Models of Science and the Role of Causation

Aaron L Bramson

June 14, 2007

> Let us chace our imagination to the heavens, or to the utmost limits of the universe; we never really advance a step beyond ourselves, nor can conceive any kind of existence, but those perceptions, which have appear'd in that narrow compass. This is the universe of the imagination, nor have we any idea but what is there produc'd.
>
> -Hume (Treatise I.2.6)

Abstract

Hume calls causation "the cement of the universe" which seems to imply that there is a universe "out there" and it needs causation to hold it together (though it seems highly unlikely that Hume himself held such a view). The position argued for here is that for some reason our perceptions admit to regularities that we encode into models and the behavior of models is entirely determined by such rules and so these rules really do hold these models together. Because we completely rely on (explicit and implicit) mental models to govern our behavior and understanding we project some of these rules into the world in the form of causes. Realizing that our concept of causation applies only to our models, and taking the skeptical thesis seriously that models are all we really have, we gain a concept of causation that does all the work it previously did to relate events but does not depend on an external world that may not be there.

1 Introduction

This paper takes seriously the realization that all that can figure into our understanding of the universe is gained through perception; or more specifically as ideas (whatever those are) in the mind that present themselves as the result of perception. Starting with this very Humean foundation we reconstruct a very Humean view of causation; what comes along the way, however, is an exploration into many current notions of ontology, epistemology and the philosophy of science grounded in terms of a current scientific techniques (Markov modeling) and reflected in ideas about the uses models (mental and formal), analogies, and the consistency of our experience.

2 Definitions and Core Concepts

My aim is to use terms of art in their most widely accepted use in order to avoid difficulties in matters of definition. Since there will remain ambiguity in what the most widely accepted use is, I begin the paper by providing a brief glossary of the terms relevant to the discussion. These are not intended to be precise, fully defended and cohesive definitions, but rather an indication of the kind of thing I mean with the hope that what is left ambiguous or underdetermined is irrelevant in the discussion. In-depth analyses of these terms exist, and I reference sources where they were explicitly consulted; otherwise all mistakes and confusions are my own.

2.1 Models and Their Domains

This paper uses a broad notion of *model* that includes any categorically structured representation of "the world".[1] The focus is on Markov model representations of systems (more below). The *domain* of a model is the set of fundamental object **types** that may appear in the model and their possible attributes (including behavior rules) in conjunction with the possible aspects of the space involved. Because the domain is a set of types, each element in the domain may or may not appear as tokens in the model, and each may appear in multiplicity. And because it is a set of **fundamental** object types, collections of elements in the domain are not themselves elements in the domain. A *specification* of a model is a set of **token** elements from the domain and is strongly related to my definition of *system*

[1] This is meant to exclude certain implicit models that some might claim exist in physical systems; e.g. that water molecules have a model of hexagonal geometry for covalent bonds as revealed by snowflake shapes. It would also exclude physical models of systems (miniature boats, plane wings, bridges, etc.) because they are not themselves abstractions and the plans to create them do not include the behavior rules or physical laws (though their designs would be informed by a model that did). The point is to include formal abstract models in any representation and mental models (including implicit mental models (cognition patterns) and folk theories).

(see below). Object tokens are sometimes referred to as *atoms* in the Democritus sense. Remember that the domain also includes aspects of the environment and these too are part of the specification. The rules defined for the included objects of the domain bind the formulation of the specification to some degree, but aggregates from some domains will be underspecified. The expression “the model” can refer to the particular specification under consideration or to its representation or to a system (see below); usually these are defined synchronously and if not context makes the usage clear.

2.1.1 Markov Models

A Markov model is a graph (aka network diagram) in which each vertex (node) represents a complete specification of a single state of the system and a weighted directed edge (arrow) represents the probability of the system changing from the tail-end vertex to the head-end vertex (along the arrow). Markov models represent the entire “space-time” of a specification but do not explicitly model processes: only their effects (i.e. the changing from one state to another). So they are useful for tracking what the model will or can do, but not particularly useful for explaining or even describing why the changes happen. In this paper the terms *Markov model*, *Markov diagram*, *Markov representation* and various combinations of those will be used interchangeably.

2.1.2 Agent Based Models

To understand the processes that the elements of a model undergo (i.e. why some state changes happen and other don’t) a model needs to be able to capture dependencies, forces, relations, and all the other things indicated by the terms ‘laws’ and ‘causes’. To do that, the best going technique is to specify the relevant parts of the domain under investigation, specify the relations among those parts, specify the environmental factors, and specify how the parts change over time through internal factors or through their interactions with each other and the environment. These parts are referred to as *agents* and are (or at least ought to be) chosen from among the domain’s fundamental objects. One important feature of these agent-based models (ABMs) is that they can represent the mereological relations among model components and between the components and their aggregates. These are discussed because they are formal models with a representation that is closest to our mental models.[2]

[2]In practice, the state of the art in agent-based models leaves them especially useful and most frequently used for intuition pumping and they are usually accompanied by more traditional analytic models for the “real” science of proving, deriving, and validating. Their proximity to mental representations allows them to better guide the construction of models and to test models’ conformity to our existing notions.

2.1.3 Linking Agent-Based and Markov Models

Building a Markov diagram representation from an agent-based model is straightforward to conceive, though practically impossible for any but the most basic models. Assuming the ABM uses discrete time (as most do) each step of an ABM is a state of the system and hence a node of the Markov diagram. The different possibilities for a next step of the ABM are obviously the possible transitions of the Markov diagram and their probabilities should match. Agent-based models, I claim, are similar in important ways to our mental models (explicit folk models and cognition itself) and so our natural ways of thinking about the relationship among events (changes in states of affairs) can analogously be converted into Markov models. This becomes more interesting once we realize that we can alter the resolution of the description of the ABM state when we encode it as a Markov state (see *resolution* below for why this is interesting).[3]

2.2 Systems and Their Scope and Resolution

Recall that a *specification* is a set of object tokens and environmental aspects from a domain. In some cases the specification will uniquely determine a configuration of these elements; but in cases where it does not then we need another term to identify a particular configuration of the elements. Defining a set of objects as well as their relationships to each other and their environment is defining a *system*. In a rich domain a specification of a model will admit to a large number of possible systems through different possible configurations of the included parts.

Aggregates of objects that exhibit a distinguishable attribute (such as coherent behavior or a higher-order property) are called *patterns*.[4] This admittedly loose definition is critically perception-dependent; the identification of patterns is limited by (and limited to) our ability to recognize attributes of aggregates. A system's or a pattern's *scope* or *scale* refers to both the set and the cardinality of the set of fundamental element tokens comprising the

[3] A note on relevance to other work. One of my research thrusts is to automate the process of generating Markov diagrams from agent-based and differential equation models for the purpose of facilitating the development and performance of system-level measures (e.g. measures of complexity, path sensitivity, robustness, etc.). The primary barrier is that due to the large number of parameter combinations the number of possible states is often greater than the number of atoms in the universe. This problem can be alleviated to some degree by defining qualitative equivalence classes in the Markov model (i.e. build in multiple realizability) and only track changes for those classes; automating the generation of such classes is a very difficult problem.

[4] Note that the terms 'system' and 'pattern' are often interchangeably and/or in ways different from how they are defined here. For example, people, including well-informed and careful scientists, often refer to hurricanes as "complex systems" but the current definitions suggests that a hurricane should be considered a pattern in a **weather system**. In some model, perhaps one for predicting the location of the hurricane eye within a hurricane, the hurricane itself would be considered a system. Whether the pattern or the system ought to be called "complex" will depend on a definition of 'complex', which I will not attempt to provide here.

aggregate and this may vary over time. A system's *resolution* is defined relative to another system as the quotient of their scopes when modeling the same thing. Hence if system X requires exactly twelve object tokens to represent a particular container of fluid and system Y requires exactly twelve hundred object tokens to represent the same (or an identical) container of fluid then system Y has one hundred times the resolution of system X. Note that if the components of one system are not parts of the components of the other system then the resolution difference may not be an integer.[5]

Systems are defined with a specific scope and resolution because those are determined to be appropriate for the modeling task at hand. If we fix the part of the world that we are modeling then there is an obvious linear relationship between resolution and scope. Systems can be nested within one another by increasing the resolution (and hence the scope) and then reducing the scope by trimming objects out of the system. Higher-resolution aggregates falling within (occurring as parts of) lower-resolution ones can form *subsystem hierarchies*. But not all series of increasingly higher-resolution systems organize into hierarchies - the scopes of the smaller aggregates may include elements outside the scopes of some of the model's lower-resolution systems. This possibility has significant implications for both reduction and emergence which will be discussed below. An immediate corollary is that hierarchy is not a sufficient notion for the analysis at hand and we must instead talk of *levels of organization* to refer to the varying resolutions of models. When comparing two levels of organization it is useful to refer to the higher-resolution level as the *microlevel* and the lower-resolution level as the *macrolevel*.

There are more distinctions to be made on these topics, but in consideration of space and what will be utilized in the substantive part of the paper I will move on to other, related areas.

2.3 Terms from Science

2.3.1 Descriptions

Models are used to gain understanding of phenomena, but communicating every detail of a system may impede that understanding rather than facilitate it (or even be practically impossible). The information in a model can therefore be compressed into a *description* (also called a *presentation*) in one of two kinds of ways. Descriptions that preserve all the information of the distinct types of elements in the domain and element tokens in the system and all their relationships with each other and the environment are *lossless compressions* (or

[5] An example of this would be modeling a human body with its organs as the components versus modeling it with its rigid segments as the components (as biokineticist might do). It might turn out that these two component sets have only slightly different cardinalities and hence the resolution difference would be a non-integer rational number near one.

lossless descriptions) of the model. Such descriptions require a one-to-one correspondence between the components of the original model and components of the description such that the original domain can be unambiguously recovered from the description. For example, if three elements of a domain were always arranged in a certain manner and always exhibited the same system properties then we can use a shorthand description of that arrangement without loss of information about the domain.[6]

Descriptions that fail to preserve all of the information in the domain or specification of a model are *lossy compressions* (or *lossy descriptions*). The advantage of such a description is that it may suffice for communicating the necessary information and have much weaker restrictions on compression than lossless compressions do. For example, "it's raining here" fails to distinguish an astronomical number of systems at the quark-lepton level but provides enough information about the world for deciding whether to bring an umbrella. Laws are often stated in terms of lossy compressions because classes of patterns or objects of the system may share behaviors and/or properties (e.g. defining bases and acids, males and females, etc.) and the behavior of an element can be determined from its inclusion in such a class.

Clearly there is a link between lossy and lossless descriptions and models at different resolutions *and* perceived levels of organization of the world. These connections will be explored further below in terms of Markov models, but making a few more distinctions will be helpful.

2.3.2 Laws, Theories, and Explanations

The world sometimes changes in what we recognize as recurrent patterns. We can often describe these patterns as conditional statements. If the regularities are consistent enough then they may be accepted as *laws* in our folk models of the world. These folk models are sometimes formalized and if the regularity persists through the formalization process and rigorous comparisons with data then these reliable conditionals may achieve the status of a scientific law. There are many distinct relations that may be called "laws" or "law-like", many of which are clearly not related to the topic at hand and will simply be ignored (e.g. anything legal). Two distinct types that **are** both included in the analysis below are *bridge laws* that relate phenomena at different levels of organization or across different domains (or different descriptions of the same domain) and *systemic laws* that relate elements of systems.

[6] A clichd example is that in certain simple atomic models of chemistry the description 'water' unambiguously refers to an oxygen atom connected to two hydrogen atoms with no extra or missing electrons (these configurations have their own names) and that pattern can therefore be recovered from the description 'water'.

Theories are taken to be a subclass of models that have the further qualities of (1) making sufficiently accurate predictions about "the world"[7] and (2) purporting to provide an *explanation* (see below) of the modeled phenomena. As models, theories **cannot** be true or false but rather useful and/or appropriate to varying degrees. Failing to match the data will likely reduce the theory's usefulness and make it seem inappropriate for the domain, but this doesn't make the theory itself false: just the claims made using the theory. There is a great deal of literature on the subject of the role of theories and the tradeoff between usefulness for (1) and (2). The distinction depends on the meaning of 'explanation', but providing an account of 'explanation' that does not collapse or make circular the distinction between (1) and (2) is not trivial.

An example may help highlight the difficulty. Neural net models are sometimes used to model systems with large parameter spaces and unknown relationships. Neural nets are nothing more than a fancy curve fitting algorithm similar to least squares regression in purpose and, as such, yield highly accurate predictions whenever the future is similar enough to the past with respect to the data fed into the model. It is generally accepted that neural net models can provide no explanation for the phenomena being modeled despite the accuracy of the predictions: nothing addresses the "why" questions. Hempel holds that explanations enable us to understand why phenomena occur by detailing the circumstances and laws that make the phenomena "expected" (Hempel 1965). But the specification of a theory (i.e. one kind of model) and a neural net model specification cannot be distinguished in this way: both provide lists of rules and conditions that produce the phenomena as a prediction.

Thus the two strongest candidates for refinements of 'explanation' that could distinguish between models with and without it is to (1) employ an empirical requirement on the observability of the states and laws or (2) tie it to causation. The argument for (2) runs thus: sure the predictions of the neural net are just as good (or better) than the theoretical model, but the rules governing its behavior do not track the causal apparatuses of the modeled system the way that the theory's system of laws does. The obvious response is that the claim that laws track the causal apparatus seems to hinge on the accuracy of their predictions and therefore merely pushes the burden of explanation down another layer. So (1) seems to be the preferable account of explanation though it requires one to bite a number of bullets regarding the nature of causation, the plausibility of physicalism and reductionism, the limits of scientific knowledge, and our folk theories of just about everything. What follows is an argument for (1) and against (2) based on models of science that reveal the mereological relationships of "the things in the world".

[7] These may be predictions about the past, that is, matching data already at hand. A new theory may suggest a new data collection which may validate or invalidate the theory (match the prediction or not); the point is that theories aim to match data in various ways and "mere models" need not make such a claim.

2.4 Physicalism, Reductionism, and Foundational Physics

2.4.1 Principle of Reduction

There may be (and most likely is) considerable dispute on the precise meanings and entailments of the doctrines of physicalism and reductionism, but I don't expect that my use of these ideas is affected by the nuances distinguishing the flavors of each one. My first step is to recognize a general *principle of reduction* relating entities, laws, behaviors, properties, etc. at one level of organization to entities, laws, behaviors, properties, etc. at a lower level of organization (i.e. a higher resolution description and, if it doesn't already exist, a higher resolution domain).

2.4.2 Bridge Laws

Level-to-level reductions are done via a collection of *bridge laws* which are laws insofar as they are reliable conditionals that state if X is the case at some level then Y is the case at some other level. Bridge relations share both this formal similarity and an empirical discovery apparatus with systemic relations. But whereas systemic relations are meant to hold only in some possible worlds, bridge relations, once discovered, are definitions of equivalence classes and hence are true in all possible worlds. Bridge relations take as antecedents system descriptions (or some part thereof) from one level and output descriptions (or some part thereof) at another level. I do not assume that these relations are symmetric; bridge laws that "go down" are *reduction relations* and bridge laws that "go up" we may call *emergence relations*. While reduction and emergence are complementary notions, we will see later that significant asymmetries exist between the two directions.

2.4.3 Multiple Realization

Under the lossy compression conception of descriptions of the world at higher levels of organization, bridge relations must be either asymmetric or "wildly disjunctive" (Fodor: 1999, Kim: 1999). Considering the directional refinement of bridge laws made above we can see that reduction relations are (generally) one-to-many and emergence relations are one-to-one. That is not to say that some microstate or microlevel behavior will not play the antecedent role in a number of different emergence relations (from molecules of water may emerge whirlpools, clouds, snowflakes, etc.). Neither this asymmetry nor the wildly disjunctive nature should worry us. The ability to describe a huge variety of particular patterns of molecules as 'dogs' makes the term useful; being able to recognize dogginess but not cattiness in any of a large number of disparate configurations is also useful.

2.4.4 Physicalism

My understanding of what is generally meant by *physicalism* is the doctrine that there exists a lowest level of organization which properly and accurately describes "the world" and that this level of organization is the domain of physics. Since there are varying levels within physics, let's call this lowest level *foundational physics*. Authors vary on whether this foundational physics refers to a further refinement of the current state of the art (quark-lepton theory with quantum effects) or some other idealized, future perfected theory. The claim can be trivial if one thinks that whatever the lowest level is like we will consider it physics; I'll use the nontrivial version. This substantive version, however, requires that we have a way to distinguish physics from non-physics in terms of how the elements of the model interact. We somehow know the difference between economics and physics, so I will not spend time worrying about that now except to identify that this may be more problematic than others have assumed.

2.4.5 Reductionism

The doctrine of *reductionism* adds to physicalism the requirement that "kinds" within the higher level are also kinds of foundational physics. The kinds include the entities, laws, behaviors, properties, states, events, relations, and anything else that one could include and identify in a model. More can (and probably should) be said here regarding interpreting reductionism within the current framework, but considering that I claim to discredit physicalism, reductionism will simply have one of its necessary components surgically eliminated.

2.4.6 Translevelism

A variety of reductionisms that holds closer to the principle of reduction rather than to a specifically physicalist reduction would be that kinds at one level translate to kinds at the level being reduced to. It makes sense to call this *translevelism*. Traditional reductionism is clearly a subclass of translevelism; and though defeating physicalism would debase reductionism, translevelism is an independent hypothesis. I will claim (and attempt to demonstrate) that it is not the case that aspects higher levels can be reduced to aspects of the lower-level even though they may be reducable to a lower level description.

2.4.7 Lower Than Physics

It is apropos at this juncture to point out that it is conceivable for there to be a level of organization with (1) a domain that is different enough from the domain of physics to be considered a different kind of science and also (2) lower than the lowest level of physics. In a hierarchal view the lower levels would be reductions of the most basic elements of physics.

We might be at the threshold of such a thing at the moment with the mysterious behavior of quantum physics. A model capable of explaining (or accounting for in a way that does not satisfy the criterion for explanation above) quantum effects may be very different from what physicists are used to and may qualify as a new kind of science.[8]

2.4.8 Physical Intermediatism

Even though it is a highly plausible scenario that we will discover, explore, and exploit lower organization levels than current physics includes, we could still maintain a form of physicalism that merely claims that all phenomena can be reduced to physical phenomena even if they may also be further reduced – I'll call this *physical intermediatism*. This can be seen as a form of organizational hierarchy in which physics plays a special role: one can skip levels up an down, but everything has an instantiation in physics. If reduction to lower levels were lossless then clearly this would be fine. But if it were lossy then it would be an empirical claim that the resolution of the domain of physics cuts the world in at least all the same places as all other special sciences. I do not expect physical intermediatism to be a viable position; it reeks of being a just so story for physicalists confronted with a coherent lower-level ontology. Already models can be developed that violate this claim, but these models are not claimed to be theories in the sense of generating empirically tested (or even testable) results.

2.5 Three Kinds of Plurality

This section attempts to tease apart and identify relationships among three kinds of plurality: ontological plurality, methodological plurality, and phenomenological plurality. *Ontological plurality* picks out the notion that there exists objects in multiple categorical levels, i.e. that parts and their wholes both exist "in the world" with the same ontological status. *Methodological plurality* will be used here for a weaker notion: techniques for measuring and analyzing the world will naturally imply or require the existence of certain objects and relations and these differ for different techniques. *Phenomenological plurality* is that we, as humans, can naturally recognize and differentiate objects at different categorical levels.

2.5.1 Phenomenological Pluralism

Phenomenological pluralism probably doesn't need much defending because it has a great deal of empirical support and that's all it needs in order to be true. All it claims is that

[8] String theory, M-theory, and other permutations on this theme are one thrust in the reductive direction - holographic theory is another – though despite their non-intuitive structure they are generally included as physics.

we *seem* to perceive (say) both a car as a whole existing and one of its wheels existing and without any conceptual confusion or cognitive dissonance.

2.5.2 Ontological Pluralism

Ontological pluralism is the view that all those things which we recognize as objects exist in the world. This claim implies that tables *and* their table legs, *and* their wood fibers, *and* their carbon atoms, *and* their quarks all exist. The operative world here is 'and'. There is no barrier to claiming all these lower-level entities exist; nor is there a barrier going up in levels of organization. The furniture in this caf exists, the businesses in this city exist, the city exists, the country exists, the Earth exists, etc. I'm not claiming that one can induce upwards (this is related to the asymmetry of bridge relations), just that objects and their proper parts can be found on the same list of things that exist.

The issue can get complicated, however, if one wants to draw a distinction between naturally occurring and unnaturally occurring ("real" and "abstract") objects. For example, in the descent through the levels of organization above we eventually reached the level of carbon atoms and quarks and in the ascension we reached an object known as *the country*. Do these things exist? Carbon atoms and quarks may be "theoretical entities" in the sense that they play a role in a theory and can enter into an explanation of observable phenomena, but there are reasons for thinking that they do not exist in the same way cars, cats, and tables do (Lewis 2001, see more below). As for countries, these are poor candidates for natural kinds and great candidates for abstract objects only existing in minds and models and not in the world at all. My point isn't that ontological pluralism must claim these as real objects, but only that it is consistent with this position.

2.5.3 Methodological Pluralism

A conspiracy theory version of methodological pluralism implies that the only reason we identify (say) biological cells as entities is because they are discernable with our techniques for investigating such things. We discern a trail in a bubble chamber and hence a textbfthing (an electron) must have generated it. Conspiratorial methodological pluralism further implies that electrons exist with the charge and energy they do only because we have the ability to measure them in a certain way that makes this the case. This is a pathological type of human perception's construction of reality – that must be true in a sense, but this thesis lacks support and is certainly not the version of methodological plurality endorsed here.

A milder version is that we posit objects so as to make our theories consistent, complete, and explanatory (which includes matching data). We can be humble and admit that the posited objects may be nothing more than placeholders for a later, more refined ontology. Those objects allow the theory to work and the theory needs to work to be useful. A closely related alternative is that the posited entities aren't placeholders *for anything*, but merely abstract components of a model. We can generalize this without changes by seeing that folk models (both explicit "commonsense" theories and implicit models realized through cognition) are models. The objects we posit for our folk ontologies are partially determined by our abilities to perceive and distinguish things and partially determined by what is useful as "variables" in our folk models. (Lewis 2001)

The position argued for below can be roughly stated in terms of the above three types of pluralism: phenomenological pluralism produces methodological pluralism within our folk theories which we confuse for ontological pluralism in the world. What follows is a tying together of the above definitions and distinctions to show that causation is not and cannot be an aspect of the world and realizing that causation is only a feature of models puts causation on a stronger foundation.

3 The Case against Causation

My next step is to differentiate two concepts that share the name 'cause'. These are *singular causation* (i.e. identifying the cause of a particular event or state) and what I will call *systemic causation*. The former is a disambiguation of which chain of events among a set of possible chains of events (series of state changes) actually brought the system to its current state. Systemic causation is revealed by entailment relations among the states of a system. Both identify precursor states to a current state; the former identifies particular actual precursor states of an instance and the second identifies all the possible precursor/successor relations for states. We will be concerned with the second concept in this paper because it fills the role of being what Hume called "the cement of the universe."

3.1 Myriad Points of Consideration against Concept

This midsection of the paper lacks a coherent structure; it is a collection of thoughts and arguments that all contain crucial material to my point, but that I could not easily structure into a standard format. Clearly this is suboptimal, but the content is there it is all tied together before the end of the paper.

3.1.1 Pushing the Level of Causation Down

There is a simple inductive argument that, given certain assumptions from physicalism, reveals a fundamental flaw in a naïve interpretation of causation's role across levels of organization - specifically in multiple causes of a single event. Some event E occurs, let's say it is A's inciting a stock market crash by making a particular trade. A stock market crash is an economic event and A's bringing about E has at least one explanation in terms of economic theory. This explanation will likely invoke economic principles such as supply and demand, expected discounted return, etc. to describe how A caused E. Whatever E's economic description, the reductive physicalist will claim that there is another description of E in terms of "lower-level" phenomena. But there isn't just one lower level, an alternative description of E can be given in terms of individual psychology, biochemistry, molecular physics, quark-lepton interactions, and so on at smaller and smaller scales. There is a widely shared intuition that given two levels of description, the lower level is closer to being causally descriptive. So given that we are looking for a causal explanation of E, it follows that the best explanation would be in terms of quarks and leptons or whatever the lowest level happens to be at the current state of the art. There are several reasons for thinking there may not be a lowest level (or that we cannot have access to it) that will not be covered here (see Kim's *Supervenience and the Mind*). Well, what if there is no lowest level? That leaves a few options: 1) there is no causation in the world, 2) causation happens at each level, 3) causation is theory-dependent or otherwise essentially subjective so that it can happen at any level but only one at a time. This paper argues that (1) and (3) are compatible and are, taken together, the best option.

To address the intuition that a causal story is better at a lower level of description I need address the relation among the ideas of levels or organization, compression, and models. As an example, consider two descriptions of a cell's behavior: one in terms of organelles (cell biology) and a lower-level one in terms of molecular chemistry. Let us suppose that the life functions of some simple unicellular life form are thoroughly developed at both levels. The biological model explains the same cell functions (e.g. osmosis, protein synthesis, etc.) as the chemical model but with very different components, behaviors, and interactions. The resolution of the chemical model is much greater than the biological one and as a result the state space is far larger. By the Pigeonhole Principle it must therefore be the case that multiple system descriptions at the chemical level map onto a singular description at the biological level. The biological description is a lossy compression of the chemical description. Since a) one cannot recover the chemical description from the biological one, b) certain interactions between components of the chemical model that are compressed in the biological model are supposed to be causal, and c) all interactions at the biological level have counterparts in the chemical model, if there is causation in the system, then the causation must be operating on the components of the lower-level model.

3.1.2 Reduction is Easier than Emergence

As Sklar points out, "Sometimes the only fruitful procedure is to use the known laws of the upper level theory to discover those connections of those laws of the foundational theory (Sklar 2003)." The analysis of phenomena from one level of organization to a lower one ties in the limitations of our theory-laden observations of phenomena and the asymmetry of emergence (Schaffer 2007b). In observing some regularity in the behavior of a system, it is a natural question to ask how the behavior emerges from its constituent parts. It is not so natural a question to ask what behavior could one observe as emerging from a collection of hypothesized entities. I don't claim that the question is unnatural because this is a strange thing to want to know, but rather because automatically detecting patterns of this sort is beyond the current state of the art. People just see the wholes and parts and make ontological claims as they see fit. But if we collect parts with the aim of discovering new higher level patterns that do not come to us readily and automatically, then humans suffer from a severe case of simultanagnosia.[9]

3.1.3 Prediction Does Not Imply Understanding Causes

Law-like connections, exemplified by Kepler's Three Laws of Planetary Motion, provide accurate predictions of behavior and properties of systems through the mechanistic (i.e. mathematical) manipulation of parameters. Because plugging in parameters derived from measurement of the world into the mathematical machinery produces the appropriate results (validated by further measurements of the world), it seems reasonable to suggest that the mathematical mechanisms encoded in the laws parallel the causal mechanisms of the world in a significant way. Yet we know that for many systems the supposed causal apparatus need not be incorporated into the predictive mechanism to yield accurate predictions; for example gravity in the case of Kepler's Laws. All that is required is that whatever is generating our perceptions be rather consistent. But if causal connection is not required for accurate prediction, then it follows that a theory's providing accurate predictions (so called "theory-confirming observations") does not reveal whether, or to what degree, a theory's mechanisms describe the corresponding causal apparatus of the world. Accurate prediction and traditional causation must be seen as completely divorced.

3.1.4 No Causes across Levels of Organization

The idea of causal priority and whether it is distinct from temporal priority in concept or in application may or may not be (or need to be) a primitive notion (Mackie 1965). But

[9]Simultanagnosia is a psychological disorder the sufferers of which cannot see elements of a scene as being related to or composing a whole scene; they can see individual trees, recognize that there are many trees, but cannot recognize that they are seeing a forest.

consider causal claims across levels of organization such as "the release of neutrons caused the mass of the block of uranium to decrease" or "the firing of the nucleus accumbens caused her to feel pleasure". Insofar as the block is constituted by parts including neutrons and the mind's function is an emergent property of brain/body function the cause and effect are synchronous in these purported cases. While the expressions do satisfy many of the identified properties of causation (counterfactualness (Lewis 1973), INUS conditions (Mackie 1965), etc.) they are actually cases of extensional substitution rather than invocations of a causal relationship. The decreasing of the block's mass is nothing more than its loss of neutrons considered at a different level of organization. The feeling of pleasure does not result from the neural stimulation – it is a manifestation of it. These cross-level bridge relations are not even candidates for causal implication for they are merely changes in the description of an event or process required by the differing resolutions of perception (Fodor 1974). As explained in more detail below, these bridge relations establish equivalence classes between wholes and sets of parts and as such do not fall under our concept of causation.

3.1.5 Some *Concepts* Cannot Be Reduced

A description of an economic phenomenon in terms of quarks and leptons would hardly be recognizable as an economic phenomenon. In part this is just the simultanagnosia problem; we are simply limited, as humans, in our ability to recognize the forest upon hears lots of individual tree descriptions. But the conceptual barrier is something more; something Bennett and Hacker call the *mereological fallacy*. This fallacy is committed whenever one uses predicates in reference to a part (the brain), although only the whole (the human being) can be the proper subject matter. This is a variety of Gilbert Ryle's wider notion of a category error where a property or concept is applied to a thing, process, etc. that could not have that property or concept (either logically or ontologically) (from Ryle's *Concept of Mind*). Certain properties apply to one level of organization rather than another. So though the quark-lepton description would describe in physical terms the same physical event, no description in terms of quarks and lepton could actually be a description of an economic phenomenon because it is a category error to have subatomic particles engage in behavior defined only for groups of people. Not all concepts fail to cross levels (atoms, rocks, and planets can all collide) yet these will fail the aggregate object perceptibility test thus demonstrating that the conceptual gap and the perceptual gap are distinct.

3.1.6 Not Non-Causal Facts Plus Laws of Nature

The account is in support of a reductivist approach to scientific inquiry consistent with the discovery of instances of mereological supervenience, but is putting forward a singularist account of causation (Tooley 1990). Causation, as it is deployed in the sciences, does not supervene on the non-causal facts plus the laws of nature. Causation is theoretical - it is invoked within a model to make salient certain relations (between events, processes, states of affairs, etc.) when describing what a model aims to explain. The laws of nature play no role in causation unless these laws are simply defined as being whatever it is that makes our perceptions consistent and reliable. This definition does not fit the spirit of "laws of nature" so that path is blocked. The above supervenience claim is supported by our folk concept of causation, but then by my argument in this paper the laws of nature are just implicit (folk) mental models that render our perceptions consistent.

3.1.7 Causation Isn't Necessary *A Priori* Anything

"Explanation is the collection of pragmatically relevant, possibly counterfactual information about causation; and causation is to be identified in a necessary a priori way with whatever physical process underwrites our explanatory claims (Braddon-Mitchell: 1993)." If one accepts reductionism about causation (rather than a weaker mere supervenience form) and that macrostates are lossy descriptions of the underlying microstates then the relations (including causal relations) posited between two macrostates is a lossy description of the actual physical process that underwrites the claim. So far so good. The problem with this version of causation is that, physical processes do not underwrite our explanatory claims and hence cannot be identified with causation. Consistencies in perception are all there is to go on and so it must be these consistencies, or perhaps predictions that depend on them, that actually underwrite our explanatory claims *if anything does.*

3.1.8 Physics is a Special Science

Physicalism also becomes a pragmatic position: whatever the lowest level of the state of the art is, that's the level at which we should restrict causal language. But then that requires physicalists to admit that, when a lower level is uncovered, they had previously been using causal language at a macrolevel. That would be contrary to the spirit of physicalism. Physicalists could instead assume that there is some most basic layer (perhaps beyond our reach) and that causation applies there. But that seems equivalent to denying physical causation completely except to save physicalism in letter. Causes are invoked to explain regularities in macro models at various levels, and so we are advised to leave them as elements of explanations and out of the perceived physical world. The elements of physics, then, no matter how basic they may seem to us now, must be considered as in principle macrolevel entities, even if we never "find" their parts.

3.1.9 Macrostates Supervene on Microstates

To illustrate causal dependence let's consider the example of a Markov diagram for some system (figure below). Let each element of the diagram be a completely described state of the system at the microlevel. So each element contains an exhaustive collection of all the properties and relations (and whatever else matters) of all the parts of the system at a high resolution. Now we can color-code these graph nodes such that all and only microstates that translate into (are a reduction of, underwrite, constitute, etc.) the same macrostate share a color. The arrows (directed edges) of the diagram represent all the ways that the system can change from one state to another. Since all the information about a state is included in a graph element, whatever results in the system changing states determines where arrows are drawn. A change in the macrostate is simply an edge that connects two nodes of different colors. It is thus clear that there can be no change in a macrostate without a corresponding changing in a microstate and so representing the universe, or some subset thereof, this way makes clear that the macrostates supervene on the microstates.

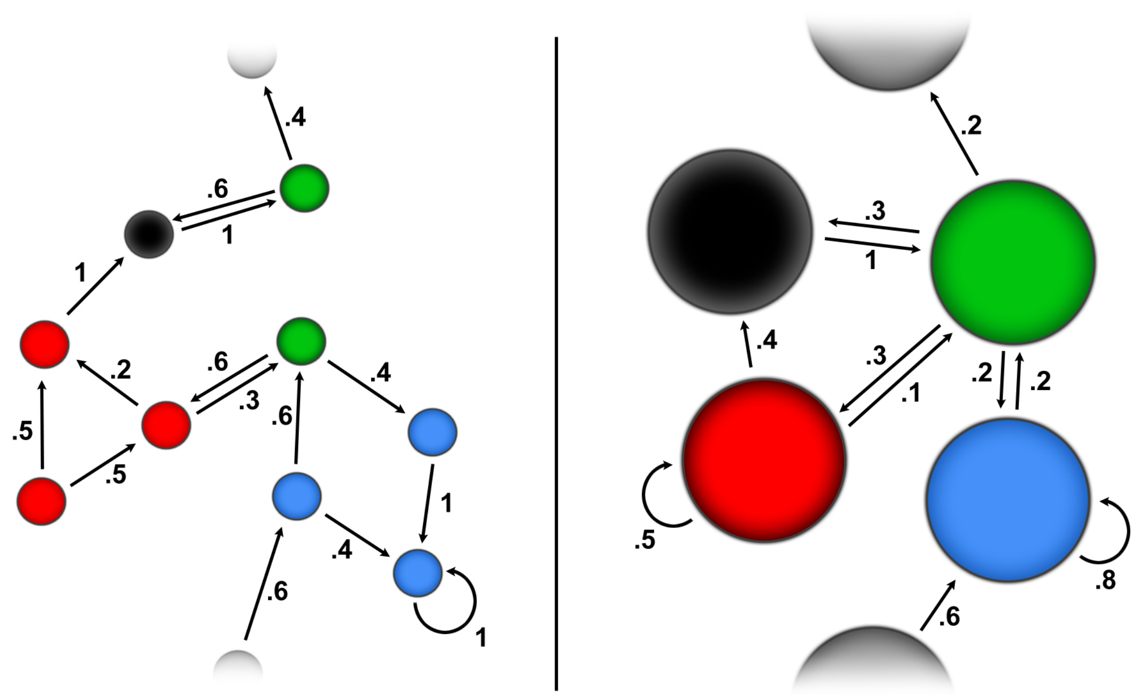

Figure 1: Microstates and Macrostates in a Markov Diagram: This diagram is a Markov representation of some model. At the macrolevel some microstates of the system are indistinguishable (green, red, blue) and others are (black). Note that the bottom-right blue microstate is an equilibrium of the model, but at the macro level this is not detectable.

In the Markov diagram the idea that macrolevel descriptions are lossy compressions of microstates is revealed by the fact that the set of nodes of a single color have a single macrolevel description but they may include multiple nodes, and so multiple descriptions, at the microlevel. This possibility bears on both whether macrolevel laws are/can be exceptionless and the multiple realization issue. A macrolevel law describing a change from macrostate A to macrostate B will be exceptionless with respect to the chosen microstate if and only if every microstate constituting A changes to some microstate constituting B. This holds regardless of whether the microlevel changes are many-to-one, one-to-one, or one-to-many.

3.1.10 Cause, Condition, or Baggage

Hart (Hart 1965) elucidates a distinction between *causes* and *conditions* of an event. The distinction, he claims, is merely a matter of one's interest. His conclusion matches my own claim; as stated above it is the salient features of a state's description that are indicated as causes for the next state change. Lewis argues that while the mind may select certain elements of the causal chain for consideration in certain context, the actual cause is always the whole causal chain (Lewis 1986). For scientific modeling I expect that conventions not too dissimilar to Grice's maxims of communication for specifying context-sensitive causal claims for our theoretical purposes.

We can create a formal distinction among causes, conditions, and irrelevant aspects of the description. Something like: causes are those aspects over which the examined variation perfectly covaries with state changes of a type, conditions are the aspects that are always fixed with respect to state changes of a type, and everything else is noise. This distinction may have some intuitive appeal or it may seem obviously false; it is admittedly an arbitrary distinction that may be useful for certain applications and incoherent for other applications. The point is that it is as good as a distinction as any other. Not only is the particular identification of causes vs conditions interest-dependent, the very conceptual distinction is just decided as appropriate for the modeling task at hand.

3.1.11 Skepticism, Solipsism, and Deceptive Demons

Regardless of the field and the tools required to obtain data, models are all build from formal idealizations of empirically revealed regularities, properties, etc. They are empirically revealed from "the world", which I have identified as simply the data against which we compare our models' predictions; this is only semi-circular. It's people who observe the world's regularities, encode them as laws, incorporate them into models, and test the models' predictions against our perceptions. The aspects included in the model may not be directly empirical verifiable (i.e. perceivable through unaided human perception), but in order to determine whether the models' output is a prediction of the world it must produce a result that can be perceived and compared to what can be perceived in the world.

The point here is the Humean point; all our thoughts supposedly about the world, including about causation in it, are all just in our minds. Our minds create models of the world to make sense of our consistent perceptions, but we cannot explain whence this consistency except within our models; and this is fully circular. Whether there is a physical world that makes our perceptions consistent or not is irrelevant, and it need be irrelevant to our concepts if we wish these concepts to be coherent and resting on a strong conceptual foundation. The solipsist may very well be correct, or the naïve realist may be, and it simply doesn't matter for causation. Causation is in our models and nobody is claiming that our models need be in (or about) any actual world.

3.2 Causal Necessity as Logical Necessity

"The goal of philosophical analysis, in dealing with such concepts [as properties, powers, and causes], should not be reductive analysis but rather the charting of internal relationships. And it is perfectly possible for a "circular" analysis to illuminate a network of internal relationships and have philosophically interesting consequences (Shoemaker 1980, p261)." There are analogies here to *my* point on the concept of causation. The reductive analysis - in my sense of uncovering relationships between levels - cannot uncover *causal relationships* because causal relationships can only be defined between aspects of states within a level and having a domain with elements at different levels of organization is an instance of the category error. Within a level of analysis it is common to find systems with circularity in their causal chains; any closed system must have this property. When a system, any system, is represented as a Markov diagram it is quite explicitly the network of internal relationships that have consequences that are scientifically, philosophically, or in any other way interesting. Mapping Shoemaker's comment onto the present analysis can link what he calls *powers* and my notion of *causes*.

Imagine two models represented as Markov diagrams with the same domain but that may differ arbitrarily in scope. Now suppose that these two models both contain a state with an identical description (i.e. both systems can be in some particular state) and hence with identical properties. "To the extent that causal laws can be viewed as propositions describing the causal potentialities of properties, it is [logically] impossible that the same properties should be governed by different causal laws (Shoemaker 1980, p261)." Seeing causation as whatever determines which state a system in some given state will go to, the implication is that two systems in the same state must go to the same state (or have the same transition probabilities). The resulting state must then also have an instantiation in both models and so its transitions must also be identical. Taking this train of thought to the limit we see that under the Shoemaker hypothesis any two models with *one identical state* must have all and only identical states and hence be *identical models*. Not only is this contrary to the assumption that we could have two models in the same domain that overlap one state and have different scopes, but it is quite contrary to ordinary intuition on the matter.

The problem is that being the same state depends on being the same description of the world at the level/resolution chosen. But we also recognize that two states described as identical at one level may vary if described at a lower level (recall that descriptions are *lossy* compressions). The transition probabilities capture an ambiguity as to which lower-level state the world is in; it is possible for $system1$ in state A to transition to B and for $system2$ in state A to transition to state C because state A covers multiple lower-level states. Models that have an identical state at one level need not be completely identical because they may vary in the included microstates and have different Markov models at different levels.

(see figure)

3.2.1 Causes Drive Changes

Though it is sensible to ask questions such as, "What causes the system to stay in the same state?", this is taken to be a different way of saying, "Why are there no causes at work?" This may be unsatisfactory in cases such as a ball remaining at the lowest point of a basin; some may claim that gravity is causing it to stay there. Yet if we take inertia as a feature of such systems then we don't need gravity to keep the ball there, but we would need something to cause it to move. We may still be unsatisfied if we consider static systems with counteracting forces. Imagine a block resting on a vertically oriented spring; gravity causes the block to compress the spring while the springs resistance keeps it from compressing completely. The two forces are constant so the system is static and the balancing of the two forces is necessary to keep it that way. At a high level of description systems having a) both forces or b) neither force are indistinguishable, but at a lower level (say, the atomic level) the balanced and inert systems have different possible transitions. The difference can be seen as the difference between a state with a strong looping transition and a state with no exiting transitions. These differences can be explicated in terms of Shoemaker's *conditional powers* of properties though the details of such an explication are somewhat tangential to the point here. We are now concerned with what aspects of models are the bearers of causation and instantiators of causal laws.

3.3 Defining Causal Laws

3.3.1 Translating Causal States into Causal Properties

Properties can also be individuated and viewed as being causally efficacious via a breakdown of the description of the states into their composing elements. Since a state description in a Markov diagram includes a specification of all the aspects of a system at that point in time we can isolate those aspects that describe the properties of interest. Then we can compare these aspects of different states and identify similarities and differences of states with respect to *these* aspects. Revealed patterns in the relationships among the identified aspects of states are precisely those relationships that can be encoded as causal laws.

The above procedure can most easily be understood via an analogy. Imagine that a state description can be encoded as a long string of characters with a specific, rigid structure: this is analogous to a DNA strand. The individual of elements of the system are smaller strings analogous to genes and their properties are analogous to codons (three base pairs that encode for a specific amino acid). We can compare strands of DNA by determining which genes and codons are shared and relate these similarities to the transition probabilities of the states. We can then catalogue which transitions are contingent on which sets of codons; these are the precisely the law-like relations among the properties of elements of the system. In some cases these relations will be unique (e.g. some set of codons is only present in states with a singular transition) and in such cases they can be defined as causal laws.

But if *only* the singular state changes (e.g. completely determined reactions) were considered to be caused then the concept of causation would have too narrow an application to do the work we put the term to do. When a state has multiple exit trajectories we would not say that the move to one of the possible next states was uncaused, but rather that we cannot uniquely specify the cause. From this we can again see from whence the intuition that causation *really* happens at the lower level derives. A state with (say) three exit trajectories may reduce to three states each of which contains only one exit trajectory. Since singular transitions are unproblematically considered causal it is natural to consider that the cause of whatever happens at the higher level is the set of salient aspects of the microstates. As intuitive as that step may be, that criterion is undefendable because a system can also have *more* singular state changes at a higher level.

What I propose is that a relation of aspects of a state to aspects of a state it can transition to is considered causal if the degree to which the two sets of aspects are correlated is sufficiently high. The threshold of correlation necessary for claiming causality is not one hundred percent, it's something lower but still quite high. Perhaps we can say that the link must hold *most of the time*; and insofar as 'most' is vague[10], so must our attributions

[10]The are some that claim that 'most' is synonymous with 'majority'. These individuals are incorrect if

of causation be. The property of being causal must be a vague property of state changes and/or of the aspects of states (properties) that are associated with the variation between states. That causation is vague is not a new idea; David Lewis in his account of causation that depends on "a relation of *comparative overall similarity* among possible worlds" also identifies this aspect because "The vagueness of similarity does infect causation, and no correct analysis can deny it (Lewis 1973)." Mackie also acknowledges that causal statements may be vague and for a reason very similar to mine; to satisfy the sufficiency part of his INUS condition the causal field must not be the entire history of the whole universe and the specification of the causal field may be vague (Mackie 1965).

3.3.2 Translating Causal Events into Causal States

In the Markov diagram of a system the events are the transitions from one state to the next. The terminology "this event caused that event" can also made sense of within the current analysis of causation (as being patterns in a system's Markov model). Instead of being patterns of particular state descriptions tending to transition into particular other state descriptions these would be patterns of transitions. If transitions were only differentiated by their probabilities then (among other things) all bifurcations would be indistinguishable and patterns in transitions would be unremarkable and unrelated to what we mean by 'causes'. If, however, transitions were typed by the states (or the salient parts of states) at the beginning and end of them then by calling an event as an event of type A we also learn that it is a transition from a state of type X to a state of type Y (or between sets of aspects of states). Let B be a transition from state Y to state Z: then A causes B translates to "states like X transition to states like Y which then transition to states like Z." This nomenclature would be particularly useful in cases where there are likely alternatives to A but not to B: i.e. "This happened because that (as a matter of *contingent* fact) happened." However this usage implies the history-based meaning of causation that this paper is explicitly *not* about (see section beginning).

Previous work has spent time ruminating on what ought to be the bearer of causation: events, properties, states, nothing, etc. The discussion is doubly difficult because all these seem to have a solid resumé for the position. If the representation of the world as a Markov model reveals the concept of causal structure with sufficient fecundity, then much of that discussion can be dissolved. Properties and states can each (i.e. both) have causal power and in the same way: as patterns in the systems' transitions. Events even share in the same

they mean to be addressing the term 'most' as it is used in colloquial English. Such a definition is analogous to claiming that 'bald' means having hair covering less than half one's head. Such artificial formalizations may be useful or even necessary for some applications, but philosophically it is more beneficial to recognize vagueness when it obtains.

sort of pattern insofar as they are the same kinds of patterns in the *dual of the graph*[11]. It is important to also remember that regardless of what one takes as cause-bearers, it is all just in the model as patterns discernable within the model. Those objects, events, etc., and the causes relating them, are not really in the world: whether there is something to refer to or not is irrelevant and hence it is also consistent, but not necessary, that *nothing in the world* is a bearer of causation. What element(s) of a model we chose to be a bearer of causation will be decided by the same pragmatic considerations that influence all the other decisions about model construction.

4 Conclusion

The paper started with the claim that our folk notion of causation as a relation between events or states of affairs of the actual world is incompatible with the principle of reduction. The problem is that reduction allows us to consider causal relations among states at different levels of organization and this implies that there would be multiple causes for the same events - contrary to our folk notion of causation. Yet referring to both (say) economic and physical causes seems coherent and is certainly useful. The solution is to realize that causation is not a relation among states or events or anything else *in the world*; it is a relation among elements of a model. The concept of causation is much stronger as a purely theoretical term: it no longer needs a world to refer to, it can apply at multiple (and sometimes across) levels of organization, and it has equal explanatory power as the folk notion. In consideration of the pragmatic justification for the inclusion of elements and relations in models and the pragmatics of communication (Grice) this conceptual change still allows causation to do most of the work we have set it to. But when we want to be careful (as scientists do) then having restrictions that keep the concept internally coherent and consistent with the other aspects of our conceptual apparatuses is an improvement.

[11] A *dual* of a graph is a representation of the nodes as edges and the edges of the original as nodes of the dual. Since the original Markov diagram is a weighted directed graph, two nodes of the dual are connected by a directed edge iff there is a path into an original node via one original edge and out via another. While it is always possible to convert in both directions, the process is not commutative: that is, you will not always recover the original graph by running the process in reverse. That isn't a problem for the point made here about events states having equivalent representations

References

[1] Braddon-Mitchell, David. "Lossy Laws" *Nous* 35:2 (2001) pp. 206-277.

[2] Braddon-Mitchell, David. "The Microstructural Causation Hypothesis" *Erkenntnis* 39: (1993) pp. 257-283.

[3] Carnap, Rudolf. "Empiricism, Semantics, and Ontology" in *Metaphysics: An Anthology.* Eds. Jaegwon Kim and Ernest Sosa. (Malden: Blackwell Publishers Ltd., 1999) pp. 13-22. Originally published in *Meaning and Necessity* (Chicago: University of Chicago Press, 1956), pp 205-21.

[4] Davidson, Donald. "Causal Relations" in *Metaphysics: An Anthology.* Eds. Jaegwon Kim and Ernest Sosa. (Malden: Blackwell Publishers Ltd., 1999) pp. 428-435. Originally published in *Journal of Philosophy* 64 (1967), pp. 691-703.

[5] Fodor, Jerry. "Special Sciences" in *Metaphysics: An Anthology.* Eds. Jaegwon Kim and Ernest Sosa. (Malden: Blackwell Publishers Ltd., 1999) pp. 504-514. Originally published in *Synthese* 28 (1974), pp. 77-115.

[6] Kim, Jaegwon. "Multiple Realization and the Metaphysics of Reduction" in *Metaphysics: An Anthology.* Eds. Jaegwon Kim and Ernest Sosa. (Malden: Blackwell Publishers Ltd., 1999) pp. 515-530. Originally published in *Philosophy and Phenomenological Research* 52 (1992), pp. 309-335.

[7] Hart, H. L. A. and A. Honor. *Causation in the Law* (1965; 2nd ed 1985).

[8] Hellman, Geoffrey and Frank Thompson. "Physicalism: Ontology, Determination, and Reduction" in *Metaphysics: An Anthology.* Eds. Jaegwon Kim and Ernest Sosa. (Malden: Blackwell Publishers Ltd., 1999) pp. 531-539. Originally published in *Journal of Philosophy* 72 (1975), pp. 551-564.

[9] Hempel, C. "Aspects of Scientific Explanation" in *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*, (New York: Free Press., 1965) pp. 331-496.

[10] Lewis, David. "Causation" in *Metaphysics: An Anthology.* Eds. Jaegwon Kim and Ernest Sosa. (Malden: Blackwell Publishers Ltd., 1999) pp. 436-443. Originally published in *Journal of Philosophy* 70 (1973), pp. 556-657.

[11] Lewis, David. "Causal Explanation" in *Philosophical Papers: Volume II.* (Oxford: Oxford University Press, 1986).

[12] Lewis, David. "Ramseyan Humility" Forthcoming in *The Canberra Plan.* Eds. David Braddon-Mitchell and Robert Nola (2001)

[13] Mackie, J.L. "Causes and Conditions" in *Metaphysics: An Anthology.* Eds. Jaegwon Kim and Ernest Sosa. (Malden: Blackwell Publishers Ltd., 1999) pp. 413-427. Originally published in *American Philosophical Quarterly* 2 (1965), pp. 245-264.

[14] Nagel, Ernest. *The Structure of Science.* (London: Routledge and Kegan Paul, 1961).

[15] Nietzche, Frederiche. " "Reason" in Philosophy" in *Twilight of the Idols.*

[16] Quine, W. V. "On What There Is" in *Metaphysics: An Anthology.* Eds. Jaegwon Kim and Ernest Sosa. (Malden: Blackwell Publishers Ltd., 1999) pp. 4-12. Originally published in *Review of Metaphysics* 2/1 (Sept. 1948).

[17] Quine W. V. "Ontological Relativity" in *Metaphysics: An Anthology.* Eds. Jaegwon Kim and Ernest Sosa. (Malden: Blackwell Publishers Ltd., 1999) pp. 45-61. Originally published in Quine, W.V. *Ontological Relativity and Other Essays* (1969), pp.26-68.

[18] Russell, Bertrand. "Existence and Description" in *Metaphysics: An Anthology.* Eds. Jaegwon Kim and Ernest Sosa. (Malden: Blackwell Publishers Ltd., 1999) pp. 23-35. Originally published as lectures V and VI of "The Philosophy of Logical Atomism" in *The Monist* (1918).

[19] Salmons, Wesley. "Causal Connections" in *Metaphysics: An Anthology.* Eds. Jaegwon Kim and Ernest Sosa. (Malden: Blackwell Publishers Ltd., 1999) pp. 444-457. Originally published in *Scientific Explanation and the and the Causal Structure of the World* (1984).

[20] Schaffer, Jonathan. "Causation and Laws of Nature: Reductionism" forthcoming in *Contemporary Debates in Metaphysics* (2007a). Available at http:\\people.umass.edu\schaffer\papers\Nihilism.pdf

[21] Schaffer, Jonathan. "From Nihilism to Monism" forthcoming in *Australian Journal of Philosophy* (2007b). Available at http:\\people.umass.edu\schaffer\papers\Nihilism.pdf

[22] Shoemaker, Sydney. "Causality and Properties" in *Metaphysics: An Anthology.* Eds. Jaegwon Kim and Ernest Sosa. (Malden: Blackwell Publishers Ltd., 1999) pp. 253-268. Originally published in van Inwagen, Peter Ed. *Time and Cause* (Dordrecht: Reidel, 1980).

[23] Sklar, Lawrence. “Types of Inter-Theoretic Reduction”, *The British Journal for the Philosophy of Science*, 18 (1967): pp. 109-124.

[24] Sklar, Lawrence. “Dappled theories in a uniform world” in *Philosophy of Science* 70.2 (2003): pp. 424-442.

[25] Teller, Paul. “How We Dapple the World” in P*hilosophy of Science*, 71.4 (2004), pp. 425-447.

[26] Tooley, Michael. “The Nature of Causation: A Singularist Account” in Metaphysics: An Anthology. Eds. Jaegwon Kim and Ernest Sosa. (Malden: Blackwell Publishers Ltd., 1999) pp. 458-482. Originally Published in *Canadian Journal of Philosophy*, suppl. vol. 16 (1990), 271-322.