

Cohesive Proportionality

ABSTRACT

Proportionality — the idea that causes are neither too general nor too specific for their effects — seems to recommend implausibly disjunctive causes (McGrath 1998, Shapiro & Sober 2012, Franklin-Hall 2016). I argue that this problem should be avoided by appeal to the notion of cohesion. I propose an account of cohesion in terms of the similarity structure of property-spaces, argue that it is not objectionably mysterious, and that alternative approaches — based on naturalness, interventionism, and contrastivism — are inadequate without appeal to it. In an appendix, I show how my proposal can be perspicuously formalized by adapting structural equation models.

1. Preliminaries

1.1. Proportionality

Some events are more general than others. Suzy's throwing the rock is more general than Suzy's throwing the rock gleefully, and less general than Suzy's throwing something. This can be understood in terms of the idea that some events determine others (Yablo 1992a, 1992b): an event A is more general than an event B just in case B determines A (in the sense that B's occurring is a way for A to occur.)¹

When it comes to causation, generality seems to matter. Take Yablo's familiar example:

(SOPHIE) Sophie the pigeon is trained to peck at red stimuli. A scarlet tile is presented. She pecks.

¹ I use 'event' throughout as a placeholder for whatever the appropriate causal relata are (possibly including facts, tropes, *etc.*). I assume a fine-grained conception of events throughout this paper, since the proportionality constraint I focus on requires it. However, given a coarse-grained conception of events,

Intuitively, the tile's being red is a cause of Sophie's pecking. The tile's being scarlet is not a cause, because it is too specific; Sophie would still have pecked if she saw some other shade of red (crimson, say). Conversely, the tile's merely being colored is also not a cause, because it is too general; not any colored tile would have led to pecking.

Such examples suggest that causes obey some form of 'proportionality' principle: at a first pass, they involve all and only those details which are relevant to their effect. Various authors have defended different versions of this idea (e.g. Yablo 1992a, 2003; Strevens 2008; List & Menzies 2009). Fleshing it out requires characterizing both what proportionality is — that is, what it is for a cause to be proportional to its effect — and the role it plays with respect to causation.

Proportionality may be understood in various ways in terms of counterfactuals, probabilities, and/or laws. For concreteness, I focus on the following simple schematic characterization:²

c is proportional to e iff: i) c is causally sufficient for e, and ii) no event which is more general than c is causally sufficient for e.

For example, the tile's being colored is too general for pecking since it fails to satisfy the first clause, and the tile's being scarlet is too specific since it fails to satisfy the second. As a first pass, causal sufficiency may be understood in terms of c's entailing e in the relevant background circumstances via some causal law (à la Strevens 2008). Of course, developing an account along these lines (for example, extending it to cover non-deterministic cases of causation) is no trivial task. However, my aim in this paper is not to defend any particular characterization of proportionality, but to address a problem which arises for any characterization. Moreover, the problem I focus on pertains to the idea that causes are general, irrespective of how their sufficiency is spelled out.

proportionality could be recast as a constraint on which descriptions of events feature in causal explanations.

² This kind of approach is developed by Strevens 2004, 2008 and Weslake 2010. Mackie 1974, Putnam 1975, and Garfinkel 1981 are important precursors.

As for proportionality's role, the most straightforward idea is that proportionality is a necessary condition: if *c* causes *e*, then *c* is proportional to *e*.³ Again, I will focus on this proposal for concreteness. Putting it together with the schematic characterization of proportionality above yields the following principle:

PROPORTIONALITY: Causes are maximally general sufficers for their effects.

Now, many have thought that this 'winner-takes-all' approach is overly demanding: in many cases, two events both seem to cause the same effect even though one is more general than the other (Yablo 1992b: 420, Bontly 2005: 340, List & Menzies 2009: §5, McDonnell 2017: §5.1, Vaassen 2022: §3.2). For example, the window's shattering may be caused both by Suzy throwing the rock and by Suzy throwing something heavy.

Various less demanding proportionality principles have been suggested. Some approaches characterize proportionality in such a way that an event and its lower-level realizer (like Suzy's throwing something heavy and Suzy's throwing the rock) may each be proportional to the same effect (e.g. List & Menzies 2009, Woodward 2018). Another idea is that proportionality must be balanced against other factors, such as naturalness (Yablo 2003): causes are then those events which constitute local maxima, or whose overall eligibility exceeds some threshold (perhaps with respect to some degreed characterization of proportionality). Suzy's throwing the rock and Suzy's throwing something heavy may each constitute local maxima, or may each exceed the relevant threshold. Another approach is to relativize proportionality — and correspondingly, causation itself — to some contextually salient domain of contrast events (Yablo 1992b, Touborg 2022). On this view, Suzy's throwing the rock may be proportional relative to some default context, whilst invoking the more general event of Suzy's throwing something heavy brings it into relevant domain, making it proportional in the new context. (This would still be 'winner takes all' within any given context, though different events 'win' in different contexts.) Finally, proportionality might not constrain causation at all, but rather causal explanation: either explanations must cite proportional causes, or those which do are (*ceteris paribus*) better (Weslake 2010, Woodward

³ Although this idea is the most commonly discussed, it is rarely endorsed — Strevens 2004, 2008 and List & Menzies 2009 are prominent exceptions.

2018). On this view, Suzy's throwing the rock and Suzy's throwing something heavy are each causes of the window's shattering, though invoking one of them may be more explanatory (depending on how proportionality is characterized).

Although the simplicity of winner-takes-all is appealing, some combination of these alternative ideas may well prove to be on the right track. For the purposes of this paper, however, it doesn't matter exactly how proportionality constrains causation(/causal explanation). As discussed shortly, some version of the problem I address arises whichever role it plays, and my proposed solution carries over.⁴

Before considering this problem, it is worth emphasizing the attractiveness of the idea that proportionality (somehow) constrains causation.

Firstly, the kind of intuition elicited by SOPHIE is not merely a feature of curious philosophical cases but extremely general: it is naturally elicited by almost any ordinary example of causation. Our overall theory of causation should capture intuitions like these if it can.

⁴ Although the central points of this paper extend to other natural ways of developing proportionality, let me briefly explain why I don't think the case against the 'winner takes all' approach is clear-cut. First, there is a lot of flexibility in naming fine-grained events, so that apparent cases of 'causal sharing' may actually be cases in which a single cause is described in different ways. (See Maslen 2017: §4 and Vaassen 2022: §3.2 for relevant discussion.) In light of the cohesion requirement introduced below, this may apply especially in cases where the more general event fails to be cohesive, yielding meta-semantic pressure on the more general description to denote its cohesive realizer. For example, if *Suzy's throwing something heavy* fails to be cohesive, 'Suzy's throwing something heavy' might refer to *Suzy's throwing the rock* in the relevant context. (This could also happen by 'scoping out', where the logical form of the causal claim is: 'For some heavy x, Suzy's throwing x caused the window's shattering'.) Second, there are arguably (at least) two notions in the vicinity of the ordinary word 'cause': a difference-making notion, which requires proportionality, and a more permissive notion of causal sufficiency. Judgments of 'causal sharing' may simply be tracking the latter, especially since events which are merely causally sufficient may nonetheless feature in good explanations, generated by chaining together difference-making causation with a non-causal determination relation. For example, the tile's being scarlet may explain Sophie's pecking by virtue of determining the tile's being red. (This raises the possibility that the debate ultimately fails to be substantive, with the two sides focusing on different notions.) Ultimately, 'causal sharing' intuitions have provided an important challenge for proponents of proportionality ever since Yablo's (1992b) original discussion; I am assuming that there is some way of dealing with them (either by accommodating them or by explaining them away), and considering how we might deal with another problem which is at least as significant.

Secondly, as Yablo (1992a:274) notes, proportionality naturally falls out from an attractive general conception of causes as *difference-makers* for their effects (Woodward & Hitchcock 2003; Strevens 2004). Difference-making relations are highly sensitive to levels of generality: an event can be too specific or too general to be the true difference-maker.

Thirdly, and most importantly, proportionality promises a significant advance with two related and long-sought holy grails in the philosophy of mind and the philosophy of science: understanding mental causation, and understanding (more broadly) the goodness of high-level scientific explanations. Indeed, Yablo (1992a) originally formulated his proportionality principle as a response to Kim's (1989) exclusion problem. The rough idea is that determinable mental events are eligible to cause behavior, despite the causal sufficiency of their more specific physical determinates, in virtue of being proportional to behavior. More generally, proportionality illuminates the distinctive goodness of the causal explanations provided by high-level sciences, despite the apparent 'causal completeness' of fundamental physics: these sciences explain high-level phenomena by citing causes that are proportional to them.

1.2 The Disjunction Problem

PROPORTIONALITY appears to recommend implausibly disjunctive causes (McGrath 1998: 171-3, Shapiro & Sober 2012: 90-1, Franklin-Hall 2016: 566-7). Consider:

(DISJUNCTIVE SOPHIE) Sophie the pigeon is trained to peck at red stimuli and at green stimuli. A scarlet tile is presented. She pecks.

It seems that PROPORTIONALITY rules out the tile's being red (call this event '*R*') as a cause of Sophie's pecking in DISJUNCTIVE SOPHIE. For the tile's being red-or-green (*R-or-G*) is a sufficer for pecking which is more general than *R*. Hence, *R* is not a maximally general sufficer. Yet, intuitively, *R* should be a cause of Sophie's pecking (no matter what we think about *R-or-G* itself).

And it gets worse: even in the case of SOPHIE, *R* seems too specific to be a cause, since there are surely alternative possible sufficers for Sophie's pecking. For example, Franklin-Hall (2016:566)

imagines that Sophie pecks whenever her chin is tickled. Now consider the event of the tile's being red or Sophie's chin being tickled (*R-or-T*). Assuming that causal sufficiency is closed under disjunction, this disjunctive event suffices for pecking — thus, again, *R* is not a maximally general sufficer.⁵ Yet, whatever we think about the causal credentials of *R-or-T*, it seems absurd to claim that its sufficiency excludes *R* from being a cause of pecking.

More broadly, PROPORTIONALITY seems to recommend as eligible cause the event constructed by disjoining all possible sufficers for the effect, since this event will always be the maximally general sufficer.⁶ This casts severe doubt on the potential benefits of PROPORTIONALITY. It renders PROPORTIONALITY unable to capture the motivating intuitions, recommending instead highly counter-intuitive causes. And it no longer seems to support the high-level causes cited by psychology and other special sciences, since these causes are not similarly repugnant disjunctions.

On less demanding versions of proportionality — those which do not embody a 'winner-takes-all' approach — one aspect of the disjunction problem may be avoided: the proportionality of repugnant disjunctions does not automatically exclude the proportionality of intuitive causes. Nonetheless, an important aspect of the problem remains: why aren't disjunctive candidates like *R-or-G* and *R-or-T* proportional — or, if they are, what distinguishes them from *bona fide* high-level causes?⁷ If proportionality is balanced against other factors, then the difficulty is to provide

⁵ There may be accounts on which causal sufficiency is not closed under disjunction: for example, if entailment by a single causal law were required. If a suitable account of 'causal law' could be provided, this would effectively be one way of imposing the kind of cohesion requirement I recommend below.

⁶ Some sufficers, e.g. Suzy's throwing the rock and the rock's striking the window, share a path of causal sufficiency. Prima facie, PROPORTIONALITY recommends disjoining these too, but perhaps this can be avoided by stipulating that determination only relates simultaneous events. In any case, I won't pursue this complication here: the challenge posed by events like *R-or-G* and *R-or-T* is different.

⁷ The consensus appears to follow Lewis 1986a: §VIII in banning disjunctive causes. Sartorio (2006) argues against a blanket ban, based on special 'switching' cases; however, this argument does not cover candidates like *R-or-G* and *R-or-T*.

A referee has pointed out that 'C1 or C2 caused E' may sound bad for Gricean reasons when all parties know that C2 did not cause E, and the disjunction is interpreted as taking wide scope (so that the claim is equivalent to 'C1 caused E or C2 caused E'). However, this does not explain the badness of all disjunctive causal claims, since (I take it) the badness remains in cases where the disjunction clearly takes narrow scope (as in: "I know that the tile was red; nonetheless, the cause of her pecking was the tile's being red-or-green), and in cases where it is an open question whether C1 or C2 caused E (as in: "I don't know whether the tile was red or green, but I know that its being red-or-green caused the pecking".)

a factor which outweighs the greater proportionality of disjunctive candidates. (In §4.1 I explain why I don't think naturalness is the required factor.) If proportionality is relativized to a domain, then the difficulty is to explain why disjunctive candidates are excluded from the relevant domain (at least in typical contexts). If proportionality constrains causal explanation rather than causation, the difficulty is to explain why invoking disjunctive candidates is not more explanatory. We can confidently expect that, on any approach, disjunctions will need to be dealt with somehow.

1.3 Cohesion

Some ways of generalizing events (e.g. from the tile's being scarlet, *S*, to *R*) seem intuitively acceptable whereas others (e.g. from *R* to *R-or-G*, and from *R* to *R-or-T*) do not. This motivates a formulation of PROPORTIONALITY which favors *R* over both *S* and *R-or-G* in DISJUNCTIVE SOPHIE, and over both *S* and *R-or-T* in SOPHIE.

A natural approach is to distinguish *cohesive* from *disjunctive* events: the acceptable ways of generalizing are those which preserve cohesion (intuitively, those which don't yield repugnantly disjunctive events).⁸ We can then reformulate PROPORTIONALITY by restricting it to cohesive events:⁹

COHESIVE PROPORTIONALITY: Causes are *maximal cohesive* sufficers for their effects.

That is: i) causes are cohesive sufficers for their effects, and ii) causes cannot be 'cohesively extended' into more general cohesive sufficers for their effects.

Thus, on the assumption that *R* is itself cohesive, and that *R-or-G*, *R-or-T*, and any other sufficers for pecking that are more general than *R* fail to be cohesive, *R* is a maximal cohesive sufficer in

⁸ A similar notion of cohesion is employed by Strevens (2004, 2008) in his 'kairetic' account of causal explanation. However, my own understanding of the notion is somewhat different from his (see fn.18).

⁹ One might think, with Lewis (1986a: §VIII), that *all* events are cohesive – disjunctive 'events' like *R-or-G* are spurious. If so, tweaking PROPORTIONALITY makes explicit a restriction already mandated by the metaphysics of events.

both SOPHIE and DISJUNCTIVE SOPHIE. Hence, COHESIVE PROPORTIONALITY does not exclude R as a cause of pecking.

Approaches which do not embody a ‘winner takes all’ assumption may also employ the notion of cohesion to address the disjunction problem. For example, one view is that causes (or those which are invoked by causal explanations) are sufficers which best balance cohesion and generality (Strevens 2004:172), or whose balance of cohesion and generality exceeds a certain threshold. Another view is that non-cohesive events are automatically excluded from the relevant domain of contrasts.

Vindicating the required assumptions about which events are cohesive requires some account of cohesion, to which I now turn. I will sketch the idea that cohesion arises from the similarity structure of property-spaces. The point of this paper is not to defend this account in detail, but to argue that the right way to solve the disjunction problem is to employ the notion of cohesion. For this purpose, it suffices to make it plausible firstly that cohesion *can* be accounted for (in a way that secures the needed verdicts), and secondly that solutions which do without cohesion are inadequate. The rest of the paper takes up these tasks in turn.¹⁰

2. Accounting for Cohesion

2.1 The general approach

Providing an account of the relevant notion of cohesion is no easy task. At least in the first instance, ‘disjunctiveness’ seems to apply to representation rather than reality. But accounting for cohesion in terms of representation is a non-starter. Any event has some non-disjunctive description: letting ‘gred’ denote the property of being either red or green, $R\text{-or-}G$ may be described as the tile’s being gred. Moreover, appealing to ‘ordinary’ language won’t help: *the tile's being either red or orange* seems to be a cohesive generalization of R despite its disjunctive ordinary description.

¹⁰ I develop and defend my account of cohesion in more depth elsewhere (‘Cohesion’, in preparation).

It is tempting to invoke the idea that events (and/or the properties they involve) have some metaphysically privileged description, which reveals their disjunctive/cohesive nature. For example, perhaps they have a disjunctive definition in ‘more natural’ terms (Langton & Lewis 1998), or a disjunctive ‘essence’ (Skiles 2016).¹¹ However, ‘metaphysical’ forms of disjunctiveness may not coincide with the ‘scientific’ (or ‘nomic’) kind at issue here. Many apparently cohesive events, like *R*, are multiply realizable and/or involve determinable properties. We might expect the privileged descriptions of multiply realizable events to be vast disjunctions of possible fundamental realizers. Similarly, on some accounts, determinable properties like redness are disjunctions of their determinates, making the events involving them metaphysically disjunctive (Bigelow & Pargetter 1990, Clapp 2001). Hence, it is unclear that this metaphysical approach draws the required distinction.

My account is consistent with the metaphysical approach, but avoids any contentious assumptions about the metaphysical analysis of high-level events. As I understand it, the cohesion of events arises from the similarity structure of property-spaces. Possible events vary along well-defined dimensions corresponding to determinable properties.¹² For example, the event *R* in SOPHIE can be compared to possible events which vary with respect to color, such as the tile’s being green, the tile’s being scarlet (*S*), and the tile’s being red-or-green (*R-or-G*). Likewise, the event of Jones’s greeting Smith loudly can be compared to possible events which vary with respect to volume, such as Jones’s greeting Smith quietly, Jones’s greeting Smith at 80 decibels, and Jones’s greeting Smith loudly-or-quietly. These dimensions of variation are causally relevant: the color of the presented stimulus makes a difference to whether Sophie pecks, and the volume of Jones’s greeting makes a difference to whether Smith replies.

When comparing possible events along dimensions like these, they are naturally represented as regions within a space whose structure is inherited from a corresponding property-space. For example, we can embed the possible events *R*, *S*, and *R-or-G* within a space structured like color

¹¹ These authors are concerned with disjunctive properties, but their accounts may be extended to events.

¹² Merely possible events are not strictly speaking events at all (in the sense of occurring particulars); they may be thought of as states of affairs instead. The idea that they correspond to properties in the manner described is most straightforward if they have properties as constituents. But I take the idea to be natural irrespective of background metaphysics, so long as they are fine-grained as required for proportionality.

space. The regions of color space correspond to color properties, such as being red, being scarlet, and being red-or-green. Color space's structure encodes the relations between color properties. These properties stand in generality relations — for example: scarlet is contained within red, and red within red-or-green. And they stand in similarity relations — for example: orange is between red and yellow; scarlet is closer to crimson than to cyan. A corresponding space of possible events inherits this structure. For example, S is contained within R , and S is closer to the tile's being crimson than to the tile's being cyan. (Likewise, we can embed possible events involving Jones greeting Smith which vary with respect to volume within a space whose structure is inherited from the structure of the space of volume properties, such as being loud and being 80 decibels.)¹³

Call a space of possible events whose structure is inherited from some corresponding property-space a *state-space*. The idea is then that state-space structure suffices to distinguish cohesive from disjunctive events within it: the cohesive events are those that correspond to 'cohesive regions' of the state-space.

Not all events seem to belong to a state-space, however. Consider the event of the tile's being red-or-round. Since redness and roundness belong to different property-spaces (color space and shape space), being red-or-round does not belong to any property-space. Or consider the event of the tile's being red or Sophie's being green. This event could not be taken to correspond to the red-or-green region of color space: what would distinguish it from the event of the tile's being green or Sophie's being red? In order, then, for an event to belong to a state-space, it must satisfy the following (somewhat rough) condition: it must correspond to the instantiation of some property(/relation) belonging to an appropriate property-space.

Let the events corresponding to cohesive regions in state-spaces be 'basic cohesive' events. Then we can say that an event is cohesive iff it is either basic cohesive or the conjunction of some basic cohesive events. Hence, there are two ways an event can be disjunctive. Events which correspond to a region in a state-space can be disjunctive by corresponding to a non-cohesive region (like R -or- G). And events which do not correspond to any region in a state-space (like the tile's being red-

¹³ For a detailed discussion of property-spaces, see Funkhouser 2006.

or-round) are disjunctive unless they are the conjunction of some basic cohesive events. *R-or-T* is disjunctive in the latter way: it does not belong to any state-space, nor is it a conjunction of basic cohesive events.¹⁴

An adequate development of this approach must answer three questions:

- i) How does the similarity structure of a property-space determine cohesion?
- ii) How is the similarity structure of property-spaces determined?
- iii) How are properties divided into property-spaces?

The rest of this section briefly takes up these questions. Providing comprehensive answers is well beyond my scope here; instead, my discussion will aim to indicate some available resources and the direction in which I expect the answers to lie.¹⁵

Before proceeding, it is worth noting that this approach focuses exclusively on the properties involved in causes. Nonetheless, it also has consequences for the involvement of objects. The basic cohesive events involve objects instantiating cohesive properties (presumably at times/time-intervals), with cohesive events formed from these by conjunction. Firstly, this disallows generalizing to object-free ‘existential events’, such as from *the tile is red* to *something is red*: such events are not conjunctions of basic cohesive events. Secondly, it allows for abstraction with respect to the object involved, such as from *the tile is red* to *the left half of the tile is red*.¹⁶

¹⁴ We might think of *R-or-T* as corresponding to the tile and Sophie standing in the unfamiliar relation of being such that the first is red or the second is chin-tickled. But this does not locate *R-or-T* within a state-space, since presumably this relation does not belong to any property-space.

¹⁵ In the appendix, I show how this understanding of cohesion may be formally implemented by adapting structural equation models. The basic idea is that instead of outfitting variables with bare sets of values as ranges, we should outfit them with structured state-spaces, and instead of assigning values to variables to represent events, we should assign them regions within these state-spaces.)

¹⁶ Which such abstractions are possible depends on i) the nature of the determination relation which holds between events and their generalizations, and ii) the metaphysics of objects. One choice point here is whether we place some restriction on the objects involved in cohesive events. A natural thought (at least, on a plenitudinous view) is that there are ‘disjunctive objects’ as well as disjunctive properties: for example, there is an object which is collocated with the fusion of Trump and Obama, and which (necessarily) exists just in case one of them exists.

2.2 Cohesion from similarity structure

How we define the notion of a cohesive region within a property-space precisely will depend on the particular structure of the space in question. But two principles serve as an initial guide. Firstly, the points of a space (corresponding to the most specific properties it represents) are cohesive. Secondly, the cohesion of a region is preserved by extending it ‘continuously’ (where the particular structure of the space determines the relevant notion of continuity.) Thus, the cohesive regions of a space are generated by continuously extending its points.

In continuous topological spaces, for example, cohesion is naturally captured by connectedness: intuitively, a region is connected when it does not consist of entirely separate regions.¹⁷ Take the space of real numbers. The connected regions in this space are all and only the intervals: for example, $(0,1)$ and $(2,3)$ are each connected regions, but their union is not. Color space is naturally conceived as a continuous topological space (sometimes depicted as a cylinder, with dimensions corresponding to hue, saturation, and brightness). Given this structure, intuitively cohesive color properties such as being scarlet and being red will count as cohesive (since they correspond to connected regions) and intuitively disjunctive color properties such as being red-or-green will count as disjunctive (since they correspond to disconnected regions). Hence, in the corresponding state-space of possible colored tile presentations in SOPHIE, S and R are both cohesive, and R -or- G is disjunctive.

The connectedness criterion is only a starting point. Something else is needed for discrete spaces, where only individual points are connected. And even in continuous spaces, we might want a degreed notion on which highly irregular and spread out regions (consisting, for example, of two islands connected by a narrow bridge) are to be counted as less cohesive. If the space has metric structure — a measure of the degree to which any two points within it are similar — then many measures of cohesion are definable. But nothing as rich as metric structure is needed: viable notions of cohesion are definable given only betweenness structure (which points of the space lie

¹⁷ More precisely: a topological space is connected iff it is not the union of disjoint open subsets; a subset is connected iff it is a connected space with respect to the subspace topology.

between which others) or graph-theoretic structure (which points of the space are directly connected i.e. nearest neighbors).¹⁸

By way of illustration, consider a discrete property-space representing a simplified spectrum of political persuasions. It has five points, representing far-left, center-left, center, center-right, and far-right. It seems implausible that this space has determinate metric structure: for example, is center-left more similar to far-left or to center? Nonetheless, it has betweenness structure — center-left is between far-left and center — and graph-theoretic structure — far-left is directly connected to center-left but not to center. This allows us to define a region as cohesive iff: i) it contains any point which lies between two of its points (corresponding to ‘convexity’); or ii) every point within it is (directly or indirectly) connected to every other (corresponding to ‘path-connectedness’). On either of these notions, for example, being far-left-or-center-left is cohesive, whereas being far-left-or-center is not.

Ultimately, I don’t expect there to be a single correct way of measuring cohesion, nor for there to always be determinate verdicts about whether, or to what extent, an event is cohesive. What I do expect — and what I think solving the disjunction problem requires — is that whenever an event is intuitively too disjunctive to be a cause, there is a suitable notion of cohesion which (when plugged into the approach I have outlined) delivers this verdict. To the extent that cohesion turns out to be indeterminate and/or relative to a context or background model, we should expect causal claims to inherit these features.

2.3 Similarity structure and property-spaces

This approach to cohesion relies on similarity structure pertaining to high-level determinable properties such as color, volume, etc. This structure is sparse in that it only pertains to an elite class of high-level properties: arbitrary gerrymandered properties do not belong to the relevant kind of property-space. It is also autonomous in that it does not flow directly from the structure of

¹⁸ Graph theory provides interesting degreed notions which measure, for example, how clustered the nodes of a set are.

fundamental physics: given the multiple realizability of the relevant high-level properties, it is unlikely that fundamental physics alone distinguishes cohesive from disjunctive properties.¹⁹

By relying on sparse, autonomous high-level similarity structure, my approach acquires the burden of accounting for this structure, and I cannot give an account of it here. However, as I argue in §3, it is very plausible that there is such structure; hence, this burden should be viewed as a research project rather than a problem. Moreover, it is important to emphasize that cohesion facts do not require particularly rich similarity structure. We certainly don't need 'global' facts about the relative similarity of utterly unrelated properties (is redness more similar to squareness than to loudness?!) All we need instead are facts about which properties belong to the same property-space (redness and blueness do, redness and squareness do not), and 'local' facts about the structure within a given space. As discussed above, facts about 'continuity' within a given property-space, as captured for example by topological and/or graph-theoretic features, suffice to exclude many properties as disjunctive.

If we wish to avoid taking property-spaces and their similarity structure as primitive, there are promising avenues for reduction to explore. To my mind, the most plausible approach is a functionalist one, according to which property-spaces are both populated and structured according to the nomic roles of their properties. On this approach, properties belong to the same property-space in virtue of sharing a schematic nomic role. For example, colors are unified by their role in robust regularities characterizing their interaction with light and visual systems, and political persuasions are unified by their role in robust regularities connecting them to voting behavior and to socio-economic factors. The fruitfulness of sciences dedicated to uncovering these nomic roles is testament to this underlying unification.

¹⁹ My approach is in the spirit of Weatherson's (2012:472) view that "what makes [high-level explanations] good is not the cohesiveness of their underlying physical mechanisms. It is, at least intuitively, the cohesiveness of the explanations from the perspective of the special science in question." By contrast, Strevens (2008, §5.4; 2012) defends a notion of cohesion based on the structure of fundamental physical possibility-space. Given multiple realizability, I am skeptical that this basis is rich enough to define the required notion, but this isn't the place to argue the point in detail.

Another significant difference between Strevens's account and my own is that he applies cohesion to trajectories from causes to effects, rather than to causes themselves. I cannot discuss this interesting proposal here, but I take it to be compatible with my general strategy of accounting for cohesion in terms of similarity structure.

Nomic roles don't only unify properties into property-spaces: they also unify the various instances of the properties themselves. For example, red objects interact with light and visual systems in similar ways, whilst red-leaning people tend to share voting behaviors and socio-economic backgrounds. This suggests that similarity structure within property-spaces should also be determined by nomic role. Roughly speaking, the similarity between red and orange is determined by the similarity in the ways their instances interact with light and visual systems, and the similarity between staunch conservativeness and moderate conservativeness is determined by the similarity in the voting behaviors and socio-economic backgrounds of their instances.²⁰

3. In Defense of Cohesion

In this section, I defend my proposed solution to the disjunction problem by arguing that cohesion is not objectionably mysterious: there are good reasons to think that there are cohesion facts, and that we have epistemic access to them when evaluating causal relations.

3.1 Cohesion isn't Naturalness

It is sometimes suggested that the disjunction problem may be solved by appeal to *naturalness*: the basic idea is that disjunctive events like *R-or-G* and *R-or-T* are excluded because they are (in a sense which may be fleshed out in various ways) unnatural. Indeed, I suspect that reactions to the disjunction problem more or less divide between two camps: naturalness-users, who think that it is straightforwardly solved by appeal to naturalness, and naturalness-sceptics, who think that appeal to any 'metaphysical' posit like naturalness is methodologically suspect. I address naturalness-users below, arguing that cohesion is a better tool for the job (§4.1). But first, I should

²⁰ How is the similarity structure in these nomically connected property-spaces itself determined? In the case of color space, for example, what determines similarity structure in the space of reflected wavelengths, or in the space of visual system states? On a 'coherentist' approach, the similarity structure of high-level property-spaces is determined all at once (perhaps via the best summary of the nomic connections between them). On a 'foundationalist' approach, similarity structure percolates upwards from the basic structure pertaining to microphysical properties (so that microphysical similarity plays an indirect, nomically mediated role in determining high-level similarity). Ultimately, some combination of these two approaches may be required.

address naturalness-sceptics — they are likely to find my appeal to cohesion instead of naturalness little more than a rebranding exercise.

The first thing to note is that cohesion and naturalness are distinct. Those who recognize naturalness facts can think of cohesion as one aspect of, or a contributing factor to, naturalness. It is standard to suppose that the disjunctiveness of a property detracts from its naturalness. For example, *being-green-or-blue* is less natural than *being-green*. There is more to naturalness than cohesion, however: conjunctions and negations are also thought to detract from naturalness. For example, although they both seem cohesive, *having-charge-and-mass* and *not-having-charge* are each less natural than *having-charge*. Whereas conjoining natural properties detracts from naturalness, conjoining cohesive properties preserves cohesion.

The question is whether the reasons to be skeptical of naturalness extend to cohesion. Franklin-Hall (2016) and Blanchard (2020) note that resorting to naturalness will strike some defenders of proportionality as an overly ‘metaphysical’ solution, where this appears to mean: offends against empiricist tendencies in the philosophy of science. To my mind, such concerns are misplaced: *pace* ostrich empiricism, avoiding any appeal to naturalness (or similar resources) in our theorizing about science is neither sustainable nor desirable.²¹ But I won’t press the point here — there is a genuine challenge for the naturalness-user in the vicinity: how are we to account for naturalness (especially, high-level naturalness) in a way which vindicates its alleged role in science? In particular, how does naturalness get its grip on theorizing about causation? In the absence of a satisfying answer, appealing to naturalness seems objectionably mysterious and *ad hoc*.²²

The question here is whether the analogous challenge can be met in the case of cohesion. I will now argue that it can be: the similarity structure of property-spaces provides a firm foundation for

²¹ It’s hard to know why naturalness should be regarded as more ‘metaphysical’ than counterfactuals, laws, or other modal notions, which are not obviously empiricist-friendly (and indeed have been thought to require ‘naturalness-infected’ analyses).

²² In this vein, Blanchard (2020:640) complains that “it is unclear how we could ever get epistemic access to natural properties and what entitles us to regard current scientific theories as latching onto them”, and Franklin-Hall (2016:574) offers the reservation that positing naturalness seems close to the mere “science-mimicry” which simply claims that “the best explanations exploit variables representing just the features to which scientists themselves appeal”.

cohesion in scientifically respectable resources. Hence, even naturalness-sceptics should be happy to appeal to it.²³

3.2 Similarity Structure isn't Mysterious

In §2, I argued that the cohesion facts needed to implement COHESIVE PROPORTIONALITY can be accounted for in terms of similarity structure on property-spaces. I will now argue that it should be uncontroversial that there is such structure, and that we have epistemic access to it. Hence, cohesion is in good metaphysical and epistemic standing.

Structured property-spaces play a crucial role in special sciences: they are posited by the sciences themselves, and exploited in discovering and describing simple generalizations. For example, the rationality of ubiquitous extrapolation practices — in which simple curves are fitted to finite data samples — crucially relies on high-level similarity structure: without it, there is nothing with respect to which these curves are genuinely simple.

For a particularly vivid illustration, consider neuronal tuning curves, widely used in psychophysics for characterizing how neurons' firing rates vary with environmental features — such as the orientation, shape, or color of a visual stimulus.²⁴ The extrapolation of tuning curves from data relies on a parametrization representing the relevant features of the stimulus. For example, we parametrize stimulus orientation by using numbers to represent angles, allowing us to mathematically describe its relationship to neural activity. This parametrization is chosen to reflect the relevant high-level similarity structure.

²³ Another kind of sceptic about naturalness holds that nothing could play *all* the roles that it has been alleged to play (Dorr & Hawthorne 2013). Cohesion avoids this worry: it can be thought of as the aspect of naturalness responsible for a limited and consistent portion of its roles (those that cluster around making for similarity).

²⁴ Butts & Goldman (2006:639) write: “Tuning curves have provided the first-order description of virtually every sensory system, from orientation columns in the vertebrate visual cortex [responding to the orientation of visual stimuli], to place cells in the hippocampus [responding to the organism's spatial location] and wind-detecting neurons in the cricket cercal system.”

Without a structured color space, for example, psychophysicists would be unable to describe how the various neurons in Sophie's visual cortex respond to the color of the presented stimulus (e.g. these neurons respond most strongly to this shade of red). Take some set of points encoding how stimulus color and firing rates covary across some number of trials. This data is evidence for a particular relationship between color and neural firing, as described by a neuronal tuning curve. But an infinite number of curves fit the data, so it can only be evidence for a particular relationship given some way of comparing the simplicity of various alternatives. This comparison relies on some background structure embedding the properties in question (as encoded by the axes used): we can make curves more or less complex by choosing suitably gerrymandered axes, which do not reflect the real similarity structure.

This particular illustration is especially powerful since it covers the rich variety of environmental features to which neurons are sensitive. But the general point extends to all high-level sciences which investigate the nomic relationships between properties — think of temperature vs reaction rates in chemistry, predator vs prey populations in biology, and supply vs demand in economics. Describing such relationships using graphs and equations relies on the high-level similarity structure of property-spaces. Moreover, the point extends beyond formal science to mundane inductive inferences: traffic vs arrival time, music volume vs anger of neighbor, time in oven vs state of dinner. The crude rough-and-ready estimations which enable us to navigate daily life can only be reliable if we track the similarity structure of the relevant properties, with respect to which certain hypotheses stand out from the many consistent with our evidence.

Thus, the vision behind COHESIVE PROPORTIONALITY is that our intuitive inferences about difference-making causation recruit the same evaluations of high-level similarity structure which are implicit in inductive inference more broadly. Our ability to routinely and reliably perform such inference provides excellent reason to think that we must be tracking the kind of structure which suffices to support cohesion facts. Even naturalness-sceptics should view this structure as unmysterious. If it turns out that cohesion cannot be accounted for without appeal to naturalness, the lesson is not that this structure is objectionable after all, but rather that naturalness-sceptics must abandon their scruples.

3.3 Cohesion is Useful

So far, I have argued that the cohesion of events can be accounted for in terms of scientifically respectable resources, and thus that cohesion is not mysterious (in the way that naturalness is alleged to be). Let me briefly add a further consideration supporting my appeal to cohesion: it is not an *ad hoc* notion wheeled in to rescue proportionality, but can (and in my view, should) be put to useful work elsewhere.

First, avoiding disjunctive explanations is an issue for *any* account of high-level explanation, whether tied to proportionality or not. It seems that something about the generality or abstractness of these explanations makes them good, where this must be distinguished from the spurious kind of generality associated with disjunctive explanations. Relatedly, avoiding disjunctive causes seems desirable for any account of causation, irrespective of one's attitude towards proportionality. Given an account in terms of counterfactuals, covering laws, or energy transfer, it is easy to construct disjunctive events from those counted as genuine causes, and hard to discount them from being causes themselves without cohesion.²⁵

Further afield, making sense of genuine similarity between objects seems to require distinguishing cohesive and disjunctive properties: sharing of disjunctive properties, like grueness and scarlet-or-cyanness, does not make for genuine similarity in the way that sharing of cohesive properties like redness does. This naturally extends to events: only cohesive events make for genuine similarity between the situations or worlds in which they occur. These similarity facts seem required in turn to make sense of counterfactuals: if we want to account for claims about what would have happened in counterfactual situations — which surely, for many reasons, we do — then we likely need to evaluate similarity between situations.

The upshot is that even an ardent naturalness-sceptic, with the least interest in defending proportionality, has good reason to recognize the notion of cohesion that suffices to insulate

²⁵ Lewis (1986a:267) and Yablo (2003:322) discuss this issue in the context of counterfactual theories.

proportionality from the disjunction problem. Independently of this particular application, accounting for cohesion is a worthwhile (and tractable) research project.

4. Doing without Cohesion?

This final section considers attempts to avoid the disjunction problem without cohesion, based on naturalness, interventionism, and contrastivism. I argue that none are adequate as they stand: to do the work of COHESIVE PROPORTIONALITY (or some similar principle), they must be supplemented by appeal to cohesion or the kind of similarity structure which, by my lights, underlies it.

4.1 Naturalness

As mentioned above, it is sometimes suggested that the disjunction problem may be solved by appeal to naturalness. Indeed, this was Yablo's own recommendation: "proportionality is not pursued at all costs but traded off against naturalness" (2003:326). The rough idea is that disjunctive events such as *R-or-G* and (presumably even more so) *R-or-T* are less natural, and this damages their eligibility as causes. So whilst proportionality pushes us towards general, and hence disjunctive, causes, naturalness pushes us back towards suitably non-disjunctive causes.

One obvious challenge for this proposal is that it requires an account of the naturalness of events (presumably via an account of the naturalness of corresponding properties) which vindicates the idea that intuitively disjunctive events are less natural. It is far from obvious how such an account would go. For example, Lewis (1986b:61) suggested that naturalness is given by length of definition in fundamental terms: the shorter the definition, the more natural the property. But this approach faces a serious challenge: multiple realizability makes it unlikely that naturalness neatly correlates with simplicity of definition. I don't want to lean on this particular criticism, however. Various alternative approaches are available to the naturalness-user.²⁶ Besides, I have not provided a complete account of cohesion, and such an account may ultimately rely on high-level naturalness facts itself.

²⁶ See Sider 2011:§7.11.1. It is worth flagging that drawing on one kind of resource — figuring in causal relations — threatens to be circular in this context.

So let's grant the naturalness-user some (suitably non-mysterious) account of high-level naturalness. Even so, this is the wrong tool for solving the disjunction problem. The reason is that naturalness pulls apart from cohesion, in both directions: relatively unnatural events may be suitably cohesive, and objectionably disjunctive events may be relatively natural.

The first kind of case arises when a sufficer may be generalized in a way which reduces naturalness but preserves cohesion. The comparative lack of naturalness does not compromise the more general event's eligibility to cause (as a comparative lack of cohesion would). For example, suppose a weighing scale flashes whenever something between 1kg and 2kg is placed on it. A 1.3kg mass is placed on the scale, and it flashes. The mass's being 1.3kg and the mass's being between 1kg and 2kg are each sufficient for the flash. Taking complexity of fundamental definition as a heuristic guide, the former event is plausibly more natural than the latter: it is defined by reference to a single plausibly perfectly natural property (being 1.3kg), whereas the latter is defined by reference to two such properties and the betweenness relation on masses. But this comparative lack of naturalness does not seem to tell against the latter's eligibility to be a cause. This is plausibly because it involves no loss of cohesion.²⁷

The second kind of case arises when an intuitively cohesive sufficer may be generalized in a way which enhances naturalness but destroys cohesion. The availability of a more general and more natural sufficer does not exclude the more specific event (as it would if the more general sufficer were cohesive). For example, suppose we discover that the property of being red-or-green is far more natural than expected: there are two fundamentally different kinds of photon, and all and only red-or-green things absorb the first kind of photon whilst reflecting the second kind.²⁸ Presumably, the newly apparent naturalness of this property would extend to the disjunctive event

²⁷ The naturalness-user might hope to accommodate such examples by arguing that the loss of naturalness is too small to outweigh the gain in generality, especially if conjunctions count less against naturalness than disjunctions. But firstly, the loss of naturalness doesn't seem to affect causal eligibility at all in this case. And secondly, in other cases the more general candidate may have a complex definition which does involve disjunctions (without being 'scientifically' disjunctive.) This will be so, for example, if determinable properties (like redness) are defined as disjunctions of their determinates; the latter would then be significantly more natural, but would not thereby make the events involving them more eligible causes.

²⁸ This example is from Gómez Sánchez (forthcoming).

R-or-G. But this wouldn't seem to make it any more palatable as a cause of Sophie's pecking: it seems that *R-or-G* remains disjunctive in the relevant sense, and hence *R* remains an eligible cause.²⁹

Thus, relatively unnatural events may be eligible causes, if they are cohesive, and relatively natural events may be ineligible, if they fail to be cohesive. Whilst disjunctive events tend to be unnatural, it is their disjunctiveness, not their unnaturalness, which prevents them from being causes.

4.2 Interventionism

According to interventionism, causal explanation aims at 'identifying interventions that would have changed the explanandum' (Blanchard 2020:634). Woodward (2018) and Blanchard (2020) have recently argued that interventionism yields a 'metaphysically lightweight' understanding of proportionality which avoids the disjunction problem. The basic idea is that explaining Sophie's pecking in terms of *R* (rather than more specific or more general events) best identifies those interventions which would have prevented pecking: namely, interventions on the tile's redness.

Interventionists use 'structural equation models' to represent causation (discussed in more detail in the appendix). These models represent events as variables taking values, and causal relations via equations connecting the variables. For example, a simple model of SOPHIE has three binary variables: *RED*, with values 1/0 representing *R*/ the tile's being some color other than red; *TICKLE*, with values 1/0 representing Sophie's chin being tickled/not being tickled; and *PECK*, with values 1/0 representing Sophie's pecking/not pecking. The equation ' $PECK = \max(RED, TICKLE)$ ' represents the fact that Sophie pecks just in case the tile is red or her chin is tickled.

²⁹ Won't *R-or-G* be cohesive relative to some more fundamental property-space, whose properties concern reflectance of photon-kinds? Perhaps, but given that Sophie's perceptual system is sensitive to color and not photon-kind, it seems to be cohesion with respect to color space (not the more fundamental space) which is relevant to causal eligibility in this case. Besides, we can modify the example so that *R-or-G* remains disjunctive relative to the more fundamental property-space too: suppose that the property of being red-or-green amounts to reflecting either photons of kind 1 or kind 2 (where there is no comparatively simple definition of redness in fundamental terms).

Intuitively, citing disjunctive events does a poor job of identifying explanandum-changing interventions. For example, suppose we combine the binary variables *RED* and *TICKLE* into a single binary variable *LUMP*, taking value $\max(\text{RED}, \text{TICKLE})$. Explaining Sophie's pecking in terms of *R-or-T*, represented by $LUMP = 1$, fails to pin down explanandum-changing interventions: if the tile is red, then the way to prevent pecking is to change the tile's color, but if Sophie's chin is tickled, then it is to stop the tickling.

This suggests that $LUMP = 1$ is too general to explain Sophie's pecking because it is equivalent to a disjunction, $RED = 1$ or $TICKLE = 1$, only one of whose disjuncts ($RED = 1$) is relevant to the actual pecking.³⁰ We can cash out *relevance* as follows: a variable *X* is relevant to an event $Y = y$ just in case there is some non-actual value of *X* such that intervening to set *X* to that value (and doing nothing else) would prevent $Y = y$.³¹ For example, in *SOPHIE*, *RED* is relevant to $PECK = 1$ since setting *RED* to 0 and doing nothing else would prevent pecking, whereas *TICKLE* is not relevant since intervening to set it to its only non-actual value, 1 (i.e. tickling Sophie's chin) would overdetermine pecking rather than preventing it.

We can then capture the interventionist idea as follows:

SPECIFICITY: $X = x$ is too general to explain $Y = y$ just in case $X = x$ is equivalent to some disjunction $X_1 = x_1$ or $X_2 = x_2$... or $X_n = x_n$, where X_1 is relevant to $Y = y$ but X_2, \dots, X_n are not.³²

³⁰ The following proposal is adapted from Blanchard (2020: §4). Woodward (2018) offers a distinct but complementary proposal, focused on type-level causation, which relies on the claim that 'there must be some basis for decisions about when it is preferable to represent a causal structure by means of distinct variables and when it is permissible (or a good strategy) to ... collapse these into a single variable'. I agree, and would add that cohesion is an important part of this basis.

³¹ In Blanchard's terminology, *X* is a 'locus of explanandum-changing interventions'.

³² Where X_1, \dots, X_n are independently manipulable variables i.e. each n-tuple of their values is metaphysically possible. Otherwise, $RED = 1$ would be too general for the explanandum $PECK = 1$ in virtue of being equivalent to the disjunction $SCARLET = 1$ or $RED^* = 1$, where $RED^* = 1$ iff the stimulus is red but not scarlet, and 0 otherwise.

However, SPECIFICITY overshoots: it only excludes disjunctive causes like *R-or-T* at the cost of excluding intuitively cohesive causes like *R*. The problem is that even cohesive events are equivalent to many disjunctions. For example, $RED = 1$ is equivalent to the disjunction $X_1 = 1$ or $X_2 = 1$, where

$X_1 = 1$ iff $RED = 1$ & it's raining in Cambridge; 0 otherwise;

$X_2 = 1$ iff $RED = 1$ & it's not raining hard in Cambridge; 0 otherwise.³³

Now imagine that, while the experiment in SOPHIE takes place, it's raining hard in Cambridge. X_1 is relevant to pecking: intervening to change its value from 1 to 0 (and doing nothing else) involves removing the red tile, which would prevent pecking. But X_2 is not relevant to pecking: its actual value is 0, and setting it to 1 would not prevent pecking. Hence, by SPECIFICITY's lights, just as citing *R-or-T* fails to identify whether *RED* or *TICKLE* is relevant to pecking, citing *R* fails to identify whether X_1 or X_2 is relevant. A more appropriately specific explanans for Sophie's pecking must mention the Cambridge weather!

It might be objected that X_1 is not relevant since some (rather convoluted) ways to intervene on it prevent the rain in Cambridge, and this kind of intervention would not prevent pecking. But they involve 'doing something else', in addition to intervening on X_1 : they would also change the value of X_2 , from 0 to 1. Indeed, the situation of *RED* vis-à-vis X_1 and X_2 perfectly parallels that of *LUMP* vis-à-vis *RED* and *TICKLE*. Intervening on *RED* to change its value from 1 to 0 (and doing nothing else) will prevent pecking. But there are (convoluted) interventions on *RED* which would not prevent pecking: changing the stimulus color whilst tickling Sophie's chin. These do not count against *RED*'s relevance, since they involve doing something else: changing the value of *TICKLE*, from 0 to 1.

Of course, there is an intuitive difference here: changing *TICKLE* from 0 to 1 in the course of changing *RED* from 1 to 0 seems like genuinely 'doing something else', whereas changing X_2 from 0 to 1 in the course of changing X_1 from 1 to 0 does not. But the challenge lies in spelling out the

³³ X_1 and X_2 are independently manipulable: all four combinations of their values are metaphysically possible.

requisite notion of *something else* (without relying on implicit appeal to cohesion or the like). If we were willing to rest on intuitive differences, we could have excluded $LUMP = 1$ from the start for being intuitively disjunctive!

The problem as I see it lies in allowing variables like X_1 and X_2 with values representing disjunctive events. For example, $X_2 = 0$ represents the event of either the tile not being red or its raining hard in Cambridge. Unsurprisingly, interventionism goes awry when applied to such variables, since the corresponding ‘interventions’ are not appropriately targeted: there are genuinely distinct ways of bringing about a disjunctive event. Although my discussion has focused on one proposal, the point clearly generalizes: it is difficult to see how interventionism itself could address the disjunction problem without relying (perhaps implicitly) on some way of restricting variables.

4.3 Contrastivism

Contrastivism is the view that causation is a four-place relation involving merely possible events as causal and effectual contrasts (e.g. Schaffer 2005). For example, in the case of SOPHIE, the following contrastive claims hold:

- (i) The tile’s being red rather than not red caused pecking rather than not pecking.
- (ii) The tile’s being scarlet rather than crimson did not cause pecking rather than not pecking.
- (iii) The tile’s being scarlet rather than cyan caused pecking rather than not pecking.
- (iv) The tile’s being red rather than Sophie’s chin being tickled did not cause pecking rather than not pecking.

Shapiro & Sober (2012) contend that the disjunction problem shows that proportionality is misguided, and that the intuitions which supported proportionality in the first place may be explained away using contrastivism. In SOPHIE, for example, we judge R to cause pecking in contexts where (i) is the salient contrastive claim, and we judge S not to cause pecking in contexts where (ii) is the salient contrastive claim. However, relative to other possible contrasts, such as those in (iii) and (iv), S is a cause and R is not. Contra proportionality, then, difference-making is

not simply a matter of how cause relates to effect, but rather a matter of the difference between the cause and some salient contrast(s) making the difference between the effect and some salient contrast(s).

Unlike the other two proposals I have considered, the idea here is not to cure disjunctivitis but to prevent it: instead of providing some countervailing pressure against grotesquely disjunctive causal claims (via naturalness or interventionist specificity), this proposal rejects the idea which created the pressure to prefer such claims in the first place. Now, this is not directly a challenge to my main thesis that proponents of proportionality should solve the disjunction problem by appeal to cohesion, since I am presupposing that some version of the proportionality idea is worth holding onto. However, if contrastivism shows that this presupposition is untenable, that would certainly make my thesis less interesting. Thus, I want to close by suggesting that, even if we adopt a contrastivist framework, the motivations for proportionality — and hence, the need to solve the disjunction problem — remain.

First, a dialectical point: avoiding disjunctive causes is everyone's problem. Granted, the issue is especially serious for proponents of proportionality (at least on a winner-takes-all approach): disjunctive causes threaten to exclude ordinary causes. But this more serious version of the problem may be solved the same way: whatever excludes disjunctive causes also reinstates ordinary causes. Moreover, for contrastivists who wish to avoid disjunctive causes, cohesion remains a valuable resource. It would be odd, then, to take the disjunction problem to motivate abandoning proportionality in favor of a view on which a closely related problem arises and is naturally solved in the same way.

Set this aside, however: let's assume that there are good reasons to adopt contrastivism, and consider whether this undermines proportionality in itself. It is important to note that, since proportionality is a constraint on a two-place relation, it is not obvious how it should even be formulated in a contrastivist framework. On a 'linguistic' approach, we might tie proportionality directly to causal claims, as the thesis that 'c causes e' is only true when c is proportional to e. However, on a natural implementation of contrastivism, 'c causes e' is true in a context C iff for all causal and effectual contrasts c^* and e^* which are salient in C, c rather than c^* causes e rather

than e^* . I grant that, as Shapiro and Sober maintain, there are contexts in which non-proportional causal claims are true (given this contrastivist semantics). For example, there are contexts in which ‘ S caused pecking’ is true, since the only salient contrastive claim is (iii) (e.g. because the question at issue is whether pecking is caused by scarlet or cyan tiles).³⁴

However, this leaves open another (more interesting) formulation of the proportionality thesis, which ties it to causal theorizing rather than ordinary language. The idea is that there is a natural, theoretically important two-place relation, in the vicinity of ordinary causal claims, which respects proportionality. In particular, there is the relation defined in terms of the objective similarity structure of corresponding state-spaces as follows: c proportionally causes e iff, for all the *nearest* contrasts c^* , c rather than c^* causes e rather than $\sim e$.³⁵ This relation deserves its name since, if c ‘proportionally causes’ e , then c is a sufficer for e which cannot be cohesively extended into a more general sufficer (since all ways of cohesively extending it would include some nearest alternative).³⁶ Thus, assuming a cohesion requirement on contrastive causation, proportional causes are maximal cohesive sufficers.

Shapiro and Sober may be ‘radical contrastivists’, who hold that the only theoretically important relation is four-place, and there is no such two-place relation corresponding to any objective restriction on contrasts. However, the motivations for a proportionality constraint in the orthodox setting also motivate a more moderate contrastivist view which recognizes proportional causation.

First, the intuitions about cases like SOPHIE show that, at least in typical contexts, nearby alternatives are salient: it is the salience of *the tile’s being crimson* which drives the intuition that S is not a cause. This suggests that proportional causation lies in the vicinity of our ordinary causal

³⁴ Proportionality may nonetheless be vindicated if it constrains causal claims in ‘default’ or ‘typical’ contexts. However, if so, this is presumably because there is some underlying proportional relation which these claims latch onto in such contexts. Hence, I prefer the more direct ‘theoretical’ approach pursued in the main text.

³⁵ There is also the following notion: c covaries with e iff c rather than $\sim c$ causes e rather than $\sim e$. If covariation is analyzed counterfactually following List & Menzies (2009), it yields their characterization of proportionality.

³⁶ I assume that if c rather than c^* causes e rather than $\sim e$, then c suffices for e and c^* does not.

thinking. Second, the popular idea that causes are difference-makers is vindicated by the failure of nearby alternatives to yield the effect, in a way that it is not by the failure of faraway alternatives. Third, and related to the last point, recognizing proportional causation allows contrastivists to capture the objective goodness of high-level explanations. Being true with respect to any old contrasts would not suffice to make the causal claims provided by special sciences explanatory; being true with respect to nearby contrasts in corresponding state-spaces does. Objective similarity-structure supports the explanatory credentials of high-level causes (for example, by privileging (i) over (iv)) whilst undermining those of their low-level realizers (for example, by privileging (ii) over (iii)).

There is, then, good reason for contrastivists to recognize proportional causation: doing so allows them to explain our intuitions, vindicate the difference-making idea, and capture the objective goodness of high-level explanations. Moreover, since similarity-structure may be outweighed or overridden by more ‘subjective’, context-sensitive factors, moderate contrastivists needn’t hold that ‘c causes e’ is only true if c proportionally causes e. They can have their cake and eat it too: they can hold that, as a linguistic matter, there are contexts in which non-proportional causal claims are true, whilst, as a theoretical matter, there is an important two-place relation of difference-making, explanation-backing causation which respects proportionality.

5. Conclusion

The idea that proportionality somehow constrains causation or causal explanation is an attractive one, but it faces the disjunction problem. I have argued that this problem should be avoided by supplementing proportionality with a cohesion constraint, understood in terms of the similarity structure of property-spaces. This notion is not objectionably mysterious, and we should not hope to avoid the disjunction problem without it.

Appendix: Cohesion in Structural Equation Models

This appendix describes how COHESIVE PROPORTIONALITY may be implemented by adapting the prominent ‘structural equation models’ framework for representing causation. (I employ this framework as a convenient formal apparatus, without importing any associated philosophical assumptions; in particular, I do not assume any account of what makes models apt.)

Structural equation models represent events by variables taking values. Each variable is associated with a range of mutually incompatible values. For example, as described above, we might represent the stimulus presentation in SOPHIE with a binary variable *RED*, where *RED* = 1 represents *R* and *RED* = 0 represents the possible event of the tile’s not being red. Or we might use a more fine-grained variable *COLOR* that can take many values, 1, ..., *n*, where *COLOR* = 1 represents *S*, *COLOR* = 2 represents the tile’s being crimson, and so on.

This formalism naturally encodes the idea that possible events vary along well-defined dimensions corresponding to properties. As it stands, however, it cannot represent the structure of a state-space: a bare set of mutually incompatible values cannot represent either generality or similarity relations between possible events. To fix this, two tweaks are required. First, to represent similarity (and hence, cohesion), ranges must be more structured than bare sets. For example, they could be spaces with topological, metric, or graph-theoretic structure. Second, instead of assigning single values to variables to represent events, we should assign sets of values (or regions). This allows us to represent generality: one event’s being more general than another is represented by its corresponding region containing the other’s as a sub-region. Hence, instead of representing an event by assigning a value from a bare set, we represent it by assigning a region from a structured state-space. As I now explain, this adaptation allows us to perspicuously implement COHESIVE PROPORTIONALITY.³⁷

³⁷ This approach departs from the standard interventionist paradigm, according to which proportionality is an aspect of the *problem of variable choice* (e.g. Blanchard 2020; Franklin-Hall 2016; Woodward 2018.) In this paradigm, proportionality is taken to be a constraint on which variables are apt to represent a given causal situation. However, this is a somewhat artificial consequence of the austerity of structural equation models: proportionality is much more naturally read as constraining the identification of causes, given some pre-existing space of possible events (i.e. choosing the *values* of given variables). This motivates enriching the formalism in the way that I suggest.

Structural equation models represent relations of causal sufficiency by ‘structural equations’ between variables.³⁸ For every ‘endogenous’ variable (one whose value is determined by factors within the model’s scope) a corresponding structural equation relates it to other variables in the model. To illustrate, a simple model of SOPHIE may pair the binary variable *RED*, representing the stimulus color, with another binary variable, *PECK*, representing whether Sophie pecks. The structural equation $PECK = RED$ encodes the fact that *R* suffices for Sophie’s pecking, and the tile’s being non-red suffices for Sophie’s not pecking.³⁹

Given some effect (such as Sophie’s pecking) represented by a region within the range of an endogenous variable (e.g. the region consisting of the single value $PECK = 1$), these structural equations define a ‘pre-image’: a set of all the ‘tuples of values of other variables which lead, via the equations, to a value within the region corresponding to the effect. For example, in our simple model of SOPHIE, the pre-image for Sophie’s pecking will be the singleton set $\{RED = 1\}$. If we adapt our simple model by replacing the binary variable *RED* with the more fine-grained variable *COLOR*, then the pre-image for Sophie’s pecking is some set $\{COLOR = 1, \dots, COLOR = k\}$, consisting of those values of *COLOR* representing the tile’s being some shade of red. And if we expand our simple model to include the variable *TICKLE* (with values 1/0 corresponding to Sophie’s chin being tickled/not being tickled), then the pre-image for Sophie’s pecking is $\{<RED = 1, TICKLE = 1>, <RED = 1, TICKLE = 0>, <RED = 0, TICKLE = 1>\}$ i.e. the set of those pairs of values representing situations in which either the tile is red or Sophie’s chin is tickled.

The pre-image of a given effect is a region within a ‘product space’. This product space is generated from the spaces corresponding to those variables which determine the variable representing the effect. In the simple model of SOPHIE, the pre-image for pecking is simply a region

³⁸ This formalism is usefully neutral on whether to understand causal sufficiency in terms of counterfactuals, robust generalisations, or in some other way.

³⁹ On the standard formalism, values are assigned to each exogenous variable, and the values of endogenous variables fixed by the structural equations. On my suggested adaptation, things work similarly: *regions* are assigned to each exogenous variable, with regions fixed for endogenous variables via the structural equations. Given an equation $Y = f(X_1, \dots, X_n)$, and some assignment of regions R_1, \dots, R_n to the variables X_1, \dots, X_n , the induced region for *Y* consists of all values $f(x_1, \dots, x_n)$ for x_1, \dots, x_n in R_1, \dots, R_n respectively.


within the space corresponding to the variable *RED*, and in the fine-grained model, it is a region within the space corresponding to *COLOR*. But in the expanded model, where *PECK* is determined by *RED* and *TICKLE*, the pre-image is a region within a product space generated by the spaces corresponding to these two variables.

The structure of this product space is determined by the structure of the spaces that generate it. In particular, a region within the product space is cohesive iff it is the product of regions which are cohesive within their respective spaces.⁴⁰ For example, suppose that all regions of the spaces corresponding to the variables *RED* and *TICKLE* are cohesive i.e. all possible assignments of regions ($\{0\}$, $\{1\}$, $\{0,1\}$) to these variables represent cohesive possible events. Then ‘*RED* = 1’ ($\{<1,0>, <1,1>\}$) corresponds to a cohesive region of the product space, since it is the product of two cohesive regions from each space: ‘*RED* = 1’ and ‘*TICKLE* = 0 or 1’. Likewise, ‘*TICKLE* = 1’ ($\{<0,1>, <1,1>\}$) and ‘*RED* = 1 and *TICKLE* = 1’ ($\{<1,1>\}$) are cohesive regions of the product space. But ‘*RED* = 1 or *TICKLE* = 1’ ($\{<0,1>, <1,0>, <1,1>\}$), i.e. the pre-image for pecking, will *not* correspond to a cohesive region, since it is not the product of two cohesive regions.

Finally, the structure of this product space allows us to divide the pre-image into ‘maximal cohesive components’: cohesive sub-regions of the pre-image which are not contained in any other cohesive sub-region of the pre-image. These regions represent maximal cohesive sufficers for the effect in question i.e. the eligible causes given COHESIVE PROPORTIONALITY.

For example, take the fine-grained models of SOPHIE and DISJUNCTIVE SOPHIE. Given the structure on the range of *COLOR* inherited from color space, the pre-image for pecking will have one maximal cohesive component in the case of SOPHIE, representing *R*, and (as illustrated on the left below) two such components in the case of DISJUNCTIVE SOPHIE, one representing *R*, and another representing the tile’s being green. Or take the expanded model of SOPHIE. Given the structure on the product space, inherited from the structure on the ranges of *RED* and *TICKLE*, the pre-image of pecking will have two overlapping maximal cohesive components (as illustrated on the right below): one representing *R*, the other representing Sophie’s chin being tickled.

⁴⁰ Cohesive regions within spaces corresponding to a single variable represent ‘basic cohesive’ events; cohesive regions within product spaces represent conjunctions of basic cohesive events.



	<i>TICKLE</i> = 0	<i>TICKLE</i> = 1
<i>RED</i> = 0	<0, 0>	<0, 1>
<i>RED</i> = 1	<1, 0>	<1, 1>

Maximal cohesive components of pre-image of pecking in DISJUNCTIVE SOPHIE (left); SOPHIE (right)

Which of these regions represents an eligible cause then depends on which represents an actual event. For example, given that the stimulus was scarlet and Sophie’s chin was not tickled, COHESIVE PROPORTIONALITY recommends *R* as eligible to cause pecking in both SOPHIE and DISJUNCTIVE SOPHIE.⁴¹ (And supposing that the stimulus was scarlet and Sophie’s chin was tickled, COHESIVE PROPORTIONALITY yields the intuitive verdict that both *R* and Sophie’s chin being tickled are eligible causes.)⁴²

REFERENCES

- Bigelow, J. & Pargetter, R. (1990). *Science and Necessity*. Cambridge University Press.
- Blanchard, T. (2020). Explanatory Abstraction and the Goldilocks Problem: Interventionism Gets Things Just Right. *British Journal for the Philosophy of Science* 71(2): 633–663.
- Butts, D. A. & M. S. Goldman (2006). Tuning Curves, Neuronal Variability, and Sensory Coding. *PLoS Biol* 4(4):e92. <https://doi.org/10.1371/journal.pbio.0040092>
- Clapp, L. J. (2001). Disjunctive properties: Multiple realizations. *Journal of Philosophy* 98 (3):111-136.

⁴¹ Of course, its *being* an actual cause requires the model’s aptness, where this is partly a matter of its structural equations satisfying certain truth-conditions (as provided by an account of causal sufficiency).

⁴² For helpful comments and discussion, I would like to thank Nathaniel Baron-Schmitt, Laura Franklin-Hall, Verónica Gómez Sánchez, Jenn McDonald, Jonathan Schaffer, Ted Sider, Michael Strevens, and several anonymous referees for the *Australasian Journal of Philosophy* and *Philosophical Studies*.

- Dorr, C. & J. Hawthorne (2013). Naturalness. In K. Bennett & D. Zimmerman (eds.), *Oxford Studies in Metaphysics: Volume 8*. Oxford University Press.
- Franklin-Hall, L. R. (2016). High-Level Explanation and the Interventionist's 'Variables Problem'. *British Journal for the Philosophy of Science* 67(2): 553–577.
- Funkhouser, E. (2006). The determinable-determinate relation. *Noûs* 40 (3):548–569.
- Garfinkel, A. (1981). *Forms of Explanation: Rethinking the Questions in Social Theory*. New Haven and London: Yale University Press.
- Gómez Sánchez, V. (forthcoming). Does Matter Mind Content? *Noûs*.
- Kim, J. (1989). Mechanism, Purpose, and Explanatory Exclusion. *Philosophical Perspectives* 3: 77–108.
- Langton, R. & Lewis, D. (1998). Defining 'intrinsic'. *Philosophy and Phenomenological Research* 58(2):333-345.
- Lewis, D. (1986a). Events. In *Philosophical Papers, Volume 2*. Oxford: Oxford University Press.
- (1986b). *On the Plurality of Worlds*. Oxford: Blackwell.
- List, C. & P. Menzies (2009). Nonreductive physicalism and the limits of the exclusion principle. *Journal of Philosophy* 106(9): 475–502.
- Mackie, J. L. (1974). *The Cement of the Universe*. Oxford: Oxford University Press.
- Maslen, C. (2017). Pragmatic explanations of the proportionality constraint on causation. In Beebe, Hitchcock & Price (eds.), *Making a Difference: Essays on the Philosophy of Causation*. Oxford: Oxford University Press.
- McGrath, Matthew (1998). Proportionality and mental causation: A fit. *Philosophical Perspectives* 12:167-176.
- Putnam, H. (1975). Philosophy and our mental life. In *Mind, Language, and Reality*. Cambridge University Press.
- Sartorio, C. (2006). Disjunctive Causes. *Journal of Philosophy* 103 (10):521-538.
- Schaffer, J. (2005). Contrastive causation. *The Philosophical Review* 114: 297–328.
- Shapiro, L. & E. Sober (2012). Against proportionality. *Analysis* 72(1): 89–93.
- Sider, T. (2011). *Writing the Book of the World*. Oxford University Press.

- Skiles, A. (2016). In defense of the disjunctive. *Inquiry: An Interdisciplinary Journal of Philosophy* 59(5): 471-487.
- Strevens, M. (2004) The causal and unification approaches to explanation unified—causally. *Noûs* 38(1): 154–176.
- (2008). *Depth: An Account of Scientific Explanation*. Cambridge, MA: Harvard University Press.
- (2012). Replies to Weatherson, Hall, and Lange. *Philosophy and Phenomenological Research* 84(2): 492–505.
- Touborg, C. (2022). Relativizing proportionality to a domain of events. *Synthese* 200 (2):1-20.
- Vaassen, B. (2022). Halfway Proportionality. *Philosophical Studies* (9):1-21.
- Weatherson, B. (2012). Explanation, Idealisation and the Goldilocks Problem. *Philosophy and Phenomenological Research* 84(2): 461–473.
- Weslake, B. (2010). Explanatory depth. *Philosophy of Science* 77: 273–294.
- Woodward, J. (2018). Explanatory autonomy: the role of proportionality, stability, and conditional irrelevance. *Synthese*: 1–29.
- Woodward, J. & C. Hitchcock (2003). Explanatory Generalizations. Part I: A Counterfactual Account. *Noûs* 37: 1–24.
- Yablo, S. (1992a). Mental Causation. *Philosophical Review* 101: 245–80.
- (1992b). Cause and essence. *Synthese* 93(3):403–449.
- (2003). Causal relevance. *Philosophical Issues* 13(1): 316–28.