question doesn't violate the causal closure of the basic physical domain as formulated in VSCC, because it doesn't consist in power bestowal. But if some of the powers of basic physical properties have irreducibly higher-level manifestation conditions, then

[18] There is another option that I do not consider here for reasons of space, which is to take geometric properties as extrinsic conditions under which some of the conditional powers bestowed by basic physical properties become powers simpliciter. It is somewhat controversial to hold that there is a principled distinction between ordinary manifestation conditions and extrinsic conditions on conditional powers. I defend that claim in order to develop a geometric version of hylomorphism in Yates (forthcoming).

there is more causal work to do even after all the causal powers have been bestowed. If you don't want to call it causal work, then call it something else. The fact remains, higher-level properties like geometric structure have a conditioning role that they don't inherit from their realizers, which role is among the determinants of the dynamic evolution of basic physical systems. That looks like downward causation to me. I turn now to the question of its source.

4. Constraint or coincidence?

The source of downward causation, as understood here, depends on what we say about the source of the law of vector composition. Lange distinguishes two potential sources for explaining why forces compose parallelogram-wise.[19] There are both dynamic and static explanations available, which is to say that the parallelogram law can be deduced either by considerations that stem from the dynamical laws (by considering how the displacements that forces produce compose) or from statics (by considering how general symmetry principles impose vector addition on the force laws). In the first case, vector composition of forces is grounded in the fact that forces produce accelerations and accelerations compose parallelogram-wise. In this case, there is no unifying principle that explains why other vectors—electric fields, gravitational fields, temperature gradients, heat flows, etc.—also compose according to the parallelogram rule. In the second case, the law of vector composition is a constraint on the dynamics and has its source in symmetry principles that are independent of the precise forms of the dynamical laws. These symmetries can also be applied to the case of other vectors, mutatis mutandis, so there is the possibility of a unifying explanation.

Lange doesn't take sides in this debate, since his primary aim is to argue that in order to make sense of the debate itself, we need to appeal to a nested hierarchy of laws. Roughly, he suggests that a law L is a constraint iff L would still have held even if the dynamical laws had been different. Consider the

[19] Lange (2017), ch.4.

counterfactual 'had the dynamical laws been different, vector fields would still have composed parallelogram-wise'. If the law of vector composition holds in virtue of the form that the dynamical laws happen to take at our world, then this counterfactual comes out false—we have no right to assume that a coincidence of the dynamics would still have held on the counterfactual supposition that the actual dynamical laws don't hold. Conversely, Lange argues, for this counterfactual to be true is for the relevant symmetry principles to be *more necessary* than the dynamical laws—to hold, that is, however the dynamical laws may be.

Lange considers dynamic vs. static explanations of the parallelogram rule as part of an argument that a particular historical debate in foundations of physics is best understood in terms of a nested hierarchy of laws. However, it should be noted that the main static account he considers relies on symmetry principles that might themselves be explained dynamically. In Poisson's account of the parallelogram rule, the main symmetry principle involved is spatial isotropy—the principle that space has no preferred direction. Poisson's explanation is complex, and I lack the space to give a full account, but the first stage of the proof relies on isotropy to establish the magnitude and direction of the resultant force of two equal forces.[20] He then generalizes the proof to unequal forces. But as Lange himself notes elsewhere, the same question—constraint or coincidence—can be raised in relation to symmetry principles such as isotropy.[21] Brown and Pooley argue[22] that the symmetries of Minkowski spacetime, including isotropy, are grounded in the dynamics. As they see it, Minkowski spacetime is just our best codification of the primitive Lorentz covariance of the dynamical laws: for Minkowski spacetime to have the metric properties and symmetries that it does, is for the dynamical laws to be Lorenz-covariant.[23] There's no unifying explanation, in Brown's view, as to why all the laws are Lorentz

[20] Lange (2017), pp. 167-169.
[21] Ch.3 of Lange (2017) is devoted to this issue.
[22] Brown & Pooley (2001); Brown (2005); Brown & Pooley (2006).
[23] Lorentz covariance of laws is the property of being invariant under the Lorentz transformations of special relativity, which tell us how to transform the co-ordinates of pointlike events between inertial frames. The laws of physics are Lorentz covariant, which is to say they are the same in all inertial frames, subject to Lorentz transformation of the relevant co-ordinates. See Lange (2017), ch.3 for a detailed account.

covariant, so the Minkowski metric arises from a coincidence of the dynamics.[24] By contrast, a constraint-based explanation of spacetime symmetries will need to locate their source somewhere other than in the dynamical laws—for instance in the structure of spacetime itself. I shall now focus on the implications of this debate for the source of downward causation.

Let's begin with the dynamical account, according to which the dynamical laws are primitively Lorentz covariant, which determines that the Minkowski metric is the best way of codifying those laws. Spacetime symmetries are fundamentally symmetries of the dynamics, so vector composition is likewise a consequence of the dynamics. It is, on this view, a coincidence that the various vector quantities all compose parallelogram-wise. The dynamical account leads to what might be termed "upward-downward causation". The symmetries of the fundamental force-laws determine that multiple forces compose according to the parallelogram rule, which is to say the laws determine that geometric structure plays a novel role in determining the resultant of multiple forces. Downward causation, on this view, is written into the fundamental dynamical laws. And if one thought, as dispositional essentialists do, that the dynamical laws themselves are grounded in the essences of basic physical properties, then those very essences would explain why certain non-basic properties like geometric structure play a novel role in determining the course of events. We can think of it in terms of the relational individuation of powers. In a pure powers ontology, it is typically held that basic physical properties are individuated solely by type-level causal relations they bear to each other. What the case of vector composition shows is that such properties are individuated by relations they bear both to each other *and* to higher-level geometric conditions on their composition.[25]

As noted above, on the dynamical explanation, it's not true that had the dynamics been different, vectors would still have composed parallelogram-wise. If the parallelogram rule is a coincidence, then

[24] Brown (2015), p. 143; see also Lange (2017), pp. 112-113. Brown's position is anti-substantivalist about spacetime and for this reason is typically seen as a form of relationism. See Knox (2019) for full discussion and for a functionalist interpretation of the dynamical conception of spacetime as applied to general relativity.
[25] I develop this idea in detail in Yates (2018) and argue that it helps pure powers ontologies avoid a regress.

we have no right to suppose that the dynamical laws at all the relevant worlds have the same symmetry properties as the actual laws, for the simple reason that coincidences of the actual dynamics are not robust under the counterfactual supposition that the actual dynamical laws don't hold. Conversely, for Lange, if the parallelogram rule still holds under the counterfactual supposition that the dynamical laws don't, then it must have a higher grade of necessity. Lange appeals to laws that hold at a broader range of possible worlds than the dynamical laws, which in Lange's system are understood in terms of a nested hierarchy of primitive subjunctive facts. In the remainder of this section, I shall offer an alternative to Lange's view, which treats vector composition as a constraint within an essentialist framework and hence has no need for primitive subjunctive facts.

The Lorentz transformations can be derived from purely kinematic principles, independently of the dynamics.[26] Lange offers such a derivation, from the principle of relativity (normally stated as the claim that the laws of physics take the same form in all inertial frames) and the invariance of the spacetime interval.[27] What's important for present purposes is that spacetime symmetries such as isotropy are here assumed by way of explaining why the dynamical laws are Lorentz covariant. As Lange notes, the principle of relativity entails spatiotemporal isotropy and homogeneity, so the kinematic explanation builds in a lot of spacetime structure at the outset. One way of interpreting this is to say that the relevant spacetime structure belongs to a substantival Minkowski spacetime and that the nature of spacetime itself is what grounds Lange's constraints. Lange himself does not commit to substantivalism, preferring instead to treat the laws as primitive subjunctive facts; for this reason, his own position is consistent with a thin conception of spacetime similar to Brown's but one which treats

[26] Janssen (2009); Lange (2017), section 3.2.

[27] Lange uses a kinematic version of the relativity principle, stated as follows: "There is a frame S, such that for any frame S' in any allowed uniform motion relative to S, the laws in S and S' take the same form," (2017, p. 104). It's more common for derivations of the Lorentz transformations to appeal to the principle of relativity together with the light postulate—the principle that, as measured in any inertial frame, the speed of light is a constant regardless of the motion of the source. The details of these derivations need not concern us here.

its structure as grounded in primitive *kinematic* laws rather than primitive dynamical laws.[28] To say that spacetime symmetries are constraints is then to say, for Lange, that the kinematic laws have a higher grade of necessity than the dynamical laws.

For those who prefer to ground laws in essences, substantivalism offers an attractive alternative. If symmetry properties like isotropy are among the essential properties of a substantival spacetime, then not only will the nature of spacetime ground the parallelogram rule, but we can also make sense of constraints without positing degrees of necessity or primitive subjunctive facts. In the arguments that follow, I assume no particular version of spacetime substantivalism and do not commit to the view that spacetime points are primitive objects. In other words, I assume that there are versions of substantivalism available that can avoid the hole argument.[29] The only aspect of substantivalism that is required in the present context is the claim that spacetime is ontologically independent of spacetime occupants. This minimal substantivalist claim is consistent with all forms of substantivalism, including metrical essentialism, according to which spacetime points are individuated by their metric properties; and spacetime structuralism, if indeed this is importantly different from metrical essentialism. We need not embrace problematic haecceitistic spacetime points in order to embrace substantival spacetime, provided we do not deny the reality of spacetime points altogether.[30]

Consider again the counterfactual: 'had the dynamical laws been different, vector fields would still have composed parallelogram-wise'. Given substantivalism, it now makes sense to suppose that the

---

[28] Heron & Knox (2019) make this point. Like Lange, Janssen (2009) also eschews any commitment to spacetime substantivalism, despite claiming that spacetime structure explains the Lorentz covariance of the dynamical laws. In Janssen's case things are more complex, as he does not have Lange's nested hierarchy of laws available to render relationism consistent with his view. See Acuña (2016) for arguments that Janssen is tacitly committed to substantivalism about Minkowski spacetime.

[29] See Norton (2019) for an introduction to the hole problem. The problem arises due to diffeomorphism invariance in general relativity (GTR). We can assign different metric properties to spacetime points within a particular region (the "hole") leaving points outside the same, preserving all observational consequences of GTR. But bodies within the hole move along different trajectories, so it looks like GTR violates determinism, because what goes on outside the hole doesn't fix what happens inside it. Whether or not hole diffeomorphisms represent genuinely different physical possibilities depends on how we conceive of spacetime points.

[30] Maudlin (1989); Bartels (1996). This is also a feature of Pooley's (2006) sophisticated substantivalism.

nearest worlds at which the dynamical laws are different are worlds at which spacetime structure is, insofar as possible, the same. For instance, if we are focused on worlds at which Coulomb's law is an inverse cube law, then we should think of worlds with one extra spatial dimension, but not worlds at which spacetime is not isotropic. Holding as much as possible of our spacetime structure fixed, we can imagine varying the dynamical laws in certain ways, since spacetime itself—as far as we know—does not suffice to fix those laws. On a dispositional essentialist account of the dynamical laws, the ontological independence of spacetime gives rise to an interesting asymmetry. If the dispositional essences of basic physical properties such as electric charge are given in terms of Lorentz covariant dynamic equations, then those properties will be ontologically dependent on spacetime. If spacetime structure doesn't depend on what occupies it, then the dependence is asymmetric.

Things are somewhat different in certain approaches to quantum gravity—in loop quantum gravity, for example, spatiotemporal localization is understood in terms of the interaction of fundamental quantum fields with a quantum gravitational field. On this view, the fundamental matter fields, the electromagnetic field, and the gravitational field might all be mutually ontologically independent, with dynamical laws arising as a result of the interaction between them.[31] Even assuming an eventual quantum theory of gravitation, what's important for my purposes is that provided spacetime is not ontologically dependent on the other quantum fields, its structure will be counterfactually robust enough to ground constraints without primitive subjunctive facts. It's not that the parallelogram rule is *more necessary* than the dynamical laws, on this approach—it's indexed to a different portion of modal reality because it has a different *source*. This in turn has the consequence that the closest possible worlds at which the dynamical laws don't hold are worlds at which a substantival spacetime

[31] For more on the potential philosophical implications of a quantum theory of gravity, see the essays in Wüthrich, Le Bihan, & Huggett (eds.) (2021). It is hard to know what will become of properties such as mass and electric charge, not to mention the idea that they bestow causal powers, in quantum gravity. Some such theories seem to have the consequence that spacetime itself is in some sense emergent rather than basic physical, and if so then presumably causal powers, understood as e.g. powers to produce accelerations, will also be emergent. In principle, however, I think the solution to the exclusion problem presented here could be reinterpreted as way of explaining why one special science—current physics—does not causally exclude those above.

with the same symmetry properties as ours exists, hence worlds at which constraints like the parallelogram rule, which follow from spacetime symmetries, also hold.

If vector composition has its source in spacetime structure, then downward causation flows from the way in which spacetime itself constrains the dynamics. However, we must be careful not to conflate the present sense of "constraint" with a causal notion according to which spacetime literally forces bodies to follow inertial paths and forces vector fields to compose parallelogram-wise. Brown objects to this suggestion on the grounds that it imports causal powers to spacetime and "spacetime feelers" to material bodies.[32] If spacetime constrains the dynamics causally, then it has the power to cause a body moving in the absence of forces to follow an inertial path; and it seems the body itself must then have the reciprocal power to be so constrained—the power to "feel" which way the inertial paths point. Spatiotemporal constraints, however, were supposed to be independent of the dynamics of spacetime occupants. Whatever it is for spacetime to constrain the motion of a body or the propagation of a field, it seems it cannot be understood as the manifestation of a power. How then should it be understood?

I do not here claim that spacetime is a cause; rather, the claim is that geometric properties have their causal roles in virtue of the essential nature of spacetime. It might seem as though there is a conflict here between a non-causal notion of constraint and a causal role for geometric structure, but it's one thing to talk about spacetime itself having causal powers and another thing entirely to talk about the causal role of geometric structure conceived as a property of pluralities of spacetime *occupants*. The kinematic explanation of why the parallelogram rule holds is not causal, but given that the rule does hold, instances of geometric structural properties play an irreducible causal role in determining the dynamic evolution of basic physical systems. This is a form of downward causation not because spacetime itself exerts a causal influence on things that are beneath it in a levels hierarchy—indeed,

[32] Brown (2005).

if spacetime itself is a basic physical structure, then there won't be anything beneath it—but because instances of higher-level geometric properties exert such an influence by acting as manifestation conditions on the causal powers of basic physical properties. Spacetime substantivalism is a potential explanation of the *source* of geometric downward causation, but that, as noted above, should not be taken to imply that spacetime itself is a *cause*.[33]

5. Conclusion: so what?

The kind of downward causation defended here may not seem sufficient to defend the autonomy of the special sciences, so let me say something in conclusion about why I think it's a promising start. Firstly, whether it's enough for a robust positive defence, it is sufficient to undermine arguments against special science autonomy based on closure principles such as VSCC. The case of vector composition leaves us with two options in relation to closure: either (i) the basic physical domain is not causally closed because geometric structure bestows novel conditional powers; or (ii) the basic physical domain is causally closed, but geometric structure still has a novel role to play as a condition on the manifestations of basic physical powers. If the causal closure of the basic physical domain is false, then it doesn't pose a problem for the special sciences; if it's true but consistent with an irreducible causal role for higher-level properties, then it is once again consistent with special science autonomy. This does not suffice to show that the special sciences *are* autonomous from physics, but it does serve to undermine the strongest argument that they are not.

Secondly, the arguments presented here can also be used to defend the claim that the special sciences are indeed autonomous. One of the core things that special sciences such as chemistry and biology do is to classify by spatiotemporal properties such as geometric structure, so their autonomy is secured by dint of the downward causal influence of such properties on the dynamics of basic physical systems.

[33] I thank Katie Robertson for pressing me on this issue.

Chemistry, for instance, classifies molecules in terms of properties such as being linear, planar, bent, cyclic, and so forth, all of which are geometric. If I am right that properties such as these are causally novel, then there is no threat to chemistry from below. Chemical properties such as molecular geometry are just the right kind of properties to condition the powers bestowed by their basic physical realizers. This simple account may not extend in a natural way to other special sciences, but I see no reason in principle why conditioning by higher level properties should not occur in sciences such as neuroscience[34] and psychology as well, even if it is not conditioning of the kind that we see in vector composition. The devil is no doubt in the details, but if at least some special science properties have a genuinely novel downward causal role in relation to physics, then that is a promising start.

Traditional accounts of special science autonomy are often framed within a functionalist conception of special science properties. Ultimately, however, the autonomy one can secure within a functionalist framework has principled limits: functional properties are defined in terms of causal roles that are occupied by *other* properties, so whatever they do, it's written into their metaphysical natures that something else is *really* doing it.[35] That something else is of course their basic physical realizers. This paper is part of an attempt to break free from functionalism. The traditional causal exclusion problem, in my view, arises not from causal closure, but from a conception of special science properties that comes with causal redundancy baked in. Geometric structure is higher-level and multiply realizable, but its realization does not consist in occupying a causal role. Ultimately, a conditioning role on the powers of basic physical properties stems either from the essential natures of those properties, or

[34] In Yates (2020) I tried to extend the simple account to neuroscience, through the idea that temporal patterns such as neural synchrony might be causally novel in the same way as the geometric structure of molecules. In a vast oversimplification, I likened the phase angle between oscillations in membrane potential to the spatial angles between atoms in a molecule. Still, I think there is something to the comparison: membrane potential is a vector quantity (the difference in potential between the inside and outside of the cell membrane) and the way in which distinct oscillating populations interact will be determined, inter alia, by the phase angle between their oscillations. This in turn will be at least partially explained by the way vectors compose.

[35] See Yates (2012) for more on the limits of functionalism. There I offer a grounding-theoretic account of the novelty of functional properties, according to which this novelty consists not in the powers they bestow, but in the distance, within a hierarchy of grounding relations, from which they bestow them. This account entails that all causal powers that special science properties bestow are ultimately bestowed by their basic physical realizers. For an attempt to circumvent these limitations of functionalism via what he calls *machresis*, see Gillett (2016).

from the essential nature of spacetime, but the conditioning role itself, on some occasion, is *occupied by geometric structure* and not by its lower-level realizers on that occasion.

In closing, let me acknowledge a potential problem for the overall approach presented here.[36] One might object that geometric properties belong in higher-order explanatory contexts, in which one explains why a cause C explains a certain effect E, but not in first-order contexts in which one explains why E happens. Geometric properties might be needed if we want to explain why a system of multiple charges explains the acceleration of a charged particle, but they do not thereby belong in the first-order explanation of why that charged particle accelerated in the way that it did. What I herein regard as a single, unified multilevel explanation of the acceleration of the particle is really two distinct explanations, at different orders, with different explananda: a first-order explanation of the acceleration of the particle, and a second-order explanation of the first-order explanation.

It's tempting to note that the above objection is consistent with the claim that geometric properties have irreducible causal-explanatory roles, it's just that these roles now consist in explaining why certain first-order causal explanations hold. However, on reflection this doesn't seem to be much help when it comes to defending special science autonomy from the threat of causal exclusion. The causal exclusion problem is that there doesn't seem to be any causal work left for any properties other than those of basic physics, which threatens the *first*-order explanations that the special sciences provide of events such as molecular motions, metabolism, neural firings, and behaviours. Those events, exclusionists say, can in principle be explained in basic physical terms alone, so special sciences provide at most a more perspicuous way of explaining the same events. My strategy here has been to argue that causal explanations even in the simplest physics cases are irreducibly multilevel and that the complete cause of the relevant basic physical effects always involves interactions between properties at multiple levels. On the alternative just mooted, by contrast, the dynamics is wholly

[36] I thank Alastair Wilson for making me aware of the alternative that follows. See Hicks & Wilson (2021) for more on the distinction that follows between first- and higher-order explanations.

driven by basic physics, with only higher-order explanatory work left over for non-basic properties. This wouldn't matter so much were it not for the fact that the ambition of the special sciences doesn't seem to have anything much to do with explaining why first-order physical explainers explain. Rather, their central aim seems to be to provide distinctive *first*-order explanations of their target phenomena.

I don't know how to refute the claim that higher-level properties belong in higher-order explanations, but I can provide some additional motivation for my multilevel approach. If a physicist wants to *predict* the acceleration of a test charge placed in the field due to two or more charges, they first need to calculate the magnitude and direction of the resultant field at the point where the test charge is to be introduced. And to do so, I argued in section 2, they must appeal to geometry. First-order predictions in such cases are simply not possible without appealing to higher-level geometric properties. But it would be very odd indeed to accept that higher-level properties are necessary for first-order predictions while at the same time denying that they feature in first-order *explanations*. Put differently, why is it necessary to appeal to geometric properties to predict the motion of a particle if such properties play no first-order causal role? If I am correct that geometric properties are among the manifestation conditions on basic physical powers, then it's easy to explain why they are necessary to predict how those powers manifest on some occasion. Such properties, on my view, are both predictively *and* explanatorily indispensable in first-order explanations of physical phenomena. Basic physical properties may bestow all the causal powers, but they don't occupy all the causal roles.[37]

[37] Versions of this paper were presented at a workshop in Birmingham on levels of explanation organized by Alastair Wilson's FramePhys project and at a joint LanCog-Complutense workshop in Madrid. My thanks to the audiences at both events for their comments and criticisms. I thank an anonymous referee for their comments. Special thanks to Amanda Bryant, Nicholas Jones, Katie Robertson and Alastair Wilson for very helpful comments on earlier drafts. Based on research funded by the Fundação para a ciência e a tecnologia (awards: IF/01736/2014 and PTDC/FER-HFC/30665/2017).

Bibliography

Acuña, P. (2016). 'Minkowski Spacetime and Lorentz Invariance: the Cart and the Horse or Two Sides of a Single Coin?', *Studies In History and Philosophy of Modern Physics* 55, pp. 1-12.

Brown, H., & Pooley, O. (2001). 'The origin of the spacetime metric: Bell's 'Lorentzian pedagogy' and its significance in general relativity', in C. Callender & N. Huggett (eds.), *Physics Meets Philosophy at the Planck Scale: Contemporary Theories in Quantum Gravity*, Cambridge: Cambridge University Press, (pp. 256–272).

Brown, H. & Pooley, O. (2006). 'Minkowski space-time: A glorious non-entity', in D. Dieks (ed.), *The Ontology of Spacetime*, Amsterdam: Elsevier, 67–89.

Brown, H. (2005). *Physical Relativity*. Oxford: Clarendon.

Crane, T. & Mellor, D. (1990). 'There is no Question of Physicalism', *Mind* 99, pp. 185-206.

Crook, S. & Gillett, C. (2001). 'Why Physics Alone Cannot Define the 'Physical', *Canadian Journal of Philosophy* 31, pp. 333-359.

Fodor, J. (1974). 'Special Sciences', *Synthese* 28, pp. 97-115.

Fodor, J. (1997). 'Special Sciences: Still Autonomous after all these Years', *Noûs* 31(S11), pp. 149-163.

Gillett, C. (2003). 'The Metaphysics of Realization, Multiple Realizability and the Special Sciences', *Journal of Philosophy* 100, pp. 591-603.

Gillett, C. (2016). *Reduction and Emergence in Science and Philosophy*. Cambridge University Press.

Heron, J. & Knox, E. (2019). 'On Constraints, Context and Spatiotemporal Explanation', *Philosophy and Phenomenological Research* 99, pp. 732-738.

Hicks, M. & Wilson, A. (2021). 'How Chance Explains', *Noûs* 57, pp. 290-315.

Jackson, F. & Pettit, P. (1990). 'Program Explanation: a General Perspective', *Analysis* 50, pp. 107-117.

Janssen, M. (2009). 'Drawing the Line between Kinematics and Dynamics in Special Relativity', *Studies in History and Philosophy of Modern Physics* 40, pp. 26-52.

Kim, J. (1992). 'Multiple Realization and the Metaphysics of Reduction', *Philosophy and Phenomenological Research* 52, pp. 1-26.

Forthcoming in A. Wilson & K. Robertson (eds.) *Levels of Explanation*, OUP 2024.

Kim, J. (1998). *Mind in a Physical World*. MIT Press.

Knox, E. (2019). 'Physical Relativity from a Functionalist Perspective', *Studies in History and Philosophy of Modern Physics* 67, pp. 118-124.

Lange, M. (2017). *Because Without Cause*. Oxford University Press.

LePore, E. & Loewer, B. (1987). 'Mind Matters', *Journal of Philosophy* 84, pp. 630-642.

List, C. & Menzies, P. (2010). 'The Causal Autonomy of the Special Sciences', in C. Macdonald and G. Macdonald (eds.), *Emergence in Mind*, Oxford University Press, pp. 108-128.

Maudlin, T. (1989). 'The Essence of Spacetime', in *Proceedings of the 1988 Biennial Meeting of the Philosophy of Science Association*, Vol. 2 (ed. A. Fine and J. Leplin), pp. 82–91.

Norton, J. (2019). 'The Hole Argument', in *The Stanford Encyclopaedia of Philosophy* (Summer 2019 Edition), Edward N. Zalta (ed.). https://plato.stanford.edu/archives/sum2019/entries/spacetime-holearg.

Papineau, D. & Spurrett, D. (1999). 'A Note on the Completeness of 'Physics'', *Analysis* 59, pp. 29-32.

Pooley, O. (2006). 'Points, Particles and Structural Realism', in D. Rickles, S. French, & J. Saatsi (eds.) *The Structural Foundations of Quantum Gravity*, Oxford University Press, pp. 83–120.

Shapiro, L. (2000). 'Multiple Realizations', *Journal of Philosophy* 97, pp. 635-654.

Shoemaker, S. (2001). 'Realization and Mental Causation', in C. Gillett & B. Loewer (eds.), *Proceedings of the Twentieth World Congress of Philosophy*, Cambridge University Press, pp. 23-33.

Wilson, J. (2006). 'On Characterising the Physical', *Philosophical Studies* 131, pp. 61-99.

Wilson, J. (2011). 'Non-reductive Realization and the Powers-based Subset Strategy', *The Monist* 94, pp. 121-154.

Wüthrich, C., Le Bihan, B. & Huggett, N. (eds.) (2021). *Philosophy Beyond Spacetime: Implications from Quantum Gravity*. Oxford University Press.

Yablo, S. (1992). 'Mental Causation', *Philosophical Review* 101, pp. 245-280.

Yates, D. (2009). 'Emergence, Downwards Causation, and the Completeness of Physics', *Philosophical Quarterly* 59, pp. 110-131.

Yates, D. (2012). 'Functionalism and the Metaphysics of Causal Exclusion', *Philosophers' Imprint* 12, pp. 1-25.

Yates, D. (2016). 'Demystifying Emergence', *Ergo* 3, pp. 809-841.

Yates, D. (2018). 'Inverse Functionalism and the Individuation of Powers', *Synthese* 195, pp. 4525-4550.

Yates, D. (2020). 'Neural Synchrony and the Causal Efficacy of Consciousness', *Topoi* 39, pp. 1057-1072.

Yates, D. (forthcoming). 'Hylomorphism, or Something Near Enough', in Yates, D. & Bryant, A. (eds.) *Rethinking Emergence*. Oxford University Press.