

Why theories of causality need *production*: an information-transmission account

Phyllis McKay Illari

Draft of November 16, 2010

Abstract In this paper, I examine the comparatively neglected intuition of *production* regarding causality. I begin by examining the weaknesses of current production accounts of causality. I then distinguish between giving a good production account of causality, and a good account of production. I argue that an account of production is needed to make sense of vital practices in causal inference. Finally, I offer an information-transmission account of production based on John Collier's work, that solves the primary weaknesses of current production accounts: applicability and absences.

Keywords Causality, production, causal inference, information, absences, Collier.

1 Introduction

Ned Hall argues that 'Causation, understood as a relation between events, comes in at least two basic and fundamentally different varieties.' [Hall, 2004, p225.] These are:

- Dependence: about dependency of E on C .
- Production: about link, biff, oomph between C and E .

To illustrate, consider Wesley Salmon's technological example: 'when I arrive at home in the evening, I press a button on my electronic door opener (cause) to open the garage door (effect). First, there is an interaction between my finger and the control device, then an electromagnetic signal transmits a causal influence from the control device to the mechanism that raises the garage door, and finally there is an interaction between the signal and that mechanism.' [Salmon, 1998a, pp17-8.]

Here, Salmon traces the link between C and E , and this is the *production* sense in which pressing the button causes the garage door to rise. There is also a dependency between the garage door rising and the button pressing. Perhaps if the button had not been pressed, the garage door would not have risen; or

pressing the button increased the probability of the garage door rising. This is Hall's *dependence*. Hall argues that there is no hope for a univocal account of causality, because *production* and *dependence* each do different jobs. He writes: 'Counterfactual dependence is causation in one sense: But in that sense of "cause", Transitivity, Locality and Intrinsicness are all false. Still, they are not false simpliciter; for there is a different concept of causation - the one I call "production" - that renders them true.' [Hall, 2004, p253.] It will be a major aim of this paper to examine these different jobs, particularly the job of *production*.

It has been recognized that persistent counterexamples to *dependence*, or what I will often call *difference-making*, come from *production* intuitions, and vice versa. To see the problem for *dependence* or *difference-making*, suppose Billy and Suzy are throwing stones at a bottle. Billy's stone hits the bottle, shattering it, while an instant later Suzy's stone whistles through the space and flying shards of glass left behind. Suzy's backup stone removes the dependence of the bottle shattering on Billy's throw. Suzy's stone would have shattered the bottle if Billy's had not been present. But since we can trace the link between Billy's throw and the shattering, we unambiguously identify Billy's throw as the cause. To see the problem for *production*, suppose Billy and Suzy are pilots in a bombing raid on an enemy city. Suzy flies the bomber to the city and drops the bombs. But Billy is her escort fighter, successfully shooting down all the enemy fighters that attempt to destroy Suzy's plane. There is no continuous link between Billy's actions and the bombing, as there is with Suzy's actions. Nevertheless, the bombing depends on Billy's actions, rendering him a cause. This suggests what Hall argues: that *dependence* and *production* answer to conflicting intuitions regarding causality, rendering a univocal account impossible. This argument has been a powerful motivation for pluralism about causality.

Accounts in terms of dependence or difference-making are clearly in the ascendancy. Consider the literature:

- **Difference-making:** Regularity; probabilistic; counterfactual; invariance; causal modelling.
- **Production:** Reichenbach-Salmon mark transmission; Salmon-Dowe's conserved quantities theory; some Anscombean pluralist primitivists, that may include Cartwright, and Machamer, Darden and Craver; Glennan's mechanistic theory.¹

The majority of work is still on difference-making accounts, with the literature on each type of difference-making account being vast. Correspondingly, the literature on production is a drop in a counterfactual ocean. Further, existing production accounts have recognized weaknesses, and few seem inclined to try to fix them.

¹I will expand on all these theories later, in section 2. There are also some less well known accounts that classify as production accounts, such as the brief account of Hall [2004]. I will set these aside here.

In this paper, I will examine the weaknesses of existing production accounts of causality in section 2. I will consider what an account of *production* might be for - what the job of *production* is - and so the continued (or not) usefulness of some intuitions regarding production, in section 3. In section 4, I will use John Collier's work to argue that an account of production in terms of *information* has some tempting features in terms of applicability, and argue that it offers a novel solution to the problem of absences (in section 5). The account it offers is particularly appropriate to causality in the domain of technology because technology often succeeds by mixing forms of information-transmission, particularly with the integration of the human agent with the technology - the extension of the human agent.

My aim here is to develop an information-transmission account of production. This account is of interest both to those interested in a univocal account of causality in terms of production, and to pluralists interested only in an account of production. I am not directly addressing the debate between pluralism and monism concerning causality. My own view is currently undecided. This is because I think we cannot complete our understanding of the distinction between difference-making and production until we understand production as well as we understand difference-making. And no good applicable account of production so far exists. So examination of the pluralism-monism issue must wait until after the project of this paper. Nonetheless, the project of this paper will further the pluralism-monism debate, by rigorously examining the comparatively neglected production intuition.

2 Current production accounts of causality

In this section I will discuss prominent production accounts of causality in turn, and their known weaknesses, to assess the current problem for such an approach to causality.

Salmon's earliest process theory, inspired by Reichenbach, says that a causal process is one that can transmit marks. (Salmon [1980], Salmon [1984], Salmon [1997], Salmon [1998b].) Salmon eventually rejected this theory on the grounds that it involves counterfactuals: causal processes are not those that actually transmit marks, but are defined in terms of what *would* happen if a mark *were* to be introduced. He moved to the account in terms of conserved quantities developed by Dowe.

For Salmon and Dowe, a process is causal if it transmits a conserved quantity, such as energy-mass, charge, or momentum. An interaction between two processes is causal if they exchange a conserved quantity. (Dowe [1992], Dowe [1993], Dowe [1996], Dowe [1999], Dowe [2000b], Dowe [2000a].) This is probably still the best-known production account of causality, in spite of four well-known problems with it.

First, if causal claims are *contextual*, then whether C causes E can alter with alterations in context. For example, whether a particular ceremony causes a legal marriage depends on the laws of that country - which presumably don't

alter whether conserved quantities are transmitted between C and E .

Secondly, we think that causes are effective in virtue of some but not all of their properties - the *relevant* ones. The movement of one billiard ball causes the movement of the next in virtue of properties such as its momentum, but not properties such as its colour. But the conserved quantity theory doesn't discriminate among properties. It tells us only that there is a genuine causal process here, but doesn't give us the resources to discriminate among genuine causal processes.

Third, the only properties the conserved quantity theory could possibly pick out as relevant are conserved quantities. These are relatively few, such as charge, mass, momentum and so on. But in the vast majority of cases of causality in the special sciences, these are not the relevant properties at all. The vast majority of the special sciences concern themselves with measurements of quite different kinds of properties. Charge, mass, and momentum seem incidental to such causal claims as 'smoking causes cancer', since the various sciences of cancer do not concern themselves with charge, mass or momentum. This is the problem of *applicability*.

Finally, there is the problem of *absences*. (See Lewis [2004].) We make claims such as 'I was late because my bus didn't turn up'. But presumably there can be no transmission of conserved quantities between an absent bus and my lateness. Dowe has a response to this which I will discuss later.

Turning to newer theories, Glennan's mechanistic theory of causality says: 'a relation between two events (other than fundamental physical events) is causal when and only when these events are connected in the appropriate way by a mechanism.' [Glennan, 1996, p56].

One major problem for this theory is that when Glennan spells out what a mechanism is, we can see that his theory is not really a new theory of causality. If mechanisms are chains of laws Glennan [1996], invariance relations Glennan [2002], or arrangements of singular determination relations Glennan [forthcoming] - which are all familiar tools used in old theories of causality - then the view merely adds some bells and whistles to existing theories. For the most part, the bells and whistles have been added to old *difference-making* theories, too. Glennan's view clearly does better than Salmon-Dowe on applicability, since mechanisms are ubiquitous in the special sciences, but he suggests that his view does not apply to fundamental physics. It also has some counterexamples in other domains. (See Illari and Williamson [forthcoming] for discussion.) Finally, Glennan offers no explicit answer to the problem of absences.

The extra bells and whistles Glennan offers matter, and there may of course be answers to offer to the old problems for production accounts. Indeed, the account I offer in section 4 can be seen as an attempt to develop more carefully the relation between causality and mechanisms.

The final view that I classify as a possible production account of causality is a second popular form of pluralism about causality. Anscombe [1975] claims that there are many thick causal concepts such as pushing, pulling, breaking, binding, and so on. Anscombe denies that there is *anything in common* to all such thick causal concepts that we can identify as causality. Recent work by philoso-

phers such as Cartwright [2004], working on capacities, and Machamer et al. [2000], working on activities, explicitly agree. In so far as they see Anscombe's thick causal concepts as primitive connectives, they seem to hold some form of production view of causality. They do go further than Anscombe, however, in offering at least *some* second-order claims about capacities and activities, respectively.

These views nicely embrace the full diversity of causal claims across domains in the sciences, cheerfully covering everything. Nevertheless they are unsatisfying in denying that there is *anything* to be said about production in general - or at least not very much. They also offer no explicit answers to the classic problems for production - including absences. Presumably we are to accept 'breaking' and 'preventing' - and any other causative verbs that might invoke absences - also as primitive causal connectives, about which we cannot say (much) more. Again, the account I offer in section 4 can be seen as an attempt to keep the general applicability of this account, but say something more about what these thick causal concepts share.

This brief survey has drawn out the major problems for current production accounts of causality, which will need to be addressed by any successful production account: context, relevance, applicability and absences. Before going on to tackle them, I will pause in the next section to consider what it is that a production account is for. Should we try to retain a production account in the face of these problems? What is it supposed to achieve? This will allow me to probe more deeply into the problems, and identify which problems an account of production should care about most.

3 What's an account of production *for*?

In this section I will examine reasons for holding on to production intuitions, disentangle the idea of a production account of causality from an account of production *tout court*, and further clarify what the aim of my account will be. I will contrast the job of production with the job of accounts of difference-making, although I will not linger on these since my primary aim is to figure out the job of production.

Production is a traditional element of our intuitions about causality. It is a persistent and remarkably stable intuition. For example, in the Billy and Suzy case above, where Billy's stone breaks the bottle an instant before Suzy's, people are in remarkably stable agreement that Billy's throw is clearly the cause, even in the absence of dependence. (See Gopnik and Schulz [2007] for a useful survey of some of the relevant literature in psychology.) But intuitions alone are not enough. We can be mistaken even in our deepest intuitions, such as the intuition that every event has a sufficient cause, that yielded in the end to quantum mechanics. We need to have a reason to hold on to intuitions, to believe that they are guiding us right. Finding such good reason can also help hone the correct content of an intuition, to make it more precise.

The non-applicability of existing substantive accounts of production, and

the unsatisfying nature of the applicable accounts, might lead us to suppose that we can safely just ignore production, at least in the special sciences. So the important question is: are there reasons coming from the sciences to hold onto an account of production? In this section I will argue that there is one primary reason, and that careful examination of this reason usefully constrains what an account of production should aim to do.

Many philosophers of science now think we use evidence of mechanisms in causal inference at least sometimes. (See Illari [under review], Weber [2009], Leuridan and Weber [forthcoming], Broadbent [forthcoming], Gillies [forthcoming], Steel [2008], Kincaid [forthcoming]) If this is true, we are committed to there being something in our accounts of causality that is at least consistent with this, and if possible helps us *understand* it. There have been several *possible* roles for mechanisms suggested, including causal explanation (Vreese [2008], Leuridan and Weber [forthcoming], Russo and Williamson [2007]); assessing the stability of established causal claims (Leuridan and Weber [forthcoming]); assisting with the problem of external validity, where you believe that C causes E in a particular population, and you wish to claim that C also causes E in a similar, but distinct, population (Leuridan and Weber [forthcoming], Russo and Williamson [2007]); and linking mechanisms (Illari [under review], Russo and Williamson [2007]).

These are all interesting ideas but I will focus on only one: the idea of linking mechanisms. The idea is that finding mechanisms helps us in causal inference by ruling *in* and ruling *out* possible causal links in a domain. This source of evidence for and against causal claims can then be matched to the difference-making relations we find - usually correlations. If nothing else, it helps constrain the possible causal claims in domains where the causal structure is underdetermined by difference-making relations alone.

It will help to examine this idea further with reference to some examples. There are some important constraints that we routinely impose even on special science mechanisms - so routinely that they are seldom mentioned explicitly. Consider:

1. Causal influence cannot travel faster than the speed of light.
2. Energy constraints are important in biochemical mechanism discovery, such as metabolic pathways.

The first is a deliverance of physics that we hardly ever consider. But it is a powerful constraint on possible causes. If an event is not in the backwards light-cone of E , it is just not a candidate possible cause of E . More prosaic constraints are more regularly important. If a posited reaction in a biochemical pathway requires more energy than is available, that mechanism is ruled out, or an alternative source of energy sought.

Dependence accounts don't explain these kinds of findings. There is nothing in the idea of counterfactual dependence, or probabilistic relations, for example, to yield these kinds of constraints. Counterfactual dependence and probabilistic relations merely characterize a relationship between cause and effect, which may

be significantly distant from each other in spacetime, and give no detail of any link. So there is no way for such empirical constraints on how such a link can or cannot be to impact on a difference-making theory. They have to be introduced ad hoc, as add-ons.

But we *need* something that answers those constraints in our account of causality, or there is nothing to explain some *vital* inferential practices. This is the best argument for not giving up on the idea of causation having something to do with linking or connection - in a way that can be constrained by empirical findings.² Note that I don't think that *any* production account will address this problem successfully. It is the primary aim of particular accounts of production that they should do this well, and I will have this in mind when examining information-transmission.

I have been writing in this section of an account of production, rather than of a production account of causality. This is because I suggest we make a distinction between:

1. A good production account *of causality*
2. A good account *of production*

As a way to make progress, I will consider only how to formulate an account *of production*. By making this switch away from giving a full-blown account of causality, I don't have to deal with the counterexamples coming from difference-making theories. I am only interested in a good account of production, which is an account that does well the particular job of an account of production, and may not do well on the jobs that fall to difference-making accounts.

- Jobs for accounts of difference-making:

Context

Relevance

Recall that these are the first two well-known problems for the Salmon-Dowe production account *of causality* that I discussed in section 2. But these are not problems for an account of production that does not claim to be a univocal account of causality. These are the jobs that difference-making accounts do well. Relevant properties of the cause are very much difference-making properties, on which features of the effect depend, as distinct from the properties that make no difference. Changing the colours of interacting billiard balls does not make a difference to their movement - their movement does not counterfactually depend on, is not probabilistically affected by, their colour, for example. Counterfactual,

²Other elements of production have been discussed by various philosophers. For Hall [2004], the core of production is transitivity, locality and intrinsicness. Many philosophers seek the Humean secret connexion: oomph, biff or glue. We might add singularism, as a core concern of Glennan's, although some difference-making accounts are also singularist, such as counterfactual dependence. Others are concerned with realism - production as the thing in the world Dowe [2004]. I suggest we let the needs of causal inference guide us in forming an account of production, and see whether any of these traditional elements of production are satisfied afterwards. I will not address them further in this paper.

probabilistic, invariance and causal modelling theories will all characterize this very nicely. Context is also a matter of difference-making. That context might alter which properties are difference-makers is no problem for difference-making theories. That patterns of counterfactual dependence or probability distribution between C and E might change with context introduces no metaphysical paradox. Woodward [2003] includes the place of context explicitly in the choice of variables. Difference-making theories do these jobs well.

Further, these jobs are not *important* to an account of production. This is because they have nothing to do with the mere possibility or impossibility of a link existing between cause and effect. But linking is what accounts of production are about. So I will leave these problems for difference-making accounts to solve, and move to consider the real problems for accounts of production.

- Remaining worries for accounts of production:

Applicability without disunity

Absences

These are the remaining two well-known problems for the Salmon-Dowe account that I discussed in section 2. But with the primary purpose of an account of production firmly in mind, we can see that these are serious problems. They *cannot* be sidestepped. To hold that production is applicable only in fundamental physics is unacceptably reductive, and means that standard constraints on causal inference cannot be applied the way they routinely are in the special sciences. The real problem of absences for production is also now revealed. If production is about a possible link between a postulated cause and an effect of interest, a gap between cause and effect ought to rule out a link, and so rule out any such cause-effect relationship. Unlike relevance and context, if these problems cannot be solved, no satisfactory account of production is available. I will further refine the problem of absences in section 5.

Once we consider the literature in this way, we can see that we have many good accounts of difference-making. They give no account of linking, but they deal well with relevance and context. We understand them well, including having a reasonably precise statement of their strengths and weaknesses. But we are still in need of a widely-applicable good account of production. To reiterate, seeking such an account is also an excellent way of rigorously examining the production intuition itself, disentangling it from difference-making intuitions, and seeing whether it can truly stand on its own. Ultimately a better understanding of production itself might either lead us back to a univocal account of causality, or establish pluralism convincingly.

Even with the aim narrowed - to give an account of production rather than of causality in general - and the job of production more clearly delimited as concerning linking, it is worth pausing to consider, finally, what an *account* of production - a philosophical account - might offer.

Dowe [2004] distinguishes between a conceptual and an empirical account of causality. The aim of a conceptual account is to tidy up our *concept* of causality, which exists in thought or language, or both. The aim of an empirical account

is to say what causality is, in the *world*. Dowe is right to raise the issue. But I do not think the two projects can succeed apart. My aim is to say what's in the world such that our concept of production tracks something, at least reasonably successfully. This is an empirically-motivated way of developing a concept.

The account I offer might be reductive. But I am not interested in eliminating causal talk, so this is not my primary aim. What's most interesting to me is the *generality* of the account. I am seeking to say what it is that is in the world, that is present most generally in cases of productive causes. I hope the resulting general account will be fruitful in theorizing and in methodological practice. If we can see what productive causes have in common, this will deepen our understanding of them, and also contribute to deepening our understanding of their use in causal inference. A background of unity also gives a nice counterfoil to understanding the plurality of productive causes undoubtedly described in various scientific fields. I will expand on this in subsection 4.3.

4 An information-transmission account of production

4.1 The applicability problem

The work of the previous sections has established that I should look for an account of production appropriate to the use of linking mechanisms to rule in and rule out possible causes in causal inference. I want an account that is as general as possible, widely applicable across the sciences, so allowing it to be illuminating about our inferential practices. Finally, I am looking for a solution to the problem of absences. In this section I am going to argue (in subsection 4.1) that the applicability problem suggests a move towards seeing production in terms of information-transmission, explain (in subsection 4.2) Collier's account of causality as information-transmission, and offer (in subsection 4.3) a development of Collier's account. I will move on, in section 5, to a solution to the abiding problem of absences.

Recall that Anscombe [1975] argues that there is no univocal conception of cause available. The mechanisms literature has taken this point on board and Machamer et al. [2000], while making some second-order claims, talk of diverse activities, which do indeed seem to be common in the language of mechanisms of the special sciences: binding, breaking, metabolising, unfurling, pushing, pulling, and so on. Descriptively, Anscombe is clearly right that these activities are ubiquitous in the special sciences, and that they are diverse. (See Illari and Williamson [under review] for further discussion.) We have come from the problem of Hall-pluralism about causality to the challenge of another form of pluralism - the diversity of thick causal concepts. It is right to look to the sciences to tell us about what constitutes linking or connection. The constraints we use to guide our causal inferences in the ways I have described are not conceptual constraints, but *empirical* ones. Perhaps the sciences themselves tell us that there is nothing much more general that can be said than: we find out

about possible causal links - the thick causal concepts, or primitive connectives - which are different in different domains.

Where should we look in the quest to find *anything* that is more unifying than Anscombe's view, more widely applicable than Salmon-Dowe's view, and a more substantial than Glennan's view? It is true, as Glennan notes, that many sciences talk about mechanisms, but I have explained why their relation to causality is not so simple to spell out as Glennan claims. Something else we find in the sciences is lots of people talking about information - right across different scientific disciplines.

4.2 Collier's Information-transmission theory

In view of this, it is not surprising that a few philosophers are also beginning to develop thinking on causality and its relation to information. Interestingly, two brands of structural realism have moved towards a view they both describe as 'information-theoretic structural realism' (Ladyman and Ross [2007], and Floridi [2008]). Crudely, these views claim that the only fundamental thing there is is information, which seems to imply that if causality is anything, it is information transfer. [Ladyman and Ross, 2007, p210-1] explicitly recognize this: the 'special sciences are incorrigibly committed to dynamic propagation of temporally asymmetric influences Reference to transfer of some (in principle) quantitatively measurable information is a highly general way of describing any process. More specifically, ... if there are causal processes, then each such process must involve the transfer of information between cause and effect'. In an apparently unrelated piece, Machamer and Bogen [forthcoming] offer the only attempt I know of to explain the relation of information to mechanistic hierarchy. I will return to this point in subsection 4.3.

This literature is small, but complex and highly technical, so I lack space to survey it all. In this paper, I will concentrate on John Collier's work in Collier [In press] and Collier [1999]. This is the earliest claim that causality is information-transfer, which influenced Ladyman and Ross. It also has the advantage of being a theory of *causality*, without the surrounding complexities of structural realism and mechanisms.

In Collier's own words: 'The basic idea is that causation is the transfer of a particular token of a quantity of information from one state of a system to another.' [Collier, 1999, p215.] He fills this out by offering an account of what information is - including what the particular case of *physical* information is - and an account of information *transfer*. I will explain these in turn.

For Collier, the information in a thing - in the first instance, in a *static* thing - is formally, and objectively, defined in terms of computational information theory:

In the static case, the information in an object or property can be derived by asking a set of canonical questions that classify the object uniquely ... with yes or no answers, giving a 1-1 mapping from the questions and object to the answers. This gives a string of 1s and

0s There are standard methods to compress these strings The compressed form is a line in a truth table, and is a generator of everything true of the thing required to classify it. There need not be a unique shortest string, but the set will be a linear space of logically equivalent propositions. The dimensionality of this space is the amount of information in the original object. [Collier, In press, p2.]

This definition is very widely applicable. It is important that it is quite independent of human minds, and so applies easily to the physical world, as well as to the special sciences. Indeed, ideas of information are being extensively developed in physics.

Physical information is a special case of the more general definition of information. To get physical information, Collier adds two further constraints. The first is that physical information in a system breaks down into what Collier calls ‘intropy’, which is broadly free energy, and ‘enformation’, which is broadly structural constraints. For example: ‘A steam engine has an intropy determined by the thermodynamic potential generated in its steam generator, due to the temperature and pressure differences between the generator and the condenser. ... The enformation of the engine is its structural design, which guides the steam and the piston the steam pushes to do work. The design confines the steam in a regular way over time and place.’ [Collier, 1999, p228.] The second constraint is the Negentropy Principle of Information (NPI), which is an interpretative heuristic that allows Collier to connect information theory to physical causation via physical entropy - which is clearly a physical item. [Collier, 1999, p226ff.] discusses the details, but these are not crucial for understanding Collier’s claim.

This yields: ‘The resulting notion of causal process is: P is a physical causal process in system over a series of states S_i from time t_0 to t_1 if and only if some part of the enformation is transferred from t_0 to t_1 , consistent with NPI’. [Collier, In press, p10.] Recall that this account is of the *static* case. What Collier’s definition means is that the steam engine is a physical causal process in so far as some part of its structural design is maintained through time. It is this structural design which allows intropy to do work. Returning to Salmon’s example of the remote control for opening the garage door, there is a physical causal process in this sense if the form of the remote, and the receiver, and the garage door mechanism, is maintained through time.

The static case above helps introduce the basics of causality as information, but the full theory requires attention to what happens in the *dynamical* case - when information is not only present, but flows. Indeed, what happens when information flows is vital to my aim of understanding causality as connection, and the kinds of constraints this imposes on possible causes, such as that causal influence cannot propagate faster than the speed of light. Illuminating such constraints is what I take to be the primary aim of an account of production.

In Collier’s early paper, information-flow is defined very simply, in terms of identity: ‘P is a causal process in system S from time t_0 to t_1 iff some particular

part of the information of S involved in stages of P is identical at t0 and t1.’ [Collier, 1999, p222.] Collier [In press] uses the theory of information flow of Barwise and Seligman [1997] to fill this out considerably in terms of information channels: ‘An information channel for a distributed system is an indexed family of infomorphisms with a common core codomain C. The infomorphisms allow information to be carried from one part of the system to another. For example, in a flashlight, the components might be a bulb, battery, switch and case. The channel is basically a connected series of infomorphisms from switch to bulb through the mediation of battery and case.’ [Collier, In press, p6.] ‘Infomorphism’ means something like the following: Consider two systems IS1 and IS2. Each system consists of a set of objects, and each object has a set of attributes. For example, an object might be a switch, and its possible attributes would be on or off. The “state” of a system means the attributes of all its objects. If, by knowing the state of IS1 you can always tell the state of IS2, then you have an infomorphism. Note that knowing the state of IS1 will not always tell you everything about the state of IS2 - some information might be lost. The rules (or functions) that let you infer the state of one system from the state of the other is what defines the infomorphism.

Collier’s final view is:

P is a causal *connection* in a system from time t0 to t1 if and only if there is an channel between s0 and s1 from t0 to t1 that preserves some part of the information in the first state. Furthermore, P is a *physical causal process* in system over a series of states si from time t0 to t1 if and only there is a channel through the states from t0 to t1, consistent with NPI, and over every intermediate state. [Collier, In press, p10-1, emphasis added.]

Note that Collier here offers an account both of causal connection, and of a physical causal process, and there is no longer a distinction between intropy and enformation. To illustrate, consider again Salmon’s electronic door opener. This is an information channel, through which we can trace infomorphisms from the button, through the signal, to the door opening mechanism.

For Collier, the causal connection in the most recent theory is still fundamentally identity, and information flow is given in terms of identity of information at various stages in an information channel. He writes: ‘the connection in this case is identity, which is perhaps the strongest connection one can have, and requires information transmission across time: it is the identical token of information.’ [Collier, In press, p11-2.] Notice finally that Collier *adds* the intermediate state constraint on physical causal processes. He says explicitly that this may be unnecessary at [Collier, 1999, p230]. I will return to this point in section 5.

Collier’s theory is very general, both in domain of application and in its relation to other theories of causality. Collier suggests that his theory: ‘applies to all forms of causation, but requires a specific interpretation of information for each category of substance (assuming there is more than one).’ [Collier, 1999, p215-6.] I will return to this idea shortly. Collier also claims: ‘my approach can be interpreted as a version of each of the others [standard approaches to

causality] with suitable additional assumptions.’ [Collier, 1999, p235.]³ As far as other production theories go, Collier claims that his theory entails the conserved quantities view in that domain. It can also be seen as an extension of the mark-transmission view. But Collier’s view avoids the problem that some causal processes cannot be marked, if we take the form of the states of a physical process to itself be a mark. They all *have* form anyway, so such a ‘mark’ doesn’t need to be introduced. This also removes the need for counterfactuals in giving the theory. [Collier, 1999, p229.]

4.3 Mechanisms as Information-channels

I have argued for the advantages of Collier’s theory as a widely applicable account of production that offers a view of causal connection and the flow of causal influence. These are advantages he does not himself stress, and I have strengthened them by explaining their value in the current debate over production versus difference-making accounts. Nevertheless, there are other clear advantages of the theory that Collier does advertise, in particular that other theories of causality come out as special cases of his more general view.

In view of these advantages, it is surprising that Collier’s theory has not been more carefully examined. I suspect the technical difficulty of understanding it has contributed to this, as has the challenge of drawing out the implications of the theory for the current concerns of the causality literature. I hope my work has helped to draw out some of these implications.

In this subsection, I will now attempt to link Collier’s theory even more intimately to existing views of production, specifically those of Glennan, and the Anscombe-inspired primitivists. I will do this by suggesting that information-transmission as an account of production should be linked to *mechanisms*. I will argue that the theory solves the problem of absences in section 5.⁴

[Machamer et al., 2000, p7] make an interesting complaint against Salmon’s theory: ‘Although we acknowledge the possibility that Salmon’s analysis may be all there is to certain fundamental types of interactions in physics, his analysis is silent as to the character of the productivity in the activities investigated by many other sciences. Mere talk of transmission of a mark or exchange of a conserved quantity does not exhaust what these scientists know about productive activities and about how activities effect regular changes in mechanisms.’ This complaint might be extended to information-transmission. It apparently identifies an analogous lack in the theory.

But this lack can be addressed by looking to mechanisms, and there are resources in Collier’s theory that invite this extension. Collier notes that on his account information-transmission is relative to a *channel* of information: ‘For example, what you consider noise on your TV might be a signal to a TV repairman. Notice that this does not imply a relativity of information to

³This suggests a univocal account of causality. As I have said, I am here interested only in an account of production, and will reserve for further work the examination of whether this account can be extended to do the jobs of difference-making accounts.

⁴These are developments that Collier is highly sympathetic to (private communication).

interests, but that interests can lead to paying attention to different channels. The information in the respective channels is objective in each case, the noise relative to the non-functioning or poorly functioning television channel, and the noise as a product of a noise producing channel - the problem for the TV repairman is to diagnose the source of the noise via its channel properties.’ [Collier, In press, p8.] By seeing *mechanisms* as information channels, we can relate Collier’s theory to Glennan’s, to Anscombe’s, and to the rapidly-growing literature on causal explanation and causal inference using mechanisms in the special sciences.

It is well known that mechanisms are relative to the phenomenon they explain. (See Bechtel [2008], Darden [2006], Illari and Williamson [2010].) In mechanistic explanation, the phenomenon is explained by decomposing it into lower-level entities and activities (or Glennan’s parts and interactions), which in turn may be further decomposed into yet lower-level entities and activities, creating a functional hierarchy of nested mechanisms with the original phenomenon at the top. When Collier says the regularities of a distributed system are relative to its analysis in terms of information-channels, this can easily be read as a claim about the functional hierarchy of mechanisms. The functional hierarchy gives you the infomorphisms, rendering a mechanism an information channel. The work of Machamer and Bogen [forthcoming], relating mechanistic hierarchy to causal continuity, is of course immensely interesting, but Collier’s view is far more general. Machamer and Bogen’s account will only capture one kind of mechanistic information.

So Collier’s work can be seen to be directly related to Glennan’s, and an explanation of how *production* is related to mechanisms. In so far as the activities in mechanisms are all Anscombe-style causative connectives, as MDC explicitly claim, they should come out as connections in the information-channel, thus relating Collier’s work to the Anscombe tradition. Collier gives us what is generally true of all cases of causative connection: they are all cases of information transfer.

And this is the point at which the domain-specific empirical work is connected to Collier’s theory. What kinds of connections exist in each domain is discovered empirically, to yield the greater informativeness that MDC claimed was lacking in the conserved quantities account. Unlike the conserved quantities theory, information-transmission admits of *modes*. Different scientific fields can be taken as studying different modes of information-transmission, and empirically discovering what constraints on information-transmission exist in their particular domain. Machamer and Bogen’s information will plausibly come out as a special case.⁵ Thus the information-transmission account naturally illuminates our successful causal inferential practices.

So my view is that causality is information-transmission across a mechanistic hierarchy. Note that a mechanistic hierarchy *might* only be relevant to the

⁵I have given a unified account of production in terms of information. If the notion of *information* splinters, I will lose the unity I set out to achieve. I intend to address this by using work by Floridi that multiple notions of information are related, with some more fundamental than others. See Floridi [2010], Floridi [2009]. I reserve this for future work.

special sciences. The difference for physics - or perhaps merely for fundamental physics - might be dropping that bit of essential context. The mechanism gives you the thing in the world that holds together disparate causal relevance conditions. It is what we call the information channels we use to track causal connections, and constraints on causal connections, in building up our picture of the causal structure of our world.⁶

5 Absences

Finally, information-transmission also offers a novel solution to the problem of absences. It does this by giving a new view of connection that is *appropriate* to the special sciences, rather than one that is reductive in the way of Salmon-Dowe processes. The old problem of absences goes with the old reductive view of connection that we have discovered is *not* always appropriate to the special sciences.

I will begin by further refining the problem of absences for an account of *production*. Dowe [2001] gives a counterfactual account of quasi-causation, holding that causation of or by an absence is not full-blown causation. It is certainly tempting to write off causal claims citing absences as non-standard, merely analogous ‘quausal’ claims. The problem with this is that absences as causes, or indeed effects, are ubiquitous in causal discourse and scientific theorizing. They can’t be ignored or got rid of by an ad hoc add-on to a theory.

In an influential paper, Schaffer discusses the example of a gunshot through the heart causing death: ‘But heart damage only causes death by negative causation: heart damage (c) causes an absence of oxygenated blood flow to the brain (\neg d), which causes the cells to starve (e).’ Later he considers the gun: ‘But trigger pullings only cause bullet firings by negative causation: pulling the trigger (c) causes the removal of the sear from the path of the spring (\neg d), which causes the spring to uncoil, thereby compressing the gunpowder and causing an explosion, which causes the bullet to fire (e).’ [Schaffer, 2004, p199.] In that paper, Schaffer says that psychologists, biologists, chemists and physicists all routinely invoke negative causation. This is absolutely right. Taking just biochemistry, examples are numerous. Cells routinely alter which enzymes they produce in response to which metabolites are available. A cell stops producing lactase, for example, in response to the absence of lactose diffusing into the cells cytoplasm. Consider [Adams et al., 1992, p577]: ‘Phosphorylation of only

⁶There is an interesting possibility that this view will offer a univocal account of causality. Information-transmission deals well with production, and mechanistic hierarchy offers the possibility of solving the problems of context and relevance. It is well known that mechanisms set the context for causal claims. They also point you to the relevant properties - properties that make a difference to various effects of interest. If accounts of these can be given in informational terms, this will yield a univocal account of causality. If accounts of these require extra work, brought in from the mechanisms literature, that cannot be incorporated in the informational account, then the resulting account will still show how production and difference-making integrate, while remaining conceptually distinct. Either way, the view will be fruitful. I reserve this possibility for future work.

approximately 20% of eIF-2 can cause complete inhibition of initiation because there are only 20-25% as many molecules of eIF-2B as eIF-2.' Finally, absences are also very common parts of mechanisms, as Craver notes: 'Absences do not bear or exchange conserved quantities. They are not processes, they are not "things", properly speaking, and they do not exhibit consistency of characteristics over time. Nonetheless, the absence of the Mg²⁺ block does seem to cause Ca²⁺ to enter the cell. At least, this is what controlled experiments suggest: when the Mg²⁺ block is in place, the Ca²⁺ does not enter the cell. When the Mg²⁺ block is removed, the Ca²⁺ current begins to flow.' [Craver, 2007, p80]. These are all serious impediments to explaining away absence-causation by a manoeuvre such as Dowe's. An account of production must engage with the problem of absences directly.

The first thing to notice is that there is not one, but two, problems of absences.

1. Metaphysics of absence
2. A particular problem for production accounts (that difference-making accounts do not share [Schaffer, 2004, p294.]

Notice that *nobody* has a solution to the first problem. The metaphysical problem of what to say is the truthmaker for absence claims is not a problem specific to accounts of production. It is not even limited to accounts of causation, but is a general metaphysical problem. I need to focus on what the problem is that production accounts are supposed to have with absences that difference-making accounts do not have. I will use the discussion of section 3 above to significantly clarify what the problem of absences really is for an account of production.

Notice that one aspect of the uses of absences in scientific discourse is about identifying a property that is highly *relevant* to an effect, and pointing out what happens when it is not there. Since oxygen is necessary for human life, its absence is highly relevant to the non-continuance of life. Further, this often occurs against a *context* - of factors usually present, or usually absent. Mechanisms often set this context. These, however, are difference-making issues. These are the things that accounts of difference-making do well, and that a mechanistic hierarchy can help with. I will set them aside as not a concern of an account of production.

A serious issue of absences remains. The serious issue is that citing an absence, by introducing a gap between *C* and *E*, seems to rule out the possibility of a link between *C* and *E*. The solution to this problem lies in the way the idea of information can modify what we think of as a *gap*.

Information can be transmitted across what are traditionally regarded as physical gaps or disconnections. For example, if Billy says to Suzy that he'll meet her for lunch at 2pm, unless he phones her by noon to cancel, the absence of a phone call by noon tells Suzy to go and meet Billy at 2pm. Information *channels* can also involve absences. A binary string is just a series of 1s and 0s, which can also be a string of positive signals, and absences of a positive signal. For Floridi, this is the essential core of data - most fundamentally, a datum is

just a difference. He writes: ‘The fact is that a genuine, complete erasure of all data can be achieved only by the elimination of all possible differences. This clarifies why a datum is ultimately reducible to a lack of uniformity.’ [Floridi, 2010, p21.] This shows that what you think is a gap depends what you think the gap is a gap *in*.

If you’re looking for gaps in *information*-transmission, you’re looking for something different than when you look for gaps in continuous space-time. Many ‘gaps’ in continuous space-time will disappear against the different background of information. [Floridi, 2010, p31] notices this: ‘This is a peculiarity of information: its absence may also be informative.’ Schaffer [2004] says causation by absences shows that causation does not involve a persisting line, or a physical connection. I am arguing that what the problem of absences shows is that our conception of *connection* is wrong. We have not updated it for the diversity of causal connections the sciences are telling us exist in the special sciences. This outdated intuition is, I think, the correct explanation for the ‘intuition of difference’ discussed by Dowe [2001].

This is why it is important that Collier has to *add* locality constraints to his account of physical causation. They are not contained in the account already. [Collier, 1999, p223-4] notes: ‘Locality, both spatial and temporal, is a common constraint on causation. Hume’s “constant conjunction” is usually interpreted this way. While it is unclear how causation could be propagated nonlocally, some recent approaches to the interpretation of quantum mechanics (e.g. Bohm, 1980) permit something like nonlocal causation by allowing the same information (in Bohm’s case “the implicate order”) to appear in spatially disparate places with no spatially continuous connection. Temporally nonlocal causation is even more difficult to understand, but following its suggestion to me (by C.B. Martin) I have been able to see no way to rule it out. Like spatially nonlocal causation, temporally nonlocal causation is possible only if the same information is transferred from one time to another without the information existing at all times in between. Any problems in applying this idea are purely epistemological: we need to know it is the same information, and not an independent chance or otherwise determined convergence. Resolving these problems, however, requires an appropriate notion of information and identity for the appropriate metaphysical category.’

So ultimately it is a *virtue* of the information-transmission view that it allows the kinds of connections that exist in different domains to be an empirical issue, not an a priori one. It allows that space-time gaps are possible, if the identical bit of information can be transmitted across a spacetime gap. For different information channels - different mechanisms - we will have to discover what they do and don’t permit in terms of connection. This will be what constitutes identity of information in that information channel. This has completely altered the traditional problem of absences by changing the very idea of what constitutes a connection, or its correlate, a gap. In conjunction with the place of difference-making as what selects relevant properties, which are sometimes absences, this solves the traditional problem of absences.

6 Conclusion

In conclusion, I have argued for a view of production as information-transmission across a mechanistic hierarchy, drawing on the existing theory of John Collier. This offers an account of production that is widely applicable across the sciences, and solves the traditional problem of absences in an entirely novel way. I have investigated the primary aim of an account of production, arguing that it is to help us better understand our use of linking or connection in causal inference, and how the results of our empirical studies delimit linking or connection. The view of production as information-transmission satisfies this by allowing for different modes of information-transmission to be discovered in different scientific domains.

Acknowledgements I would like to thank the Leverhulme Trust for funding this work on mechanisms. I am also grateful to numerous colleagues and two anonymous referees for extensive discussions leading to improvement of the work. These colleagues include, but are not limited to: John Collier, Luciano Floridi, James Ladyman, Bert Leuridan, Federica Russo, Erik Weber and Jon Williamson. Remaining errors are, of course, my own.

References

- RL Adams, DP Knowler, and JT Leader. *The Biochemistry of the Nucleic Acids*. Chapman and Hall, 11 edition, 1992.
- GEM Anscombe. Causality and determination. In *E. Sosa, (ed.) Causation and Conditionals*. OUP, 1975.
- J Barwise and J Seligman. *Information Flow: The Logic of Distributed Systems*. CUP, 1997.
- William Bechtel. *Mental Mechanisms: Philosophical perspectives on cognitive neuroscience*. Routledge, Oxford, 2008.
- Alexander Broadbent. Inferring causation in epidemiology: mechanisms, black boxes, and contrasts. In Phyllis McKay Illari, Federica Russo, and Jon Williamson, editors, *Causality in the Sciences*. OUP, forthcoming.
- Nancy Cartwright. Causation: one word, many things. *Philosophy of Science*, 71:805–819, 2004.
- John Collier. Causation is the transfer of information. In Howard Sankey, editor, *Causation, natural laws, and explanation*, pages 215–263. Kluwer, Dordrecht, 1999.
- John Collier. Information, causation and computation. In Gordana Dodig Crnkovic and Mark Burgin, editors, *Information and Computation: Essays*

on Scientific and Philosophical Understanding of Foundations of Information and Computationc. World Scientific, In press.

Carl Craver. *Explaining the Brain.* Clarendon Press, Oxford, 2007.

Lindley Darden. *Reasoning in Biological Discoveries.* Cambridge University Press, Cambridge, 2006.

Phil Dowe. Wesley salmon's process theory of causality and the conserved quantity theory. *Philosophy of Science*, 59(2):195–216, 1992.

Phil Dowe. On the reduction of process causality to statistical relations. *British Journal for the Philosophy of Science*, 44:325–327, 1993.

Phil Dowe. Backwards causation and the direction of causal processes. *Mind*, 105:227–248, 1996.

Phil Dowe. The conserved quantity theory of causation and chance raising. *Philosophy of Science (Proceedings)*, 66:S486–S501, 1999.

Phil Dowe. Causality and explanation: review of Salmon. *British Journal for the Philosophy of Science*, 51:165–174, 2000a.

Phil Dowe. *Physical causation.* Cambridge University Press, Cambridge, 2000b.

Phil Dowe. A counterfactual theory of prevention and 'causation' by omission. *Australasian Journal of Philosophy*, 79(2):216–226, 2001.

Phil Dowe. Causes are physically connected to their effects: why preventers and omissions are not causes. In Christopher Hitchcock, editor, *Contemporary debates in Philosophy of Science.* Blackwell, 2004.

Luciano Floridi. A defence of informational structural realism. *Synthese*, 161: 219–253, 2008.

Luciano Floridi. Semantic conceptions of information. Online, Summer 2009. URL <http://plato.stanford.edu/archives/sum2009/entries/information-semantic/>. Edward N. Zalta (ed.).

Luciano Floridi. *Information: A Very Short Introduction.* OUP, 2010.

Donald Gillies. The Russo-Williamson Thesis and the question of whether smoking causes heart disease. In Phyllis McKay Illari, Federica Russo, and Jon Williamson, editors, *Causality in the Sciences.* OUP, forthcoming.

Stuart Glennan. Rethinking mechanistic explanation. *Philosophy of Science. Supplement: Proceedings of the 2000 Biennial Meeting of the Philosophy of Science Association. Part II: Symposia Papers (Sep., 2002)*, 69(3):S342–S353, 2002.

- Stuart Glennan. Singular and general causal relations: A mechanist perspective. In Phyllis McKay Illari, Federica Russo, and Jon Williamson, editors, *Causality in the Sciences*. OUP, forthcoming.
- Stuart S. Glennan. Mechanisms and the nature of causation. *Erkenntnis*, 44: 49–71, 1996.
- Alison Gopnik and Laura Schulz, editors. *Causal learning: psychology, philosophy, and computation*. Oxford University Press, New York, 2007.
- Ned Hall. Two concepts of causation. In LA Paul, EJ Hall, and J Collins, editors, *Causation and Counterfactuals*, pages 225–76. MIT Press, 2004.
- Phyllis McKay Illari. Mechanistic evidence: Disambiguating the Russo-Williamson Thesis. under review.
- Phyllis McKay Illari and Jon Williamson. Function and organization: comparing the mechanisms of protein synthesis and natural selection. *Studies in the History and Philosophy of the Biological and Biomedical Sciences*, 41:279–91, 2010.
- Phyllis McKay Illari and Jon Williamson. What is a mechanism?: Thinking about mechanisms across the sciences. forthcoming.
- Phyllis McKay Illari and Jon Williamson. In defense of activities. under review.
- Harold Kincaid. Causal modeling, mechanism, and probability in epidemiology. In Phyllis McKay Illari, Federica Russo, and Jon Williamson, editors, *Causality in the Sciences*. OUP, forthcoming.
- James Ladyman and Don Ross. *Every Thing Must Go*. OUP, 2007.
- Bert Leuridan and Erik Weber. The IARC and mechanistic evidence. In Phyllis McKay Illari, Federica Russo, and Jon Williamson, editors, *Causality in the Sciences*. OUP, forthcoming.
- David Lewis. Void and object. In EJ Hall LA Paul and J Collins, editors, *Causation and Counterfactuals*, pages 227–90. MIT Press, 2004.
- P. Machamer, L. Darden, and C. Craver. Thinking about mechanisms. *Philosophy of Science*, 67:1–25, 2000.
- Peter Machamer and Jim Bogen. Mechanistic information and causal continuity. In Phyllis McKay Illari, Federica Russo, and Jon Williamson, editors, *Causality in the Sciences*. OUP, forthcoming.
- Federica Russo and Jon Williamson. Interpreting causality in the health sciences. *International Studies in the Philosophy of Science*, 21(2):157–170, 2007.

- Wesley C. Salmon. Causality: production and propagation. In Ernest Sosa and Michael Tooley, editors, *Causation*, pages 154–171. Oxford University Press, Oxford, 1980.
- Wesley C. Salmon. *Scientific explanation and the causal structure of the world*. Princeton University Press, Princeton NJ, 1984.
- Wesley C. Salmon. Causality and explanation: a reply to two critiques. *Philosophy of Science*, 64(3):461–477, 1997.
- Wesley C. Salmon. *Causality and explanation*. Oxford University Press, Oxford, 1998a.
- Wesley C. Salmon. Scientific explanation: Three basic conceptions. In *Causality and Explanation*, pages 320–32. OUP, 1998b.
- Jonathan Schaffer. Causes need not be physically connected to their effects: the case for negative causation. In Christopher Hitchcock, editor, *Contemporary debates in Philosophy of Science*. Blackwell, 2004.
- Daniel Steel. *Across the boundaries. Extrapolation in biology and social science*. Oxford University Press, 2008.
- Leen De Vreese. Causal (mis)understanding and the search for scientific explanations: a case study from the history of medicine. *Studies in the History and Philosophy of Biological and Biomedical Sciences*, 39:1424, 2008.
- Erik Weber. How probabilistic causation can account for the use of mechanistic evidence. *International Studies in the Philosophy of Science*, 23(3):277–295, 2009.
- James Woodward. *Making things happen: A theory of causal explanation*. OUP, 2003.