

Replies to Weatherson, Hall, and Lange

Michael Strevens

Draft of January 2012

It is a great pleasure and privilege to have such acute, insightful, and probing commentaries to wrestle with! In these responses I hope to illustrate how, in certain difficult and interesting cases, my “kairitic” account of explanation goes about its business, and in so doing to develop and clarify some of the lines of thought in *Depth*.

The question of the individuation of causal-explanatory mechanisms—the question of what I call cohesion—runs through my commentators’ critiques, and so it will run through these responses. In part I want, of course, to defend the kairitic account from objections; in this respect my aim is narrow. But I have in addition a broader goal: to emphasize that cohesion is a question for everyone taking a causal approach to scientific explanation, and perhaps everyone working on any kind of explanation at all. In that respect, it is hardly necessary to rebut my commentators, since they have done a magnificent job of identifying and articulating precisely the kind of worries that drew me through the thorns of causality and into the thickets of cohesion. Indeed, Weatherson’s beautiful essay could stand alone without any further remarks; in what follows, I hope to show that Hall’s and Lange’s contributions as well amount not only to worthy critiques of some doctrines in *Depth*, but also identify pressing questions that anyone interested in explanation should feel bound to answer.

1. Reply to Weatherson

Explanatory models in economics and in many other sciences—perhaps the great majority of high-level sciences—are radically multiply realizable, by which I mean that they lack cohesion in the kairetic sense. Weatherson’s example, the explanation of the outcomes of auctions of oil drilling rights in the Gulf of Mexico, provides a perfect illustration of this pervasive fact.

So: the kairetic account demands that explanatory models satisfy the requirement of cohesion. Many legitimate explanatory models in economics and elsewhere violate this constraint. Game over?

As Weatherson points out, *Depth* develops several lines of defense against objections such as these. Some defenses offer ways to see that the cohesion requirement is not violated after all. But these do not apply in the present case, for the reasons that Weatherson sketches. One possible defense remains: economists and others are making use of explanatory frameworks.

Let me elaborate on Weatherson’s succinct description of frameworking. Sometimes we seek not to explain a fact simpliciter, but to explain why it holds *given that* some other fact holds. For example, if I ask for an explanation of why a certain window broke, you might tell me that it was hit by a large rock, and perhaps describe the causal process that led to the rock’s arrival (a local riot or a meteor storm?). But if I am an engineer at the company that made the supposedly unbreakable glass, I might ask why the window broke *given that* it was hit by a rock. To this explanatory question, information about the cause of the rock’s hitting the window is irrelevant. That the rock hit the window is taken as given; the explanation will focus on the properties of the window in virtue of which it shattered when subjected to such an impact.

This shows that “*e* given that *f*” is a special kind of explanandum to be distinguished from *e* simpliciter, and requiring a special kind of explanation with its own rules of relevance. (Compare the sense in which the explanandum “*e* rather than *f*” demands a different explanation than does *e*.)

On the kairetic account, all explanation notionally proceeds by taking

a complete description of some chunk of causal reality and progressively abstracting until only difference-makers for the explanandum remain. In a “given that” explanation, the procedure is amended as follows. First, the element of the story that is taken to be “given” must appear in the description and may not be removed. Second, anything that causally explains that element (and that plays no other causal role in the story) must be removed. In the example above, the rock’s hitting the window is taken to be a fixed part of the explanatory model, then, but no explainers of the rock’s trajectory and impact are admitted into the model. In such explanations, I say that the “given” element is moved to the *explanatory framework*, where the term “framework” is intended to suggest something that is necessary for the explanation without itself being explanatorily active.

When the element that is frameworked is a causal mechanism, its explanation is a description of the workings of the mechanism itself. To remove the explanation of what is given from a causal model, then, is to remove from the model all mention of these workings. What is left behind is a mere black box, a functional specification of the mechanism that maps inputs to outputs while saying nothing of the structure of the internal causal paths. If you are asked, for example, to explain the crash of Air France flight 447, you might well give a detailed description of the causal processes that led to the malfunctioning of the aircraft’s Pitot tubes (resulting in faulty airspeed readings). But if you are asked to explain the crash *given that* the Pitot tubes were malfunctioning, the same description will be explanatorily irrelevant. The tubes will enter into your explanatory model only in the guise of a functional description specifying what (faulty) readings would be yielded in what circumstances, without any hint of the way in which the circumstances cause the readings.

You will see that the effect of explanatory frameworking is to focus the explanatory spotlight on some parts of the relevant causal mechanism while excluding others. Frameworking the rock in the window-breaking case shines the spotlight on the structure of the window; frameworking the malfunction-

ing Pitot tubes in the Air France crash shines the spotlight on the action in the cockpit. This manipulation of the explanatory spotlight is entirely analogous, I think—perhaps identical with—what is called causal selection in the literature on causation.

What does all of this have to do with cohesion? Simple: moving a mechanism to the explanatory framework creates explanatory models containing black boxes, and black boxes are radically multiply realizable. When frameworking is in the offing, the kairetic account allows for incohesion—provided that the incohesive part of the model, the multiply realizable part, is confined to the framework.

This, I propose, is what is happening in multiply realizable models in economics and elsewhere. For practical reasons derived principally from the need for an efficient division of cognitive labor, researchers in different areas framework different aspects of causal processes, so as to examine more minutely and under brighter light some particular part of a process. Most important, researchers at different levels of inquiry tend to framework mechanisms that are the proper object of study at the next level down. Economists framework psychological processes; psychologists framework neuroscientific processes; neuroscientists framework some chemical or physical processes (e.g., mechanisms responsible for chemical bonding); and so on. As a consequence, at every level but the lowest—everywhere but in the basement of the physics department—explanatory models teem with black boxes, and so multiple realizability is rife.

Applying this analysis to auctions of drilling rights, I will try to persuade you that the creators of Weatherson's explanatory model are frameworking Nash equilibrium-seeking behavior. Observe, on the one hand, that the mechanisms responsible for the oil companies' seeking a Nash equilibrium are clearly explanatorily relevant to the outcome of the auction since, had the companies made their bidding decisions in a way that did not seek such an equilibrium, the auctions' outcome would very likely have been different. On

the other hand, as Weatherson makes clear, the model tendered to explain the auctions' outcomes in no way explains the participants' tendency to seek a Nash equilibrium. This is precisely what the kairetic account predicts if the equilibrium-seeking mechanism or mechanisms are frameworked: though it is crucial for obtaining the observed outcome that the bidders seek a Nash equilibrium, scientists accept an explanation that makes no attempt to explain this equilibrium-seeking behavior. So I conclude that the model has a "given that" explanandum—it explains, and is intended to explain, why the auction ends the way it does given the equilibrium-seeking behavior of the bidders.

Since incohesion enters this model only in the framework, there is no violation of the kairetic criteria for a good explanation.

Consider two objections to this handling of the auction example. The first is voiced by Weatherson, who writes of the frameworking approach:

It seems too defeatist to me. Part of the appeal of game theoretic explanations is that they are supposed to explain why we get to, and stay at, equilibrium. I don't think a practicing economist would say that they are merely presupposing that players in a game reach equilibrium.

What Weatherson says in this passage about economic explanation is, I think, completely correct. But in saying it, he underestimates the power of the frameworking explanation. The aim of the frameworking explanation is to explain a certain outcome by describing the causal process that leads to the outcome. That process comprises several parts: what you might call the initial conditions, such as the bidders' knowledge of the quality of the potential drilling sites; the mechanism by which the auction is run; and the mechanisms by which the bidders make their decisions. Only the last of these is frameworked. Thus the frameworking explanation describes in substantive terms a significant portion of the causal process that brings about the outcome to be explained.

The second objection is that the frameworking maneuver is ad hoc: what real content is there to cohesion, if a violation of the requirement can invariably be explained away by positing an explanatory framework? As with any putatively ad hoc maneuver, what is needed in defense of frameworking is some independent handle on what typically goes into a framework and what does not, against which the apparent violations of cohesion can be assessed.

There are many motivations for frameworking, but the most important, foreshadowed above, is the division of cognitive labor. Frameworking is a way of dividing a causal process into pieces or strata that may be fruitfully investigated by independent groups of scientists. In the case of drilling rights, the division effected by the frameworking in this and all similar auctions is into, on the one hand, the structure of the auction, and on the other hand, the strategies of the individual actors (executives, boards of directors, companies) given that structure. If this is an especially fruitful division—if the tools for studying auctions, such as game theory, are sufficiently different from the tools for studying boardroom cognition, such as embedded observers and cognitive psychology—then frameworking along these lines is to be expected on independent grounds, and so my treatment of Weatherson’s case cannot be considered objectionably ad hoc.

All of this must of course be regarded as rather provisional: there is far more to be said about the structure of game-theoretic explanations in economics, about the division of explanatory labor, and about the uses of frameworking, in particular. But I hope that it gives the reader some sense of my reasons for thinking that explanatory practice is governed by a cohesion requirement that might appear to be violated as a matter of course.

2. Reply to Hall

Hall asks whether the kairetic account can handle a simple account of pre-emption; more exactly, he asks whether the account has what it takes to reject

a certain model that, if counted as genuinely explanatory, would qualify a mere backup cause (event *a* in his neuron diagram) as an actual cause. He concludes, correctly, that I would reject this model on the grounds that it is not cohesive, that is, that it fails to describe a single causal mechanism but rather has among its instantiations two distinct mechanisms.

To motivate this treatment, let me try to convince you, invoking only everyday explanatory intuitions and making no appeal to the kairetic account, that his model—his purported counterexample—fails as a causal explanation for precisely this reason, that is, because it is not a single explanation but a compilation of possible explanations.

Put some flesh on Hall's neuron diagram: suppose that the event to be explained (the event of neuron *e*'s firing) is the destruction of a bottle of rare single malt whisky. In the diagram, there are two possible causes of this effect, the actual cause *c* and the backup cause *a*. Suppose that *c* represents a rogue meteor heading toward the bottle, while *a* represents a rioter hurling a brick that will hit and break the bottle if no piece of space debris does the job first. In Hall's figure 1 the meteor appears and the rioter's projectile arrives only after the bottle's destruction; in figure 2 no meteor happens along and it is the rioter who breaks the bottle.¹

Suppose, with Hall, that the meteor does the job. Then we do not want to count the rioter's brick-hurling as explanatorily relevant to the breaking of the bottle. As Hall notes, the rioter will count as relevant if there is a cohesive, veridical causal model for the breaking from which the rioter's hurling cannot be removed (in the kairetic account's proprietary sense of removal). He suggests that there is such a model, namely, the model that describes the situation along with the relevant laws or generalizations, that specifies that the rioter throws their brick, but that does not specify whether or not the meteor appears.

1. Strictly speaking, this is a case of late preemption, not the early preemption shown in Hall's neuron diagrams. But it will make no difference.

I claim that the model is not cohesive. It does not represent a single causal mechanism, a single way that the bottle might be broken, but rather two distinct mechanisms and thus two distinct explanations, one in which the bottle is pulverized by the meteor and the other in which it is shattered by the brick. It says: either a meteor broke the bottle or a rioter's brick did. That is not a single explanation, however; it is a disjunction of possible explanations including both the actual explanation and another, irrelevant explanation. For this reason, the model should be rejected by any right-thinking explainer, whether or not they endorse the kairetic account.

I conclude that there is no problem with the kairetic account's handling of cases of preemption, since it appeals to a criterion for explanatory goodness for whose existence there is much independent evidence and that clearly judges Hall's case as the kairetic account wants and needs it to.

But perhaps—as Hall suggests in his closing remarks—there is something wrong with the kairetic account's treatment of cohesion? Perhaps the account's criterion for individuating causal mechanisms is flawed? That would be a problem quite independent of any worries about preemption.

Hall suggests that the correct criterion for individuating causal mechanisms is one that captures “a cognitively effective means for organizing our information about the world”. That sounds sensible to me. The question is: which of many possible schemes for efficiently organizing causal information do the mechanism individuation criteria employed by the norms of explanation implement? And why those ones? (Technically that is two questions, but I have no plans to answer the second in this short piece.)

Hall suggests that the taxonomy of mechanisms relevant to explanation will be based on a similarity relation. He also hints that the taxonomic principles will differ from domain to domain. My individuation principle, the cohesion requirement, has neither property. Let me use the rest of this reply to explain why.

The cohesion requirement turns on contiguity rather than similarity. The

realizers of a causal model are contiguous if, roughly speaking, you can get from any realizer to any other realizer by a process of minimal tweaks. That is, there should be a path through “realizer space” (a space of causal processes) connecting any two members of the set of realizers.

This differs from a similarity criterion for mechanism individuation in three ways. First, it requires a weaker structure on the space of possible causal processes: it requires a neighborhood or topological structure (very loosely, for any process there should be a set of processes minimally different from that process), whereas similarity requires a metric structure (for any two processes, there must be a degree to which they resemble one another). The former structure is far easier to come by: it seems very hard to imagine a precise quantification of the degree of difference between any two processes, but not nearly so hard to imagine, for any process, which other possible processes would count as “almost the very same process”.

Second, to implement a similarity scheme for individuation, it is not enough to supply a similarity metric. You must also supply a function determining, for any given explanandum, exactly how similar the realizers in the explanatory model must be to one another. There is no need of such a thing on the contiguity approach: contiguity is all or nothing.²

Third, a contiguity criterion is in one important respect much weaker than a similarity criterion: a contiguous set may contain realizers that are radically dissimilar.

Let me give two arguments for preferring a contiguity to a similarity implementation of mechanism individuation. One of these arguments turns on the remarks made above, and is only a half-sentence long: a contiguity criterion for individuation is much easier to come by, much simpler to implement, than a similarity criterion.

To set up the other argument, return to the bottle and its breaking. Suppose

2. That said, I play around with the idea of “nearly contiguous sets” in section 5.43 of *Depth*.

this time that there is no meteor, but only Rebecca the Rioter and her identical twin sister Ruby throwing bricks at the bottle. A causal model that contains this information and the relevant laws, and lacks only the fact about who threw first that determines whether it was Rebecca's or Ruby's brick that broke the bottle, has the same defect as the model above: it is not a single causal explanation but two explanations disjoined, the actual explanation and a possible but non-actual explanation alongside it. (To add to the inconvenience, the actual explanation is not labeled as such.) Call this the *bad model*.

Now consider a different model. This one models Ruby's breaking the bottle but does not specify the position from which she launched her brick. Rather, it specifies certain facts about the relative positions of Ruby and the bottle, together with the relative velocity of the brick. These determine a direct hit and so a shattering, but without pinning down a specific trajectory. This, I claim in *Depth*, is a better explanation of the breaking than is provided by a model that does specify a precise trajectory, since the precise trajectory does not make a difference to the bottle's breaking—what matters is only that the trajectory, the thrower, and the target were related so (see Lange's explanation of the case of the cannonball's exact mass and velocity not making a difference). Call this the *good model*.³

How is a similarity theorist to make sense of the good model's goodness while acknowledging the bad model's badness? The most dissimilar realizations of the bad model are no more dissimilar than the most dissimilar realizations of the good model. (In both cases, the principal difference between the realizations is the angle, in absolute terms, from which the fateful brick is thrown.)

It seems to me that the similarity theorist will have to borrow the contiguity theorist's central idea. The problem with the bad model is not that it contains

3. The bad model describes only cases where preemption of some sort is guaranteed. The good model describes some cases like this, and others as well. This is unimportant. It is facts about cohesion that settle facts about preemption, not vice versa.

dissimilar realizers; it is that it has a hole in the middle—its realizers are not contiguous. This reveals the true nature of the similarity theory: it is not an alternative to the contiguity theory, but rather an augmentation of that theory. It takes the contiguity criterion, and adds an additional criterion of similarity: realizers must not be too dissimilar. If there is a good reason to impose this additional criterion, I don't know it. So I suggest that we stick with what the contiguity and similarity theorists agree on: the contiguity criterion. Let us see if we can make an adequate criterion for mechanism individuation out of that.

On to the other way in which Hall's proposal differs from my own: the domain-specificity of individuation. I envisage a single space of causal processes, corresponding more or less to the level at which those processes are realized in fundamental physics, contiguity with respect to which determines cohesion in every explanatory domain, from physics itself through economics (even if, as suggested in my reply to Weatherson, cohesion has a limited role to play in the high-level sciences). Hall suggests that the relevant notions of contiguity and similarity may change from domain to domain. Or if he does not mean to suggest it, I will suggest it, as it is a reasonable and plausible alternative to my own view.

Why go with a unitary, fundamentalist conception of contiguity? Let me begin by putting to rest a bad argument against fundamentalism. To adhere to an explanatory norm that invokes facts about the fundamental physical level, it is not necessary to have a thorough theoretical knowledge of fundamental physics, any more than to be a law-abiding citizen, it is necessary to have a thorough knowledge of the statute book. If the norm is to have any bite, it had better be possible to follow it loosely without such knowledge, but I think that in the case of physical contiguity, this requirement is easily met. You do not need to know quantum physics, or even its Newtonian approximation, to see what counts as a minor physical variation on Ruby's rock-throwing or a drilling-rights (or any other kind of) auction. Without expertise you may

make errors in applying the explanatory norms, but we can hardly require these or any other norms to insulate us from that possibility.

What can be said in favor of fundamentalism, then? Two things. First, it offers the prospect of a relatively simple, relatively observer-independent, universal individuation scheme, one that is ready and waiting to guide our explanatory judgments whenever we encounter a new realm of facts. Perhaps it is this rationale that moves Hall to worry that I am “hopelessly optimistic”. But surely it is about time that something turned out to be okay?

A second reason for favoring fundamentalist contiguity is as follows. Suppose we have settled on contiguity; the key notion underlying judgments of cohesion, then, is that of a tweak, a small change in the unfolding of a causal process. If a domain-specific view of contiguity is correct, then we ought to have different conceptions of a tweak, or at least different conceptions of the appropriate level of grain for a tweak, in every domain. But as far as I can see, in any domain, there is no tweak so small that it is inappropriate in type or level of grain. When I ask “Do these processes make up a cohesive whole?”, I am quite happy to trace their continuity through minute variations in physical parameters. This suggests to me that it is the fundamental level that ultimately adjudicates such matters in every case.

Neither of these considerations is anywhere close to decisive, needless to say. As I wrote in *Depth*, the cohesion criterion for mechanism individuation is offered in a spirit that is more exploratory than apodictic. Its best test is use. Are there kinds of processes that strike us as unified causal processes, but whose set of realizers is discontinuous? (And for which the “given that” interpretation is inapplicable?) If so, we may need domain-specific bridges to bring together these causal islands into an explanatory whole. But I will hold on to fundamentalist contiguity until I see the gaps, and a theory of bridges capable of spanning them.

3. Reply to Lange

Explanatory Abstracta The kairetic account frequently attributes difference-making power not so much to individual variables—mass, velocity, charge—as to relations among those variables. What makes a difference to a brick’s hitting a window is the relation between the thrower, the angle of the throw, and the window, rather than the absolute values of any of these quantities.

Lange’s first challenge is to make a distinction between those abstracta—those relational or functional facts about quantities—that seem genuinely to have their own explanatory power, and those that are merely stepping stones from what has real explanatory power, the individual variables themselves, to the explanandum. Citing the explanation proffered by the kairetic account for the existence of a certain force between two charged bodies, Lange writes that “the force’s magnitude is not explained by the value of q_1q_2/r^2 . Rather, it is explained by the values of q_1 , q_2 , and r ”. By contrast, a static equilibrium (for example) might be explained by certain objects’ centers of mass rather than by the positions and masses of the particles that determine the centers of mass.

One way to read the kairetic theory is as giving the wrong account of this distinction between genuinely explanatory abstracta and abstracta that are merely middle terms in an explanatory derivation—wrong because it places both center of mass and the non-explanatory abstractum q_1q_2/r^2 on the same side of the divide, attributing the same explanatory status to both. It is better, however, to see the kairetic theory as giving no account of the distinction at all, that is, of failing to find a meaningful explanatory distinction where Lange sees one.

Let me explain. Though the kairetic account cites the abstractum q_1q_2/r^2 as a difference-maker, it does not thereby fail to do justice to the causal roles of the individual quantities, that is, to the causal roles of the two charges and the distance between them, for the following reason. It is not enough, on the kairetic approach, merely to cite the abstractum; its difference-making power

must be derived from the relevant aspects of the fundamental laws. In the case at hand, the derivation will attribute the abstractum's explanatory significance to what the fundamental laws have to say about charge and distance. The individual quantities will therefore be recognized as explainers (though their particular values will not be specified). Of course, the same is true for an explanation in terms of centers of mass. A center of mass will be cited as a difference-maker, but in demonstrating its difference-making power, the explanation will attribute it to the positions and masses of the individual particles (without specifying those positions and masses).

More generally, a kairetic explanation, however much it abstracts, will invariably derive the difference-making power of its abstracta from their fundamental-level parts. The abstracta and the parts are each given their due, then, but for that reason, Lange's distinction between abstracta that are explanatory in their own right and those that are mere middle terms is not made.

Must I hold, then, that the distinction is illusory? Not necessarily. An account of the distinction might be given in terms of the division of explanatory labor discussed in my reply to Weatherson: the division of labor motivates the adoption of certain default explanatory frameworks in a discipline or sub-discipline, and relative to those frameworks certain explanatory abstracta—such as centers of mass, Mendelian genes, or representative agents (in economics)—are so persistently useful that they become a part of the discipline's semi-official explanatory inventory. Then, insofar as each discipline has an explanatory practice that is legitimate “in its own right” (that is, relative to its assigned frameworks), the components of the corresponding explanatory inventory also come to be seen as explanatory in their own right. To attribute Lange's distinction to practical interests is more a matter of expedience than philosophical insight, I readily admit; I find the practical explanation to be plausible in the present case, but I look forward to further reflecting on the matter.

Conservation and Meta-Laws Lange's second challenge: the cohesion criterion's individuation of explanatory models is too fine-grained when it comes to explaining conservation properties. Electrostatic interactions conserve momentum because of the translational symmetry of the electrostatic force law. This explanation attributes momentum conservation to a property of the relevant fundamental laws that is radically multiply realizable: it is a property that may be possessed by laws governing a wide range of utterly different causal processes. But the kairetic account cannot countenance such a property as explanatory, because like all multiply realizable properties it violates the cohesion requirement.

In reply, I might argue that the laws having translational symmetry, though they are myriad and heterogeneous, form a contiguous set. But rather than putting so much weight on the cohesion requirement, let me explore a different defense that I think casts considerable light on the structure of the explanation of conservation properties.

A symmetry-based explanation of momentum conservation under electrostatic interaction has two parts:

1. A demonstration that interactions under any force law with translational symmetry will conserve momentum, and
2. A demonstration that the law or laws governing electrostatic interactions have translational symmetry.

It is clear, I hope, that only the first part is shared with the explanation of momentum conservation under other laws, and only the first part makes no mention of electrostatic interactions. To the extent that there is, as Lange claims, a single explanation of momentum conservation that has no specific connection to electrostatic interactions, it is at most a partial explanation, or better, what I call in *Depth* and in my reply to Weatherson a "given that" or "frameworking" explanation. More specifically, it is an explanation of an interaction's momentum conservation *given that* the interaction law has

translational symmetry. As such, as also set out in the Weatherson reply, the explanation is excused from conforming to the demands of the cohesion requirement.

In short, there is, as Lange claims, an explanation of momentum conservation in electrostatic interactions that abstracts away entirely from electrostatics, that is identical to the explanation of momentum conservation in a large class of other interactions, and that is (perhaps) incohesive. But it is the sort of frameworking explanation in which the kairetic account tolerates incohesion. Further, the complete, non-frameworking explanation of electrostatic momentum conservation is cohesive but does not abstract away from electrostatics.

What if distinct, momentum-conserving laws have the same mathematical form (think Coulomb's law and Newtonian gravitation, bracketing the direction of the force)? Will they not have the same part (2) as well as part (1), and so the same incohesive complete explanation? I don't think so. They may share their mathematical form, but the basis of that form is a different aspect of nomological reality, and so explanations of facts that follow from the form will themselves be grounded in different aspects of nomological reality. This is a sufficient reason, I suggest, to consider the explanations, despite their formal similarity, to be themselves distinct, as cohesion requires.

Lange concludes this section of his discussion by suggesting that principles of symmetry constitute "meta-laws" that constrain, and so explain the symmetries of, the first-order dynamical laws. Assuming he is right (his own work contains enticing arguments for this view), should the kairetic account capture this sort of explanation? That depends, I think, on the nature of the dependence relation between the dynamical laws and the symmetry principles. If it is a relation of causal dependence—if the dynamical laws themselves come to hold by way of a causal process, perhaps in the gods' nomological foundry or perhaps in the early evolution of a universe initially pregnant with many nomological possibilities—then the answer is affirmative, and I would hope

that the kairetic account would rise to the occasion. But if the dependence is more metaphysical than causal, as I think is true on Lange's (2009) account, then the job belongs to an account of metaphysical explanation, which the kairetic account, in spite of my imperialist inclinations, is not.

Dimensional Explanation Dimensional explanation is a fascinating topic. *Depth* does it no justice whatsoever—Lange's third challenge. Let me conclude with a brief sketch of the way in which the kairetic account might understand the example developed in Lange's critique.

Lange poses two challenges. First, dimensional explanations do not appear to be causal. Second, they do not appear to be cohesive. Let me deal with these in turn.

Dimensional explanation's non-causal aspect is shared by many varieties of explanation in which the explananda are regularities, generalizations, or laws. Such explanations typically have the form of a derivation from more fundamental laws, something that looks quite different from the description of a causal mechanism.

On the kairetic account, these derivations are to be understood as exhibiting the causal mechanism (or mechanisms) in virtue of which the explanandum obtains. To explain the correlation between drinking red wine and cardiac health, for example, is to exhibit the mechanism by which drinking red wine improves cardiac health (or prevents cardiac disease). The operation of the mechanism will be governed by more fundamental laws or regularities; to exhibit the relevant causal mechanism is in part to show how these more fundamental entities play their governing role. That is what is accomplished by an explanatory derivation from the laws.

Adapting *Depth's* story about regularity explanation to Lange's example: To explain a law or regularity is, as I have said, to exhibit the causal mechanism in virtue of which the explanandum obtains. To explain some feature of the law is to exhibit those facts about the explanatory causal mechanism in virtue

of which the feature exists. Thus, to explain why a planet's orbital period is proportional to the $3/2$ -power of its orbital radius you should, first, find the causal mechanism responsible for the orbit, and second, exhibit whatever features of the mechanism explain the $3/2$ proportionality. The dimensional explanation, I think, does precisely this. If so, dimensional explanation fits well into the kairetic account of regularity explanation.

What of the dimensional explanation's lack of cohesion? Here I propose to make the same move as I did above for the explanation of momentum conservation. There is a dimensional explanation of the $3/2$ proportionality that is the same for a wide range of possible physical implementations of gravitation—and perhaps other orbit-generating mechanisms—and which is therefore quite possibly incohesive. But it is a frameworking explanation with a “given that” explanandum, and is therefore permitted to ignore the cohesion requirement.

More specifically, the universal dimensional derivation explains why the $3/2$ proportionality exists *given that* the relevant gravitational mechanics' L and T lines are as shown in Lange's formulation. But it does not explain why the relevant gravitational mechanics' L and T lines have the form they do. This latter explanation—required for a complete account of the $3/2$ proportionality in our universe—will have to respect the cohesion requirement, and so will derive the form of the L and T lines from the facts about the kind of gravitational mechanism that we have in the actual world (with cohesion determining, then, how broadly that mechanism will be individuated). Which, I suggest, is exactly the way that things should be.

Let me stop there. I have not offered an account of dimensional explanation, of course, but I hope I have persuaded you that the kairetic account might provide a suitable framework for doing so. And for much else besides!

References

Lange, M. (2009). *Laws and Lawmakers: Science, Metaphysics, and the Laws of Nature*. Oxford University Press, Oxford.