On explanations from "geometry of motion"[*]

Juha Saatsi
University of Leeds
J.T.Saatsi@leeds.ac.uk[†]

29th May 2015

Abstract

This paper examines explanations that turn on non-local geometrical facts about the space of possible configurations a system can occupy. I argue that it makes sense to contrast such explanations from "geometry of motion" with causal explanations. I also explore how my analysis of these explanations cuts across the distinction between kinematics and dynamics.

1 Introduction

Most philosophers nowadays grant that not all explanations in science are causal, but some still stubbornly defend the hegemony of causal explanations with respect to explanations of individual events, for example (Skow 2013). This hegemony, or what remains of it, is partly maintained by the opaqueness of the nature of non-causal explanations: there is little agreement how to demarcate between causal and non-causal explanations. At the same time, many people intuitively recognise a difference between causal and non-causal explanations at least in some cases. (I will give an example shortly.) This provides a prima facie motivation for thinking that there is a genuine, worthwhile distinction here to be clarified. The debate can be advanced by examining in detail some intuitively non-causal explanations

[*]Forthcoming in *The British Journal for the Philosophy of Science*

[†]**Acknowledgements:** Thanks to Steven French, Marc Lange, Alexander Reutlinger, anonymous reviewers, and audiences at BSPS 2014 and workshops on non-causal explanations in Cologne and Munich. I would also like to acknowledge that this paper was partly written during my visiting Fellowship at Durham Institute of Advanced Studies (IAS).

in attempt to demarcate their explanatory credentials from causal explanations. This paper does exactly that.

There is no reason to presuppose, prima facie, that all intuitively non-causal explanations are of a piece. Although many non-causal explanations are broadly speaking geometrical in character, some non-causal explanations only concern regularities, for example, as opposed to particular events.[1] Such a difference may, or may not, turn out to be relevant for demarcating between causal and non-causal explanations (cf. Saatsi and Pexton, 2013). Be that as it may, in this paper I focus on explanations of particular events, since the hegemony of causal explanations is least contentious here. I also focus on a class explanations that exhibits a common, unifying character.

I take as a point of inspiration a somewhat amorphous division between *kinematic* and *dynamic* explanations in physics. This division resonates in interesting ways with the non-causal vs. causal distinction. Consider, for example, a typical (but by no means definitive) 'textbook' presentation of the physicists' amorphous distinction in the context of classical mechanics.[2]

Kinematics: the study of the *geometry of motion*; kinematics is used to relate displacement, velocity, acceleration, and time, without reference to the cause of the motion.

Dynamics: the study of the relation existing between the forces acting on a body, the mass of the body, and the motion of the body; dynamics is used to predict the motion caused by given forces or to determine the forces required to produce a given motion.

If taken at face value this suggests that while dynamics is explicitly concerned with causes, kinematic explanations are non-causal. This face-value reading turns out to be misleading: many kinematic explanations are causal. But further reflection supports the notion that *some* explanations from "geometry of motion" are indeed non-causal. We can begin to grasp the nature of these explanations by reflecting on typical kinematic explanations from mechanics, identifying some as non-causal. These can then serve as *exemplars* of a broader notion of non-causal explanation that is not restricted

[1]For instance, it is not very clear whether there is a particular event that can be explained by the seemingly non-causal explanation involving the 'graph-structure' of Königsberg's bridges. The explanation of Kleiber's law discussed in Saatsi and Pexton (2013) only concerns a regularity.

[2]Cf. Beer (2010).

to kinematics, and is also applicable outside classical mechanics. I will illustrate this broader notion with reference to other non-causal explanations that turn on "geometry of motion" in the way brought out by the exemplars.

As an entry point I will look at a simple toy example of a familiar, paradigmatic kinematic explanation that is intuitively non-causal. By analysing this example I argue that some kinematical explanations really do furnish genuinely non-causal explanations by virtue of involving irreducibly non-local geometrical features. I characterise the non-causal nature of this exemplar by defending it in detail against some potential objections (§2). After that I will explore explanations that turn on such "geometry of motion" in more general terms, appropriating the exemplar to explanations from different walks of physics (§3 and §4). The upshot is an account of non-causal explanations that are: (a) *non-local*, by virtue of turning on non-local geometrical facts about the space of possible configurations a system can occupy, and (b) *robust*, by virtue of being independent of the details of the underlying dynamics. Such explanations, although arguably non-causal, are actually not radically different from familiar causal explanations; both can be accommodated in similar counterfactual terms. Nevertheless, the distinction between causal and non-causal explanations is a worthwhile distinction, given its value in physics (§5).

2 Toy example

Take a straight stick and parallel-transport it over a closed path on the surface of a sphere of radius r (see Figure 1).[3] Returning to the starting point A, the stick has rotated an angle α. Whence the angle α?

Figure 1: Parallel transport of a stick on a spherical surface.

The explanandum is an individual event, or (to be precise) an aspect thereof, with a natural contrast class: why α, as opposed to some other angle α'? The event in question takes place when the stick returns to the point A. We are interested in a why-question that presupposes that the stick

[3] Mathematically speaking, a vector can be parallel transported along a smooth curve if the manifold in question is equipped with an affine connection which gives a sense in which the vector must locally 'stay straight' along the way.

started its journey earlier from this very point: why α, *given that* the stick pointed in a different direction at this point earlier? This question concerns the stick's direction at A, but only in as far as this differs from its earlier direction.[4]

You are probably familiar with the explanation, hinging on a simple regularity that holds between the following variables characterising the set-up: α, r =radius, and A =the area enclosed by the path:

$$\alpha = r^{-2} \times A \tag{1}$$

The stick has rotated an angle α because (i) the path traversed (without rotating the stick along the way) encloses area A over a sphere of radius r, and (ii) α depends on variables A and r as per the regularity above.[5]

Intuitively this strikes me as a purely geometrical, non-causal explanation, exhibiting an asymmetric dependence of α on the curvature of the sphere and the area enclosed by the loop. The explanation seems non-causal, because the variables in question do not seem to qualify as causes of the outcome stated in the explanandum. The curvature does not 'twist' the stick at any point along the path, and the path on the whole determines α in a way that only depends on its non-local feature (the area A).[6] This is—as I

[4]This why-question does not ask why the stick, returning to A, points in some *particular* direction in space. Explaining the stick's particular direction at P at the end of the journey would require reference to its particular direction at the start of the journey. This complicates matters by introducing a further local and causal explanatory variable into an explanation that is (I will argue) otherwise non-causal.

[5]This dependence of α on A and r^{-2} is easiest to establish for a triangular path as in Figure 1, but the results generalises to arbitrary closed paths that can be approximated with some complex polygon that can be divided into triangles.

[6]The distinction between local and non-local properties of geometrical objects (such as a curve) is standardly drawn both in mathematics and physics. (See e.g. Banchoff and Lovett, 2010.) Local properties, such as local curvature, can be attributed to a small neighbourhood of a point on a curve. Non-local properties, by contrast, are only applicable to the whole curve (without being mere sums over local properties, in the way a curve's overall length is, for instance). The area enclosed by a closed curve is a paradigmatic example of a non-local property of the curve. (The total area of a sheet of paper is not a non-local property *of the sheet* in this sense, however, for it is just a sum of the areas of its parts.) For another example, consider the non-local topological property of a rubber band: being a closed loop. Admittedly being a closed loop *depends* on local facts: e.g. one can cut open a closed loop at a single point. So this sense of non-locality is clearly different from the non-locality associated with quantum mechanics, for example, and the non-local features supervene on the local ones. But the topological feature remains non-local in an interesting sense, however: the point of the present distinction is to bring out a way in which one has to take the *whole* rubber band into account in determining whether it is closed or not. Being a closed loop is not a feature applicable at all to any proper part

will presently explain—a paradigmatic example of a kinematic explanation in terms of "geometry of motion".

The intuition from this toy example has to be defended against two challenges that aim to massage it to fit the causal hegemony. In the course of vindicating the intuition I will identify the key features of the exemplar that we can associate more generally with the kind of non-causal explanations "from geometry of motion" that I am interested in.

Here is the first challenge. One might argue—in the spirit of Skow (2013), perhaps—that the explanation is causal after all, since it does provide *some* information about relevant causes. This line of thought is part and parcel of a tradition crystallised in David Lewis's influential 'Causal Explanation', according to which "to explain an event is to provide some information about its causal history" (Lewis 1986, 217). The explanandum at stake has a causal history involving whatever causes performed the parallel transport as described, and the geometrical explanation seems to provide *some* information about it, to boot. After all, it is explicitly said that *no* causes rotated the stick en route, and it is implicitly said that *some* causes (perhaps relating to your intentions and actions) were responsible for taking the stick around the loop. Hence, the explanation allegedly is causal after all, albeit it is so in a somewhat roundabout way.[7]

I rebut the first challenge as follows. The explanation is not a causal one, since none of its explanatory power is due to the causal information it provides.[8] Following Woodward (2003) and others, I take a paradigmatic form of scientific explanation to consist in showing *how* the explanandum depends on the explanans. The explanation given above of the value of α does exactly this by virtue of relating α to r and A in a way that brings out a precise and informative functional dependence between these relevant

of the loop, and it is not a 'sum' of the local features. If some explanandum depends only on this feature of the rubber band, then any feature attributed to any proper part of the band (location, local connectedness, curvature, what-have-you) says nothing about whether the explanandum holds or not. (Thanks to a referee for pressing me on this.)

[7]Cf. Skow's criterion for causal explanation:

> A body of fact partially causally explains [event] E iff it is a body of fact about what causes, if any, E had; or if it is a body of facts about what it would have taken for some specific alternative or range of alternatives to E to have occurred instead. (213, p. X)

[8]Lange (2013, pp. 495-6) presses this point in connection with different non-causal explanations.

variables. Armed with the regularity expressed in equation (1) we know precisely how α would have been different had r and/or A had different values. By contrast, the (minimal and vague) causal information provided by the explanation does not amount to anything like an informative functional dependence of α on the existence of some causes (or absences of others). They are better regarded as contingent features that are presupposed by the why question: whence α, *assuming that* the stick somehow performs the parallel transport? Moreover, any details of the specific causes and dynamics involved are simply supererogatory, since the explanandum is actually completely independent from the actual laws of motion and the forces involved (as long as the presuppositions that are part of the why-question are satisfied). (Cf. Lange 2013)

So, arguably all the explanatory power springs from the functional dependence of α on the sphere's curvature and the area enclosed by the loop. Grasping this dependence in sufficient detail gives one something to *work with*: if you wish the angle to shift only half the original amount, for example, you can accordingly 'manipulate' (broadly speaking) the sphere's radius or the area enclosed by the path to make it happen. This complies with the *spirit* of Woodward's 'manipulationist' account of causal explanation:

> [Causal explanations] furnish information that is potentially relevant to manipulation and control: they tell us how, if we were able to change the value of one or more variables, we could change the value of other variables. (2003: 6)

You would, of course, have to perform the parallel transport anew to do this, because the dependence of α on these variables is non-local. So, even though you may be able to thus broadly speaking 'manipulate' the explanans variables, this cannot be achieved by a surgical intervention on any relatively local part of the path. Furthermore, changing of the value of an explanans variable (e.g. radius) need not involve anything like an *intervention* in Woodward's sense at all: whether or not a change in these variables is caused by something (e.g. air pressure inside a balloon) is neither here nor there for the explanation at stake. Thus, the explanatory dependence does not appear causal, as the explanatory geometrical features—the sphere's curvature combined with the area enclosed by the loop—do not boil down to any localisable or intervenable features that 'make a difference' in a difference-making sense associated with causation.

The last claim requires careful elaboration, as this is where the second challenge comes in. For one might argue—in the spirit of Strevens (2008), say—that the pertinent geometrical variables can be construed as causal after all, despite not being localisable. This line of thought hinges on the notion that causes qua *difference-makers* typically involve an abstraction from some ultra-specific causal facts. For example: in causally explaining the breaking of a window that happened to lie along a rock's path, we appeal to those causal features that made a difference to the breaking (as opposed to not breaking): the rock hitting the window in *some* 'sufficiently impactful' manner, which can be realized by countless specific combinations of the rock's hardness; a specific location of the impact; a specific momentum that is high enough for that location, etc. We typically explain with causes thus identified by a higher-level description of the system that abstracts away from the unnecessary specifics that a true lower-level causal description would attribute to the system.[9]

Applying this line of thought to a simple geometrical example shows how non-local geometrical features can be causal qua difference-makers. Consider a paradigmatic equilibrium explanation, for instance: why does a ball eventually end up at the bottom of a concave bowl, when released just inside the bowl's rim? An ultra-specific causal description involves the specific forces acting on the ball over time, due to air resistance and the specific parts of the bowl that are in contact with it before reaching the energy-minimising equilibrium state. But clearly those very causes and parts of the bowl are not the difference-makers with respect to the explanandum at stake, for the outcome would not be any different if the ball was released at some other point along the bowl's rim. (Ditto regarding the specific details concerning the amount of air resistance and friction, and the strength of gravity, for example—as long as these causal factors are large enough.) What makes a difference, rather, is the bowl's *concave shape*. Although this difference-maker is clearly a geometrical, non-local feature of the system, it is also a clearly identifiable causal feature that makes a difference for the explanandum at stake, given the fundamental causal facts in play. The fact that an explanation in terms of this difference-maker abstracts away from the un-

[9]See Strevens (2008) for a philosophy of causation and causal explanation that makes the most of this idea.

necessary specifics of a lower-level causal description does not make it any less causal—according to the advocates of the difference-making concept of causation at least.[10] So, surely, the challenge goes, we can similarly construe as causal the dependence of α on r and A.

But not so: there are critical disanalogies between the two cases. In particular, in the case of the bowl there is always an ultra-specific causal description that can be associated with the high-level 'geometrical' explanation. Indeed, the explanatoriness of the 'geometrical' explanation in the case of the bowl requires the existence of *some* such specific causal description: the 'geometrical' explanation breaks down if it is not taken to imply something about the explanatory relevance of the gravitational and frictional forces involved. By contrast, in our toy example this is not so: the validity of this geometrical explanation actually does not imply *anything at all* about some specific lower-level causes involved. The explanation would hold just the same—it would explain equally well and for the exact same reasons—even if no forces were in play at all. Thus, the fact that some non-localisable geometrical feature (such as *being concave*) may count as a causal qua difference-maker does not imply that all geometrical explanations are causal in this sense. After all, higher-level causal difference-makers are obtained by abstraction from lower-level causal facts, and while such an abstraction takes place in the case of the equilibrium explanation, it does not

[10]Woodward also regards the equilibrium explanation causal, saying that this is an upshot of his "broad notion of causal explanation according to which, roughly, any explanation that proceeds by showing how an outcome depends (where the dependence in question is not logical or conceptual) on other variables or factors counts as causal." (2003, 6). The qualification 'roughly' is critical here; without this qualification the notion of causation would be too liberal, including automatically too many explanations under the causal heading (as Woodward clearly recognises elsewhere). The qualification relates to the interventionist aspect of Woodward's account, the idea being that the notion of intervention makes sense with respect to the shape of the bowl, for instance. This judgment seems right on the basis of our pretheoretic folk conception of causation, which takes causation to be a relatively local matter. We must allow our pretheoretic folk conception of causation to inform Woodward's account to some extent, at least, due to the oft-noted conceptual circularity in Woodward's account: causation is characterised partly in terms of intervention, which itself is a causal notion. We need some idea of causation to get the account started. The fact that the folk conception of causation takes it to be a relatively local matter renders some variables non-intervenable, as Woodward notes for example in admitting dimensionality of spacetime as a variable on which the stability of planetary orbits non-causally depends. The vagueness in our folk conception of causation leads to a degree of vagueness in the distinction between intervenable and non-intervenable variables, and correspondingly in the demarcation between causal and non-causal explanations. In unclear borderline cases we can appeal to other theoretical benefits of drawing the line in a particular way.

take place in the toy example. The toy example thus has explanatory *robustness* that distinguishes it from explanations such as the bowl case: it is fully independent of the details of the underlying dynamics. The explanation would function just the same, even if there is no underlying causal dynamics at all. (By contrast, Strevens's analysis of higher-level causal explanations is premised on causal fundamentalism: the idea that there is a network of causal influences at the fundamental level that ultimately underwrites the truths about higher-level causal relevance.)[11]

We can appeal to this kind of distinctive robustness in demarcating explanations akin to the toy example from causal explanations that appeal to geometrical features. Thus, I conclude that the second challenge need not compel us to give up the intuition that we have a non-causal explanatory dependence at stake in the toy example. This vindication of the toy example as non-causal may not be watertight; perhaps intuitions also differ in this regard at this point. Still, I believe to have shown that there is a principled line that can be drawn between the toy example, on the one hand, and explanations—geometrical or otherwise—that come out as straightforwardly causal by the lights of the leading counterfactual accounts causal explanation. The rest of the paper amounts to a case for regarding this line as a *worthwhile* one one to draw. (Given my pre-theoretic grasp of causation I also regard it completely natural to signify this distinction by the contrast between 'causal' vs. 'non-causal', although ultimately I care less about framing the distinction in this way, and more about recognising its philosophical importance.)

Having characterised nature of the explanation furnished by the toy example, let us summarise its key features. The explanation of the specific angle α is **geometrical** by virtue of turning on the dependence of α on geo-

[11] Admittedly the bowl case has considerable modal robustness as well, but this is largely due to a high degree of abstraction in the explanandum itself. That is, the abstraction from 'unnecessary specifics' in causal difference-making explanations is licensed by a degree of abstraction in the explanandum itself. To wit, there are countless ways for a window to *end up broken*, and there are countless of ways for a ball to *end up at the bottom of a bowl*. Instead of the very broad question, 'Why did the ball eventually end up at the bottom when released just inside the bowl's rim?', we can ask a much more specific one: 'Why did the ball end up at the bottom in 17.5435 seconds, given $g = 9.80665 m/s^2$, such-and-such air pressure, temperature, etc.?' *Having a concave* shape no longer functions as an explanatory difference-maker, because the specific concave shape now *does* make a difference—for almost all concave shapes—to the time it takes for the ball to reach the equilibrium.

metrical variables: the area A inside the loop, and the radius r quantifying the curvature of the sphere. The explanans variables are **non-local**: the explanation boils down to non-local geometrical facts about the (relevant features of the) sphere and the area enclosed by the whole path.[12] The explanatory value of these non-local geometrical facts is not dependent on any assumptions about underlying dynamics or causal structure of the world (nor does it have anything to do with the degree of abstractness of the explanandum). Indeed, the explanation is fully **robust** in the sense that it is completely independent from any laws of motion or forces involved.[13] The explanatory dependence at stake is **asymmetric**: α depends on A and r, not vice versa, in a way that supports a broadly 'manipulationist' reading and provides corresponding **counterfactual information** (cf. *what-if-things-had-been-different* questions in Woodward, 2003). Finally, given the set-up of the toy example, there are **constraints on possible states** of the stick, imposed by the spherical surface one the one hand, and by the requirement of parallel transport on the other.

As said, I also take the toy example to provide an exemplar of a type of **non-causal** explanation. (As such it provides a straightforward counter example to the idea that all explanations of particular events are causal, cf. Skow, 2013.)[14] But what exactly makes the explanation non-causal? In responding to the two challenges above I suggested that this is due to *both* (i) the non-local character of the explanans variables, as well as (ii) robustness of the explanation with respect to variation in the dynamical facts (including laws) involved. It is these two features that demarcate this type of non-causal explanation, since the other features are typical to causal explanations nicely captured by popular counterfactual accounts (e.g. Woodward 2003). I

[12]In as far as the explanatory geometrical facts are faithfully represented by mathematics, the explanation is a **distinctly mathematical** one (Lange, 2013).

[13]The explanatory geometrical facts are in this sense fully explanatory of the explanandum, unless there are causal facts ('twisting' the stick en route) also involved, in which case the explanation of α requires the geometrical part as a distinct and indispensable explanatory component.

[14]To be exact, Skow defends the idea that all explanations of particular events are causal—*except* for those that clearly are not, namely 'in-virtue-of' explanations. For example, a chemist arguably thus explains the fragility of a piece of glass in terms of its molecular structure; more generally, such explanations "explain why some fact obtains by citing some other fact or facts that 'ground' the target fact, that are the 'deeper' facts 'in virtue of which' the target fact obtains" (Skow 2013, p. XX). I see no reason to think that the toy example should be classified as a 'grounding' explanation, as opposed to a counterfactual explanation based on non-causal dependence.

will further analyse the nature and role of these demarcating features below by first arguing that neither non-locality nor robustness are singly sufficient for rendering an explanation non-causal. Then, the recognition that the demarcating features are not specific to kinematics points towards a more general appreciation of non-causal explanations from "geometry of motion".

3 "Geometry of motion" in classical mechanics

Having described and defended an exemplar of non-causal explanation, I now wish to explore in broader terms other explanations that it can be assimilated with. This section will locate the key features of the toy example in some actual explanations in classical mechanics. The next section goes beyond mechanics with illustrations from other areas of physics.

Recall the initial characterisation of kinematics as a "study of geometry of motion". Another classic treatise in theoretical kinematics begins by presenting its subject matter in similar terms:

> Formally, kinematics is that branch of mechanics which treats the phenomenon of motion without regard to the cause of the motion. In kinematics there is no reference to mass or force; the concern is only with relative positions and their changes. (Bottema and Roth, 1979)

If kinematics (thus construed) furnishes explanations, these have the potential to be non-causal by virtue of being completely independent of the cause of motion. The toy example exemplifies how this potential can be realised: the toy example is a kinematical explanation that concerns the "geometry of motion" that provides only information about relative positions and their changes with respect to a stick constrained to parallel transport on a spherical surface. This information brings out explanatory, counterfactual dependencies by reference to the relevant geometrical features of the space of possible states. Furthermore, these can be explanatorily independent of the system's dynamics or causes of motion, making the explanation robust. The geometrical features need not be independent from dynamics and forces in any fundamental metaphysical sense, of course; the independence can be only relative to a level of description. The spherical surface in the toy example, for instance, is assumed to be rigid so as to support the weight of a straight stick; at a lower level of description there are forces in play corresponding to this rigidity. But the kinematic explanation is *robust* so that the details about those forces—whatever they are—are completely 'screened off' from explanatory relevance by the assumptions regarding the sphere's rigid shape (presupposed by the why-question at stake). That is, we can 'vary' the forces and laws responsible for rigidity and for the geometrical features of the surface and the stick as much as we like without affecting the kinematic explanation in any way. The non-causal explanatory dependencies have such *modal* independence, even if some of the features of the explanans physically depend on underlying causal mechanisms or laws. (cf. Lange 2013)

Kinematic explanations are robust by virtue of being independent of dynamical features. I have argued that *some* kinematic explanations are non-causal due to incorporating explanatory geometrical features that are furthermore non-local. It is the combination of non-locality and robustness that is needed to drive a wedge between non-causal explanations from geometry of motion, and explanations that appeal to non-local geometrical features as causally difference makers (as in the case of the bowl). Clearly many kinematic explanations do *not* tick the box of non-locality, despite being robust. I would not argue that such explanations are non-causal.

Consider the following elementary kinematic explanation, for instance. A motor bike starts from rest and has a constant acceleration. Why does it travel 100 meters in 4.47 seconds? Because its rate of acceleration $a = 10m/s^2$ and the time it takes to travel distance s depends on the rate of constant acceleration as per $t = \sqrt{2s/a}$. This equation provides explanatory counterfactual information: had the (constant) acceleration been different, the time would have been different accordingly. This explanation is kinematic since it says nothing about the forces in play or the laws of evolution. But despite being kinematic, the explanation is straightforwardly causal: the explanans variable a attributes a local feature to the motorbike, namely its acceleration at any given point, that makes a difference to how swiftly it covers the distance.[15]

So, a face value reading of textbook characterisation of kinematics as

[15]Think of the dependence between the variables t, s, and a in terms of a graph plotting the velocity v as a function of t. (Acceleration a gives the slope of the graph.) The area under the line quantifies s, the distance travelled, and it clearly depends additively on local features of the line.

“the study of classical mechanics which describes the motion without consideration of the causes of motion” is misleading, since ‘cause’ (as philosophers understand the term) cannot be simply equated with ‘force’, ‘energy’, or some other notion that clearly falls on the dynamics’ side of physicists’ division.

Nevertheless, there are various explanations in kinematics that turn on “geometry of motion” in a non-causal way by virtue of satisfying both nonlocality and robustness, as in the case of the toy example. There are interesting, much less toy-ish real-life explanations to exemplify this, too. Consider, for instance, the following simple device consisting of two rigid 10cm long arms XY and YZ joined together by a hinge (Y), and a wheel (W) rotating around a fixed point on the arm YZ. The device can freely rotate around a pivot point (X). (Figure 2)

Figure 2: Polar planimeter. From Jennings (1994, 39).

Assume the pointer end (Z) traces exactly once around a simple closed curve in the plane.[16] As it does so, the two arms move and the wheel rotates. Assume that at the end the wheel has rolled 1cm from its starting position. The wheel, let’s imagine, is set to mechanically give you a reading which is simply ten times the distance rolled by the wheel: viz. 10. This (in cm^2 units) turns out to be exactly the area enclosed by the curve. Why?

The explanation of the reliable functioning of such *polar planimeter* turns on the fairly easily provable fact that the area (A) inside the closed curve is equal to the length (L) of the arm YZ times the distance (D) rolled by the wheel.[17]

$$A = LD \tag{2}$$

This simple equation brings out explanatory dependencies between the relevant geometrical variables pertaining to the set-up’s geometry of motion: the distance D rolled by the wheel depends on A (and a fixed L). This dependence of the rolled distance D on the non-local explanans variable is

[16] Simple closed curve does not cross itself like, say, the figure of eight.

[17] As said, for our imaginary planimeter L=10cm, of course. The regularity in equation (2) has an interesting, purely geometrical derivation, which furnishes a further non-causal, ‘distinctly mathematical’ explanation of the regularity itself. See Jennings (1994).

non-causal, for clearly the area as a non-local variable does not cause the wheel to roll, or cause 'the final reading'. Nor is the dependence in any relevant way based on the laws of motion or forces involved in the process. This dependence is nevertheless explanatory, since we know precisely how D would have been different, had the area been smaller or larger, and we know that D had nothing to do with any specific local feature of the path around A or the causes of the planimeter's motion. This dependence is furthermore asymmetric, since the area only depends on the path taken by the pointer (Z). Finally, there are geometrical constraints to the system given its set-up (fixed pivot point, single hinge, the length of the rigid arms, etc.): these constraints limit the possible (kinematic) states the system can occupy. The explanation of polar planimeter's correct reading shares all the key features identified in connection with the initial exemplar. In particular, as we have seen, there are non-local geometrical facts about the space of possible states associated with a polar planimeter that are explanatory without being causal.

So, some, but not all, kinematic explanations are non-causal. But is it the case that all non-causal explanations that thus turn on non-local "geometry of motion" are exclusive to kinematics? The answer is no. There are significant non-causal explanations in physics that, despite falling on the dynamics' side of physicists' division, turn on "geometry of motion" in a way exemplified by our kinematic exemplar. One paradigmatic geometrical explanation in mechanics pertains to the so-called 'geometrical phase' of a dynamical system. This is broad-ranging phenomenon, but some of its instances are closely related to our toy example. Consider the familiar Foucault's pendulum, for instance. It exhibits a phase shift in the pendulum's plane of motion that is explained in terms of "geometry of motion". The turning of Foucault's pendulum is too often falsely 'explained' in terms of the idea that the pendulum's plane of motion remains fixed as Earth rotates 'underneath it'. This idea completely fails to account for the *precession* of the pendulum, however: the fact that the plane of motion fails to return to its original orientation after a full (24h) revolution of the Earth—unless the pendulum is located at a pole or the equator. In reality the phenomenon can be explained in purely geometrical terms.

The daily precession α' experienced by a Foucault pendulum, as a function of the latitude λ, is given by $\alpha' = -2\pi \sin \lambda$. This is known as a geometrical phase (or (an)holonomy, or Hannay's angle) of a classical dynamical system that travels a closed loop in a parameter space without returning to its original state. As suggested by its name, geometrical phase depends on geometrical features. As said, in some cases it does so in a way that is closely related to our toy example. The overall phase shift of Foucault's pendulum, for example, is equal to the solid angle subtended by the circle of latitude, i.e. $\alpha = 2\pi(1 - \sin\lambda) = r^{-2} \times A$.[18] The geometrical nature of the phenomenon allows the phase shift to be derived from geometrical considerations that indicate how the dependence of the phase shift on latitude, as captured by this equation, has "a deep geometric meaning and is independent of the local properties of the path—it depends only on the enclosed area." (von Bergmann et al, 2007, 891)

A geometrical explanation of Hannay's angle requires the use of dynamical concepts, as the phenomenon depends on the motion exhibiting conservation of momentum, for instance. More generally, it depends on the Hamiltonian, which must be integrable, and the requirement that the Hamiltonian's excursion in the parameter space is sufficiently slow, i.e. adiabatic.[19] Nevertheless, despite this dynamic context, the explanation has a geometrical character that is robust and non-local in the sense exemplified by kinematic explanations from "geometry of motion": the explanation is fully independent of the specific laws and causes in the sense that the geometrical phase only depends on non-local geometrical variables involved in a closed loop in the system's parameter space. For Foucault's pendulum, for instance, it does not matter how fast or slowly we complete the rotation of the globe (as long as we do it adiabatically, of course). Neither does it matter at all what the causes (if any) of the rotation are. This explanation from "geometry of motion" is thus robust and non-local as per our exemplar, and thus non-causal in this sense.

I will bring this section to a close by precisifying further the present viewpoint by contrasting it to an explanation discussed in Lange (2013). The why-question at stake involves a double pendulum in homogeneous gravitational field, which (as a matter of necessity) has at least four equilibrium configurations. Why is that?

[18] A is the area of the Earth's surface enclosed by the circle of latitude; r is the radius of the Earth.

[19] See Hannay (1985) or Arnold (1989) for details and other examples.

Figure 3: Double pendulum. From Lange (2013).

Lange documents an explanation of the number of equilibrium configurations that is 'distinctly mathematical', and in his view also non-causal. As is well known, the configuration space of a double pendulum is a torus parameterised by the two angles α and β (Figure 3). The explanation of interest turns on a general geometrical fact about compact surfaces such as this configuration space: most differentiable functions are Morse functions, which on a compact surface of genus g have at least $2g + 2$ critical points. A torus ($g = 1$) thus has at least four critical points (viz. saddle points, maxima, or minima). The potential energy function of the system is given by a Morse function; it thus has at least four critical points. Each critical point of the potential energy function is also an equilibrium point: at critical points the gradient of the potential energy vanishes, and this amounts to there being no force, viz. an equilibrium.

Figure 4: Pendulum's configuration space with a potential energy function.

This explanation, too, clearly belongs to dynamics (as opposed to kinematics), depending on a potential energy function and forces deriving from it. As Lange puts it, the explanation appeals to "a natural law: a system is at equilibrium exactly when the net force on each of its parts is zero (i.e. when its potential energy is stationary)—a particular case of Newton's second law." (2013, 503). Lange nevertheless argues that the explanation is a non-causal one, since it does not include any *specific* information about the way in which the system evolves under force.

> Any causal explanation in terms of forces (or energy) must go beyond Newton's second law to describe the particular forces at work (or energy function in play)—if not specifying them fully, then at least giving their relevant features (such as their proportionality to the inverse-square of the distance). That is why the distinctively mathematical double-pendulum explanation I have just described is non-causal despite including Newton's second law. (Ibid.)

But I am not sure why an explanation should thus involve any more specific features of the forces at work (or their effect on motion) in order to be causal. Admittedly the causal information provided by (the pared-down version of) Newton's second law is rather *abstract*, but arguably this is all the causal information that is relevant for the (relatively abstract) explanandum at stake. The explanation brings out the way in which for almost any configuration of the system there are local differences in the potential energy that cause the system to move, apart from a few (but at least four) equilibrium configurations. By virtue of saying all this about the causes within a double pendulum system we might well be inclined to regard this as a causal explanation. (Note that with respect to its high degree of abstractness this explanation is not that different from the causal equilibrium explanation discussed in section §2—the equilibrium explanation did not require specific information about particular forces at work either.)[20]

It may nevertheless seem that there is a significant non-causal aspect to the double pendulum explanation, turning on the fact that the configuration space is a torus—a fact about the global "geometry of motion". Given this, is there still a close parallel between this explanation and the non-causal explanations explored above? I do not think so. For while it is true that there is a non-causal *aspect* to this (in my view causal) explanation in a sense, this can be taken to be merely due to the specific why-question at stake. The explanation for the equilibrium nature of each of the critical points, taken individually, turns on relatively local matters: there are no forces at *this* configuration, but forces would be in play were the system to shift to any nearby configuration allowed by the kinematic constraints. The fact that there are more than one—or, as it turns out, more than three—such equilibrium points altogether is, unsurprisingly, a non-local feature of the system. Thus, the explanation turns on non-local matters regarding the configuration space purely by virtue of asking a question about such non-particular fact. This is quite a different explanatory input from the cases of "geometry of motion" explored above.

[20]The two explanations are quite different with respect to their explanans, however, for while the explanandum of the equilibrium explanation is an individual event, in the present case the explanandum is a more abstract regularity without any corresponding individual event.

4 Beyond classical mechanics

Thus far we have discussed non-causal explanations firmly in the context of classical mechanics. But the key idea applies naturally also to other areas of physics where we can conceptually distinguish between pertinent non-local explanatory facts regarding the space of possible physical states of a system, and its laws of evolution.[21] In the light of this general distinction, the leading idea of this paper—that some non-causal explanations may completely turn on non-local geometrical facts about a system's space of possible physical states, viz. its "geometry of motion"—is not at all wedded or native to classical mechanics. Rather, such non-causal explanations can take place whenever a theory posits a space of possible physical states with sufficiently rich structure to ground robust explanatory non-causal dependencies of the sort we have explored above. Here I provide a further example to illustrate this.

The discussion of geometrical phase above anticipates a natural extension of our analysis beyond classical mechanics. A huge range of physical phenomena is accountable in terms of an anholonomy in a system's parameter space; Hannay's phase in classical mechanics is but a special case.[22] For example, in quantum mechanics the so-called Berry phase relates to the fact that a system's wavefunction often exhibits a phase difference after its parameters cycle slowly (adiabatically) around a circuit in an abstract parameter space that specifies the system in relation to its environment. In general terms, it follows from Schrödinger's equation that for a slowly changing Hamiltonian $H(\mathbf{R})$ that depends on parameters $R_1(t)$, $R_2(t)$,..., $R_n(t)$, and for a system in a discrete eigenstate ψ_n of $H(\mathbf{R})$, the geometric phase shift over a closed circuit $\mathbf{R}_0 \to \mathbf{R}_T$ in parameter space is

$$\gamma_n(T) = i \oint \langle \psi_n | \nabla_R \psi_n \rangle$$

This phase difference again depends only on non-local geometrical properties of the circuit in parameter space, and it is independent of the time taken to complete the circuit (apart from the requirement that the process is slow

[21] The conceptual distinction between kinematics and dynamics is also completely independent from mechanics.

[22] There are various known examples of geometrical phase, coming e.g. from classical optics, nuclear magnetic resonance, hydrodynamics, and quantum field theory. See Wilczek et al. (1989).

enough to constitute an adiabatic chance). In the simplest cases the phase difference is proportional to the area enclosed by the circuit, in close analogy to the classical Hannay's angle of Foucault's pendulum.

For a concrete example (adapted from Griffiths (1995)), consider a spin $S = \frac{\hbar}{2}$ neutron in a magnetic field $\mathbf{B}(t)$ that is constant in its magnitude but changes direction. In particular, assume that $\mathbf{B}(t)$ makes a fixed angle α with the z-axis, precessing around it at constant angular velocity ω. Assume that the neutron starts in a spin-up eigenstate in direction of $\mathbf{B}(0)$. When the direction of $\mathbf{B}(t)$ slowly shifts, the neutron remains in a spin-up eigenstate in the direction of the magnetic field. Returning to the starting point in the parameter space, the neutron is in the original eigenstate but it has acquired a geometric phase which, for such an adiabatic excursion around the parameter space of the Hamiltonian $H(t) = \frac{e}{m}\mathbf{B} \cdot \mathbf{S}$, is

$$\gamma(T) = \pi(\cos\alpha - 1).$$

That is, the geometric phase is proportional to the solid angle subtended by the magnetic field vector from the origin. Thus this equation brings out the precise functional dependence of the Berry phase only on the non-local geometrical variable determined by α. Furthermore, this result generalises to arbitrary closed excursions of the Hamiltonian in the parameter space: if the shifting direction of the magnetic field leads the neutron's spin adiabatically around an arbitrary closed path, the geometric phase shift is equal to minus one half the solid angle swept out by the magnetic field vector. (See e.g. Griffiths 1995) This dependence of the phase shift on non-local features of such a cyclic change is a purely geometric consequence of the assumption that the particle's quantum state obeys Schrödinger's equation: the phase shift is given by a closed line integral in parameter space that is independent not only of any local features determining the shape of the excursion, but also of the rate at which the excursion is done and the energy of the system.[23]

[23]The reality of this kind of phase shift has been well documented in neutron interferometry, for instance, where the spin polarisation of neutrons travelling through a magnetic field is rotated 'around a loop'. The geometric phase shift produced by the experiment contributes to the total phase difference with the other interferometer path, revealed in the diffraction pattern. (See Werner (2012) for a recent review.) But did I not just use causal language to describe all this, saying that the experimental set-up *produces* a given phase-shift to one of the neutron currents in the interferometer? Does the magnetic field

In Berry's treatment of quantum geometric phase he required that the system remains in the same eigenstate of the Hamiltonian at every instant throughout this evolution. This adiabaticity requirement has been relaxed in the subsequent work, e.g. in Aharonov and Anandan (1987), and Anandan and Stodolsky (1987). In their more general treatment of quantum geometric phase Anandan et al. also associated the phenomenon directly with the 'motion' of the quantum system, as opposed to tracking the particular Hamiltonian used to achieve this motion. Thus, the features of the geometry of motion at play can be associated with the quantum mechanical state itself, and this furthermore does not depend on the requirement of adiabaticity for the varying environment (the Hamiltonian). Furthermore, Mukunda and Simon (1993) show how the geometric phase can be decoupled from Hamiltonian dynamics altogether by construing it in even more fundamental terms as a purely *kinematic* consequence of the structure of Hilbert space, being a non-local property of curves connecting quantum state vectors. This brings out even more forcefully its purely geometrical nature.

In the philosophical literature Batterman (2003) has rightly emphasised the importance of non-causal explanations from geometrical phase. According to Batterman the "most remarkable features" of the Berry phase, for example, is "the fact that topological and geometric structures of an abstract space of parameters can have observable, physical, but obviously noncausal 'effects'." (2003, 553) I have been writing very much in the same spirit, offering an analysis the non-causal character of these explanations: the geometrical structures pertaining to the system's "geometry of motion" are explanatory by virtue of providing counterfactual information about an asymmetric dependence of the explanandum on non-local geometrical features of the space of possible states of the system. The prevalence of such explanations remains to be further examined; there are various explanations that physicists characterise as 'purely kinematic' or 'purely geometrical', and many of these, I suspect, bear the hallmarks of non-causal explanations from

causes the neutron wavefunction phase to shift, after all? Yes, and no. When we describe an experimental set-up in higher-level terms, or explain its diffraction pattern outcome in terms of the regularities pertaining to the relevant variables, we can, of course, appeal to causal notions in as far as those variables are amenable to a causal reading by virtue of being intervenable. At a more fundamental level, when describing what happens when the neutrons travel through the magnetic field, the relevant explanatory variables are no longer amenable to such causal reading, however.

"geometry of motion".[24]

5 Conclusion: a worthwhile distinction

This paper has engaged with an issue of demarcation: drawing a line between causal and (some) non-causal explanations. How worthwhile is this demarcation issue? One might worry that the issue ultimately turns out to be largely terminological: "So you think it is *worth calling* that a 'cause'? Well, I do not". All the more so, perhaps, given that I have argued that non-causal explanation from "geometry of motion" only require a natural extension of already broadly applicable counterfactual framework for analysing explanations.

There are actually two distinct issues here, one concerning the value of drawing the distinction *in causal terms*, the other concerning the value of drawing the distinction *at all*. I care more profoundly about the latter issue, and I think the significance of the distinction at stake is ultimately underwritten by its value to science. It really matters to science that we can recognise and are able to conceive of different kinds of scientific explanations: some straightforwardly dynamical and causal; others geometrical, non-local, and (explanatorily) independent of the dynamics.[25] The recognition of the existence of explanation-supporting dependences of the latter ilk is a hugely significant scientific achievement, and it marks a natural 'joint' in the nature of scientific explanation. It is a joint that furthermore naturally supports philosophers' renouncement of the hegemony of causal explanation. Insisting on the hegemony, it seems to me, only works to satisfy an outdated philosophical prejudice for the price of weakening our intuitive grasp on causation and causal explanation.

At the same time, the significance of the distinction at hand should not be taken signify the need for a radically different account of explanation. Rather, a counterfactual framework along the lines of Woodward

[24]For example, I am inclined to think (*pace* Skow, 2013) that various explanations turning on the Pauli exclusion principle, fit this mould, although I will not argue for this here.

[25]Often such explanations furthermore contribute in tandem to a complete explanation of a particular effect, as neither type of explanation fully account for the explanandum on its own. Consider, for example, a Foucault's pendulum that is also subject to dynamical torsion, so that there is both a causal and a non-local geometrical contribution to the end result. See also footnote 4.

(2003) still functions well as a philosophical framework for analysing non-causal explanations from “geometry of motion”, for these explanations also turn on informative functional regularities that are invariant under changes in clearly identifiable explanans variables that lead to answers to “what-if-things-had-been-different” questions. It is only the robustly non-local geometrical nature of the explanans variables that demarcate these explanations from otherwise profoundly similar causal explanations.