

Chapter 18

Explaining Causal Selection with Explanatory Causal Economy: Biology and Beyond

Laura R. Franklin-Hall

Abstract Among the factors necessary for the occurrence of some event, which of these are selectively highlighted in its explanation and labeled as causes—and which are explanatorily omitted, or relegated to the status of background conditions? Following J. S. Mill, most have thought that only a pragmatic answer to this question was possible. In this paper I suggest we understand this ‘causal selection problem’ in causal-explanatory terms, and propose that explanatory trade-offs between abstraction and stability can provide a principled solution to it. After sketching that solution, it is applied to a few biological examples.

Keywords Causal selection • Causal explanation • Explanatory trade-offs • Robustness • Abstraction • Causal democracy

1 Introduction: Explanatory Sparseness and Systematicity

Our universe is dizzyingly complex, and everything that happens within it causally depends on innumerable other things. The living world in particular can appear almost horrifically complicated. Though some of this complexity remains beyond our grasp, scientists have unraveled ever-larger portions of it. The combination of this complex world and our increasingly sophisticated theories accounting for it should make two features of our causal-explanatory practice appear surprising: its *sparseness* and its *systematicity*.

Explanatory practice is *sparse* in that many apparently legitimate causal explanations are rather thin affairs, in which happenings are accounted for with only the tiniest sliver of information, and not by citing all, or even very many, of an event’s

L.R. Franklin-Hall (✉)
Department of Philosophy, New York University, 5 Washington Place, New York,
NY 10003, USA
e-mail: lrf217@nyu.edu

causal influences.¹ Explanatory practice is *systematic* in that those few morsels that sparse explanations feed to us do not seem to emerge higgledy-piggledy, as if they were the output of some ‘explanatory lottery’ in which all causal factors enjoyed equal odds. Instead, the features scientists judge explanatorily relevant follow regular contours, perhaps indicating that hidden principles govern their selection. More specifically, I will distinguish two dimensions of explanatory sparseness and systematicity: horizontal and vertical.²

Along the ‘horizontal’ dimension, we do not normally explain the occurrence of an event by citing all of the conditions *necessary* for its occurrence, that is, by describing what J. S. Mill called its ‘total cause.’ Instead, one or a few features are given special priority, with other factors relegated—for good or ill—to the status of ‘background’ or ‘enabling’ conditions. For instance, the life-threatening sickling of an individual’s red blood cells during metabolic stress may be explained by citing a particular gene sequence—that coding for the hemoglobin protein—and not the equally necessary features of intra- and extra-cellular environments. More complicated, though still excruciatingly simple, explanations of biological development appeal to very spare gene regulatory architectures—like the double-repression network—in accounting for cell differentiation. These accounts also elide a profusion of essential cellular machinery. And moving from scientific to folk explanations, Hart and Honoré (1959: 10) provide a simple, and more famous, example of what is arguably the same phenomenon: it would be customary to explain the occurrence of a fire by citing the dropping of a lighted cigarette, but not by mentioning the environmental oxygen, even though its presence may have been equally necessary for the conflagration.

Along a ‘vertical’ dimension, those few features that do make the explanatory cut are often themselves quite ‘high-level,’ abstracting from sundry micro-details. For instance, a county’s high fox population may be explained by the springtime boom in bunnies—a favorite prey item of the medium-sized canid—with the account remaining silent on the particular hoppings-about, matings, and eatings of bunnies across the rural landscape (activities still very consequential for the foxes’ flourishing). The character of gene regulatory explanations in cell and developmental biology is identical; there, the production of all-important transcription factors is characterized in terms of coarse-grained concentrations, not by describing spatio-temporal details of molecular location. And, as above, Hart and Honoré’s (1959) fire example can provide a folk window on the same phenomenon: even when circumstances conspire to make the presence of oxygen explanatorily

¹The distinction between explanations as *communicative acts* and explanations construed in an ‘ontic’ mode as *sets of facts* will not loom large in this paper; throughout, I will presume that the content of communicative acts are the explanatorily relevant facts.

²Though many explanations are strikingly sparse, others—among them some so-called ‘mechanistic’ accounts—are less so. Though both sorts will be dealt with in this paper, as both are recommended by the explanatory theory that I will articulate, I begin by emphasizing explanatory sparseness because it is comparatively puzzling and in need of philosophical elucidation.

relevant to a fire's ignition—as when the inferno erupted in a manufacturing plant from which oxygen was normally evacuated—the explanation will invariably cite a rather high-level feature, such as the oxygen's non-trivial quantity, while omitting finer points, such as which particular oxygen molecules directly contributed to the blaze (even were such information miraculously available).

What accounts for these two varieties of selection—the ('horizontal') omission of background conditions and the ('vertical') omission of low-level physical detail? There is no consensus answer to this question in the philosophical literature on causation or causal explanation, it being widely believed that they represent independent dimensions and require different sorts of treatments.

On the one hand, vertical selection—that of the proper 'level of explanation'—is usually, though not invariably, understood as an objective matter. Thus, it is thought that an explanation might err in describing causes at the *incorrect* level, considering the explanatory target. For instance, defective explanations might be too low-level, including putatively irrelevant "gory details" (Kitcher 1984: 370); or, they might be too high-level, omitting objectively relevant details, as when a black-boxing explanation fails completely to "reveal underlying mechanisms" (Kaplan and Bechtel 2011: 442). Though there is agreement neither on just what the right level is, nor in virtue of what it would be correct, the hunch that there is a real phenomenon here, one amenable to a systematic treatment, is widespread.

Horizontal selection, on the other hand, is more often considered to be a purely pragmatic matter, and the omission of 'background factors' accounted for on quasi-Gricean principles. Mill, for instance, emphasized the "capricious" nature of causal selection, suggesting that factors are omitted "because some of them will in most cases be understood without being expressed" (1882: book 3, chapter 5) rather than on "any scientific ground."³ Along the same lines, both Lewis (1986) and Hall (2004) put horizontal causal selection down to "invidious distinctions" between causes of equal ontological, and presumably also explanatory, mettle. Summarizing the general philosophical mood, Schaffer (2007) states that "selection is now generally dismissed as groundless."

2 Aim: To Explain Horizontal Sparseness and Selectivity

In the face of such skepticism, the central aim of this paper is to offer a causal-explanatory analysis of the *horizontal* dimension of selection (often called the 'causal selection problem,' a label I use interchangeably). My proposal will, I hope, cut a pleasing path between two unsightly extremes: first, the implausibly strong metaphysical claim—and one with appropriately few advocates—that background factors and those cited as causes in explanations are of entirely distinct

³I am somewhat simplifying Mill's discussion, in which he floats a number of more concrete proposals concerning causal selection, though in each case emphasizing their haphazard application.

ontological genera; second, the unsatisfyingly shallow proposal—one belied by a close look at our practice—that there is no interesting, objective structure to the cause/background conditions distinction, and that context-dependent pragmatics always reigns.

While the horizontal dimension of selection will be my special focus, to deal with it I must sketch an account of causal explanation—one general in aspiration, though designed with the biological sciences centrally in mind—from which my treatment of horizontal selection derives. According to this *Causal Economy Account*—described in more detail elsewhere (Franklin-Hall [forthcoming](#))—complete explanations describe packages of causal factors that ‘cost less’ and ‘deliver more,’ making them, in metaphorical terms, maximally *economical*. Very briefly, ‘cost’ is equivalent to *total content of the explanation*—understood in terms of the number of ways the world might be that it excludes—while ‘delivery’ is equivalent to *stability that the explanans bestows on the explanatory target*—a factor tracking the ability of the factors cited in the explanans to make the explanandum event robust.

As will be detailed below, not only does this explanatory account offer a solution to the causal selection problem, but it also promises, in contrast with other proposals, to dispose of both sorts of causal selection—horizontal and vertical—at once, with background conditions and micro-details falling short in precisely the same way. Furthermore, the standard determinative of explanatory worth on the Causal Economy account—the *bang-for-your-buck principle*—processes these dimensions in concert, such that whether a factor is relegated to the status of background condition partly depends on its explanatory level. In consequence, my handling of horizontal selection cannot be completely insulated from the problem of vertical selection; I maintain that the common practice of separating them has been a mistake, as consideration of *background* and *level* must be kept in view simultaneously for either one to be seen clearly.

Beyond its solution to the causal selection problem, the Causal Economy account can address an issue more central to this volume: the well-advertised fact that scientists in different fields—ecology, evolutionary biology, and molecular biology, for instance—sometimes offer different explanations of the same event. While it has become customary to maintain that such differences result from the existence of substantially different explanatory norms governing the construction of explanations in different fields, the Causal Economy account can offer a more unified account of such differences, maintaining that a *single principle* guides explanation construction across the sciences, with context simply highlighting one or another equally acceptable accounts.⁴ It is able to do this by embracing a variety of explanatory pluralism according to which there sometimes exist multiple, distinct yet correct explanations of particular individual events. As will be explained below, this pluralism is organic to the explanatory account itself, in particular a consequence of the fact that different candidate explanations of a single event can

⁴See Press (2015) (Chap. 16, this volume) for alternative views on explanatory unification.

be equally ‘economical’—some costing and delivering much, others costing and delivering little, still others costing and delivering the same amount, but doing so in different ways.

The discussion unfolds as follows: in the next section I offer a brief overview of the causal selection problem and consider two responses to it. Following that, I sketch the Causal Economy explanatory account. Finally, I highlight its solution to the selection problem, while applying it to a few folk and biological examples, including to a dispute concerning the explanatory status of genetic and environmental causes of ontogeny.

3 Candidate Solutions to the Horizontal (or ‘Causal’) Selection Problem

The causal selection problem springs from two observations influentially highlighted by Mill (1882: book 3, chapter 5):

1. For any event (understood broadly to include states of affairs), there is an enormous set of factors—both positive and negative—necessary for its occurrence. We can think of these factors as those *on which the event depended*, or (perhaps equivalently) *which determined the event*. Call these an event’s *determinants*.
2. For any event, a comparatively small set of factors—both positive and negative—is given special causal or explanatory priority, being labeled *causes* of the event (in the extreme, ‘the cause’), and (perhaps equivalently) which are selectively cited in its explanation. Call these an event’s *narrow causes*. Determinants that aren’t narrow causes are *background factors*.

In its standard formulation, the problem of causal selection is the challenge of explaining the gap between the *determinants* of any event—which will all, technically speaking, be causes of it, on most any account of causation you may prefer—and its *narrow causes*. To do this in a substantive way requires offering a *selection principle* capable of telling the difference between narrow causes and background factors, as well as a rationale for the principle, in case the practice is thought justified. This principle might involve almost anything under the sun, possibly in complex combination: norms of communication, scientific convention, a causal metaphysics, or an account of explanation.⁵

⁵Should this principle be considered as part of the *semantics* or the *pragmatics* of causal-explanatory claims? That is, in cases in which a particular factor that is usually back-grounded (e.g., oxygen with the fire) is (apparently illegitimately) claimed to be *a* or *the cause* of the event, does this involve saying something *false* (the semantic view) or saying something strictly true, but falling short in some other respect (inappropriate, uninformative, irrelevant, etc.)? I prefer a semantic approach, but don’t think the choice here makes any difference to the substance of my analysis; you should feel free to reconstruct the discussion on your preferred picture.

To set the stage for my own explanation-based proposal, I offer a whistle-stop tour of alternative takes on the causal selection problem, covering populational and frame-working accounts. These stops are chosen for their proximity to our final destination; like my account, they address the causal selection problem at a kind of ‘intermediate level’ of philosophical depth, neither via a deep causal metaphysics nor a shallow conversational pragmatics.

3.1 *The Populational Account*

On the *populational* approach, the narrow causes of an event—or phenomenon more generally—are distinguished from background conditions by their special status *relative to a population*. This view can be spelled out in a number of ways, depending on whether the population is actual, or merely hypothetical. Here I focus on a version of Waters’ (2007) strategy, which appeals to an actual population, ignoring the many nuances irrelevant to my aims.

Waters’ picture has two principal parts. The first is a counterfactual difference-making account of causation borrowed from Woodward (2003). This serves to define what I call the determinants. Broadly speaking, these are the factors that, had they been different, the phenomenon would have been different (or would have had a different probability).⁶ Since the clash between accounts of causal selection does not concern the precise specification of this set—pre-emption puzzles aside, causal accounts agree enough on its composition—I will not explore this further.

The second part is a principle that picks the narrow causes—which Waters suggests we understand as *actual difference-makers*—from among the determinants. It is here that the populational element plays its part. Actual difference-makers are only definable relative to a population of entities, which might be “different actual entities” or “the same entity at different actual times” (2007: 566). They are the features that actually made a difference to the target event, across members of the population. Naturally, this population must be a mixed one with respect to the phenomenon in question, with some members of the population taking one state of it, and the other members some other. Relative to this population, some determinants (as they pertain to each member of the population) will be uniform, and others not. Actual difference-makers—if they exist, which isn’t invariably the case—are those determinants whose actual variation in the population at least partly accounts for the actual difference in the target phenomenon in members of that population.

An example will illustrate. Consider a particular fruit fly with red eyes. What is the narrow cause of this redness? *Prima facie*, there are many factors—both intrinsic and extrinsic to the fly—that might be cited, since there are many factors

⁶I am glossing over numerous features of Woodward’s account: how these differences are affected, that is, via interventions; what the causal relata themselves are, that is, variables, etc. Though important in other contexts, rehearsing these features would be only a confusing distraction here.

such that, had they been different, the eyes would not have been red. Yet Waters claims that biologists will often focus on a particular factor or factors in accounting for redness; it is the gap between the many determinants and the particular factors actually highlighted that makes this a problem of causal selection. This winnowing happens, on Waters' view, by conceptually embedding the red-eyed fly in a broader population, and considering what made a difference to eye color between *members of that population*.⁷

For instance, let us embed the red-eyed fly, which has genotype (+, pr), in a population of flies, some also with red eyes and that same genotype, and some lacking red eyes (possessing purple eyes, instead), with an alternative genotype (pr, pr). Virtually all other features of these flies are identical: they vary *just with respect to* whether they have a pr allele, or a wild-type (i.e., +) allele, at a particular locus. This is also among the many determinants of eye color. In this case, Waters says that the actual difference-maker for red eyes (in my language, its *narrow cause*) is the allele type (+ or pr) at that locus; all other determinants are mere background factors.

What if all the facts about flies and eyes had been just the same, but the red-eyed fly had been conceptualized as part of another actual population, one in which all flies shared the same (+, pr) genotype, but were cultivated at different temperatures? Importantly, this involves changing nothing about the causal order with respect to flies—it only changes *how the scientists think about the situation*. (Water says that he means 'population' in a statistical sense, so for the population to change, nothing about how flies interact, for instance, need be modified.) Assume that cultivation temperature is also a determinant of eye color. In this case, the narrow cause of red eyes—its actual difference maker—will change: now cultivation temperature, rather than allele-type. Thus, it is *our selection of a population*, on Waters' scheme, which ultimately explains causal selection.

3.2 The Frameworking Account

On a *frameworking* account, narrow causes differ from background factors relative to, not a population, but an 'explanatory framework' ('framework' henceforth). I will explore a version of this strategy from Strevens (2008: §§5.3 and 6.1). As above, this account has two parts. First comes a characterization of the determinants, the

⁷This embedding can be understood in one of two ways. On what appears to be Waters' preferred formulation and which I do not follow, when properly understood the real explanandum in this case is not the red eyes in a particular fly, but instead a difference in eye color in a particular population. Put in just this way, Waters would not actually be addressing the causal selection problem, since that is the problem of accounting for why particular determinants are cited in explanations of events, not in explaining differences in types of events across populations. (Were his solution to *that problem* it would be completely un-controversial.) So that Waters can be addressing the causal selection problem itself, I am articulating a version of his strategy that maintains the same explanandum, but makes the causal-explanatory claim relative to a population.

total set of an event's causes. These are difference-makers, though Strevens defines them via an optimization procedure rather than counterfactually. Subtleties aside, this procedure takes a fine-grained vertical model of the causal influences on an event and makes it as abstract as possible, while still requiring it be usable to derive, following the causal order, a statement of the occurrence of the explanatory target. Anything left in this model, post-optimization, is a difference-making cause. This model may look a sleek machine beside the more fine-grained one from which it was produced—but bulky it will still be, generously endowed with determinants.

The second element in this account of causal selection is the *framework*. This instrument, which Strevens modifies from Mackie's (1974: chapter 2) notion of a *causal field*, has many applications beyond causal selection, but its use there will be my exclusive focus. The framework specifies a state of affairs—one that must be veridical—that is assumed or held fixed in an explanation, and can be thought of as a distinct part of the explanatory request itself. When an explanation involves a framework, the explanation is not for some event—full stop—but rather for that event *given* whatever the framework specifies. The presence of the framework interacts with the optimization procedure described above, such that it only evaluates models in which this state of affairs holds. Furthermore, the state of affairs is not considered a difference-maker for (or explainer of) the target phenomenon, nor not one; it is simply a fact that the explanation takes for granted and evaluates difference-making relative to.

To illustrate, consider a frame-working approach to the fly example. On this view, if the particular allele is to be a narrow cause of red eyes, and temperature not, this will be because temperature, and not allele type, has been placed in the explanatory framework. In such a case, the explanatory target is not the *red eyes*, but, *red eyes, given that the fly was raised at such-and-such a temperature*. By the same token the reverse conclusion might be reached, that it is the *temperature* that is responsible for eye redness. For this result, simply place the fly's genotype in the explanatory framework (as well as, possibly, other difference-makers), but do not do this with respect to temperature. This picture generalizes to any example of causal selection (including those to which a populational account would not apply). So, while for Waters it was the explainer's selection of *an actual population* that did the work of causal selection, for Strevens it is the explainer's selection of a *framework*.

4 Evaluating Candidate Solutions to the Causal Selection Problem

Populational and frame-working accounts of causal selection are different on the surface, but share an essential commonality: *relativization*. On either view, an extra factor has been added to the explanatory request over and above the explanandum event. It is the selection of this factor that distinguishes narrow causes from background factors.

In relativizing in this way, the views just sketched are hardly unique. For instance, Schaffer (2012) offers a structurally similar account of causal selection (and causal semantics more generally) in which the relativizing utensils are *cause-and-effect-event-contrasts*. Similarly, Hesslow (1988) suggests that an *object of comparison* may be the proper relativizing device. Whatever the particular incarnation, interest in relativization is unsurprising in light of the fact that some explanations are explicitly relativized. It is in play, for instance, when sociologists ask outright what made a difference to educational achievement *in an actual population*, such as between students in a particular Los Angeles middle-school (Rumberger and Larson 1998). It is also in evidence when biologists ask why some individuals experience brain or liver damage, *given* (i.e., frameworking) that they were exposed to environmental lead (Onalaja and Claudio 2000). Accounting for relativization is thus a necessary part of any complete chronicle of our causal-explanatory practice, including in the account I will eventually recommend.

Nevertheless, if offered as the exclusive explanation for horizontal causal selection, relativizing strategies should not inspire you, for two reasons. First, there is the uncomfortable fact that many explanatory requests don't (apparently) make *any* reference to populations or frameworks at all, but despite this, selection seems to go off without a hitch. To illustrate with the simplest example: in the explanation of the breaking of a window, the throwing of the brick, not the absence of a wall between the brick and the glass (equally a determinant), will usually be prioritized. This selection is completely straightforward, even though in the normal circumstance neither population nor framework (nor any other relativizing implement) will be in view. This arguably indicates that such paraphernalia are not, at least invariably, responsible for selection.

Second, even in cases where the relativizer's apparatus is front-and-center, relativization-based solutions to the problem of causal selection provide no explanation for what is arguably the central issue: why certain kinds of causal factors are treated, by either the scientists or the folk, as narrow causes, and others as background factors. After all, *pace* Lewis and others who parade the capriciousness of selection, a good deal of actual selective practice is rather systematic. For instance, as noted in the introduction, transcriptional machinery is usually backgrounded in developmental explanations, and gas concentrations in explanations for forest fires. An illuminating account of causal selection will account for these patterns (as well as making sense of their exceptions). But relativizers don't do this. Instead, they out-source the explanatory task by suggesting that just what is selected is a consequence of which actual populations, frameworks, etc., are included in explanatory requests. And on the selection of *these* they remain silent.

Let me anticipate two responses. To the first point, the relativizer may counter that—while the relativizing apparatus isn't always noted out-right—in any given case a relativizing tool was *implicitly* included in the explanatory request, reaching out to do the selective work from just below the surface. To the second point, the relativizer may simply reject the explanatory demand. After all, on her view selection follows from *what explanatory requests are made*. And for the *origins*

of our questions, perhaps no informative account—at least, none that it is the philosopher of science’s job to provide—is possible.

These replies are difficult to decisively parry, but in a way revealing of their deeper deficiency: they combine to make the relativizer’s theory, if not completely immune from potential counter-example, certainly very close to it. Let me explain. First, at the core of her account the relativizer offers a variable—the chosen relativizing instrument, of whatever nature—whose value, in any given case, establishes which determinants are narrow causes, and which not. Next, the relativizer states that the value of this variable will, often if not usually, be left implicit. Finally, she declines to provide a theory that might give us access to its setting, claiming that beyond her bailiwick.

In total, this makes the relativizer’s solutions exceedingly prone to ad hoc maneuvers and just as suspect as a scientific theory of the same character. In any given case the relativizer can claim that the implicit variable’s setting is—rather conveniently—fitting to the actual selective behaviors we observe. That such an account is incorrect, of course, doesn’t follow. But cognizance of this characteristic can certainly motivate the search for a more systematic alternative, in particular one that is (1) more explicit, and (2) able to account for patterns in our causal-selective practice.

I hope my positive proposal can provide satisfaction on both of these counts. As noted already, it derives from an account of causal explanation whose attractions extend beyond a solution to the causal selection problem itself; for instance, as noted in the introduction, it is able to make sense of vertical selection as well as the pluralism of our causal-explanatory practice. In the next section I outline the heart of that account. Then the spotlight returns to the causal selection problem proper, eventually focusing on selective instances in biology in particular.

5 The Causal Economy Account

The philosopher of explanation—at least one without a revisionist bent—will take our actual explanatory practice as her datum and will devise a theory that both explains and rationalizes the principal features of that practice. Such an account will usually have two connected parts: (1) an articulation of just what gives an explanation its explanatory force (e.g., unifying disparate phenomena, answering what-if-things-had-been-different questions),⁸ (2) a description of explanatory *form*, that is, of what a complete explanation consists in (e.g., a derivation of a statement of the occurrence of an event, a veridical causal model of a certain kind, etc.).

Over the last few decades, philosophical opinion has coalesced around a broadly causal explanatory picture, one according to which it is an event’s *causes*, either some or all, that explain its occurrence. Yet, even with this constraint the explanatory

⁸See Issad and Malaterre (2015, Chap. 12, this volume) for a discussion of explanatory force.

tent must still be a large one, as causal enthusiasts harbor disagreements respecting both motivation and form. After all, while some see explanations as concerned with difference-making (Lewis 1986; Woodward 2003), with mechanistic models (Bechtel 2008; Glennan 2002; Machamer et al. 2000), or with a combination of the two (Craver 2007; Strevens 2008), others suggest that an ideal explanation will trace most or all of an event's causal influences—whether difference-making or no (Railton 1981; Salmon 1984). My own view, the *Causal Economy* account (Franklin-Hall forthcoming), is also a member of this causal-explanatory crowd. It has two principal features.

First, in contrast with many other theories, it is not tied to any particular causal metaphysics, but might be combined with alternative accounts of the causal relation, e.g., counterfactual (Lewisian or interventionist), regularity, conserved quantity, etc. This is possible because a metaphysics of the causal relation is not “where the action is” on the Causal Economy view, that is, the place from whence its proprietary explanatory constraints originate. Better then to leave metaphysical edification to the metaphysicians, and offer an explanatory account welcoming to all prospective partners.⁹

Second, Causal Economy offers a selection principle used to characterize the components of a complete explanation. According to this principle, good explanations—those deemed ‘complete’ and capable of providing understanding—are special for their *economy*: they ‘cost less,’ in virtue of being abstract, and ‘deliver more,’ in virtue of citing causal influences that make the event to be explained stable or robust. The motivation for this biggest bang-for-your-buck standard—and for the explanatory account more generally—is to isolate an event's *most important* causal factors, those in virtue of which it was, to a rather large degree, ‘bound to happen.’

Before applying it to the causal selection problem, I elaborate on this account's two key aspects—the ecumenical causal metaphysics and the sectarian selection principle. To aid in doing so, bring to mind any candidate explanandum event. It may be a ‘high-level’ event, like the simple death of an animal, possessing coarse-grained identity conditions, or, less commonly, it may be perfectly concrete—a death individuated in terms of all of its intrinsic properties, down to the horrible particulars of a final gasping breath. Whatever the target event, many of features of our universe, both positive and negative, will be responsible for its happening *just as it did*. In fact, given the structure of at least two of the fundamental physical forces (gravitation and electromagnetism), which diminish rapidly with distance but never go to zero, virtually any other event in the backwards light cone of an explanatory target will be so implicated.

All such minimally responsible events—as well as the laws in virtue of which they are related—are, in my language, *causal influences*.¹⁰ Granting the existence of this influential fabric, just what is its warp and woof? It is here that my metaphysical

⁹Here I am following Strevens (2008).

¹⁰This is a term of art, and should not be assimilated to Lewis' views in “Causation as Influence” (2000).

ecumenism rears its amiable head. I make no commitment respecting the real nature of causal influence, for instance, on whether it should be understood in terms of the *transfer of conserved quantities* via a relation of *in principle manipulability*, or something else entirely. I insist only on this: causal influence is fully physical.¹¹

But why demand even this? Here are two reasons. The first reason involves our need to compare different candidate explanations—at different levels of abstraction, for instance—on the basis of their economy. To do this, they must be in some sense commensurate, something that is only straightforward if all explanations are constructed from the same physical starting materials. The second reason is tied to my commitment to physicalism itself. However significant high-level *explanatory* relations may be—and I strongly believe in their import and thus design my explanatory account to make room for them—there is ever more reason to think that there is at least *some sense* in which it is physics, if any science, that describes our universe's basic movers-and-shakers. So, though no *explanatory* fundamentalist, I will grant physics a precedence of some kind, one located in the character of causal influence itself.

It is the second element of the Causal Economy account, the selection principle, that stops this physicalist approach to causal influence from effecting an explanatory fundamentalism. Rather than an event's explanation offering up a total, fine-grained causal-influential chronicle—as would be provided by one of Railton's (1981) 'ideal explanatory texts'—a complete explanation should cite just a special part of that saga. The precise size and shape of this part—just which features it includes, and which it does not—will depend on the architecture of the run-up to the event in question. Yet, at least in many cases, the part judged explanatorily relevant will be *sparse* in both vertical and horizontal dimensions, omitting small influences, low-level details, and (what are often considered) background factors.

How does the selection principle do this? Space will only permit a sketch of its treatment of the simplest kind of explanation, what I will call a *direct explanation*. In a direct explanation, an event is explained by reference to other events or states of affairs, in concert with a causal law connecting these with the target event,¹² but without mention of intermediate events. This focus is apt because the causal selection problem is always posed in terms of direct explanations. Furthermore,

¹¹Most transference accounts of the causal relation—such as Dowe's process theory—are already physically constrained, so my requirement is, on them, without effect. It has the most impact on counterfactual accounts, such as on a causal interventionism of the Woodward (2003) variety. If causal influence is cashed out in interventionist terms, I insist that the causal model to which interventionist causal claims are relativized be a fine-grained physical model, not a 'high-level' one (even though those would otherwise be kosher) (Here I follow Strevens (2008)).

¹²For lack of space, I am equally unable to say much about the content of these causal laws, and will focus on characterizing the *states of affairs* are to be explanatorily cited. But in brief, the content of such laws is determined—not by some independent 'high-level' account of causation, such as that provided by Woodward (2003)—but by the explanatory selection procedure itself. The causal law connecting state of affairs A and target event B asserts that A is a winner of the causal economy competition with respect to B.

though many of the most interesting scientific explanations are not of this kind, I see them as assembled from direct-explanation building blocks. So, in accounting for the nature of those blocks, I will still have the opportunity to describe a key component of the Causal Economy picture.

There are two conceptual steps in explanation assembly for direct explanations: *production* and *selection*. *Production* is a process in which an exceedingly large number of candidate explanations are manufactured from the complete causal-influential tale for a target event. That tale will be tellable in physical locution, given the nature of causal influence itself. Candidate explanations are produced by censoring the complete account via omission and abstraction. In the case of omission, particular causal influences are completely deleted. In the case of abstraction, causal influences are described in less detail, through *coarse-graining*, *amalgamation*, or *populational* transformation. In coarse-graining, a particular feature is described as falling in some range or exceeding some threshold (*over 30 miles per hour* replacing *35 miles per hour*). In amalgamation, multiple lower-level features applicable to a particular individual are combined in a more complex parameter, and the particular values of the components are thereby lost (*15 kgm/s momentum* replacing *5 kg mass at 3 m/s velocity*). In populational transformation, a population-level feature is cited rather than a set of parameters applicable to the individuals constituting the population (*temperature* of a gas, e.g., mean kinetic energy, replacing its constituent molecule's particular *kinetic energies* (themselves products of amalgamation, since kinetic energy $\approx \frac{1}{2}$ mass * velocity²)).¹³

When all of these kinds of transformation are applied in different orders and degrees to the complete story for a particular event, there will result exceedingly many candidate explanations, each one a separate 'package of causal influence.' Some will include a good deal of the total chronicle, others small slices of it. Needless to say, most packages will in no way resemble the actual explanations scientists offer up. But within the rubbish nestle explanatory gems, special packages of causal influence that do appear explanatorily suitable.

These gems are identified in the selection step. As noted already, relative to a target event, an explanatory package maximizes the ratio of delivery to cost (or equivalently, maximizes the product of delivery and cheapness). The *cost* of a package is its *total content*, which I understand to be the number of ways that the world might be that it rules out. Other things equal, good explanations say very little about the catastrophically complicated run-up to a target event. The *delivery* of a package reflects the *stability boost* that it provides the explanatory target; this tracks the extent to which the package makes it the case that the target event would still have happened, even had circumstances been in various ways different.

¹³Both amalgamation and populational transformation are species of what is sometimes called *variable reduction*. Variable reduction isn't always considered a kind of abstraction, but I class it thus; it equally involves moving to a representation that leaves information out in contrast to the original.

I will amplify on both the cost and delivery aspects of the selection principle, in order. For *the number of ways that the world might have been that the explanation rules out*—my take on *cost*—to have any significance, it must be tied to scheme of world individuation. To maximize expressive potential, I assume a scheme that is fine-grained and physical.¹⁴ Whether this scheme is itself objective—that is, whether it is rationally required for any explainer, rather than being in some way ultimately *dependent on us*—is something on which I remain agnostic. Whatever the metaphysical status of the individuation scheme at its heart, what pressure does cost put on complete explanations? In short: less is better. Other things equal, an explanation should be very abstract, trimming away as much of the causal-influential run-up as possible. This might involve deleting what we normally consider to be small causal influences—factors that aren’t even among the *determinants* of an event. It might also involve removing (what we usually consider) *background factors* or *low-level details*, those extracted via omission or any of the kinds of abstraction noted above, thereby bringing about horizontal and vertical sparseness. Indeed, if we were considering cost alone, we would even omit the factors that *do* appear explanatorily relevant.

While cost considerations favor paring causal influences away as much as possible, delivery favors including factors in an explanatory account to the degree to which those factors make the target event stable or robust, providing what I will call a *stability boost*. Intuitively, a big booster of an event’s stability is a factor that makes it the case that the event would have happened *even had many other things been different*. For instance, the impact of the Chicxulub asteroid provides a stability boost for the dinosaurs’ demise, since, just so long as the impact itself occurred, many features—grazing patterns, immunological status—might have been different, yet the animals would still have perished.

To be somewhat more precise, the stability boost offered by a candidate explanans reflects the *additional stability* that the factors cited in the explanans contribute to the event, over and above the stability of the event simpliciter (the *baseline stability*). To measure this additional stability, we subtract the baseline stability from what I will call the *construct stability*. Thus,

$$\text{construct stability} - \text{baseline stability} = \text{stability boost}$$

Though it might be understood in a number of ways, I measure stability using a possible-worlds framework, which I can only sketch briefly here. On this approach, the baseline stability of an event is equivalent to the number of a privileged set of nearby possible worlds—worlds differing from our own via one or a number of simple, physical perturbations (effected via Lewisian small miracles)—

¹⁴Though not important to my task here, the cost measure should also be relativized to the size of the causal-influential fabric for the event at the time of the candidate influence (e.g., the span of the event’s backwards light cone at the relevant time). After all, it is only as between properties of that material—one that gets rapidly larger earlier in time—that the selection principle must pick.

in which the target event nevertheless takes place.¹⁵ More specifically, in each world experiencing a perturbation, let events unfold according to the actual laws until the time of the target event. Events are more stable simpliciter to the degree to which they occur across a greater range of these possible worlds. For instance, World War I might turn out surprisingly stable simpliciter: perhaps lots of things could have been different—among them the assassination that actually triggered the cataclysm—and yet a European war in that period still would have eventuated.

The *boost* provided by a candidate explanans is determined by how much stability, over and above the baseline stability, is due to factors cited in the explanans. To measure this, we must construct a new set of worlds, taking as our starting material the set of nearby possible worlds appealed to above. Consider those worlds at a time just after the perturbations that distinguish them from our own, and let them unfold, according to the laws, until the time at which the causal influences cited in the explanans should appear. Now in some of these worlds, the causal influences cited in the candidate explanans will probably be absent due to the particular perturbations found in those worlds. For every such world, let a second miracle occur, making it the case that the causal influences cited by the candidate explanans are in fact present.¹⁶ (This second miracle will permit us to probe the relationship between the factors cited in the explanans, and the target event, in worlds somewhat different from our own.) Then, let all worlds continue to unfold according to the physical laws. The ‘construct stability’ is gauged by the number of worlds in which the target event (the explanandum) takes place. Generally, this value reflects both the degree to which target events are stable simpliciter, and the degree to which events become stable once we’ve fixed the influences cited in the explanans.

Having characterized baseline stability and construct stability, the stability boost is just the difference between the two. Broadly speaking, this factor reflects the stability that the explanans contributes to the target event.¹⁷ One important consequence of the measure is this: packages of causal influence that offer substantial stability boosts—greater than that provided by competing packages—are those that are themselves unstable simpliciter, but such that, given their occurrence, the event to

¹⁵This set of worlds is one that I describe exclusively to make sense of our explanatory practice. I take no stand here on whether it is in any way an objective set, one that is special from a metaphysical point-of-view. See my (forthcoming) for more on the privileged set of nearby possible worlds, produced by a basic set of perturbations, that is used to define the stability boost.

¹⁶Crucially, this will not usually just reverse the initial perturbation that had the consequence of disrupting the influences cited in the explanans. Many consequences of that original perturbation will persist, despite the re-enactment, since the disruption of the explanans influences will usually be but *one of many* down-stream effects of the initial perturbation.

¹⁷Because the baseline stability of any target event is a constant, subtracting it will not change the ordinal ranking of candidate explanans. The subtraction is nevertheless useful by allowing us to conceptually distinguish target events with high baseline stability from those only stable with the aid of factors cited in the explanans.

be explained is very stable.¹⁸ Such influences are distinctive in two ways. First, they will be (what we would usually consider) difference-makers for the occurrence of the target event, setting them apart from the abundance of factors that make a difference to *how* it occurred, but not to *whether* it occurred.¹⁹ Second, and more discriminately, they will, in most cases, constitute just a subset of the difference-makers—a feature that forms the heart of my solution to the causal selection problem. Excluded will be those factors that cannot contribute to an event’s stability because they were themselves so stable. These factors cannot offer much of a stability boost because, in virtue of being so stable, they will remain present in many of the worlds in the privileged set (i.e., those that experience the original basic physical perturbations). And since they are so pervasively present, there is no role for a second miracle to institute them and in that way impact the occurrence of the target event. As I see it, such highly stable factors are mere channels through which the work of the real stabilizers has been transferred from elsewhere. Naturally, it is the originators of stability—not their envoys—that must be explanatorily relevant on any picture on which explanations tell *in virtue of what* an event was ‘bound to happen.’²⁰

Putting the pieces together: when the selection principle evaluates an event’s candidate packages of causal influence—the output of the step producing the candidate explanans—it will extract those that jointly maximize abstractness and stability-boosting. There may, of course, be multiple packages that do this, some that are more concrete or even mechanistic in flavor, but which also make it the case that the target event is virtually bound to happen, and others that are more abstract while still offering a substantial stability boost. I submit that any of these optimal

¹⁸Predictably, it will be possible to contrive counterexamples to this gloss on my procedure, cases in which a baroque causal architecture thwarts my measure’s ability to isolate factors that do have this property. When considering such cases, it is important to keep the spirit of the Causal Economy proposal in mind: its selection principle aims to informatively describe the broad principles influencing our explanatory practices, practices that I believe to be best defined for the causal systems that scientists normally encounter and on which their norms have been trained. The most important task for an account is to deal with those systems; after all, it is challenging enough to provide an informative description—that is, one that doesn’t simply appeal to pervasive relativization—of principles guiding horizontal and vertical selection in those central cases.

¹⁹Fans of counterfactual difference-making accounts of causation may wonder about the nature of the base or contrast state relative to which difference-making is being implicitly evaluated. My algorithm, in effect, does not pick one ‘default,’ but instead surveys a large range of states, checks for difference-making relative to each of these states, and integrates over those results. In particular, big stability boosters are difference-makers relative to many or all of these alternative states, not just those present in the collection of worlds produced from our own via simple, physical perturbations at some prior time. This strikes me as better solution to the ‘default problem’ than privileging one such state, perhaps the one deemed (on subjective grounds) ‘normal.’

²⁰Though my procedure does, I hope, give special preference to the unstable difference-makers, those that seem particularly explanatorily relevant, I do not deny that even stable difference-makers might, in some contexts, have an explanatory role. In particular, if we expand the set of nearby possible worlds appealed to in the construction procedure—by allowing for more and more radical perturbations—even seemingly stable difference-makers will sometimes be absent in the privileged set, thus making them at least potentially explanatorily relevant.

packages, in concert with the causal laws connecting them with target events, can constitute *complete explanations*, and are scientifically acceptable accounts that yield real understanding.²¹ But do such packages indeed resemble those direct explanations that we actually offer up? Yes, or so I will suggest in the next section in the process of addressing our central topic: an explanation-based solution to the problem of causal selection.

6 The Causal Economy Treatment of the Causal Selection Problem

Defending a substantive account of causal selection requires establishing the extensional adequacy, as well as rational defensibility, of its selection principle. Most importantly, it must be shown to pick *narrow causes* from *determinants* in a way at least broadly consistent with how those cuts are actually made. Three data streams might point us to the location of such cuts: (1) arm-chair philosophical intuition; (2) psychological studies, such as those probing the folk distinction between causes and ‘background’ or ‘enabling’ conditions (see Cheng and Novick 1991; McGill and Tenbrunsel 2000; N’gbala and Branscombe 1995); (3) the actual explanatory annals, that is, the sum total of extant explanatory texts. Absent space constraints, I would explore all three sources of evidence. Given constraints, establishing Causal Economy’s consistency with the actual explanatory annals is my first priority. Yet, for the sake of getting out the basic move before bringing in scientific complexities, I start by treating a simple datum from stream (1): philosophical intuition.

6.1 An Intuitive Example

The most rehearsed case of causal selection concerns the explanation of a forest fire, an example mentioned already in the introduction. The determinants of such a fire extend to all the conditions, positive and negative, on which the fire depended, including: the fact that there was wood around; that the wood wasn’t so wet as to be non-flammable; that there was oxygen present; that a lighted cigarette was dropped. Though all of these are in some sense *causes* of the fire, intuition suggests that only the last will invariably be mentioned in its explanation. It is the fire’s *narrow cause*.

²¹Though not directly relevant to the causal selection problem—and thus not worth detailing here—Causal Economy requires a further constraint on abstraction to prevent a preference for disjunctive explanations that are contrived to be both abstract and stability boosting simultaneously. The constraint I prefer is a cohesion requirement—modeled on a standard from Strevens (2008)—on which a particular feature of the influential nexus cannot be made so abstract that it is impossible to move, in physical state space, from one possible realizer of it to another without moving through a realizer that is not an instance of it.

Agreeing with intuition, Causal Economy recommends an explanation in terms of the cigarette-dropping because that event is the feature of the fire's causal-influential run-up that maximizes causal economy: being a simple, local event, it is comparatively cheap to characterize; and it provides the target event, the fire, with a large stability boost. This is in virtue of the fact that the drop of the cigarette is a relatively unstable event, but one given which the fire is stable. In contrast, 'background factors'—the presence of oxygen, the fuel in the vicinity—are omitted because they are so stable that they can't much contribute to the stability of the fire.

Next, consider a variant case. Here the fire breaks out in a manufacturing plant in which facilities have been constructed to eliminate oxygen. Such a fire still required oxygen's presence. Thus, one feature of the run-up to the conflagration will have been a failure of the oxygen-evacuation mechanism. Many determinants of this fire will resemble those of the woodland blaze: the presence of a particular igniter (not a cigarette, perhaps, but some other spark source); that there was oxygen around; that flammable materials were not overly damp, etc. Yet intuition suggests that, in this case, the presence of oxygen *will be* (among) the fire's narrow cause(s).

How to make sense of oxygen's distinctive relevance? Crucially, on the Causal Economy view relevance is neither traced to some difference in explanatory presuppositions—as relativizers might have it—nor to the mere fact that oxygen is not 'normal' to the factory (though that is indeed true). Instead, its relevance is tied to an objective feature of the causal-influential architecture of the factory fire's run-up. Most notably, in the factory—and not in the forest—the presence of the oxygen was rather unstable. Thus, oxygen's presence provides the fire a stability boost that the oxygen in the forest could not provide to the forest fire.²² In combination with the fact that oxygen's specification is itself cheap—given facts about fire propagation, oxygen gas must just exceed some concentration, but need not (for instance) have its constituent molecules arranged 'just so'—oxygen offers rather good economy.

I submit that the cigarette-dropping in the forest, oxygen's presence in the factory, and all other systematically selected features, share this in common: they are located at 'sweet spots' in the architecture of the run-up to the events that they explain. These are features at which two characteristics—cheapness and stability-boosting—somehow converge. A universe might have lacked sweet spots, given that these characteristics often trade-off: a sure way to maximize an event's stability is to give a full—and thus very costly—specification of its antecedents. And the reverse is also true: very cheap specifications usually deliver only minimal stability.

²²A nuance: given that oxygen will infiltrate the factory following a failure of whatever mechanism was responsible for its evacuation, *that failure* itself may be an economical part of the causal-influential run-up. Thus the question arises as to whether it—and *not* the oxygen—might constitute the most economical—and thus complete—explanation. As far as I can tell, on the causal economy view both of these would make for good explanations, as both are unstable events that make the target event stable. They differ, of course, in their place in the temporal sequence, but causal economy will often find multiple co-equal (and equally complete) explanations that cite influences at different times (as well, as those at the same time). This, I hope, is such an instance.

Nevertheless, sometimes you get more than you pay for. And for this, cognitively limited creatures like ourselves should be thankful: it is on these grounds that complete explanations are, with any frequency, within our grasp.

6.2 *Scientific Cases*

With this core strategy in view, next I consider some selective examples in scientific explanatory practice. Saying anything punchy about something so vast is not easy, so I further focus attention on causal selection in biology. Being a science of exceedingly baroque systems, causal selection is pervasive in the life sciences, even among sophisticates. This makes it a gold mine for those wanting to identify the selective patterns to which a philosophical account is responsible. Even more narrowly, consider just two (somewhat overlapping) domains: *signaling systems* and *organismal development*.

6.2.1 Signaling Systems

Biologists use the term ‘signaling system’ to refer a variety of phenomena: visual and auditory animal communication mechanisms; immune system coordination; interactions between molecules within a single cell (Weng et al. 1999). Put aside the complex question of just what these systems have in common and consider particular explanations for events within them.

In explaining the hiding of a vervet monkey in a nearby bush its neighbor’s short, grouped vocalization between 0.2 and 1 kHz—the “eagle alarm” (Seyfarth et al. 1980)—may be cited. In accounting for the growth of the shmoo (a membrane protrusion) in an alpha-type yeast cell, the membrane binding of **a**-factor pheromone, a molecule produced by a compatible **a**-type cell, will be named (Bardwell 2005). Or, to make sense of the breakdown of glycogen during your daily exercise, the release of epinephrine will be mentioned. It is this protein whose binding to 7TM receptors on liver cells begins a molecular cascade that eventuates in the glycogen’s demise.

What do these accounts have in common? They omit any description of the inner workings of the signaling process and speak just of a system’s production of a certain output in response to single, narrow input: the feature often labeled as the signal. This narrow focus is not a result of any target event depending exclusively on just such inputs, as in each case innumerable other conditions were necessary for the output to be explained; the vervet’s hiding required his being awake, unrestrained, possessing sufficient oxygen to react, etc., and the shmoo’s growth required an equally complex set of molecular requirements. But in these and other cases, the signal’s presence will, at least judging by actual practice, be the explanatory relevant determinant.

This is just what we’d expect were the Causal Economy account correct: signaling events possess the properties that make for good explanations. First, they

are cheap. Though presumably not a conceptual requirement on something's being a signal, each actual signal's presence is low in content (e.g., a relatively high-level phonological pattern, or a single molecule type's coarse-grained concentration). Second, they deliver, providing stability-boosts, in virtue of the following: (1) their presence is a difference-maker for the target event (e.g., vervet hiding, shmoo growing, glycogen destruction); (2) their presence is itself unstable, making both 'on' and 'off' states among the nearby possibilities; (3) other in-principle difference-makers for the target event *are* (comparatively) stable.²³ Given this combination, signaling states usually win the explanatory competition.

6.2.2 Organismal Development

Next consider organismal development. For most multicellular creatures that reproduce by way of a single-celled bottleneck, development is a process by which that single cell—the zygote—transforms into a multi-cellular and functionally integrated organism via coordinated growth, morphological change, and cellular differentiation. Though complicated, development is an important phenomenon for a theory of causal selection to treat in the light of an ongoing debate in the philosophy of biology concerning the developmental importance of 'genetic causes,' the nucleotide sequences that constitute an organism's DNA. Though controversy over the status of genetic causes is not always understood against the backdrop of the problem of causal selection, it often is, and I will take it in that light here.

Trivially, every aspect of an organism's developmental trajectory will depend on (or be determined by) a variety of factors, positive and negative, near and far. Some of these features are genetic; others are not. Non-genetic developmental determinants include non-nucleic organismal properties, as well as properly 'environmental' features, those that originate in processes fully extrinsic to the organism. Though there is broad agreement on the simple fact that genetic and non-genetic features, both necessary for development, interact in complex ways to produce organismal traits, talk of 'a genetic blue-print' and 'genes-for' particular traits is common enough in scientific and (even more) in popular discussions that some have wondered whether there might be reason to judge genetic factors *more important* or *relevant* to developmental outcomes than other factors. Ignoring all sorts of subtleties, there are two families of reactions to such a suggestion. Some, like Weber ([forthcoming](#)) and Rosenberg (2006), believe that genetic factors are,

²³To spell out in more detail how the stability boost measure applies to a particular case, consider a specific signaling event, such as a particular monkey's flight to the bush. To consider the stability boost of a signal from a neighboring animal, rewind the tape to some time before the signal was sent and perturb the actual world in a variety of ways, thus producing the set of privileged nearby possible worlds. Some number of these perturbations would prevent the signaling animal from emitting the call. And since the call itself is required for the hiding to occur (and because other factors equally required are themselves very stable), an explanans that includes the call's production will substantially augment the stability of the hiding event.

at least to a first approximation, privileged developmental explainers, features that scientists very properly prioritize. The project of these ‘gene-centric’ philosophers is to say in just what respect genes are so causally special. Others, like Gannett (1999), follow a broadly Millian line and deny that there is a principled and across-the-board difference between the roles of different kinds of developmental determinants, whether genetic, environmental, or otherwise. Their burden is to provide an ‘error theory’ of the genetic preference one sometimes finds in scientific explanatory practice; Gannett, for instance, points here to the scientists’ financial motives.

The degree to which Causal Economy is in tension with these alternatives so briefly sketched will depend on precisely how they are spelled out (Stegmann 2012). Yet, at least when viewed in their broadest outlines, Causal Economy steers between the parity and gene-centric views, rejecting one aspect of each. First, it straightforwardly discards the Millian idea: because of the existence of sweet spots in the causal architectures of developing organisms, for many developmental explananda there will be principled explanatory—and not merely pragmatic—differences between developmental determinants, elevating some as narrow causes, while relegating others to the background. Second, Causal Economy rejects an across-the-board gene-centrism.²⁴ Though in *some cases* genetic features might be uniquely explanatory of developmental events, in other cases environmental features offer the best pay-off, and in still other circumstances features falling neatly in neither category (e.g., non-nucleic, molecular features, such as methylation patterns) should be (and are) prioritized on explanatory grounds.²⁵ In this respect the Causal Economy approach is consistent with critiques of gene-centrism, like that offered by Stotz (2006), which claim that pervasive prioritization of genetic features as developmental explainers is unmotivated by the causal complexities of living systems.

Let me illustrate this middle position via an example: the explanation of phenotypic sex, that is, for why particular organisms are either phenotypically male or phenotypically female. While in mammals and birds sex is usually given a genetic explanation—mammal sex is explained by an aspect of the father’s chromosomal contribution and avian sex by the mother’s—in turtles it is explained environmentally. During a particular period (the aptly named *thermosensitive period*), the temperature of an egg-bound turtle embryo is said to determine its sex, with higher temperatures making females, and lower temperatures, males. Of course, as per a biological application of the Millian insight, the fact that any particular organism—tortoise, cat, or cockatoo—develops the phenotype (typical) of one sex or the other will strictly depend on *both* environmental and genetic

²⁴To be fair, it is unclear whether anyone really holds an across-the-board gene-centrism; even broadly ‘pro-gene’ philosophers, such as Weber and Rosenberg, articulate positions that are considerably more nuanced. If they too want to reject gene-centrism, our disagreement would concern the subtler issue of the *standard* by which a subset of an event’s necessary conditions—sometimes genetic ones and sometimes not—might be explanatorily privileged.

²⁵There might also be cases in which multiple factors—some genetic and some not—are equally causally economical, and would each constitute complete explanations.

factors, among others. To see this, note that there are changes to a developing mammal's intra-uterine environment that would lead even a genetically male fetus to develop a characteristically female phenotype. And similarly, certain genetic features of a developing turtle are necessary for low temperatures to eventuate in male offspring; after all, in other amniotes (such as crocodiles), the relationship between temperature and sex is different, with moderate temperatures yielding males, and extremes, females.

In light of this pervasive multi-causal determination of sex, why is it said to be 'genetic' in mammals and birds, and 'environmental' in turtles? This difference in selection follows neatly from the different economy of the genetic vs. environmental explanations across these cases, something well defined even absent relativizing maneuvers. The economy of such patterns vary in light of the fact that—just as the causal run-up to the forest fire was objectively different from the run-up to the factory's blaze—turtles and mammals (among others), as embedded in their environments, are complex systems with different causal architectures. In the case of the developing turtle, low temperature is the cheap and stability-boosting determinant of male phenotype, while in mammals, presence of the Y chromosome is. After all, a low incubation temperature is itself an unstable event—had the egg been just somewhat differently positioned in the subterranean nest, that temperature would have been different—but one in virtue of which the turtle's sex is stable. Yet this particular kind of temperature sensitivity *does not hold* of a mammalian juvenile, though the presence of the sex-determining chromosome has an equivalent character. This difference in architecture exists in spite of the fact that, as emphasized already, every animal's (phenotypic) sex strictly speaking has genetic, environmental, and other determinants besides.

7 Conclusion: Relativization Revisited

Though I hope that the inherent plausibility of Causal Economy's selection principle speaks somewhat in its favor, the theory's ability to account for selective patterns in scientific explanatory practice is the ground on which I recommend it here. Beyond the few examples that space permitted me to explore, the ambition is that Causal Economy could negotiate *all clear instances* of selection. It is thus vulnerable, in principle, to counterexample. And just what sorts of cases would be most threatening? Naturally, those in which systematically explanatory cited factors were not causally economical. This would be because, in contrast to alternatives, they were comparatively (1) costly, or (2) not stability-boosting of the events they were offered to explain. For instance, what a ticklish situation it might have been for Causal Economy had herpetologists accounted for the sex of Lonesome George—the last Galapagos tortoise from the Isle of Pinta—on genetic grounds!

I know of no such clear counterexamples to Causal Economy's recommendations. Yet a different kind of case may appear to make trouble more obliquely. By way of conclusion I will respond to an instance of this kind: the car crash, as

presented by (among others) Hanson (1958: 54) and Carnap (1995, 1966: 191–92). In the run-up to this crash, an angry driver, depressing the gas pedal in a fit of furry, speeds down a country lane in a rainstorm. Hitting a bump, the car spins out of control, colliding with a barn wall. What accounts for the crash? Naturally, this event has many determinants, including: any property of the mechanism connecting pedal depression with vehicular acceleration, the car's speed, the lack of steer stabilizing equipment, the road's wetness, and the driver's anger. Carnap claims that—rather than expecting any consensus on which of these was *the cause* of the crash—different factors will be highlighted by different individuals, each one “looking at the total picture from his point of view” (192). The policeman, for instance, may put the crash down to the speed; the psychologist to the anger of the driver; the engineer to the road's wetness.

Of course, the mere suggestion that this event may have multiple explanations, which may be cited by different parties, is not in itself a problem for the Causal Economy view. The selection principle allows for multiplicity in virtue of simply requiring an explanation to maximize the ratio of delivery to cost, and, as noted already, there may often be multiple ways of achieving this maximum. In particular, whenever there are multiple sweet spots we should expect there to be multiple maxima, and thus multiple co-equal explanations. And though the focus of this paper has been on competition between ‘horizontal’ alternatives, all of which have been relatively high-level, legitimate alternatives can also vary vertically: low-level explanations, such as those often called ‘mechanistic,’ may cost and deliver a lot, but offer the same economy as high-level ones that cost little and deliver proportionality. Alternatives may even vary along both horizontal and vertical dimensions simultaneously—as Potochnik (2010) has noted is characteristic of alternative scientific explanations for the same phenomenon.

However, in the case at hand it seems that merely noting the possibility of multiple explanations will not provide a satisfactory response, as none of the multiple explanations of the car crash that Carnap mentions will win the Causal Economy competition. After all, in order to make the crash stable, some confluence of these factors, each independently unstable, must occur. Thus, Causal Economy will require a complete explanation to cite that collection of unstable difference-makers as a package. Why then does it seem that the policeman and the psychologist, for instance, will just state part of the full explanation for the crash, and a part that seems guided by their unique circumstances? Granting this datum for the purposes of argument, my response is simple: there are some contexts—including extremely practical ones—in which partial explanations may be offered, and pragmatics (by which I mean either conversational pragmatics, or implicit relativization of some kind) then do determine which part is cited, and thus determine which cause is selected.

In exploiting the relativizer's line here, have I sapped the Causal Economy account of its advertised explanatory pay-off? I don't believe so. First, even if there are instances, such as the case of the crash, in which relativization is called for, the Causal Economist uses this strategy far less frequently than does the full-blooded relativizer. After all, there are many circumstances—such as respecting

the biological events examined above—in which the architecture of the systems involved yields winners in the explanatory competition that are singular factors, rather than a combination of factors. Whenever this is the case, the Causal Economist can move from causal architecture to selection on objective, explanatory grounds. Second, even when the Causal Economy does bring relativization into play, it does less heavy lifting than it does for alternative views. This is because the selection principle has still wildly narrowed the space of possible determinants that might be explanatorily cited, and thus the role of pragmatics in accounting for the selected feature is proportionality diminished.

Even granting that relativization does not eradicate Causal Economy's explanatory offerings, it should still be acknowledged to limit them. There will now be a variety of selective questions on which it has nothing to say except the following: A was selected, and not B, because B was (for any of a variety of unspecified reasons) frameworked. Yet perhaps this is fitting, given the data before us. Let me explain. The discussion opened by emphasizing two features of explanatory practice: its sparseness and its systematicity. Over the course of the paper, I suggested that these were the offspring of two parents: first, an account of causal explanation according to which uniform principles determine selection, along both horizontal and vertical dimensions—principles that prefer sparse explanations; second, the fact that the kinds of causal structures responsible for the events that scientists target—particularly in the life sciences—have features that make for clear explanatory winners across a range of cases, accounting for the systematicity that we observe.

Yet this systematicity is not pervasive, and those stressing the haphazard and interest-relative nature of selection are on to something; at times causal selection is influenced by practical considerations. Best for an account of selection to acknowledge this, and to neither shoe-horn the data to suite an overly systematic theory, nor stipulate that penumbral cases should be awarded as “spoils to the victor.” Instead, the victor should be she who both accounts for those patterns in selective practice that actually exist—as I hope the Causal Economy account does—and makes room for an absence of systematicity elsewhere.

Acknowledgements For helpful comments on this paper, thanks to David Frank, Maria Kronfeldner, Michael Strevens, David Velleman, two referees, and to (this volume's editors) Pierre-Alain Braillard and Christophe Malaterre. Though I was not able to directly address all of the useful suggestions I received, they uniformly aided me in the development and presentation of my argument.

References

- Bardwell, L. (2005). A walk-through of the yeast mating pheromone response pathway. *Peptides*, 26(2), 339–350.
- Bechtel, W. (2008). *Mental mechanisms: Philosophical perspectives on cognitive neuroscience*. New York/London: Routledge.
- Carnap, R. (1995, 1966). *An introduction to the philosophy of science*. Mineola, NY: Dover Publications.

- Cheng, P. W., & Novick, L. R. (1991). Causes versus enabling conditions. *Cognition*, 40(1–2), 83–120.
- Craver, C. F. (2007). *Explaining the brain: Mechanisms and the mosaic unity of neuroscience*. Oxford: Clarendon.
- Franklin-Hall, L. R. (forthcoming). The causal economy approach to scientific explanation. In C. K. Waters (Ed.), *Minnesota studies in the philosophy of science*. Minneapolis: University of Minnesota Press.
- Gannett, L. (1999). What's in a cause? The pragmatic dimensions of genetic explanations. *Biology and Philosophy*, 14(3), 349–373.
- Glennan, S. (2002). Rethinking mechanistic explanation. *Philosophy of Science*, 69, S342–S353.
- Hall, N. (2004). Two concepts of causation. In J. Collins, N. Hall, & L. A. Paul (Eds.), *Causation and counterfactuals* (pp. 225–276). Cambridge, MA: MIT Press.
- Hanson, N. R. (1958). *Patterns of discovery: An inquiry into the conceptual foundations of science*. Cambridge: CUP Archive.
- Hart, H. L. A., & Honoré, T. (1959). *Causation in the law*. Oxford: Clarendon Press.
- Hesslow, G. (1988). The problem of causal selection. In D. J. Hilton (Ed.), *Contemporary science and natural explanation: Commonsense conceptions of causality* (pp. 11–32). New York: New York University Press.
- Issad, T., & Malaterre, C. (2015). Are dynamic mechanistic explanations still mechanistic? In P.-A. Braillard & C. Malaterre (Eds.), *Explanation in biology. An enquiry into the diversity of explanatory patterns in the life sciences* (pp. 265–292). Dordrecht: Springer.
- Kaplan, D. M., & Bechtel, W. (2011). Dynamical models: An alternative or complement to mechanistic explanations? *Topics in Cognitive Science*, 3(2), 438–444.
- Kitcher, P. (1984). 1953 and all that. A tale of two sciences. *The Philosophical Review*, 93(3), 335–373.
- Lewis, D. (1986). Causal explanation. *Philosophical Papers*, 2, 214–240.
- Lewis, D. (2000). Causation as influence. *The Journal of Philosophy*, 97(4), 182–197.
- Machamer, P., Darden, L., & Craver, C. F. (2000). Thinking about mechanisms. *Philosophy of Science*, 67, 1–25.
- Mackie, J. L. (1974). *The cement of the universe*. Oxford: Clarendon Press.
- McGill, A. L., & Tenbrunsel, A. E. (2000). Mutability and propensity in causal selection. *Journal of Personality and Social Psychology*, 79(5), 677.
- Mill, J. S. (1882). *A system of logic, ratiocinative and inductive: Being a connected view of the principles of evidence and the methods of scientific investigation*. London: Longmans, Green.
- N'gbala, A., & Branscombe, N. R. (1995). Mental simulation and causal attribution: When simulating an event does not affect fault assignment. *Journal of Experimental Social Psychology*, 31(2), 139–162.
- Onalaja, A. O., & Claudio, L. (2000). Genetic susceptibility to lead poisoning. *Environmental Health Perspectives*, 108(Suppl 1), 23.
- Potochnik, A. (2010). Levels of explanation reconceived. *Philosophy of Science*, 77(1), 59–72.
- Press, J. (2015). Biological explanations as cursory covering law explanations. In P.-A. Braillard & C. Malaterre (Eds.), *Explanation in biology. An enquiry into the diversity of explanatory patterns in the life sciences* (pp. 367–385). Dordrecht: Springer.
- Railton, P. (1981). Probability, explanation, and information. *Synthese*, 48, 233–256.
- Rosenberg, A. (2006). *Darwinian reductionism, or, how to stop worrying and love molecular biology*. Chicago: University of Chicago Press.
- Rumberger, R. W., & Larson, K. A. (1998). Toward explaining differences in educational achievement among Mexican American language-minority students. *Sociology of Education*, 91, 68–92.
- Salmon, W. (1984). *Scientific explanation and the causal structure of the world*. Princeton: Princeton University Press.
- Schaffer, J. (2007). The metaphysics of causation. In E. N. Zalta (Ed.), *Stanford encyclopedia of philosophy*. Stanford, CA: The Metaphysics Research Lab.

- Schaffer, J. (2012). Causal contextualisms: Contrast, default, and model. In M. Blaauw (Ed.), *Contrastivism in philosophy*. New York, NY: Routledge.
- Seyfarth, R. M., Cheney, D. L., & Marler, P. (1980). Vervet monkey alarm calls: Semantic communication in a free-ranging primate. *Animal Behaviour*, 28(4), 1070–1094.
- Stegmann, U. (2012). Varieties of parity. *Biology and Philosophy*, 27(6), 903–918.
- Stotz, K. (2006). Molecular epigenesis: Distributed specificity as a break in the central dogma. *History and Philosophy of the Life Sciences*, 28(4), 533.
- Strevens, M. (2008). *Depth*. Cambridge, MA: Harvard University Press.
- Waters, C. K. (2007). Causes that make a difference. *Journal of Philosophy*, 104, 551–579.
- Weng, G., Bhalla, U. S., & Iyengar, R. (1999). Complexity in biological signaling systems. *Science*, 284, 92–96.
- Woodward, J. (2003). *Making things happen: A theory of causal explanation*. Oxford: Oxford University Press.