Causes with material continuity

Lauren N. Ross

Recent philosophical work on causation has focused on distinctions across types of causal relationships. Of the various ways that causal relationships can differ, differences that have been discussed in this literature include stability, specificity, and proportionality (Woodward 2010; Lombrozo et al. 2018). This paper argues for another distinction among causation that has yet to receive attention in this work. This distinction has to do with whether causal relationships have "material continuity" or not, which refers to the reliable movement of material from cause to effect. I provide an analysis of material continuity and argue that causes with this feature (1) are associated with a unique explanatory perspective, (2) are studied with distinct causal investigative methods, and (3) provide different types of causal control over the world.

1 Introduction. A significant amount of the philosophical literature on causation has focused on distinguishing causal from non-causal relationships. This has generated interest in various causal criteria, including those that appeal to statistical relations, counterfactual analyses, and notions of conserved quantities, to name a few. Despite this dominant focus, recent work in this area has taken a different aim. Instead of distinguishing causal from non-causal relationships, this work aims to clarify differences across types of causal relationships. These projects capture "distinctions among causation" in the sense of clarifying how legitimate causal relationships differ from each other (Woodward 2010). Of the various ways that causal relationships can differ, differences that have been discussed in this work include those having to do with stability, specificity, and proportionality (Woodward 2010; Lombrozo et al. 2018).

In this paper, I argue for another distinction among causation that has yet to receive attention in this literature. This distinction has to do with whether causal relationships have "material continuity," which refers to the reliable movement of material from cause to effect. Why does this distinction matter? According to Woodward (2010), a significant motivation for appreciating distinctions among causation is that they identify causal relationships that figure in deeper and better explanations. In this work, I suggest an additional motivation for appreciating various causal distinctions, including material continuity. Causes with material continuity are not important because they provide explanations that are deeper, better, or more paradigmatic than others. Instead, these causes are important because they capture how causal systems in the world are different from each other and how these differences influence our causal reasoning. Different causal systems are associated with unique explanatory perspectives, they require study with distinct scientific methods, and they provide different types of control over the world. Appreciating this diversity is necessary for understanding the varied causal structure of the world and how we successfully navigate it.

This paper provides an analysis of material continuity as a distinction among causation. A guiding motivation for this work is that knowing that a relationship is causal only tells you so much.

[†]To contact the author, please write to: Lauren N. Ross, Department of Logic and Philosophy of Science, University of California, Irvine, 3151 Social Science Plaza A, University of California, Irvine 92697-5100; email: rossl@uci.edu.

[*]Acknowledgements: I would like to thank Jim Woodward and audience members at the Munich Center for Mathematical Philosophy (MCMP) for helpful feedback on this paper.

There are many rich, dynamic, and colorful differences across causal relationships that matter for how we study causation, engage in causal reasoning, and use causal knowledge to control outcomes in the world. This analysis focuses on three main ways that causes with material continuity differ from causes that lack this feature. In particular, causes with material continuity (1) are associated with a unique explanatory perspective, (2) are studied with distinct causal investigative methods, and (3) provide different types of causal control over the world. Discussion of these features will clarify the importance of appreciating causal diversity and how this diversity can be captured within an interventionist framework.

2 Background: Commonly discussed distinctions among causation Recent philosophical work on causation has acknowledged various distinctions among causal relationships, including stability, specificity, and proportionality (Woodward 2010; Lombrozo et al. 2018). While this work appreciates different types of causal relationships, it does not advocate for different causal criteria–it does not claim that what makes these relationships "causal" differs from one to the other. In these projects, and the analysis in this paper, causes are understood with an interventionist framework. Within this framework, the relata of causal relationships are variables (X, Y, Z,...), which represent different properties (a gene, a light switch, and so on). These variables can take on different values (0, 1, 2, ...), which correspond to different states of these properties (such as absent 0, and present 1) (Woodward 2003). This background allows us to specify a minimal interventionist criterion (I), that all causal relationships are expected to meet:

(I) X is a cause of Y if there is an ideal intervention on X (and only X) in background circumstances B, where this intervention produces changes in the values of Y.

In this sense, to say that X is a cause of Y simply means that there are some values of X such that intervening on X and changing its values (in some specified set of background conditions) produces changes in the values of Y. This account is counterfactual in the sense that determining whether X is a cause of Y depends on counterfactuals about what would have occurred to Y under different interventions on X. Thus, a cause "makes a difference to" its effect and an effect "counterfactually depends" on its cause (Woodward 2003).

A first distinction among causation has to do with the stability of causal relationships (Woodward 2010). Stability concerns the degree to which the generalization between X and Y would break down if the background conditions were changed in various ways (Woodward 2010). For example, the causal relationship between a gene and a disease phenotype is stable when the gene causes the disease in a wide variety of "background" conditions. An example of this is the causal relationship between the *huntingtin* mutation and Hungtinton's disease. As Kendler states "[i]f you have one copy of the pathogenic gene for Huntington's disease, it does not matter what your diet is, whether your parents were loving or harsh, or if your peer group in adolescence were boy scouts or petty criminals. If you have the mutated gene and you live long enough, you will develop the disease" (Kendler 2005, 394). Changes in other genetic, environmental, or social factors will not influence or alter this cause-effect relationship–it remains stable across these changes (Kendler 2005, 397). Compare this with the claim that a car crash (Y) was caused by the birth of the driver's paternal grandmother (X) (Lewis 1986, 216). Although X is in the "causal history" of Y, we view this causal claim as questionable, because we can imagine many situations and (different likely background conditions) in which X takes place, without reliably leading to Y. For example,

this relationship breaks down in situations where X does not have children, Y is a better driver, or safer road conditions exist. Both cases reveal how stability is assessed by considering "relevant" background conditions, where these are specified in a context-dependent manner. Unsurprisingly, causal relationships that are more stable are viewed as more paradigmatically causal, more useful, and more supportive of deeper explanations.[1]

A second distinction among causation involves specificity, which comes in two varieties–value specificity (specificity$_1$) and variable specificity (specificity$_2$). Value specificity concerns the number of values of the effect that the cause can produce. In this case, a cause is highly specific when it provides fine-grained control over many values of its effect. An example of this is the fine grained-control that a radio dial provides over the speaker's volume and that a dimmer switch provides over the light from a light bulb. This form of specificity is often associated with the explanatory importance of DNA and the "information" it contains (Waters 2007; Weber 2017; Stegmann 2014). In fact, this form of causal influence is said to explain why genes are "privileged" causes in biology, namely, because they provide more specific, fine-grained control than other causes. A cause that is less-specific provides coarse-grained, binary control over its effect, such as a standard light switch controlling whether a light is "on" or "off".

Variable specificity, on the other hand, has to do with the number of causes and effects in a causal relationship. The paradigm of variable specificity is a one-cause-one-effect relationship, in which the cause produces a single, specific effect and the effect has a single, specific cause. An example of this is the "monocausal model" in medicine, in which a single gene variant, bacteria, or vitamin deficiency produces a single, unique disease (as is the case with Huntington's disease, tuberculosis, scurvy, etc.). Non-specificity of effect refers to situations in which a particular cause has many different effects such as a single earthquake causing various outcomes in a city (water pipes to break, bridges to crumble, and fires to erupt) and a single gene producing many different effects (as in cases of pleiotrophy). Non-specificity of cause refers to situations in which different causes can all produce the exact same effect. An example of this emerged in research on Parkinson's disease when it was discovered that completely different causes were all individually sufficient to produce this same disease.[2] Non-specificity of cause captures an important form of causal complexity that is found in many sciences (Ross 2020a,b).

A third distinction among causation is proportionality, which concerns the level of description of cause and effect. This distinction is well-captured with Yablo's pigeon-pecking example Yablo (1992). In this example a pigeon is trained to peck at red targets. One day, this pigeon is presented with a scarlet target and, unsurprisingly, this causes it to peck. This leads us to consider two different ways of describing this causal relationship: (i) Did the presentation of a *red* target cause the pigeon to peck? or (ii) Did the presentation of a *scarlet* target cause the pigeon to peck? Description (ii) is less correct because it describes the cause in a way that is less proportional to the effect–it suggests that the scarlet shade caused the pigeon to peck, when any shade of red would have caused this. The sense in which the cause in (ii) is less proportional relates to its inclusion of irrelevant detail (namely, the particular shade of red). Thus, (i) describes a more proportional relationship–red better captures the proper contrast because if a non-red color had been presented, the pigeon would not have pecked. While this case considers whether the cause is proportional

[1] A further set of distinctions are discussed by (Lombrozo et al. 2018), which concern different types of stabilty.

[2] These different causes included single gene variants, single environmental factors, and combinations of each (Ross 2020a).

to some effect, other examples consider whether the effect is proportional to the cause Woodward (2010).

These distinctions are claimed to capture causal relationships that are more "useful" and that figure in "better" explanations Woodward (2010). While these features do appear valued in various contexts of causal reasoning, there are interesting exceptions to this. Unlike some claims (Waters 2007; Lynch et al. 2019), causes with value specificity do not always provide better explanations. Reasons for this include the fact that coarse-grained control can be preferred and that fine-grained control is impossible if the effect variable is strictly binary. Similar points can be made for variable specificity.[3] Nevertheless, the claim that these causes support higher quality explanations is a significant motivation for studying and clarifying these distinctions.

3 Causes with material continuity: A new distinction Another distinction among causation, which has yet to receive attention in this literature, has to do with whether causal relationships have material continuity or not. Material continuity refers to a situation in which, the change in X that produces a change in Y is accompanied by the reliable transfer of material from X to Y. Examples of this include a ball rolling down an incline, blood traveling along a blood vessel, subatomic particles moving along a decay series, a carbon skeleton moving along a metabolic pathway, and an organism progressing through the steps of a developmental pathway.

As a more detailed example, consider the metamorphosis of a butterfly. The stages of this developmental process include the egg, larvae, pupa, and adult. While there are a number of factors involved in the transformation of the egg into an adult butterfly one important cause of this is the egg itself. The egg's presence in the right environment is a cause of its transformation into the most immediately downstream step and so on. This causal sequence involves counterfactual dependencies between the different forms that the organism takes on. The egg's presence is a cause of its transformation into a larvae, which is a cause of its transformation into a pupa, which is a cause of its transformation into an adult butterfly. Relatedly, controlling whether the egg is present or not, controls whether a downstream butterfly is produced or not–if the organism gets held up at any step, complete metamorphosis will not occur. This sequence of causal steps involves the flow of material, associated with an organism, from one step to the next. This is similar to metabolic pathways and factory assembly lines, in which some upstream substrate is converted into a series of downstream products. The substrate at each upstream step is causally relevant to creation of the most immediately downstream product and there is some material that reliably flows along this causal process.

In clarifying what is meant by material continuity a few additional points should be mentioned. First, causes with material continuity often involve some significant amount of material carryover from cause to effect–it does not count to simply have microscopic particles or traces of dust moving from one to the other.[4] Second, in these cases the material "reliably" moves from cause to effect without being a matter of sheer luck or coincidence. The movement of material is, in a sense, "built-in" to the causal process. Third, the flow of material is largely constrained by and limited to the causal process. A large majority of the material moves from cause to effect, without moving

[3]With respect to variable specificity, heterogenous effects can be desirable in some cases (in cascade-like processes) and heterogeneous causes can also have their advantages (as they are more available options to produce an outcome).

[4]What counts as a significant amount of material is likely to depend on the context and the amount of material an object can loose or gain at each step, while still being considered the same object.

to other external properties that are non-effects.

While some causal processes involve the flow of material, not all causal relationships have this feature. As an example, consider a sequence of dominos that fall over in series. This captures a process that we view as genuinely causal, although there is no material that starts with the first domino and moves down to the final domino as the causal process unfolds. In this case, the causes and effects are physically interacting, but these interactions do not reliably transmit material from one to the other. Similar examples of this are colliding billiard balls, the moving gears of a watch, and causal mechanisms with physically (or mechanically) interacting parts. Other types of causal relationships that lack material continuity are situations in which causes and effects lack physical connection. Examples of these include "absence causation" and "double prevention," in which the absence of an entity can produce some effect and in which prevention of an action can produce (or allow for) an effect, respectively (Beebee 2004; Schaffer 2016).[5] A final set of examples are the cases Woodward (2010) uses to illustrate his distinctions among causation. There is not any material that moves from the *huntington* gene to the disease phenotype, from the radio dial to the speaker's volume, or from the presented target to the pigeon's pecking behavior.

Why does it matter whether causes have material continuity or not? According to Woodward (2010), a significant motivation for appreciating distinctions among causation is that they identify causal relationships that figure in explanations that are deeper and better than others. I am going to suggest a different motivation for appreciating the material continuity distinction. Causes with material continuity are not important because they provide causal explanations that are deeper, better, or more paradigmatic than causal relationships that lack this feature. Instead, these causes are important because they help capture how causal relationships and causal systems in the world are *different*. These differences motivate distinct explanatory questions, they are associated with unique explanatory perspectives, they require study with distinct scientific methods, and they provide different types of control over the world. Appreciating this diversity is necessary for understanding of the varied causal structure of the world and how we successfully study it. In the rest of this paper, I support these claims by clarifying how causes with material continuity (1) are associated with a unique explanatory perspective, (2) are studied with distinct causal investigative methods, and (3) provide different types of causal control over the world. An analysis of these points will show the importance of appreciating the diverse causal structure of the world and how this can be captured within an interventionist framework.

4 Explanatory perspective: Object-oriented causation Causes with material continuity are associated with a unique explanatory perspective. Causes with this feature are likely to surface when we are focused on "object-oriented causation," in which we are interested in a particular object and how it changes over time. As the objects in these cases are typically associated with material, we view this material as "moving" through a sequence of causal steps. This is similar to what Russell refers to as the "persistence of something" through a causal relationship, such as "a person, table, a photon, or what not" (Russell 1948, 404). In these cases, there is a "constancy of quality" and "constancy of structure" through the steps of the causal process, where this can take the form of constancy of physical material. This is seen in how we track an object as it develops, matures, or transforms over time, or in how we track an object as it moves through a system. We

[5]For instance, if my houseplants die because I fail to water them, I have caused their death without coming into physical contact with them and, thus, without transmitting any material to them. Examples of absence causation and double prevention are commonly found in science and everyday life.

can track an organism through developmental stages, a product through an assembly line, blood through the vessels of an organism, and water through the pipes of a city. In contrast, notice how these examples of "object-oriented" causation differ from mere "causal influence," in which an effect ricochets through the world without involving the reliable flow of material. For example, the spread of a fascinating news story through a community is clearly a causal process, although it need not involve the movement of material or any "object" through these steps. Similarly, Woodward's causal distinction examples also lack this object-oriented perspective Woodward (2010). We do not identify any object moving from the *huntington* gene to the disease phenotype, from the radio dial to the volume output, or from the red target to the pigeon's pecking behavior.

This object-oriented perspective is further illustrated with two main categories of causes with material continuity: (i) manufacture and (ii) movement cases.[6] Manufacture cases involve changes in an object's form or constitution, while movement examples involve changes in an object's spatial location. Manufacture cases involve tracking some starting material as it undergoes a sequence of transformations, which culminate in a final product. These are associated with processes such as metamorphosis, synthesis, development, growth, transformation, maturation, and so on. Movement cases, on the other hand, involve tracking an object's change in location, such as movement of blood through vascular pathways, water flowing along channels, and a toy car moving along a race track. These cases are associated with "object-oriented" explanatory why-questions, such as "What does X develop into over time?" and "Where does X move to over time?". While these questions are "locked-in" or "tethered" to a particular object (from cause to effect) other common causal-explanatory questions lack this feature.

These manufacturing and movement cases both meet a minimal interventionist criterion. In these examples each upstream change in the constitution or spatial location of an object "makes a difference" to an immediately downstream change.[7] While these cases primarily focus on changes to some single object of interest, they include many other causal factors that guide, channel, and direct the object as it undergoes these changes. In order to see this, consider manufacturing examples in which some starting material is used to produce a final product. These examples bear similarity to Aristotle's notion of "material causation" in which some material is the substance "out of which" a particular product is made (Lennox 2001a). For example, a block of marble that is sculpted into a statue is a material cause of the statue.[8] The material cause is often associated with an

[6] I do not claim that these categories are exhaustive of all types of causes with material continuity.

[7] These examples may seem in tension with interventionism in the sense that cause and effect appear to reference the same object. Within the interventionist framework cause and effect must refer to sufficiently distinct properties such that it is possible to (hypothetically or conceptually) intervene on X without, at the same time, also intervening on Y. If this is not possible it suggests that, instead of X being a cause of Y, that X might just be Y (and vice versa). For example, consider the claim that turning a door handle to the clockwise position causes a "dextrorotatory" handle, or that introducing a disease gene into an organism causes a "mutated" genome, or that opening a refrigerator causes the door to be ajar. In each of these examples, instead of the purported cause producing the effect, both merely re-describe the same phenomenon. Examples of object-oriented causation do not have this issue because, although cause and effect are associated with the same object, they refer to sufficiently distinct properties of the object. For example, in movement cases X and Y refer to different spatial locations, as there is interest in knowing whether the object's presence at one is causally relevant to its later presence at another. In manufacture cases X and Y refer to distinct forms–there is interest in knowing whether taking on one form X is causally relevant to the object then taking on another Y.

[8] While my analysis of "material continuity" bears similarity to the notion of "material causation" discussed in the philosophical literature, they should not be viewed as compatible in all respects. There are many

"object" that is followed through the causal process as the marble "is not only the material out of which the statue is made; it is also the subject of change, that is, the thing that undergoes the change and results in a statue" (Falcon 2019). Of course, the starting material is just one cause of the downstream outcome, as other causes convert these materials into the final product. So, for example, just as a sculptor converts the bare marble into a statue, an enzyme converts a metabolite into a chemical product and developmental factors transform the egg into a butterfly. These "guiding" causes–sometimes referred to as "structuring causes" or "constraints"–are responsible for converting, channeling, or guiding the upstream starting material into a particular downstream outcome (Dretske 1988; Haslanger 2016). Part of what this shows is that causes with material continuity are one of many causes in multicausal situations.[9]

In these examples, we tend downplay the role of causes with material continuity, while emphasizing the role of guiding (or structuring) causes. This is evident in the fact that we are less likely to say that the building materials caused the house to be built, or that the block of marble caused the creation of the statue, or that the butterfly egg caused the butterfly to develop. This is associated with the view that starting materials alone do not produce the product as they remain motionless and inert unless acted upon by something else. As mentioned by Aristotle, "a statue does not come to be spontaneously" (Lennox 2001a, 640a, 30). This de-emphasis on the causal role of starting materials is seen in the fact that they are characterized as a "passive cause," "passive matter," and as having "passive power" compared with the more "active" role of guiding causes (Barnouw 2007; Noble 2017; Crombie 2020).[10] While we admit that starting materials are necessary for the outcome to occur, we do not always highlight their causal role in bringing about the outcome.

Why are starting materials often viewed as passive causes, while guiding factors are viewed as more active and causally responsible? There are at least two main reasons for this. First, while the presence of the material is causally relevant to whether production takes place or not, it rarely determines which particular effect (of many potential effects) is produced. As the guiding cause is responsible for this it tends to be viewed as actively determining the specific outcome. For example, while the block of marble can be converted into a variety of different products (a column, statue, tile, etc.) the sculptor controls which one is produced. In this case, causal responsibility and whether a cause is "active" is associated with the fact that the guiding cause has specific$_2$ causal control over the final product formed, while the starting materials do not.[11] A second reason why causes with material continuity are viewed as less causally responsible is that we often assume that they are present in a given situation, while other factors change in various ways. If we want to know how a house is built or how a butterfly a develops, we often assume that the starting materials are

discussions of material causation that differ from the interventionist-oriented picture that I support in this work (Sosa 1980).

[9]Notice again that this differs from the examples used to illustrate Woodward's distinctions, as in these cases we tend to identify a single main cause.

[10]This is also related to the view that there is some "insufficiency of explanation wholly in terms of material causes" (Gotthelf 1987, 221).

[11]The same rationale holds for movement cases, in which we are interested in the changing location of an object over time. In these cases, guiding causes determine the particular location that the starting materials move to. For example, blood vessels dictate where blood will flow, water channels determine where water will move, the race track determines where the toy race car will move, and so on. The upstream presence of material is causally relevant to its downstream movement, but it does not control the specific location of this movement, which guiding causes are responsible for. So causes with material continuity often lack causal specificity$_2$, although they interact with other causes that do have this specificity.

available.

While the explanatory role of material causes is often minimized, there are some exceptions to this. One main set of exceptions have to do with cases in which we assume that the guiding causes are fixed and unchanging. In many situations these guiding factors are viewed as "structuring" causes, not just because they channel the system to a particular outcome, but because they are fixed and less likely to change, relative to other causes. For example, given a set of water channels or blood vessels, we commonly assume that these structures are fixed and we focus on whether the material cause (the liquid) is introduced into this system or not. In this set up, availability of the object or starting material is what causes the process to unfold.[12] In fact, in biological contexts, organisms exploit this control of material causes–when they want production to stop they sequester starting materials and when they want it to ensue they make them available. This is similar to having a prepared sculptor but varying whether the marble starting material is available or not. When guiding causes are an assumed part of the background, or held fixed, causes with material continuity are viewed as having more causal responsibility.

Causes with material continuity are associated with a unique explanatory perspective, which focuses on some material object and how it changes over time. These causes are more likely to arise in domains in which: there are material objects of interest, these objects have some stability within a relevant time-frame, and these objects undergo changes that are reliable, repeatable, and generalizable. As illustrated by the examples discussed in this section, causes with this feature arise in many scientific and ordinary life contexts.

5 Investigative strategy: Tracer techniques Causes with material continuity are often studied with a unique investigative strategy. These causes are often studied with physical tracer and tagging techniques, which exploit the flow of material through these causal processes. These techniques involve attaching a physical tag to the upstream cause and then following the tag as it moves through the steps of the causal process. Following this tag allows scientists to discover these downstream steps, where these steps can reveal different stages in an object's development, different locations it can travel to, and different ways that causal links are (and are not) connected up in some domain.

This basic approach is found in many scientific domains, including biochemistry, ecology, and neuroscience. In the context of biochemistry, scientists use tracers to identify metabolic pathways, which capture the sequence of chemical conversions that transform a metabolic substrate into a product. These cases involve attaching a radioactive tracer to a metabolite and following its journey through a sequence of chemical conversions. This reveals the "fate" of the original metabolite and the different chemical substances that it is transformed into on its way to becoming some final product. A second example comes from radioecology, in which ecologists use radioactive tracers to identify ecological pathways, which capture the prey-predator connections between species linked in food chains (Odum 1971). While it is often obvious which species are located in a particular area, it is much harder to determine the particular feeding relations among them. One way to identify this is to introduce a radioactive marker into an upstream prey (such as a species of grass) and then follow its movement into a sequence of downstream species at increasing time intervals.

[12]Similarly, given a developmental or biochemical environment replete with standard guiding causes, it is the presence of a butterfly's egg that will trigger development and the presence of metabolic starting material that will trigger biochemical conversion. This is similar to having a factory flush with workers and tools, but varying whether the starting materials have been delivered or not.

This allows ecologists to "unravel" and "untangle" these causal connections in ecosystems (Odum 1971). A final example involves the use of tracers to study neural pathways, which constrain the flow of nerve signals in brain and other nervous tissue. In this context, neuroscientists use dyes, radioactive materials, and viruses to chart the flow of material along and across neurons, as this reveals the routes along which signals flow (Morecraft et al. 2014; Saleeba et al. 2019). Identifying this structural and anatomical connectivity provides insight into the function of different nervous tissues (Ross 2020c).

In scientific contexts, tracers are expected to meet particular criteria in order to provide reliable information about a causal system.[13] A first criterion is that the tracer should not alter the causal process being studied, as this would provide information about a perturbed system rather than the "natural" process of interest. Due to this, scientists will often claim that an ideal tracer should be "relatively harmless" and have minimal toxicity to the system (Hu et al. 2013; Huh et al. 2010; Mi et al. 2019). Second, tracers should be sufficiently recognizable in the causal process such that they can be properly identified and followed. The tracer should be distinguishable from surrounding materials and it should remain in high enough concentration in later steps of the causal process. Third, the tracer is often expected to be "specific" in the sense that it tags and marks exactly and only material that moves through the causal process (Martin and Dolivo 1983). One exception to this are tracers that reliably flow with material moving along a causal process, without the tracer binding to this material.[14]

Tracer techniques are a unique causal investigative strategy in the sense that they are not used to establish or prove causality, but to get further information about systems that are already viewed as causal. This is supported by two main points. First, tracer methods require background information about the causal process, which is clear from tracer criteria. In order to use a tracer it needs to be clear what should be tagged, which requires knowing what (if anything) reliably flows through the causal process. Before using tracers in the biochemical, ecological, and neurological contexts above, scientists had already identified materials that reliably flowed through the causal systems.[15] This background causal knowledge was required for the tracer methods they employed. What scientists did not yet know is where the downstream steps of these pathways would lead and what would be found at each step. You can know that caloric energy flows through ecosystems and that cellular materials flow along neurons withou knowing that the meadow frog eats the grasshopper and that the hypoglossal nerve innervates the tongue. Second, that tracer methodology requires causal information is also suggested by the fact that movement of material alone is not sufficient for causation. We can identify and track the movement of some physical matter, without this outlining the steps of a causal process. For example, we can track the flow of material from an organisms's diet to the microconstituents of its cells, without this material being causally or explanatorily relevant to all of the cell's behaviors.[16] As tracking material alone is insufficient to

[13]For further discussion of this see: (Ross 2020c).

[14]This is present in radiotracers used to study blood vessels and the gastrointestinal tract–when these tracers are introduced into these locations they reliably flow in the direction of relevant material (blood and food, respectively) without tagging this material.

[15]For example, they had identified that carbon and other atoms move through metabolic pathways, that caloric energy moves through prey-predator connections, and that axonal transport delivers cellular materials along neurons.

[16]For example, a cancerous skin cell can share a connection to material consumed in the diet, without this material explaining its cancerous state (suppose it is malignant from overexposure to the sun). In other words, explanatory targets can be made up of material that is causally and explanatorily irrelevant to the

provide reliable information about a causal system, something more is needed to ensure that tracer methods meet this standard.

A main suggestion of this analysis is that these tracer techniques can be used for some, but not all, causal systems. As these techniques involve tagging material and watching it flow through a causal process, they cannot be used for causal processes that lack this feature. Can tracer techniques be used to study causal systems that lack material continuity? If so, what makes this investigative strategy meaningful? To the extent that all causal relationships involve the flow of "causal influence" there is a loose (and less meaningful) sense in which "tracing" this influence can provide information about the system. A more compelling account of tracers for causal systems that lack material continuity would involve tracing "immaterial" entities such as the flow of "information" or "signals." There are a number of likely examples of these cases, including the use of genetic markers in biological systems and "traitor tracing" in social organizations.[17] While these signal or informational tracers are likely to share various features with tracers that tag material, they still have a number of important differences. Tagging material is a different technique than tagging information or material–many systems will be better studied with one technique over another. This difference is likely associated with different explanatory perspectives–the focus on a persisting object in one case, and a persisting signal in another. Finally, to the extent that signal tracers appear less common than physical tracers, it is worth considering why this is the case. It may be that signal tracers are more difficult to implement for reasons having to do with complications in defining what counts as information, challenges in identifying this, and distinguishing information from causality.

6 Causal control: Microstructure of effect Material continuity differs from Woodward's causal distinctions in an interesting way. In particular, material continuity concerns the microstructure of cause and effect, while stability, specificity, and proportionality are completely silent on this aspect of causation. This can be appreciated by examining the type of "causal control" that each distinction among causation is associated with. For example, the control that a cause has over its effect can be stable, it can be fine-grained versus coarse-grained (specificity$_1$), or it can be specific in the sense of producing a single outcome, versus non-specific in the sense of producing many (specificity$_2$). Importantly, causes with material continuity provide control over the microstructure of their effects. For example, while the sculptor controls which particular product is created–column, tile, or statue–the starting material controls its microstructure as it controls the physical material it is made of. Changing the starting materials provided to the sculptor–whether they are marble, granite, alabaster, limestone, soap, or plastic–changes and controls the microstructure of the final product. This captured in an Aristotelian conception of material causation, as "[t]he material cause explains out of what kind of matter the effect comes," that is to say, what the effect is made of (Dodig-Crnkovic 2003, 531).

Consider other causes that appear to have microstructural control over their effects, but that lack material continuity. Heat can change the microstructure of water by converting it from liquid to gas. Chemicals can change the microstructure of metal by causing it to rust. While these causes exert control over an effect's microstructure, they can do so without exhibiting material continuity

target, despite the fact that this material can be traced to some earlier place (or point) in time. What this means is that, if some material is followed, the mere fact that this material ends up in some property is no guarantee that it explains some feature of the property that we are interested in.

[17]Thanks to Daniel Dennett for bringing some of these examples to my attention.

as they need not transfer material to the effect. Does this show that microstructural control is not unique to causes with material continuity? The important point here is that causes with material continuity exert control over the microstructure of some effect by transferring material to it–they "share" microstructure with their effect, which is not the case in the above examples. The block of marble shares microstructure with the statue in a way that heat does not share microstructure with water and that chemicals do not share microstructure with the metals they oxidize.

The fact that causes with material continuity have this type of causal control is important for a number of reasons. First, microstructure matters because it is often explanatorily relevant to behaviors of the object of interest. Whether a final product is radioactive, magnetic, or buoyant depends on the material it is made up of, which is often determined by the starting materials used. In these cases, choice of starting materials will be explanatorily relevant to features of the final product. This helps clarify how properties of the final product are explained by the starting materials and not the craftsman–for example, that the marble statue will sink in a body of water is explained by the chosen starting materials, as opposed to the skill of the craftsman. Changing the starting materials (for example, to styrofoam) would change whether the statue exhibits this sinking property or not.

Second, in the context of manufacture cases, this microstructural control relates to the functionality of the final product. If the final product is supposed to serve some function–e.g. the feature of being waterproof–the chosen starting material will impact this and this will restrict which materials can be used. In this sense, "the nature of the craft product constraints the choice of appropriate materials" (Lennox 2001b, 289). This is due to the fact that, "whatever material is used has a nature of its own: the craftsman cannot do anything he likes with his material, but only what it is capable of being compelled to do" (Lennox 2001b, 289). This is relevant to creating products that are functional and it shows how the craftsman is limited by the starting materials as they provide causal control over the final product.

Third, this microstructural control provides a way to "import" material into the final product, in ways that can be accidental or deliberate. Examples of accidental cases include situations in which some toxic material in the cause delivers toxic material to the effect. Examples of this include radioactive materials that enter the food supply, moving from atomic constituents of food product to the organisms that ingest it.[18] While these cases involve accidental contamination, the same control can be used intentionally to move material to a particular location. This is exploited in some forms of drug delivery in which drugs are incorporated into materials that are then assimilated by particular areas of the body. If these drugs are toxic, they can destroy the specific cells that engulf them (which is valuable in cancer treatments). Some of these delivery methods are said to implement a "Trojan horse" approach, as products that are normally welcomed into various cells of the body are loaded with various drugs, allowing them to be smuggled into the cell.

Finally, microstructural control can also be used to get information about a cause by studying its effect. The fact that cause and effect share microstructure means that there are physical "traces" of the cause in the effect such that studying the latter provides information about the former. For example, studying the materials that a statue is made of can identify the region in which it was created or where the materials were sourced. Additionally, radiocarbon dating can be used to date the time period of the cause and when it produced the effect, as is done with various archeological finds, including artifacts and fossils. Of course, other features of an effect, besides its microstructure, can reveal information about its cause. A sculptor might have special carving

[18]This happened after the nuclear disasters at Chernobyl and Fukushima Daiichi, for example.

technique that is used to identify their pieces. These features are likely to provide different types of information about the causal process–one about the craftsman and the other about their materials.

7 Conclusion This paper has considered a novel distinction among causation, referred to as "material continuity," which concerns the reliable transfer of material from cause to effect. This distinction has been compared to commonly accepted causal divisions, including stability, specificity, and proportionality (Woodward 2010). While study of these divisions is commonly motivated by an interest in identifying causes that provide deeper and better explanations, this paper suggests another motivation–namely, capturing the diversity of causal structures in the world. Capturing this diversity is important, because it reveals how different types of causes involve unique explanatory perspectives, are studied with distinct causal investigative strategies, and provide different types of control over the world. Why exactly is capturing this causal diversity important? What turns on appreciating material continuity and the unique features associated with it?

We take interest in causation because it helps clarify how we navigate the world, provide scientific explanations, and control outcomes and bring about change. Appreciating the diverse causal structure of the world matters because different structures teach us different lessons for reaching each of these goals. Navigating causal systems that are unstable requires a different strategy set than systems that are stable. Explaining outcomes with non-specific, heterogenous causes involves various challenges that are not present when specific, single main causes are at play. The control we have available to us changes depending on whether causes have fine-grained specificity or are material causes with microstructural control. In each of these cases, the manner in which we select relevant causes, frame an explanatory-why question, provide an explanation, and exert control on the world will differ depending on the causal structure we are presented with. Appreciating distinctions among causation provides a fuller picture of the causal complexity of the world and the nuanced ways in which we study, navigate, and explain it.

References

Barnouw, J. (2007). Hobbes's causal account of sensation. *Journal of the History of Philosophy*, 118:115–130.

Beebee, H. (2004). Causing and nothingness. In Hall, N. and Paul, L. A., editors, *Causation and counterfactuals*, pages 291–308. The MIT Press, Cambridge.

Crombie, A. C. (2020). *The history of science from Augustine to Galileo*. Dover Publications.

Dodig-Crnkovic, G. (2003). Shifting the paradigm of philosophy of science: Philosophy of information and a new Renaissance. *Minds and Machines*, pages 521–536.

Dretske, F. (1988). *Explaining behavior*. The MIT Press.

Falcon, A. (2019). Aristotle on Causality. *Stanford Encyclopedia of Philosophy*.

Gotthelf, A. (1987). Aristotle's conception of final causality. In Gotthelf, A. and Lennox, J. G., editors, *Philosophical issues in Aristotle's biology*, pages 204–242. Cambridge University Press.

Haslanger, S. (2016). What is a (social) structural explanation? *Philosophical Studies*, pages 113–130.

Hu, W., Liu, D., Gu, J., Zhang, Y., Gu, X., and Gu, T. (2013). Neurological function following intra-neural injection of fluorescent neuronal tracers in rats . *Neural Regeneration Research*, pages 1253–1261.

Huh, Y., Oh, M. S., Leblanc, P., and Kim, K.-S. (2010). Gene transfer in the nervous system and implications for transsynaptic neuronal tracing. *Expert Opinion on Biological Therapy*, 10(5):763–772.

Kendler, K. S. (2005). "A Gene for. . .": The Nature of Gene Action in Psychiatric Disorders. *American Journal of Psychiatry*, pages 1–10.

Lennox, J. G. (2001a). *Aristotle: On the parts of animals I-IV*. Oxford University Press, Oxford.

Lennox, J. G. (2001b). *Aristotle's philosophy of Biology*. Cambridge University Press, Cambridge.

Lewis, D. A. (1986). Causal Explanation. In *Philosophical Papers Volume II*. Oxford University Press, Oxford.

Lombrozo, T., Blanchard, T., and Vasilyeva, N. (2018). Stability, breadth and guidance. *Philosophical Studies*, 175(9):2263–2283.

Lynch, K. E., Parke, E. C., and O'Malley, M. A. (2019). How causal are microbiomes? A comparison with the Helicobacter pylori explanation of ulcers. *Biology & Philosophy*, 34(6):1–24.

Martin, X. and Dolivo, M. (1983). Neuronal and transneuronal tracing in the trigeminal system of the rat using the herpes virus suis. *Brain Research*, 273:253–276.

Mi, D., Yuan, Y., Zhang, Y., Niu, J., Wang, Y., Yan, J., Yang, Y., and Hu, W. (2019). Injection of Fluoro-Gold into the tibial nerve leads to prolonged but reversible functional deficits in rats. *Scientific Reports*, pages 1–8.

Morecraft, R. J., Ugolini, G., Lanciego, J. L., Wouterlood, F. G., and Pandya, D. N. (2014). Classic and Contemporary Neural Tract-Tracing Techniques. In Johansen-Berg, H. and Behrens, T. E. J., editors, *Diffusion MRI*, pages 359–399. Elsevier.

Noble, D. (2017). *Dance to the tune of life*. Cambridge University Press.

Odum, E. (1971). *Fundamentals of Ecology*. Sauders, Philadelphia, PA, 3 edition.

Ross, L. N. (2020a). Explanation in contexts of causal complexity: Lessons from psychiatric genetics (Forthcoming). . In *From scientific metaphysics to biological practice*. Minnesota studies in the philosophy of science.

Ross, L. N. (2020b). Multiple realizability from a causal perspective. *Philosophy of Science*, pages 640–662.

Ross, L. N. (2020c). Tracers in neuroscience: Causation, constraints, and connectivity (Forthcoming). *Synthese*.

Russell, B. (1948). *Human knowledge: Its scope and limits*. Taylor & Francis Routledge.

Saleeba, C., Dempsey, B., Le, S., Goodchild, A., and McMullan, S. (2019). A Student's Guide to Neural Circuit Tracing. *Frontiers in Neuroscience*, 13:1–19.

Schaffer, J. (2016). Metaphysics of causation. *Stanford Encyclopedia of Philosophy*.

Sosa, E. (1980). Varieties of causation. *Grazer Philosophische Studien*, 11(1):93–104.

Stegmann, U. (2014). Causal Control and Genetic Causation. *Noûs*, 48(3):450–465.

Waters, C. K. (2007). Causes that make a difference. *The Journal of Philosophy*, pages 1–30.

Weber, M. (2017). Causal selection versus causal parity in biology: Relevant counterfactuals and biologically normal interventions (Forthcoming). In *Causation in biology*. Minnesota studies in the philosophy of science.

Woodward, J. (2003). *Making things happen*. Oxford University Press, Oxford.

Woodward, J. (2010). Causation in biology: Stability, specificity, and the choice of levels of explanation. *Biology & Philosophy*, 25(3):287–318.

Yablo, S. (1992). Mental Causation. *The Philosophical Review*, 101:245–280.