

CAUSAL STRUCTURALISM¹

John Hawthorne
Syracuse University

How should we think as metaphysicians about the nature of properties? A promising place to begin, perhaps, is with the following observation: Properties confer causal powers upon the things that have them. Being spherical confers a capacity to roll. Being hot confers a capacity to make ice melt. Being possessed with mass confers a capacity for gravitationally attracting other massy things. But what is the relationship between properties and the causal powers that they confer?

On one view, it is no part of the essence of a property that it confers the causal powers that it does. Properties have a causal role all right, but the role is utterly contingent. God could have bestowed a very different role upon, say, the property of being negatively charged. Such a view has been embraced, notably, by David Lewis. According to his conception of reality, the causal powers of a property are constituted by its patterned relations to other properties in the particular Humean mosaic that is the actual world.² In other worlds, that property will be combined differently. In those worlds, then, the laws of nature governing it and thus the causal role it enjoys may be utterly different. In this regard, Lewis is largely following David Hume. Hume clearly thinks of the nature of simple impressions as given by how they are intrinsically and not by their patterned connection to other elements in reality: hence the denial of necessary connections between simple existences. What is missing in Hume is a robust realism about properties.³ Simple impressions of a certain shade of red do not, literally, have some shade in common. Let us then call Lewis' view the 'neo-Humean view'.

On a second view, some or all of the causal powers of a property are essential to it. It does not go so far, though, as to insist that the causal powers of a thing exhaust its nature. Two possible properties may be distinct and yet the same causal powers be essential to each. Those powers do not, then, constitute an individual essence (where an 'individual essence' is a profile that is necessary and sufficient for some particular thing). Like the first view, this view holds that there is something more to the nature of a property than the causal

powers that it confers—the intrinsic nature of the property if you will—holding that two different internal natures might necessitate the same causal profile.⁴ Call this the ‘double aspect view’.

On a third view there is, for each fundamental property, a causal profile that constitutes the individual essence of a property. That is, the profile is both necessary and sufficient for each property. (The relevant profile, we should note, may include facts about how a property figures as an effect as well as how it figures as a cause.) Let us call this view ‘causal structuralism’. The first two views but not the third holds that there is something to a property—call it its quiddity—over and above its causal profile.⁵ Causal structuralism holds that quiddities are a will-o’the-wisp.⁶

Note that causal structuralism is not intended to be applied directly to properties that are neither vehicles of genuine change nor serious causal explanation. As such, it does not offer a theory of mathematical and logical properties—such as the successor relation and the identity relation. (Also excluded, perhaps are such properties as overlap and part of, codified by the mereologist.) Nor does it, in the first instance, offer a theory of mere Cambridge properties, whose change does not intuitively amount to real change and which are not, intuitively, the conferrers of bona fide causal powers.⁷ Call the candidate domain of properties for which causal structuralism is intended the ‘natural properties’. I shall for the purposes of this paper assume that this category is reasonably well understood.

Note also that the causal structuralist needn’t hold that all of a thing’s causal profile is essential to it. She could consistently claim that some property A has the power to bring about both B and C, but that only the power to bring about B is part of its individual essence. But, *prima facie*, it would seem that to distinguish the essential from the accidental in this way would be to draw lines in an arbitrary way. For now, then, let’s focus on a version of causal structuralism according to which the whole causal profile of a property is essential to it.

In what follows, I shall be looking at arguments for and against the following two key theses of causal structuralism: (i) That for any given natural property, there is some causal profile such that having that profile is sufficient for being that property and (ii) that for any given natural property, there is some causal profile such that having that profile is necessary for being that property. I can’t imagine anyone liking the first thesis but not the second.⁸ Meanwhile, if one likes the second thesis but not the first, one will likely find the double aspect theory attractive. Finally, if one eschews both theses, one will likely find the neo-Humean position attractive.

Views on quiddities have rough analogues in debates about haecceities. The neo-Humean view on properties resembles a radical haecceitism according to which all the qualities of a particular are contingent to it: only its haecceity is essential. (I could have been a poached egg so long as my haecceity was present.) The double aspect view also has an analogue in the metaphysics of individuals: There are those who believe that certain kind-properties (and perhaps other

origin-theoretic properties) are essential to individuals but that one cannot construct an individual essence out of such properties. Any such list would leave out the thisness or haecceity which distinguishes a particular individual from other possibilia of the same kind (and origin). Meanwhile, causal structuralism about properties is analogous to a view according to which some qualitative profile is both necessary and sufficient for being a particular individual and hence that haecceities are a fiction.⁹ I shall be exploiting these analogies in due course.

In part one, I shall briefly explore various arguments for causal structuralism, considering semantic, metaphysical, epistemological and methodological considerations that offer prima facie support for one or both of its key theses.

In part two, I shall consider five kinds of considerations that apparently militate against causal structuralism, suggesting that none of them are decisive.

In part three, I shall consider an especially worrying kind of objection to causal structuralism, in response to which I shall offer a fresh way to think about that view.

Part One: Arguments for Causal Structuralism

1 Semantic: Referring to Properties

One might think that if a pair of properties are alike in causal role, then one would be unable to single one of them out. And if that is right, then it will be incoherent to think of, say, 'redness' as singling out a property whose causal profile is exactly like some other actual property. In his 'Causal and Metaphysical Necessity', Sydney Shoemaker endorses something like this line of thought, viz: "And what epistemological considerations show, in the first instance, is that if there are sets of properties whose members are identical with respect to their causal features, we necessarily lack the resources for referring to particular members of these sets."¹⁰

This doesn't secure causal structuralism. But it is fair to say that it at least goes part of the way, by undermining the idea that, say, the property of being negatively charged might be distinct from some other property whose actual causal profile is just the same.

The argument's plausibility trades on an equivocation on 'causal feature'. Suppose A and B have the same causal powers. To simplify, let's suppose the following profile is exhaustive: 'A' and 'B' are both 50 percent probable on C and both, if instantiated, bring about D. Suppose John has C. It is compatible with the profile of C that John has A and not B. Supposing John has A, can John uniquely refer to A? Surely he can. He can do this not by exploiting an asymmetry in the laws governing A and B but rather by an asymmetry in the pattern of instantiation. Thus he can say: "Let 'Jones' pick out the property that C just caused to be instantiated in me."

Response: But in the case described, A and B do differ in causal features, viz: A but not B caused the state of affairs ‘John’s being D’.

Reply: But this isn’t the sense of ‘causal feature’ in question. The causal structuralist does not think it essential to A that it causes John to be D. What is essential to A, according to the causal structuralist, are the causal powers it confer upon whatever thing instantiates it. As the example shows, even if two properties have the same causal profile in the sense relevant to causal structuralism, we may have the capacity to single one of them out. The key point is that a pair of properties may have equivalent causal profiles and yet asymmetrical patterns of instantiation, the latter affording a perfectly good basis for unique reference to one of the pair. The line of thought we are considering is not all that much better than one which claims that if I have the same qualitative profile (relational and non-relational) as some other guy, then I can’t single myself out.

2 *Metaphysical: Intra-World and Inter-World Variation*

It seems to me a general feature of our thought about possibility that how we think that something could have differed from how it in fact is [is] closely related to how we think that the way something is at one time could differ from the way that same thing is at a different time. In possible worlds jargon, the ways one and the same thing of a given sort can differ across worlds correspond to the ways one and the same thing of that sort can differ at different times in the same world. Could I have been a plumber or an accountant instead of a philosopher? The answer seems to be yes—and this goes with the fact that we acknowledge the possibility of a scenario in which something who was exactly as I was at some point in my life undergoes a series of changes resulting in his eventually being a plumber or accountant. Could I have been a poached egg? Pace Lewis, the answer seems to be no—and this goes with the fact that our principles of trans-temporal identity rule out the possibility of a scenario in which something starts off as a human being of a certain description and ends up a poached egg.¹¹

This line of thought, taken from Shoemaker’s ‘Causal and Metaphysical Necessity’ and aired earlier in his ‘Causality and Properties’ takes off from the idea that in the case of particulars, the possibility of something’s being F is of a piece with the possibility of a world which branches from the actual world where that thing is F. It then generalizes that thought to all things. Combine this generalization with the thesis that a property cannot change its causal powers over the course of time and then it follows that there is no world where some property enjoys some power that it actually lacks—since there is no world branching from the actual world where that property enjoys that power.

Here’s the motivating principle—call it the ‘Branch Principle’:

If x exists at a world w^1 and x could have been F then there is a world w^2 which is exactly like w^1 up to some point in time where x is F (call such a world a ‘branch world’).¹²

Given that properties can’t change their causal powers over the course of time, we can use the Branch Principle to secure at least one of the two key causal structuralist theses (that the causal powers of a property are essential to it).

Should we agree with Shoemaker that, at least in the case of particulars, the Branch Principle is right? If determinism is true, there is no branch world with the same history and the laws of nature but where I am a plumber. So, according to Shoemaker, if determinism is true, things could not have been such that I was a plumber. Even if one thinks that the laws of nature are metaphysically necessary, it does not seem that one wants to say that it is impossible that I be a plumber: for why not admit a world which is a bit different to ours all along where I am a plumber? To focus our intuitions, we might do well to consider a simple world, one where two particles have always existed (either since the beginning of time or for all eternity). Call them A and B . Intuitively, it seems that A could have existed alone. But there is no branch world where A exists alone. So Shoemaker is committed to the thesis that A could not have existed alone. At the very best, such modal commitments are extremely tendentious. It thus does not seem that a promising way to motivate causal essences for properties is by the Branch Principle.

3 *Epistemological: An Argument from Recognition*

Suppose a property is something over and above its causal profile. We then seem to have conceptual space for something like the following: There is negative charge 1 and negative charge 2 that have exactly the same causal powers. What we call an instance of negative charge 1 is sometimes an instance of negative charge 1, sometimes an instance of negative charge 2. Since 1 and 2 have the same propensities to affect all possible detection mechanisms, there is no way of discriminating 1 and 2. We would now be unable to tell, it seems, whether two groups of particles that we call “negatively charged” had the same property or else distinct but indistinguishable properties. But this is absurd: We can recognize property sharing. So we had better not allow properties to have an individual essence that transcends causal features. Here is Sydney Shoemaker:

...if two property can have exactly the same potential for contributing to causal powers, then it is impossible for us even to know (or have any reason for believing) that two things resemble one another by sharing a single property.¹³

I myself am not much moved by arguments of this form. We are all familiar with arguments with the following structure: If metaphysics M is right, then there are unlucky worlds where our judgments are way off with respect to sub-

ject matter S. Further, there is a natural sense in which, if M is right, we can't tell whether we are in an unlucky world. But we are very knowledgeable about subject matter S. If we can't tell whether we are in an unlucky world, we are not knowledgeable about subject matter S. So metaphysics M is all wrong.

Examples: (i) If possible worlds are Humean mosaics then there are unlucky worlds where everything is just like this and then the world stops altogether at some point in the very near future. Supposing the world is a Humean mosaic, we can't tell whether we are in a world like that. But we know a whole lot about the future. So the world is not a Humean mosaic.

(ii) If metaphysical realism about physical objects is true, then there are unlucky worlds where we are brains in vats. Supposing realism is true, we can't tell whether we are in a brain in a vat world. But we know a whole lot about tables. So we should not be metaphysical realists about physical objects.

(iii) If qualia are superadded features of the world then there are unlucky worlds where everyone else but me is a zombie. ...

Now it's true that there is no stable consensus concerning how to react to these kinds of arguments. We haven't come as far as we'd like when it comes to epistemology. Some of us will go reliabilist and say that so long as we in fact don't live in an unlucky world and unlucky worlds are not nearby, then our knowledge is not under threat. Applied to the case at hand, we will say that so long as there is not in fact any real risk of pairs of properties being presented to us that are role indiscernible, then our knowledge that certain things share certain properties is safe. Some of us will go contextualist and say that once we start doing serious philosophy then we should all be sceptics when it comes to the deployment of 'know' but that the standards are far lower in ordinary contexts—when truth plus moderate credentials are sufficient for the proper applicability of 'know'. Applied to the case at hand, we will say that so long as scientists are lucky enough to be right when they say that the same feature is being presented over and over again and so long as they rule out alternatives that are relevant to them, they can say 'I know that is the same feature as that', but that we philosophers, having gotten all worried about duplication of role by different quiddities should hesitate to claim to know.¹⁴ Some of us will go abductionist and say the simpler and more elegant hypothesis is evidence enough for knowledge, so long as there is belief and truth. Applied to the case at hand, we will say that it is reasonable to believe in the simpler hypothesis—that there is only one quiddity there, not several—when presented with the negative charge role.

I don't know what to say of these various epistemological reactions. But I do admire a common thread—namely, "Don't throw out a metaphysical hypothesis on the basis of unlucky world arguments".¹⁵

There is one further thing to be said. Once one takes unlucky world arguments seriously then one sees soon enough that causal structuralism doesn't help all that much. Consider the following sceptical scenario. There is negative charge 1 and negative charge 2 that are exactly alike except that were property

P which is in fact uninstantiated to be instantiated, charge 1 would interact with it in a different way than charge 2.¹⁶ Such an unlucky world is perfectly possible even if causal structuralism is true. By hypothesis, charge 1 and charge 2 would engage with our sensory mechanisms and other detection devices in the same way. How then can we *know* that there is one property, negative charge, rather than a pair of properties whose causal differences are inscrutable? Insofar as causal structuralism is designed to offer respite against scepticism, its comforts may be largely illusory.

4 *Epistemological: Getting to Know a Property.*¹⁷

We distinguish, intuitively, between being able to refer to a property and knowing what that property is. Only in the latter case is the nature of the property genuinely revealed to us. One way to try to flesh out this intuitive contrast is via George Bealer's distinction between semantically stable and semantically unstable expressions:

An expression is semantically stable iff, necessarily, in any language group in an epistemic situation qualitatively identical to ours, the expression would mean the same thing. An expression is semantically unstable iff it is possible for it to mean something different in some language group whose epistemic situation is qualitatively identical to ours. ...

Semantic instability has to do with the effects of the external environment. An expression is semantically unstable iff the external environment makes some contribution to its meaning. Natural kind terms are paradigmatic—'water', 'gold', 'heat', 'beech,' 'elm', etc. Logical, mathematical, and a great many philosophical terms, by contrast, are semantically stable: the external environment makes no such contribution. For example, 'some', 'all', 'and', 'if', 'is identical to', 'is' 'necessarily', 'possibly', 'true', 'valid'; '0' '1' '+' '÷', '∈; 'property', 'quality', 'quantity', 'relation', 'proposition', 'state of affairs', 'object', 'category', etc. It seems clear that all these are semantically stable: any language group in an epistemic situation qualitatively identical to ours would mean what we mean by these "formal" expressions."¹⁸

The following seems rather natural as a necessary condition on knowing what a property is:

K1: One knows what a property is only if one has a semantically stable way of referring to it.

On this gloss, users of the term 'water' that are altogether ignorant of chemistry succeed in referring to the property of being H₂O, but, lacking a semantically stable way of referring to that property, do not know what property it is that they are referring to. It is tempting, meanwhile, to suppose that many of us are now in a position to know what property it is that 'water' refers to—thanks

to modern chemistry. Indeed, it is tempting in general to suppose that the advance of science has enabled us in many cases not merely to refer to causally efficacious properties but to know what they are. Thus:

K2: For a large range of causally efficacious properties we are in a position to know what those properties are.

Accept K1 and K2 and we are in a good position to argue against quidditism. For suppose, as the quidditist admits, that a different quiddity could have played the hydrogen role. Then there would seem to be a community whose epistemic situation is qualitatively identical to ours but who refer to a different property by ‘hydrogen’. Moreover, there seems to be no imaginable way of achieving a semantically stable way of referring to the property of being hydrogen, given quidditism. Any sort of role-description (whether rigidified or not¹⁹) will, in a qualitatively identical situation, be used by a counterpart community to pick out a different property. Does this spell real trouble for quidditism? I doubt it. Notice that we do not appear to have a semantically stable way of referring to any given person, say Saul Kripke. I can’t take a cognitive photograph of his haecceity. My ability to refer to him will exploit either intrinsic features that he happens to possess, or else relations that he stands to myself and others, or both. Whatever the means I have of referring to him, it seems clear that there will be a counterpart community that uses the same reference-fixing devices to refer to a different individual. So ‘Saul Kripke’ is not semantically stable in the relevant sense. Nor is a semantically stable way of referring to Saul Kripke available, even in principle. Does this mean that in the ordinary sense, people do not know Saul Kripke? Surely not. By analogy, the preceding reflections should not be taken to show that, in the ordinary sense, people do not know which property is picked out by ‘hydrogen’. It remains very unclear why it should count as a cost of a theory that it reckons most or all causally efficacious properties to fail the demanding standards of knowability set by K1.²⁰

5 Methodological: Don’t Invoke What You Don’t Need

The best case for thinking that the causal profile of a property exhausts its nature proceeds not via the thought “Well otherwise we wouldn’t know a whole lot of what we do know” but rather via the thought “We don’t need quidditative extras in order to make sense of the world.” Let us return to negative charge. All scientific knowledge about negative charge is knowledge about the causal role it plays. Science seems to offer no conception of negative charge as something over and above “the thing that plays the charge role”. If there were a quiddity that were, so to speak, the role filler, it would not be something that science had any direct cognitive access to, except via the reference fixer “the quiddity that actually plays the charge role”. Why invoke what you don’t need?

Unless certain logical considerations forced one to suppose that properties are individuated by something over and above their causal role, then why posit mysterious quiddities?

Bertrand Russell was certainly onto the idea that science reveals nothing beyond causal structure when it comes to fundamental properties when he tells us that while introspection reveals the intrinsic quality of percepts, "... we know nothing of the intrinsic quality of the physical world," adding that "We know the laws of the physical world, in so far as these are mathematical, pretty well, but we know nothing else about it."²¹ Supposing we are unpersuaded by Russell's idea that role transcendent qualities are required to accommodate the facts of introspection, we can easily generate an argument for causal structuralism: Why posit from the armchair distinctions that are never needed by science?

I recommend the current line of thought as the most promising for the causal structuralist to pursue. But one should be aware of a trade off: Perhaps science doesn't need a robust conception of causation and can get by with thinking of causal laws in a Humean way, as the simplest generalizations over the mosaic. If so, it seems that one needs an independent characterization of the mosaic's pixels. It hardly seems plausible to be a deflationary Humean about causation and yet a causal structuralist about properties. To eschew quiddities on the basis of considerations of scientific economy may serve to saddle us with a view of causality that is far from economical.

Part Two: Arguments Against Causal Structuralism

1 *The Circularity Argument*

Suppose we define A in terms of its capacity to bring about B and then go on to define B in terms of its capacity to be brought about by A. Isn't the definition circular? Perhaps Bertrand Russell was moved by something like this concern when he wrote:

There are many possible ways of turning some things hitherto regarded as "real" into mere laws concerning the other things. Obviously there must be a limit to this process, or else all the things in the world will merely be each other's washing.²²

Its not so clear that the problem is very serious. We avoided circularities in the functional analysis of belief and desire in the philosophy of mind by defining belief and desire together. Perhaps we can do the same in the functional analysis of properties. Suppose with David Armstrong that causal laws are relations between universals.²³ Assume, then, that a causal necessitation relation N that holds between certain universals. At the risk of oversimplification, let us look a simple world where the lawbook for properties instantiated in that world is very small. There are five properties A, B, C, D, E. Here are the laws in the lawbook: $A \mathcal{N} B$, $A \mathcal{N} C$, $B \mathcal{N} D$ and $D \mathcal{N} E$. Just as functionalism in the

philosophy of mind was helped by aid of the Ramsey-Lewis technique for functional analysis, so the same is true of causal structuralism about properties. Take the laws of the lawbook and conjoin them. Replace each property name by a distinct variable ($F^1 \dots F^n$) and prefix each variable by a quantifier. So we have $\exists F^1 \exists F^2 \exists F^3 \exists F^4 \exists F^5 (F^1 \supset F^2 \wedge F^1 \supset F^3 \wedge F^3 \supset F^4 \wedge F^4 \supset F^5)$. The lawbook was supposed to be exhaustive. To capture this, we can add such clauses as $\forall F^6 (F^4 \supset F^6 \supset F^6 = F^5)$. Call this the Ramsified lawbook. We can now articulate causal structuralism very easily, and whatever its merits, we cannot be accused of vicious circularity. Since the variable ‘ F^1 ’ replaced A, we can give a theory of the individual essence of A by the open sentence you get by dropping the existential quantifier prefixing ‘ F^1 ’. According to causal structuralism, it is a necessary truth that anything that satisfies that open sentence is identical to A. Generalizing, the causal structuralist will say that any natural property can be defined by a suitable open sentence delivered by the Ramsified lawbook for that property.

(Note that, according to the causal structuralist, one Ramsifies the lawbook and not the worldbook to get the essence of a property. If I am tall, that does not go into the definition of tallness. What goes into the definition are the laws concerning how tallness relates causally—as cause and effect—to other properties. Note also that the approach assumes what the Humean abhors—that the causal relation is fundamental. But that is to be expected of causal structuralism in any case. Note finally that the laws in the lawbook may be considerably more complicated than the simple single property to single property laws discussed above: but this makes little difference to the viability of the Ramsification strategy.)

2 *Combinatorialism.*

On one attractive picture of modality, you get possible worlds by mixing and matching the fundamental natural properties. Call this Humean combinatorialism. Here is David Lewis’ version of that view: “To express the plenitude of possible worlds, I require a principle of recombination according to which patching together parts of different possible worlds yields another possible world.”²⁴

Lewis goes on to deploy this principle against the thesis that is common to causal structuralism and double aspect theory:

Another use of my principle is to settle—or as opponents might say, to beg—the question whether laws of nature are strictly necessary. They are not; or at least laws that constrain what can coexist in different positions are not. Episodes of bread-eating are possible because actual; as are episodes of starvation. Juxtapose duplicates of the two, on the grounds that anything can follow anything; here is a possible world to violate the law that bread nourishes. So likewise against the necessity of more serious candidates for fundamental laws of nature.²⁵

The principle of recombination offers a very elegant theory of possibility space. Pity to abandon it if there is nothing elegant to put in its place by way of expressing the plenitude of possibilities. But there is.

Let us replace Humean combinatorialism by Structural Combinatorialism. I sketched a simple lawbook above, containing five properties and four laws. Consider the Ramsified lawbook corresponding to that simple lawbook, the former sufficing to capture the content of the latter according to the causal structuralist. That Ramsified lawbook expresses five possible properties. What other properties are possible? According to the Structural Combinatorialist, any logically consistent Ramsified lawbook expresses a possible set of properties. The structural combinatorialist can work with his own fundamental principle of plenitude—one that governs properties: it corresponds to the plenitude of consistent lawbooks. He can then build a theory of possible worlds upon that principle of plenitude. I see no reason for thinking that the result will be inelegant or unsatisfying.

3 Experience

One might instinctively react to causal structuralism by worrying that it fails to capture the color of the world. Taken literally, the problem is not so pressing. After all, it is not so unreasonable to suppose that a property counts as redness so long as it has the appropriate capacity to cause sensations of the right sort. But how about the sensations themselves? One might think that, say, pain (or phenomenal red), is constituted by an intrinsic *thusness* and not simply by its causal role. On this way of thinking, something could duplicate the causal role of pain in its entirety and yet not *be* pain.

Clearly, part of what held Russell back from full blooded causal structuralism are considerations such as these, as is evidenced by his view that while “we do not know their laws so well as we would wish,” we do know the intrinsic quality of percepts (which are in turn “part of the physical world”) by introspection.²⁶ One way to push this line of thought is via zombie thought experiments: There are, it seems, possible worlds with zombies that are structurally isomorphic to us. The zombies possess states with the right causal role but which lack any qualitative character. The difference between a zombie and us, on the current conception, is that the quiddities have been switched. Quiddity switching is what turns the light on and off, so to speak.²⁷ If that is right, only the neo-Humean and double aspect views remain as serious options.

It would be silly to try and engage at length with this objection in a survey paper such as this one. It is clear enough, though, what the causal structuralist should say in broad outline: How is one so sure that it is some intrinsic, role-transcendent quiddity that is responsible for consciousness? Consider me, and suppose I am in pain. Assume that I do have evidence-transcendent quiddities and that God were now to switch the quiddity that underlies my pain for another one—makes the quiddities dance, as it were. I would still believe that I

am in pain, as the causal propensities with respect to my belief system would be unchanged. There are two perspectives here. On one perspective, the experiences dance with the quiddities—its just that we don't notice that pain has been replaced by ersatz pain. On another perspective, the experiences do not dance with the quiddities: we remain in pain and hence do not start making introspective mistakes when the quiddity switch occurs. Is it so clear that the first perspective is correct? I believe that the jury is still very much out on this issue.

(Note that even if one does believe in role-transcendent properties—as, say, Lewis does—this needn't be because one thinks that experience is role-transcendent. Lewis' perfectly natural properties are role-transcendent—but Lewis' reason for positing role transcendent properties has far more to do with his combinatorialism than with any deep intuition that introspection brings us into acquaintance with something role-transcendent. Analogy—there are plenty of people who believe in haecceities but who don't believe that self-consciousness brings us into special acquaintance with a peculiar sort of meanness that serves as the individual differentium.)

4 *Categorical and Dispositional Properties*

We've all learned that dispositional properties have a categorical basis. But isn't this violated by causal structuralism? Frank Jackson rejects the view that properties are defined by causal role for exactly this reason: "This, to my way of thinking, is too close to holding that the nature of everything is relational cum causal, which makes a mystery of what it is that stands *in* the causal relations."²⁸

I find it hard to see how to make the worry compelling. Suppose, to parody Jackson, we worried about a necessity of origins thesis for substances in the following way: "This to my way of thinking, is too close to holding that the nature of particulars is relational cum causal, which makes a mystery of what it is that stands *in* the causal relations." That objection doesn't seem very compelling at all.

Its true enough that we wish to distinguish between the second order property of, say, being a property that causes headaches and those first order properties which instantiate that second order property. But the causal structuralist has no reason to deny distinctions like that. Consider, for example, the open sentence from a Ramsified lawbook that defines a property. There is a property that the open sentence expresses. And there is a different property that is the unique realizer of that open formula. No need for collapse.

Perhaps the worry is that categorical bases are supposed to be without any relational features essentially. It is true enough that if this is a requirement, then both the double aspect theorist and the causal structuralist are in trouble. But it seems that the requirement needs some sort of motivation rather than being put forward as an unargued axiom. Note also that the categorical basis

would, *prima facie*, be a poor explanans for the disposition as explanandum, if the categorical basis did not drag any causal powers along with it.²⁹

5 *Argument Five: Hyperstructuralism*

According to the hyperstructuralist, one defines a property by taking a lawbook and Ramsifying through the whole thing entirely, so that every relational and non-relational predicate is replaced by a variable. This means that the relation of causal necessitation is itself replaced by a variable. The hyperstructuralist claims that the resulting Ramsified sentence (which, note, says nothing explicitly to the effect that its subject matter is causality) is sufficient to provide definitions of each predicate in the original vocabulary that figured in the lawbook, including the relation of causal necessitation.

I take it that none of us are hyperstructuralists.³⁰ None of us think that the truth conditions of the original lawbook are given by the purely formal sentence that the hyperstructuralist obtains by comprehensive Ramsification. So it looks like the content of the predicate ‘causes’ cannot be recovered by Ramsifying through the theory in which ‘causes’ figures. Why then believe—as the causal structuralist believes—that other predicates can be defined by the Ramsey-Lewis technique? What’s so special about causality?

The question is fair enough. But it is not clear that the causal structuralist is devoid of an answer. After all, it is part of the metaphysical picture of causal structuralism that the structural properties of a system have a different status than the nodes of the structure. Causality is not the only property that turns out to be primitive and indefinable—probably the same goes for identity and for various mathematical and mereological properties. You may not like the picture. But it does not seem altogether arbitrary or ad hoc to treat the structure of the world (the “form” of the world) in a different way to the nodes in the structure (the “matter” of the world). It thus does not seem to me then to be altogether ad hoc or arbitrary to endorse causal structuralism but to resist hyperstructuralism.

Part Three. Symmetrical Roles: Two Varieties of Causal Structuralism

The following seems to me to be a perfectly possible causal structure: There are four properties, call them A, B, C, D. Here are the laws governing them: $A \supset C$, $B \supset C$, A and $B \supset D$. It is crucial to this structure, note, that A and B are distinct. Their coinstantiation has different effects (the addition of D to the world) than is produced by either being instantiated alone. Not only are such causal structures intuitively possible; the structural combinatorialism that I sketched and which seems to me to be an elegant principle of plenitude for properties militates in favor of their possibility.³¹

Suppose such a structure is possible. That spells trouble for causal structuralism. According to that view, the individual essence of a property can be

given via the Ramsified lawbook. But there is no way for that to work here. Ramsify the lawbook and there is nothing to distinguish the pair of properties corresponding to 'A' and 'B' even though the structure does require that there *be* a pair and not just one. The situation is similar to that besetting the most straightforward version of anti-haecceitism, according to which some purely qualitative profile provides the individual essence of each individual substance. That view famously got into trouble when confronted with Max Black's symmetrical world³² (and Scotus' duplicate angels³³), which seemed to require that there be two substances whose qualitative profile is exactly the same. Anti-haecceitism of that sort is troubled by intra-world duplication of qualitative profiles. Causal structuralism of Shoemaker's sort is troubled by intra-world duplication of causal profiles. If causal structuralism is the view that each property has a unique individual essence consisting of a causal profile, then that view seems to be wrong.

David Lewis has provided an alternative way of thinking about anti-haecceitism than the one adumbrated above: Here is his favored conception of haecceitism:

If two worlds differ in what they represent *de re* concerning some individual, but do not differ qualitatively in any way, I shall call that a *haecceitistic difference*. Haecceitism, as I propose to use the word, is the doctrine that there are at least some cases of haecceitistic difference between worlds.³⁴

Lewis is clear that on this gloss, anti-haecceitism can allow intra-world duplication of things with the same qualitative profile. Max Black's world is perfectly possible according to this brand of anti-haecceitism. What is not possible is the existence of two qualitatively indiscernible worlds such that the *de re* truths concerning one are different from the *de re* truths concerning the other.

One might think that once the possibility of Black's world has been conceded, haecceitism quickly follows. For isn't there one world where one of the balls exists alone and a different world where the other ball exists alone in a qualitatively duplicate state? Here is Lewis:

I might have been one of a pair of twins. I might have been the first-born one, or the second-born one. These two possibilities involve no qualitative difference in the way the world is. Imagine them specified more fully: there is the possibility of being the first-born twin in a world of such-and-such maximally specific qualitative character. And there is the possibility of being the second-born twin in exactly such a world. The haecceitist says: two possibilities, two worlds. They seem just alike, but they must differ somehow. They differ in respect of 'cross-identification'...I say: two possibilities, sure enough. And they do indeed differ in representation *de re*: according to one, I am the first-born twin, according to the other I am the second-born. But they are not two worlds. They are two possibilities within a single world. The world contains twin counterparts of me, under a counterpart relation determined by intrinsic and extrinsic qualitative similarities (especially, match of origins). Each twin is a possible way for a person to be, and in fact is a possible way for me to be.³⁵

Lesson: Utilizing Lewisian counterpart theory we can resist multiplying worlds whenever we have multiple possibilities for a thing to be. In so resisting we can allow intra-world qualitative duplication without taking on haecceitism.

What is interesting is that there is a version of causal structuralism that exactly mimics Lewis' approach. Let a structural description of a world be a description which describes the world using certain structural primitives—like part/whole and causal necessitation—and which otherwise uses merely the resources of logic (if you want to be a haecceitist about things but causal structuralist about properties, then throw in all the individual constants corresponding to each thisness). The causal structuralist can map out his position in the following way:

If two worlds differ in what they represent *de re* concerning some property, but do not differ structurally in any way (i.e. have the same structural description), I shall call that a quidditistic difference. Anti-structuralism is the doctrine that there are at least some cases of quidditistic difference between worlds. I say that there are no quidditistic differences between worlds.

Lewis' anti-haecceitist can allow that there is intra-world duplication of some qualitative profile. My causal structuralist can allow that there is intra-world duplication of a causal profile. One might think that one can generate quidditistic differences out of such duplication. Take the symmetric structure that I described earlier. Isn't there the possibility that there are two things that instantiate A and one thing instantiates B and another possibility where two things instantiate B and one thing instantiates A? The lesson learned earlier can be applied here: don't multiply possible worlds whenever one has a multiplication of possibilities.

Lewis' anti-haecceitism does not either require a deep answer to the question: "Which elements of a thing's qualitative profile are essential to it?" His counterpart theory allows context-sensitive flexibility when responding to such questions. Neither does my causal structuralist need treat as deep the question "Which element's of a property's causal profile are essential to it?" A counterpart theory can allow context-sensitive flexibility here too. What is crucial to this brand of causal structuralism is that it does not allow that worlds can be alike structurally and yet different concerning what is true *de re* of the properties in them.

Suppose one doesn't like counterpart theory and yet does believe in the possibility of Black's world of two numerically distinct but qualitatively duplicates balls. As I see it, there is then no resisting haecceitism. Suppose one doesn't like counterpart theory and yet does believe in the possibility of pairs of properties with symmetrical causal roles. As I see it, there is then no resisting anti-structuralism. But symmetric structures do seem eminently possible. So counterpart theory combined with a rejection of quidditistic differences between worlds is certainly the best way to pursue the causal structuralist's vision of reality. I leave it to the reader to judge whether that vision is worth pursuing.

Appendix: Modest Structuralism

In replying to a version of this paper, Sydney Shoemaker noted the availability of a more modest version of causal structuralism than any discussed in the body of this paper (evincing some temptation to embrace it rather than any of the more radical versions of causal structuralism). Lest readers assume that his writings are unequivocally committed to full-blooded causal structuralism, let me sketch the modest view. Let the Shoemaker sentence for a target property P be obtained as follows: Take the lawbook, conjoin it, and replace each occurrence of P by a variable v , (leaving the other property names as they are). The result is an open sentence. The modest structuralist claims that the open sentence produced by this procedure individuates the target property. Notice that the symmetric causal structure that I described is no problem at all for the modest structuralist. After all, the Shoemaker sentence for B (as it figures in that structure), will include some such clause as ‘in combination with A , suffices for D ’, while the Shoemaker sentence for A will not say that. Thus we have a kind of “anti-reductive” structuralism that is much more tolerant in that it can make room for symmetric structures. This modest version of causal structuralism will still put some constraints on possibility space: It will not allow the following pairs of lawbooks: (1) $A \text{ } N \text{ } B$, $B \text{ } N \text{ } C$ (and that’s all) (2) $A \text{ } N \text{ } D$, $D \text{ } N \text{ } C$ (and that’s all), since the Shoemaker sentence associated with B and D would be the same. Insist that these are different and genuinely possible lawbooks and the thesis of modest structuralism—that the Shoemaker sentence for a property individuates it—is violated. From the perspective of the causal structuralism examined in this paper, modest structuralism risks being far too permissive with regard to possible structures of properties. After all, modest structuralism can allow a distinction between the following pairs of lawbooks: $F \text{ } N \text{ } G$ (and that’s all), $H \text{ } N \text{ } I$ (and that’s all), since each of the properties F , G , H and I will have a different Shoemaker sentence associated with them. The radical structuralist will think that there is no such multiplicity of possibilities.

If one were to put an intuitive gloss on what modest structuralism amounts to, I would suggest it is this: There are quiddities, though one can get a necessary and sufficient condition for being a particular quiddity in terms of its causal relations to other quiddities. Whether this attempt to find a middle ground between the merits of causal structuralism and a metaphysic of quiddities has much going for it, I leave as a question for another time.

Notes

1. Thanks to Tamar Gendler, David Lewis, Ted Sider, Zoltan Szabo, Dean Zimmerman and audiences at Cornell and Notre Dame for helpful discussion.
2. See, for example, his introduction to *Philosophical Papers Volume II*, Oxford University Press, 1986.
3. Lewis himself wavers on whether to endorse a full-blooded commitment to universals, conceived as more than set-theoretic entities. In this paper, I shall not be call-

- ing the existence of universals into question. I shall leave the reader to judge to what extent the issues are significantly affected by a shift to a set-theoretic conception of properties.
4. Such a view was not uncommon, I believe, in scholastic philosophy. For example, while William of Ockham shows little willingness to think that a thing's nature is exhausted by some causal profile, he claims in Reportatio IV q2 that it is part of the very nature of an effect that it can be produced by one kind of efficient cause and not another.
 5. I am not requiring of the 'quidditist' that he be ontologically serious about quiddities, considered as something metaphysically distinct from the properties themselves.
 6. The causal structuralist will thus see something deeply misleading in Hilary Putnam's distinction between 'a causal description' and 'a canonical description' of a property, where a canonical description is of the form 'the property of being F' while the causal description picks out a property via its causal role. For the causal structuralist, it is the causal description that most deserves the label 'canonical'. See 'On Properties' in *Mathematics, Matter and Method, Philosophical Papers Volume 1* (Cambridge University Press, 1975), p. 316.
 7. Cf. Shoemaker 'Causality and Properties,' in *Identity, Cause and Mind* (Cambridge University Press, 1984), p. 207.
 8. For one thing, if a number of profiles were sufficient, it would seem that the disjunction of them would be necessary. For another, I wouldn't know how to motivate the thesis that a pair of distinct possible properties could not share some profile once one has denied that the profile is necessary to either.
 9. Cf. Thomas Aquinas, who thought that form in the case of angels, and that form plus a certain originating quantity of matter in the case of corporeal substances (where 'quantity of matter' was not conceived of haecceitistically) was sufficient for individuation. See his *On Being and Essence*.
 10. 'Causal and Metaphysical Necessity,' *Pacific Philosophical Quarterly* 1998, Volume 79 p. 66.
 11. 'Causal and Metaphysical Necessity,' p. 69–70.
 12. Cf 'Causality and properties,' p. 218: 'the assertion that a certain particular might have had different properties than it does in the actual world...implies that there is a possible history 'branching off' from the history of the actual world in which it acquires those properties...'
 13. 'Causality and Properties,' p. 215.
 14. Cf David Lewis 'Elusive Knowledge,' *Australian Journal of Philosophy* 74 (1996), p. 549–567.
 15. Similar considerations apply to the spectre that there are lots of inert properties that bear heavily on what really resembles what, falsifying most of our overall resemblance judgments as between particulars. See p. 215 of 'Causality and Properties'.
 16. Alternatively: negative charge 1 and negative charge 2 are exactly alike except that negative charge 1 causes epiphenomenon 1 (which is itself causally inert viz a viz our detection devices) and negative charge 2 causes epiphenomenon 2 (which is similarly inert).
 17. I am grateful to David Lewis here, who proposed something like the argument that follows in discussion at Notre Dame, 1999.
 18. 'On The Possibility of Philosophical Knowledge', *Philosophical Perspectives 10: Metaphysics*, 1996, p.23.

19. In the lingo made popular by David Chalmers' *The Conscious Mind* (Oxford University Press, 1986), the issue is whether there is a primary intension that always delivers the same property for any world, considered as actual, not whether there is a secondary intension that delivers the same property for any world considered as counterfactual.
20. Though it may make trouble for the philosophical idea, in circulation since antiquity, that forms are more fully knowable than particulars, at least where the forms concerned are the causally efficacious properties that are found in the natural world.
21. *The Analysis of Matter* (Kegan Paul, 1927), p. 264.
22. Russell's concern, taken from *The Analysis of Matter*, p. 325, is voiced by Simon Blackburn ('Filling in Space' in *Essays in Quasi-Realism*, pp. 255–258), who adds that the point is particularly pressing if one uses a possible worlds account of powers).
23. See his *What is a Law of Nature?* (Cambridge University Press, 1983).
24. *On The Plurality of Worlds* (Blackwell, 1986), p. 87–88.
25. *On The Plurality of Worlds* p. 91.
26. *The Analysis of Matter*, p. 264.
27. Note that if one wishes to suppose that the quiddity of a property is always something like experiential character, then one will be driven to a sort of panpsychism and will not tolerate the possibility of zombies after all. Quiddity switching may make for spectrum inversion but cannot, on this view, turn the lights off.
28. *From Metaphysics to Ethics* (Clarendon Press, 1998), pg 24.
29. On this last point, see Blackburn's 'Filling in Space'. It is of course not decisive against the Humean—what the latter thinks as the appropriate explanation for some dispositional property is a combination of one or more perfectly natural properties *together with* some suitable set of laws.
30. Especially if identity, mathematical relations (such as the successor relation) and mereological relations (such as being a part of) disappear into variables.
31. Admittedly, what would be nicer still would be a sketch of how circumstances in the actual practice of science might compel the positing of a symmetrical structure such as the one above (embedded, presumably, in a more complex structure that preserves the symmetry between A and B.) Mill's methods of agreement and difference certainly have no straightforward bite here, since the reidentification of circumstances is in this context epistemologically problematic. We normally discriminate properties by their differential impact on our sensory organs or on some detection instrument. But in the case described there is no straightforward basis for such discrimination. Where I posit a structure like the one above, you may posit a structure whereby there are four properties such that $A \ N \ C$ and $B \ N \ C$ and $B \ N \ D$. Nevertheless, the symmetrical structure does seem perfectly possible and well-motivated by proper principles of plenitude. It strikes me as ad hoc to deny its possibility.
32. 'The Identity of Indiscernibles,' in *Problems of Analysis*, Routledge and Kegan Paul, 1954, p. 80–92.
33. See *Ordinatio* II d3 p1 q7
34. *On The Plurality of Worlds*, p. 221.
35. *On The Plurality Of Worlds*, p. 231.