

The Principle of Least Action and Teleological Explanation in Physics*

David Glick †

Abstract

The Principle of Least Action (PLA) has often been cited as a counterexample to the dominant mode of causal explanation in physics. In particular, PLA seems to involve an appeal to *final causes* or some other teleological ideology. However, Ben-Menahem (2018) argues that such implications no longer apply given that PLA can be recovered as limiting case from quantum theory. In this paper, I argue that the metaphysical implications of PLA-based explanations are not undermined by its status as derivative. However, I contend that PLA functions as a *diachronic constraint* that licenses explanations by constraint (Lange, 2016). PLA-based explanations, on this account, are non-causal but also differ from familiar cases of teleological explanations in several respects.

1 Introduction

The Principle of Least Action (PLA) presents a potential counterexample to the dominant mode of causal explanation in physics. The principle, which explains a range of phenomena in the classical domain, constrains the evolution of a physical system to the path through configuration space for which the action takes an extremal value. This seems to suggest that how a system evolves is constrained by features of its total evolution from an initial state to a final state that has yet to occur. This mode of explanation hearkens back to Aristotelian *final causes* and related notions of teleology considered out of place in contemporary physics. However, in a wide range of contexts the PLA is equivalent to

*Forthcoming in *Synthese*.

†Department of Philosophy, University of California, Davis, daglick@ucdavis.edu.

the ordinary dynamical laws of classical physics and, moreover, can be shown to be a limiting case of quantum theory. This allows for an understanding of PLA as a consequence of these dynamical laws rather than an independent explanatory principle. Such an attitude avoids the implication of teleology while preserving the role of PLA as a useful heuristic in classical physics.

My aim in this paper is to argue that, notwithstanding these correspondence results, PLA may plausibly be regarded as an independent explanatory principle.¹ Explanations based in PLA are legitimate and autonomous from the underlying dynamical laws. This means that there remains a role for variational principles in physics. Where this leaves teleology is less clear. Variational principles typically appeal to *global* features of a system, which may include future facts. This does not require that systems evolve toward some *goal state* but rather that their total evolution must adhere to certain constraints. So, while the legitimacy of PLA-based explanation does put pressure on the dominant mode of causal explanation, it doesn't require returning to Aristotelean notions of final cause.² In sum, PLA provides an example of a diachronic non-causal explanation. This is a feature shared with other *meta-laws* in physics, such as conservation laws. Lange (2016) argues that conservation laws can be understood as *constraints* that license non-causal *explanations by constraint*. Below, I extend this analysis to PLA.

The paper will proceed as follows. In section 2, I discuss Fermat's principle of least time, which provides a simple example of a variational principle in physics. Next, in section 3, I turn to PLA, which may be seen as a more

¹Two sets of laws may be empirically equivalent without being theoretically equivalent. The claim being advanced here is that PLA is a more general principle that agrees with (the dynamical formulation of) Newtonian mechanics in many contexts. I use the term "correspondence" to reflect agreement between laws that is less than full equivalence—for example, agreement in certain contexts, or when taking the limit. One important theoretical difference between PLA and Newton's laws concerns the greater generality of the former. See section 3 below.

²While this is perhaps the most familiar the conception, it is an oversimplification to limit teleology to the pursuit of a goal state. I address this point in more detail in section 5.1 below.

general principle of which Fermat’s principle is an illustration. Both of these principles admit of alternative interpretations that differ on the question of whether they function as constraints on the dynamical laws, or are merely consequences of them. Section 4 focuses on Feynman’s (1963) recovery of PLA as a limiting case from his path-integral formulation of quantum theory, which Ben-Menahem (2018) alleges dispels any teleological implications of PLA.³ In section 5, I take up Ben-Menahem’s argument more directly, and find that it relies on problematic assumptions about Feynman’s path integral formulation and the status of derivative principles in physics. Finally, in section 6, I conclude that PLA grounds legitimate non-causal explanations in physics, but they are not necessarily teleological explanations.

2 Fermat’s Principle

Before turning to PLA, it’s worth considering a simpler case: Fermat’s principle of least time. In 1662, Fermat demonstrated that light travels along paths that take the least time. Fermat’s principle allows for the derivation of Snell’s law, describing the angle of refraction of light as it travels through different media.

As an illustration, consider the path of light from through a region of air and one of water as depicted in figure 1. Because light travels slower in water than air, the shortest-time path from origin (A) to destination (B) may differ from a straight line. To use an analogy deployed by Feynman (1985, 51–52), a lifeguard trying to get to a victim in the ocean as quickly as possible should run along the beach for a while before getting into the water, because running is faster than swimming. Similarly, the shortest-time path will be for light to

³There is some reason to think this may have been Feynman’s view as well. On the other hand, in his lecture on PLA, he also uses teleological language in explaining the quantum behavior of electrons (the same “smelling” language used in discussing light). I focus on Ben-Menahem’s argument here because it provides a more clear and explicit defence of the claim that quantum theory provides a non-teleological explanation of PLA.

remain in the faster medium longer to minimize time in the slower medium.




Figure 1: Snell's law from Fermat's principle (image: Wikipedia)

From Fermat's Principle, we can derive *Snell's Law*, according to which when light is refracted, the ratio between the sines of the angles of incidence and refraction is a constant that depends only on the respective natures of the two media:

$$\frac{\sin \alpha}{\sin \beta} = \frac{v_1}{v_2} = \frac{n_2}{n_1}, \quad (1)$$

where each α and β are the angles measured from a line normal to the boundary, v_i is the velocity of light in the respective medium, and n_i is the refractive index of the respective medium.

2.1 Alternative explanations

Intuitively, Fermat's principle provides a compelling explanation of Snell's law. From the fact that light travels along the path of least time, we can derive Snell's law. However, this offends against the mode of causal-mechanistic explanation favored by many at the time. For example, Descartes produced a derivation of Snell's law that appealed to an analogy with colliding balls in his 1637 essay *Dioptrique*. Others, notably Leibniz, argued that Fermat's principle provided a more satisfying explanation:

... the demonstration Descartes attempted to give of [Snell's] theorem by way of efficient causes is not nearly as good. At least there is room for suspicion that he would never have found the law in this way, if he had learned nothing in Holland of Snell's discovery. Leibniz (1686, p.55)

Another alternative explanation is provided by Huygens's wave theory of light. In 1678, Huygens proposed that every point that a luminous disturbance meets turns into a *source* of a spherical wave. The sum of the secondary waves, which result from the disturbance, determines the path of the light in the new medium. If we suppose that light travels slower in the secondary medium, this allows for a derivation of Fermat's principle as depicted in figure 2 below.




Figure 2: Huygens' construction (image: Wikipedia)

Thus, the situation with respect to Fermat's principle allows for alternative interpretations. One may regard the principle as a consequence of the ordinary dynamical laws involving light—either understood as a particle (Descartes) or wave (Huygens). Alternatively, Fermat's principle could be regarded as a more fundamental explanation of the phenomena those theories must reproduce (following Leibniz). The chief problem with the latter view, is that it's hard to imagine how Fermat's principle could operate as a metaphysical principle. Feynman highlights the mystery as follows:

The idea of causality, that it goes from one point to another, and

another, and so on, is easy to understand. But the principle of least time is a completely different philosophical principle about the way nature works. Instead of saying it is a causal thing, that when we do one thing, some thing else happens, and so on, it says this: we set up the situation, and light decides which is the shortest time, or the extreme one, and chooses that path. But what does it do, how does it find out? Does it smell the nearby paths, and check them against each other? Feynman (1963, Vol. I, § 26-5)

Fermat's principle is an example of a *variational principle*—it asks us to consider varying the path, check the time required to traverse it, then select the path with the least time. While variational principles certainly have heuristic and pedagogical value, it's hard to see how they could be principles of nature itself. As Feynman highlights, the subjects of these principles—light, in this case—aren't capable of considering the relevant variations, which concern their own future evolution. Below, I will argue that variational principles can be understood as diachronic *constraints*, in the sense of Lange (2016), but first, let's consider the most famous variational principle in physics: the principle of least action.

3 The Principle of Least Action

Roughly speaking, the principle of least action (PLA) may be thought of as a more general version of Fermat's principle.⁴ Rather than considering the path of light through physical space, we consider the path of a physical system's state through *configuration space* and we seek to find an extremal value of a particular

⁴While Fermat's principle helps to illustrate PLA, the latter is not simply a generalization of the former. Rather, PLA was the result of figures such as Bernoulli, Euler, and Maupertuis applying an argument from optics to mechanics. Some of these figures (e.g., Euler and Maupertuis) had a more general understanding of PLA, but it was not a generalization in the usual sense. Thanks to Michael Stöltzner for this point.

feature of that path. In particular, in its modern formulation (due to Hamilton), PLA asserts that a system moving under the influence of conservative forces takes the path for which the *action* has an extremal value.

When first proposed by Maupertuis in 1744, the action was thought to be closely associated with the kinetic energy of the system. The modern formulation takes the action (S) to be the difference between the kinetic and potential energies ($T - V$), in other words, the Lagrangian L , integrated over time:

$$S = \int_{t_1}^{t_2} (T - V) dt = \int_{t_1}^{t_2} L dt. \quad (2)$$

PLA asserts that for a small variation of the path, the first-order change in the action vanishes:

$$\delta S = \delta \int_{t_1}^{t_2} L dt = 0. \quad (3)$$

To get a feel for the PLA, consider a simple application (Feynman, 1963, Vol. II, § 19-1). Consider a system that moves in one spatial dimension. The system begins at position x_0 at time 0 and ends at x_f at time f , so we can ask which path between these points minimizes the action ($T - V$). First imagine there are no forces acting on the system. Then the potential energy $V = 0$, so we should just minimize the kinetic energy T . This means having the system move at a constant velocity—anything else would require putting more energy into the system than needed to get it from x_0 to x_f . Next, suppose the system is in a uniform gravitational field and consider its motion in the vertical direction (opposite to the field orientation). If we suppose that the system travels from a lower point y_0 up to a higher point y_f , what is the path of least action? If the object travels higher than y_f it will be subject to greater potential energy V , which decreases the value of the action, but this must be balanced against the

kinetic energy T required to reach that height. The path which strikes the ideal balance of these components—the path of *least action*—is the familiar parabola shape of projectile trajectories in classical mechanics.

Indeed, it can be shown that, under a wide range of conditions, PLA is *equivalent* to Newton’s laws of motion. This leads to a similar situation as the one we encountered with Fermat’s principle. We have alternative explanations available for the same physical phenomena. One may appeal to PLA or, alternatively, one may regard the dynamical laws—in this case, Newton’s laws—as the ultimate source of the explanation of a system’s evolution over time.

In other words, the laws of Newton could be stated not in the form $F = ma$ but in the form: the average kinetic energy less the average potential energy is as little as possible for the path of an object going from one point to another. (Feynman, 1963, §19-1)

This equivalence (or better, *correspondence*) allows for a deflationary attitude toward PLA and its potential implications for the nature of physical explanations. For instance, Lagrange, whose seminal 1788 work *Mécanique Analytique* put PLA at the center of mechanics, said the following of the principle:

... I view [PLA] not as a metaphysical principle but as a simple and general result of the laws of mechanics. Lagrange (1811, p.183)

As a contemporary example, consider Yourgrau and Mandelstam (1968), who note that we should expect that, in general, physical laws can be stated in either a differential or integral form (e.g., $F = ma = m \frac{d^2x(t)}{dt^2}$ or PLA). This leads them to claim that, “whether the differential or integral formulation is employed depends on convenience alone” (Yourgrau and Mandelstam, 1968, p.176) and therefore:

On account of these critical observations, it is not too much to say that, if certain scientists are constrained to indulge in metaphysical daydreams about the action principle, their meditations, though perhaps interesting, are nevertheless scientifically unintelligible. (Yourgrau and Mandelstam, 1968, p.176)

Such deflationary attitudes toward PLA are not uncommon, but there are also those like Planck who view the principle as more than a mere accidental consequence of the underlying dynamical laws.⁵ And, while some may take correspondence results to suggest that there is no physically-significant difference between the two kinds of physical laws, this also may be resisted. For those concerned with the metaphysical implications of physics, it is often necessary to distinguish between what working physicists may regard as “equivalent formulations.”⁶

4 Feynman’s Path-Integral Formalism

Historically, variational principles have been viewed both as laws of nature and as mere consequences of the true (dynamical) laws of nature. Some have proposed alternative explanations of target phenomena that bypass variational principles, while others have taken the explanations they provide to be superior. If we suppose (for now) that variational principles engender some sort of teleology, then taking them seriously would seem to allow teleology to play a role in physics. Some have taken developments in contemporary physics to foreclose

⁵“Among the more or less general laws that mark the achievements of physical science in the development of the last centuries, the principle of least action today is probably that which by form and content may claim to come closest to the ideal goal of theoretical research” (Planck, 1934). Planck distinguished the general form of PLA from the specific instances of it that result from the specification of a particular Lagrangian. This understanding of PLA as a general scheme is the one that Planck has in mind in this passage. For more on Planck’s views on PLA, see Stöltzner (2003, pp. 292–296).

⁶For example, North (2021) argues that among different formulations of physical theories, those attributing to reality the least structure are metaphysically privileged.

this possibility. In particular, Ben-Menahem (2018) argues that quantum theory provides a setting in which, “for the first time. . . the principle of least action is given an explanation that completely defeats the teleological interpretation that had accompanied it for more than two centuries.”⁷

Ben-Menahem’s argument centers on the approach to quantum theory provided by Feynman’s Path-Integral Formalism (FPI). The core idea of FPI is to sum over *all the possible paths* of a system through configuration space from an initial state A to a final state B (see figure 3). Associated with each path is a probability amplitude, and these are added together. The resultant amplitude allows us to determine the probability of the state B . FPI takes the probability amplitude of a given path to be represented by a number proportional to $e^{iS/\hbar}$, where S is the *action* and \hbar is Planck’s constant. If we associate the phase of a path with a complex number, then $\frac{S}{\hbar}$ is the phase angle associated with a given path. Now, in the classical regime, the action will be much larger than Planck’s constant ($S \gg \hbar$), so the phase angle will be dominated by the action. Those paths whose phase angles differ significantly will cancel out (deconstructively interfere) leaving only the paths where the action is relatively stable. In the classical limit where $\hbar \rightarrow 0$, there will be only one path with non-zero amplitude: the path where the action is stationary to first approximation ($\delta S = 0$). Thus, we recover PLA as a limiting case from FPI.⁸

⁷Stöltzner (1994) also claims that a teleological interpretation of PLA is no longer viable in the context of Feynman’s path integral formalism: “Looking at the path integral, the interpretation that a particle moves goal-directed, in the sense that it behaves as if it acted intentionally, breaks down” (Stöltzner, 1994, p.53). Ultimately, Stöltzner’s (1994) interpretation of PLA as revealing *global structure* is in line with the interpretation of PLA I defend below, but we disagree about the bearing of quantum theory on the issue (at least as expressed in the quoted passage).

⁸See Feynman (1963, §19-1). This derivation of PLA was first proposed by (Dirac, 1933, p.69).




Figure 3: Feynman's Path Integral Formalism. (image: Wikipedia)

4.1 The Case Against Teleology

So, what exactly is the argument that contemporary physics—in particular, FPI—sanitizes PLA of its teleological implications? Ben-Menahem isn't explicit, but the reasoning seems to go roughly as follows:

1. We can derive PLA as a limiting case of FPI.
2. So, we should regard PLA as a consequence of FPI.
3. FPI is free from teleological implications.
4. So, PLA is free from teleological implications.

There are several issues with an argument of this kind. The first set of issues concerns recovery at the limit, or *general correspondence* as Post (1971) calls it. Even granting that a Feynman-style derivation gives us the result that $\delta S = 0$ in the limit where $\hbar \rightarrow 0$, we might worry (a) that the notion of the action is different in the context of quantum theory and/or (b) that the limit is at best an idealization. Regarding the first point, there is a worry that goes back at least to Kuhn (1962) that even if we can achieve a formally-equivalent equation via some limiting procedure, we may still not have recovered the desired scientific

law. The reason is that scientific laws, unlike mathematical equations, require *interpretation* and this is tied to the scientific theory in which the equation is embedded. In the present case, in order for the derivation to succeed, S must refer to *the action* under an interpretation that's neutral between quantum and classical theories. Otherwise, the resulting formula will be isomorphic to PLA, but S will remain a *quantum* quantity. The second point is that the limit described is physically unrealistic. Planck's constant is small but not zero ($\hbar = 1.054571817\dots \times 10^{-34} \text{J}\cdot\text{s}$), so it is an idealization to take its value to be strictly zero.⁹¹⁰ Perhaps, such an idealization is unproblematic in the classical regime, but is strictly false. This would seem to suggest that PLA itself is strictly false as well, which undercuts the aim of the derivation—namely, to show how quantum theory supports PLA.

All of this being said, correspondence results play an important role in physics and are widely regarded as explanatory of the theories and laws they concern. However, notice that equivalence is a symmetric relation; agreement of PLA and FPI at the limit fails to determine which theory explains which. Moreover, it is unsurprising that such an equivalence exists. Given that PLA is empirically adequate in the classical regime, FPI must agree with its predictions lest it be empirically disconfirmed. Now, empirical equivalence does not require theoretical equivalence, but it is hard to see how a theory could be in empirical agreement with a principle as general as PLA in the classical regime without recovering it as a limit case. So, why should we think that FPI explains PLA rather than the other way around? One potential basis for this claim is that FPI is more *fundamental* than PLA. But there are several notions of fundamentality in the literature and it's not clear which is relevant here. If it just means that we

⁹ \hbar is the *reduced* version of Planck's constant: $\hbar = \frac{h}{2\pi}$.

¹⁰It is an oversimplification to identify taking the classical limit with setting $\hbar = 0$. In a variety of contexts, more subtle methods are required to achieve correspondence with classical physics and this is the case for FPI as well. For a numerical example (and references to more comprehensive mathematical treatments), see Forgione (2020, pp. 802–803).

take FPI to be more explanatorily-basic, then it's question-begging to assume that FPI is more fundamental. FPI is a quantum theory, and quantum theories typically apply at higher energy levels (or shorter distance scales) than classical theories, so that is one sense in which FPI is more fundamental, but its relevance here is unclear. There is a traditional view according to which high-energy physics provides the most fundamental scientific principles, and everything else is (ultimately) to be explained in their terms (Weinberg, 1987), but that view has been challenged (Morrison, 2006). If a principle, such as spontaneous symmetry breaking in condensed matter physics, figures in the explanation of a wide range of phenomena, then one might argue that it is fundamental in the sense relevant to explanation even though it doesn't appear at the highest energy level. One argument used to support this attitude appeals to the autonomy of certain mid-level principles (again, symmetry breaking provides an example). When a principle is robust under a range of different underlying dynamical laws, this gives the principle a kind of autonomy that is relevant to explanation.¹¹

Even if one grants that FPI is more fundamental, and hence, the direction of explanation runs from FPI to PLA, this is insufficient to undermine the putative metaphysical implications of PLA. Few today would limit their metaphysical commitments to only fundamental principles, so even if PLA is a derivative principle, it may still reveal real dependence relations in nature. And if those dependence relations have a teleological character, then this may be seen as a part of physical reality, irrespective of the character of the more fundamental laws from which it emerges.

We might appeal to FPI's *generality*. After all, FPI is true even in cases where \hbar is non-negligible, while PLA is not. However, there are other dimensions along which PLA may be plausibly understood as more general. In particular,

¹¹Perhaps this captures Planck's understanding of PLA as a more general principle that guides physical theorizing. See section 5.2 and footnote 18 below.

FPI is an approach to quantum theory, which attributes a particular dynamical evolution to physical systems (as represented by Schrödinger's equation or its relativistic analogue) while PLA applies to *any* possible law of dynamical evolution. In other words, PLA concerns a broader notion of physical modality than FPI, or any other theory that is committed to a particular dynamics.¹²

Finally, let's turn to the third claim: that FPI is non-teleological. Recall that FPI computes the probability of an end state by considering all paths to it from the initial state. The key move is to associate an amplitude and a *phase* with each path so that their sum is a function of both their amplitudes and the phase angle between them. That is, one treats each path as a wave which allows them to *interfere* with each other. The interference of waves propagating in space is a familiar phenomenon that admits of a relatively straightforward causal-mechanistic analysis. When the ripples created by rock thrown into a pond intersect, they undergo interference which causes a change in the ripples. But the connection between paths in FPI and physical waves is much more tenuous. First, the paths traverse *configuration space* rather than ordinary three dimensional space. This immediately problematizes the idea of physical interference between paths. Unless one is willing to regard configuration space as a physical space (rather than an abstract one used to represent physical systems), talk of "interference" between paths has to be read metaphorically. Moreover, the amplitudes and phases associated with each path are not easily understood as physical waves. For his part, Feynman seems to have viewed the paths instrumentally as a means for making accurate predictions of measurement outcomes. Clearly such a view doesn't provide a causal-mechanistic account of FPI. This is not to say that there's reason to regard FPI as teleological in character, but rather, to challenge the idea it sanitizes of teleology PLA by

¹²See Butterfield (2005) for detailed discussion of the role of modality in PLA and other variational principles.

grounding it in purely causal-mechanistic story. At best, what FPI seems to offer is an explanation of PLA that is non-metaphysical, and hence, it does nothing to replace the metaphysical implications of PLA with a non-teleological alternative.¹³

In summary, the case against the teleological implication of PLA based in FPI is problematic at almost every step. Aspects of the alleged derivation of PLA from FPI may be challenged. Even if the derivation is successful, this does little more than establish that the two theories agree in the classical domain, which is to be expected on empirical grounds. It is an open question whether one should regard FPI as more fundamental than PLA in a sense that justifies the claim that the latter is a consequence of the former. Finally, the interpretation of FPI is far from clear, which makes it impossible to assess whether it condemns or vindicates teleology. Feynman's instrumentalist leanings suggest a view of FPI which fails to provide a reduction of PLA to a purely causal-mechanistic metaphysics.

5 Teleology and the PLA

My aim thus far has been to argue that PLA may be regarded as a legitimate explanatory principle, and not merely an accidental consequence of the dynamical laws. If PLA engenders a teleological metaphysics, then this would seem to imply that physics ascribes teleology to our world after all. However, the proclamations of prominent physicists notwithstanding, there are reasons to be skeptical of PLA's teleological character. In this section, I will argue that PLA

¹³To be clear, I am not contesting the claim that Feynman's derivation provides some kind of explanation of PLA. Rather, the claim is that it fails to provide a causal-mechanistic underpinning for the principle. FPI makes predictions that agree with PLA (at the classical limit), but these predictions are not based in a causal-mechanistic story, despite the suggestiveness of Feynman diagrams and the analogy with the interference of (classical) light waves. In other words, the explanation provided is best understood as *epistemic* rather than *ontic*, but the latter is what's required to dispel the threat of teleology.

is best understood as a diachronic constraint.

5.1 Teleology

Teleology involves the invocation of final causes, or ends, in the explanation of some target explanandum. Standard examples come from organisms capable of intentional states. I might take a drink of water *because* I'm thirsty and have the goal of not being thirsty. The application of teleology outside intentional contexts is much more controversial. For his part, Aristotle finds final causes for natural phenomena that are non-intentional, and indeed, non-biological. Famously, Aristotle's explanation for the falling of massive bodies centers on the tendency of earth to return to its natural place at the center of the universe. This is not to attribute to such bodies intentions or a mind, but simply to appeal to a goal state in the explanation of behavior in a non-ideal state. Another thread in Aristotle's teleological thinking is that only good things can serve as the final causes (goal states).

“not every thing that is last claims to be an end (*telos*), but only that which is best” (Phys. 194a 32–33, quoted in Falcon (2022))

Turning to the PLA, do we find in it an explanation based in a final cause that is best? Initially, it may seem so. The path of least action is privileged in that it minimizes a certain kind of energy term. But there is an immediate problem. The modern formulation of PLA doesn't minimize the action, rather, it requires that the action be *stationary*: $\delta S = 0$. But, this holds at the *maximum* as well as the minimum. This undercuts an understanding of PLA as a kind of efficiency of nature, in which the goal state is one that requires the least energy (action) to get from its initial state to its final.¹⁴ In this respect, PLA differs

¹⁴This seems to have been a key part of Maupertius' original conception of PLA. As he says, it is "... a principle so wise, so worthy of the Supreme Being, a principle, to which nature

from Fermat’s principle of least time. In the latter case, there is a clear sense in which that path is optimized—it allows the light to get from source to target as quickly as possible. But, given that PLA cannot be understood as getting from the source to the target using the least amount of energy, there is no obvious respect in which the reasoning can be carried over to this case. Absent some new argument to replace the “efficiency of nature,” there is no reason to regard PLA as selecting the *optimal* path.

There is another respect in which PLA diverges from familiar cases of teleological explanation. Typically, we think of the explanandum to be some activity done in the service of arriving at a final goal state. In applications of PLA, however, both boundary conditions enter as premises. Consider the simpler case of Fermat’s principle. It’s not as though the light has some target it’s trying to hit as it passes through different media. Instead, we are asking why the path of the light from *A* to *B* is the path it is. Saying the light had to take this path to arrive at *B* suggests there is something special about this location, which needn’t be the case. In other words, the goal state (*telos*) of the system is not its final state.

If PLA is to be understood teleologically, the goal state must be an extremal (stationary) value of the action in the evolution from initial to final conditions. But this is a feature of the *evolution*, not of the system that is evolving. So, even if we grant that a non-intentional system like a beam of light can have a *telos*, here its *telos* is not a later state, but rather, a feature of its evolution across time. At the very least, this differs from the most familiar kind of *telos*—it isn’t a goal in the future, but a global feature of its entire timeline.

There are two possible replies to these worries.¹⁵ First, we may think of

seems to be constantly attached” (quoted in Stöltzner (1994, p.33)). Note that Maupertius’ formulation of PLA differs from the modern version discussed here, so the present point doesn’t straightforwardly apply to him.

¹⁵I thank two very helpful referees for raising these points (among others).

each step along the path as an intermediate goal state, which is “best” with respect to arriving at the final state. Second, we may employ a broader sense of teleology which isn’t committed to the pursuit of manifest goals. Sticking with the simpler case of Fermat’s principle, the first reply would suggest that every point on the path is selected to be the best way for the light beam to get to the target. However, it remains the case that each point is best only in reference to an endpoint (target) which has no inherent value. In other words, if there is nothing optimal about the endpoint, and the points along the way are only optimal in reference to the endpoint, then there is nothing inherently “best” about these points either. At most, this allow us to regard each intermediate point as optimal *given the final state*, so the sense in which the path of a beam of light is “best” is ultimately without a ground in a final goal state that is inherently good.

It may be argued that this conception of teleology is too limited—not all cases of teleology involve a tendency toward a manifest goal which is inherently best. So, we can still regard PLA (and Fermat’s principle) as a teleological principle in virtue of the fact that it requires the optimization of a feature of the entire path. While I am open to thinking of a constraint on a temporally-extended process as teleological, there are two points worth emphasizing here. First, this differs from the most familiar instances of teleology, which involve manifest goals directed at some inherently good end state. This isn’t to rule out other conceptions of teleology, but just to note that PLA would have a different character from the most familiar cases. Second, as discussed above, there is no clear basis for regarding PLA as singling out the *optimal* the path. Absent some further argument, constraining the action to be stationary does not seem to select the “best” path in any obvious sense.

I don’t take these considerations to rule out a teleological interpretation of

PLA. Rather, my aim is to raise some concerns which show that the link to teleology is not as straightforward as it may initially seem. At the very least, PLA differs from familiar cases of teleology in important respects.¹⁶ Ultimately, these considerations suggest an understanding of PLA as comparable to other principles in physics which impose constraints on the total evolution of a physical system (e.g., conservation laws). These principles are typically regarded as non-teleological, which suggests we should take the same attitude toward PLA.

5.2 Constraints

Even if PLA-based explanations aren't teleological for the reasons just provided, they still provide an example of a non-causal form of explanation in physics. PLA-based explanations may be understood as what Lange (2016) calls *explanations by constraint*. The key distinction here is between constraint and coincidence. A principle like PLA may be regarded as a coincidentally true generalization, or alternatively, as a limitation on what dynamical laws are possible. Lange regards it as an empirical matter whether a principle functions as a coincidence or a constraint, and says that PLA “has sometimes been interpreted as a constraint but more often as a coincidence” (Lange, 2016, p.68).

However, when we look at the considerations Lange appeals to when discussing conservation laws, many of those same points apply to PLA as well. In particular, Lange claims that energy conservation is *resilient*, in that it applies to a wide-range of different dynamical laws. Regarding a principle as possessing resilience indicates that it's functioning as a constraint because a mere coincidence wouldn't be expected to hold if the underlying dynamical laws were different.

So, in assessing whether PLA is a constraint or coincidence, we should ask

¹⁶For this reason, PLA and other variational principles may be helpful in developing an account of formal teleology. See Stöltzner and Weingartner (2005).

whether it's resilient: even if the laws of dynamical evolution were different, would we expect PLA to be respected? Of course, quantum theory tells us that PLA isn't strictly true, so we might say that the *actual* laws of dynamical evolution are inconsistent with PLA.¹⁷ But, if we restrict ourselves to the classical context where PLA is applicable, then there is reason to think that PLA is robust under different dynamical laws. Indeed, after he shows the equivalence of PLA and Newton's second law of motion, Feynman shows that PLA can be extended to the relativistic analog of $F = ma$ as well. Indeed, PLA's resilience in the face of the changes prompted by relativity motivates Planck to view it as a fundamental principle:

The most brilliant achievement of the principle of least action is shown by the fact that Einstein's theory of relativity, which has robbed so many theorems of their universality, has not disproved it, but has shown that it occupies the highest position among physical laws. The reason for this is that Hamilton's "action" . . . is an invariant with respect to all Lorentz transformations, that is, it is independent of the system of reference of the observers. Planck (1934, pp.77–78)¹⁸

Regardless of such cases, it remains an open possibility to regard PLA as a mere coincidence—one that just happens to follow from many dynamical laws in the classical regime—but, this leaves us with a puzzling coincidence.

¹⁷Recall that Feynman's derivation of PLA from quantum theory requires one to take the limit where $\hbar \rightarrow 0$, but this limit doesn't describe reality, so PLA is strictly inconsistent with the full quantum theory.

¹⁸Planck distinguishes between the specific version of PLA used in mechanics and the understanding of it as a more general scheme. It is the latter that is resilient under changes of the dynamical laws. There are difficult questions about how to make this notion of resilience precise and its connection to invariance in the sense discussed in relativity. The claim here is just that there is some serviceable notion of resilience according to which PLA is plausibly understood as resilient. Of course, the devil is in the details when it comes to specifying the precise form of PLA and the nature of its resilience. For some discussion of the technical difficulties involved, see Butterfield (2005).

The most important reason we should be open to thinking of PLA as a constraint is that it enters into explanations that are naturally understood as explanations by constraint. When we use PLA to explain, this differs from ordinary causal explanation: we do not appeal to earlier causes to explain a later effect. I’ve also argued that it differs from familiar cases of teleological explanation: there is nothing for which minimizing (or making stationary) the action is a final state to which it aims (and is “good” for it). So, PLA-based explanations, if they exist at all, are non-causal but also not teleological in the usual sense. Regarding PLA-based explanations as instances of explanation by constraint allows us to vindicate them as genuine explanations by taking PLA to be a genuine constraint.

Lange explicitly adopts a methodology according to which it is a matter of empirical investigation whether a generalization is a coincidence or constraint—if scientists regard the generalization as having the kind of resilience discussed above, then it should be regarded as a constraint. But, as we’ve seen, there are supporters of both views among prominent physicists.

So where does that leave us? I submit that we shouldn’t foreclose the possibility that PLA is genuinely explanatory. It is possible that those who regard PLA as a mere accidental coincidence fear that giving it any explanatory status is inconsistent with it being only true in the classical regime or would reintroduce teleology into physics. But both of these concerns have been addressed above, leaving it open to take PLA-based explanations at face value.

6 Conclusion

PLA provides an interesting case of non-causal explanation in physics. Notwithstanding Feynman’s recovery of PLA as a limiting case of quantum theory, we may regard it as a genuinely explanatory principle in the classical regime.

This means that if PLA carries with it the implication that nature is teleological, then this remains true even if it isn't *fundamental* in the sense of particle physics. However, there are reasons to question this implication of PLA-based explanations. Unlike familiar instances of appeals to final causes, PLA-based explanations do not posit a final state of a system that is “good” for it. PLA is a variational principle that explains the evolution of a system from an initial state to a final state. Moreover, the path of development it singles out isn't “best” in any obvious sense—it is only the path where the action is stationary. Ultimately, successful PLA-based explanations should be understood as explanations by *constraint*, along the same lines as those grounded in conservation laws and symmetry principles. While I agree with Lange that it is ultimately a matter for physics to decide the status of PLA, we should allow for the possibility that it functions as a constraint and provides a basis for successful non-causal explanations. This is contrary to the standard causal-mechanistic form of explanation in physics, but needn't amount to a reintroduction of teleology into physics.

Conflicts of interest

The author has no competing interests to declare that are relevant to the content of this article.

Acknowledgements

Thanks to Daniel Kodaj and the Teloi Project at the Central European University in Budapest for inviting me to give a talk on this subject. I am grateful to the audience of that talk and two anonymous referees for helpful feedback on previous versions of this paper. I am also thankful to participants of the MEN-

TaL reading group at UC Davis and Alejandro Naranjo Sandoval for discussion of this material. Finally, I am especially indebted to Michael Stöltzner for his thoughtful comments and suggestions.

References

- Ben-Menahem, Y. (2018). *Causation in Science*. Princeton: Princeton University Press.
- Butterfield, J. (2005). Some aspects of modality in analytical mechanics. In M. Stöltzner and P. Weingartner (Eds.), *Formal teleology and causality in physics*. Paderborn: Mentis.
- Dirac, P. A. (1933). The lagrangian in quantum mechanics. *Physikalische Zeitschrift der Sowjetunion* 3, 64–72.
- Falcon, A. (2022). Aristotle on Causality. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Spring 2022 ed.). Metaphysics Research Lab, Stanford University.
- Feynman, R. (2011 [1963]). *The Feynman lectures on physics*. New York: Basic books.
- Feynman, R. P. (2006 [1985]). *QED: The strange theory of light and matter*. Princeton: Princeton University Press.
- Forgione, M. (2020). Path integrals and holism. *Foundations of Physics* 50, 799–827.
- Kuhn, T. (1996 [1962]). *The Structure of Scientific Revolutions* (Third ed.). Chicago: The University of Chicago Press.
- Lagrange, J. L. (1997 [1811]). *Analytical mechanics*. Number 191 in Boston Studies in the Philosophy of Science. Dordrecht: Springer Science & Business Media.
- Lange, M. (2016). *Because Without Cause: Non-Casual Explanations In Science and Mathematics*. Oxford University Press.
- Leibniz, G. W. (1989 [1686]). Discourse on metaphysics. In R. Ariew and D. Garber (Eds.), *Philosophical Essays*. Indianapolis and Cambridge: Hackett Publishing Company.
- Morrison, M. (2006). Emergence, reduction, and theoretical principles: Rethinking fundamentalism. *Philosophy of Science* 73(5), 876–887.
- North, J. (2021). *Physics, Structure, and Reality*. Oxford University Press.

- Planck, M. (1934). *Wege zur physikalischen Erkenntnis*. Verlag von S. Hirzel.
- Post, H. (1971). Correspondence, invariance and heuristics: In praise of conservative induction. *Studies in History and Philosophy of Science Part A* 2(3), 213 – 255.
- Stöltzner, M. (1994). Action principles and teleology. In H. Atmanspacher and G. Dalenoort (Eds.), *Inside Versus Outside*, Volume 63 of *Springer Series in Synergetics*. Berlin and Heidelberg: Springer-Verlag.
- Stöltzner, M. (2003). The principle of least action as the logical empiricist's *Shibboleth*. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 34 (2), 285–318.
- Stöltzner, M. and P. Weingartner (Eds.) (2005). *Formal teleology and causality in physics*. Paderborn: Mentis.
- Weinberg, S. (1987). Newtonianism, reductionism and the art of congressional testimony. *Nature* 330(6147), 433–437.
- Yourgrau, W. and S. Mandelstam (1968). *Variational Principles in Dynamics and Quantum Theory* (Third ed.). London: Sir Isaac Pitman & Sons Ltd.