CAUSAL PRIORITY AT THE SINGULAR LEVEL: FORK ASYMMETRY-BASED ACCOUNTS MEET THE GENERALIZATION STRATEGY

by

LINDA LEANNE HABEEB

A dissertation submitted to the Graduate Faculty in Philosophy in partial fulfillment of the requirements for the degree of Doctor of Philosophy, The City University of New York

1998

UMI Number: 9820537

UMI
300 North Zeeb Road
Ann Arbor, MI 48103

Approval

This manuscript has been read and accepted for the Graduate Faculty in Philosophy in satisfaction of the dissertation requirement for the degree of Doctor of Philosophy.

12/2/97
Date

Chair of Examining Committee

12/2/97
Date

Executive Officer

Professor Hartry Field

Professor John Greenwood

Professor Michael Levin

Professor Richard Mendelsohn

Professor Stephen Schiffer

Supervisory Committee

THE CITY UNIVERSITY OF NEW YORK

Abstract

CAUSAL PRIORITY AT THE SINGULAR LEVEL: FORK ASYMMETRY - BASED ACCOUNTS MEET THE GENERALIZATION STRATEGY

by

L. Leanne Habeeb

Advisor: Professor Hartry Field

I dissect one particular strategy for explaining causal priority: the fork asymmetry strategy. I argue that the best objections that have been raised are objections to versions which are less than optimal. I develop the strongest version of the principle and argue that this stronger version is fully equipped to explain causal priority at the type level. So far so good. But, I discover and attempt to deal with one remaining and perhaps insurmountable problem -- Fork asymmetry accounts cannot explain causal priority at the token level. I suggest using Daniel Hausman's notion of "causal connection" to explain causal priority at the singular level.

Acknowledgments

I am fortunate to have many caring and understanding people in my life who have supported me in diverse ways through the course of this project. Though they are too numerous to list, I am beholden to them all. I list my principal champions here.

I am indebted to Hartry Field for his constructive criticism, encouragement and patience. I thank Stephen Schiffer for his helpful advice and kind disposition. I am especially grateful for the suggestions and compassion generously bestowed by Margaret Cuonzo, Beth Hassrick and Ligia Römer.

My charming and intelligent sisters, Alisa Williams and Cindy Habeeb, have a faith in me that knows no bounds. Also a source of great comfort are my brothers, Tony Habeeb, Joe Habeeb, Jr., and Scott Williams, and my father, Joe Habeeb. My mother, Patsy Habeeb, inspires me with her inherent goodness and determination. The attentive and caring nature of my mother-in-law, Margaret Coady, father-in-law, Tony Coady, and brother-in-law, Ben Coady touches me deeply.

Last, and most importantly, I could not have completed this project without the loving reassurance, as well as the cogent philosophical input, of my cherished husband, David Coady. His tenderness, optimism and devotion fill my life with happiness. I dedicate this work to him.

Table Of Contents

Causal Priority at the Singular Level: Fork Asymmetry-Based Accounts Meet the Generalization Strategy
by Leanne Habeeb

Chapter One: Overview

1:1: What is Causal Priority?[1]

Among many other things a viable account of causation must explain why causes typically precede their effects. This phenomenon is sometimes referred to as *causal priority*. Unfortunately, there is a tendency in the literature to conflate the expressions 'causal asymmetry' and 'causal priority'. And there is a question about what we mean by each. Huw Price uses "causal asymmetry" to refer to two concerns about the causal relation: "The causal relation is asymmetric, so that if A is a cause of B then B is not a cause of A; and effects never (or almost never) occur *before* their causes."(Price 1992, p. 253) In contrast, Daniel Hausman uses "causal priority" in the following manner:

> "The relation between cause and effect is asymme ric. Causes determine their effects in a way in which effects do not determine their causes. One reliable way to tell causes from effects is by examining their temporal order. Effects seem not to begin before their causes. Yet there appears to be more to causal priority than just temporal priority." (Hausman 1984, p. 261)

Thus, Hausman seems to be more concerned with distinguishing between cause and effect than with the fact that causes typically precede their effects.

[1] I use the following notation throughout: A capital letter (X) denotes an occurrence-type of which the phenomenon represented by the corresponding lowercase letter (x) is an instance.

Presumably "causal asymmetry" means asymmetry in the logical sense (i.e., causation is not symmetric). In fact, our standard assumption about the logical relation 'is a cause of' can be strengthened to *anti*symmetry. A relation, R, is *antisymmetric* if and only if whenever Rab it is not the case that Rba. So, I use *causal antisymmetry* to refer to the fact that the causal relation is such that, given two phenomena, a and b, whenever a causes b, b does not cause a.

Part of what we mean by "causal priority" is that causes and effects are distinguishable in a way that is independent of any temporal ordering. I call this aspect of the causal relation *causal distinguishability.* A different aspect involves the fact that (at least in most cases) effects occur after their causes. I call this *causal direction.* I use "causal priority" to refer to the conjunction of causal distinguishability and causal direction. Any account of causal priority must explain each of these two features of the causal relation. This leaves us with four expressions representing different aspects of the causal relation:

1. **Casual Antisymmetry:** *If x causes y, y does not cause x.*
2. **Causal Distinguishability:** *Causes and effects are distinguishable in a way that is independent of any temporal ordering.*
3. **Causal (temporal) Direction:** *Effects rarely (if ever) precede their causes.*
4. **Causal Priority:** *causal distinguishability* plus *causal direction.*

While (1) and (3) are uncontroversial, (2) and (4) are not clearly true. A Humean cannot accept either (2) or (4). Since there is a tendency in the literature to conflate these expressions, the relationship between them (and what might hang on this relationship) is unclear. Throughout I try to sort out this relationship and its impact on current accounts of causal priority.

1:1:1: The Problem With Conventionalist Accounts

David Hume gives the following account of why causes typically precede their effects:

> "The appearance of a cause always conveys the mind, by a customary transition, to the idea of the effect. We may, therefore, suitably to this experience, form another definition of cause and call it *an object followed by another, and whose appearance always conveys the thought to that other.*" (Hume 1953, p.)

Thus, Hume stipulates that causal priority coincides with temporal priority. That is, Hume maintains that, by definition, a cause occurs earlier than any of its effects. This *conventionalist* approach has three undesirable consequences. First, backward causation (i.e., an effect occurring before its cause) becomes impossible, by definition. While everyone agrees that backward causation rarely or never occurs, it seems, at least initially, that we have no trouble making sense of the notion, because we can tell stories in which backward causation occurs. It may be the case that upon closer inspection we find we cannot make sense of the notion of backward causation, as when we try to make sense of the ıotion of a barber who shaves everyone except those who shave themselves. But, until we determine that backward causation is logically impossible, we should assume our stories involving the notion make sense. Furthermore, some physicists postulate the existence of particles called "tachyons" that travel faster than light. A particle traveling faster than light according to one frame of reference travels backward in time according to another frame of reference. Thus, in the latter frame of reference the particle brings about past events. Since we appear to b able to make sense of such things, the idea of "backward causation" is conceivable in a way that the idea of "married bachelor" is not. So, it would be better to have an account of causal priority that permits the possibility of backward causation.

The conventionalist account of causal priority also renders simultaneous causation impossible. As in the case of backward causation, however, we can conceive of cases where the cause and effect occur simultaneously. For example, according to Newton's law of gravitation, if a mass suddenly comes into existence it instantaneously exerts a force on all other masses. So, we should also look for an explanation of causal priority that does not rule out the possibility of simultaneous causation.

Finally, we make certain assumptions about the causal relation that demand a deeper explanation of why causes precede their effects. I discuss these assumptions shortly. For example, one can manipulate causes to bring about effects, but cannot manipulate effects so that they bring about causes. This suggests a distinction between cause and effect that does more than define causal priority as temporal priority. Thus, defining causal priority as temporal priority lacks explanatory power in the sense that it cannot account for some fundamental assumptions we make about the causal relation.

In response to these undesirable consequences many philosophers have turned their attention to providing an account of causal priority that does not define causal priority as temporal priority. On such accounts causal priority rests on an objective asymmetry in the world that is (usually) aligned with temporal ordering. One such objective asymmetry in the world is the *fork asymmetry*, roughly, that for any pair of correlated events, where one event does not cause the other, there is usually a third event which causes them both. In short, my project centers around attempts to explain causal priority using the fork asymmetry. I contend that such approaches have trouble dealing with a certain fact about causation, namely, that we normally describe causation as occurring at both the general (e.g., causal claims like "Smoking causes lung cancer") and singular level (e.g., claims like "Cindy's smoking causes her to develop lung cancer").

Fork Asymmetry accounts of causal priority rely on correlations between event-types to explain causal priority. The fork asymmetry principle itself distinguishes between cause and effect by pointing out that causes sometimes function as screeners-off, while effects do not. To be more precise, given causal phenomena, A, B, C and D, if C is a common cause of A and B, when we hold C fixed the correlation between A and B disappears. But, rarely is it the case when D is a common effect of A and B that the correlation between A and B disappears given D. In fact, when D is a common effect of A and B something like the opposite of screening-off occurs. Holding D fixed *creates* correlations between A and B.

Screeners-off, because they depend on correlations, exist only at the type level; that is, a screener-off is a phenomena type. We cannot distinguish between cause and effect at the token level using the fork asymmetry principle, because there are no token level screeners. For example, consider the following two claims:

(1) Smoking causes lung cancer and heart disease.

(2) Cindy's smoking caused her lung cancer and her heart disease.

If smoking screens off the correlation between lung cancer and heart disease we can determine that (1) is true. But, we cannot do this for (2). Cindy's smoking, by definition, is not a screener-off. So, fork asymmetry-based accounts explain causal priority only at the level of general causal claims.

Most proponents of fork asymmetry accounts do not examine the manner in which causal priority moves from the general to the singular level. Perhaps this is because they take a Humean approach to explaining the relation between singular and general causal claims. Humean theorists hold that a singular causal claim is best characterized as an instantiation of some general causal claim. If we analyze singular causal claims in terms of general causal claims we would expect causal priority to fall naturally from the general to the singular level. That is, the singular causal claim would derive its causal priority from that of the general causal claim that instantiates it. If, on the other hand, we analyze general in terms of singular (as I think we have good reason to do), we need some explanation of how the fork asymmetry principle is relevant to singular level causal priority.

Nancy Cartwright puts forth an analysis of general causal claims as generalizations over singular ones (i.e., the *generalization strategy*). She provides a means of generating causal laws from causal instances. (Cartwright 1989, p. 134) In Chapter Four I explore the possibility that causal priority somehow "filters down" from the general to the singular level. In the process I examine whether the use of laws or probabilities in some version of the generalization strategy will help to explain how this filtering might occur. Since I reject taking general causal claims as basic, opting to analyze general causal claims in terms of singular causal claims, I submit that proponents of the fork asymmetry account have some explaining to do. Though fork asymmetry accounts explain causal priority at the general level, they have yet to show how we explain causal priority at the singular level. I submit a suggestion about how fork asymmetry accounts might explain causal priority at the singular level.

To this end I address the following prefatory queries:

1. What precisely is causal priority?
2. What is the strongest formulation of the fork asymmetry?
3. What is the relationship between singular and general causal claims?

1:1:2: Objective and Subjective Aspects of the Causal Relation

There are certain epistemological and pragmatic, as well as metaphysical questions to address when clarifying the notion of "causal priority". From the epistemological point of view, we observe that choices are made for the sake of what they might cause. With respect to phenomena occurring in the present we can observe records of their causes but not their effects. We understand there to be some means of distinguishing between cause and effect. Metaphysically we are concerned with why it is the nature of the world that effects rarely, if ever precede their causes. We concern ourselves with the fact that there is some feature of a cause that distinguishes it from its effect and vice versa. This leaves us with the four following questions that an adequate theory of causal priority must answer:

(1a) What is the difference between cause and effect?

(1b) Why do we think there is a difference between cause and effect?

(2a) Why do causes happen before their effects?

(2b) Why do we think causes happen before their effects?

(1a) and (1b) are questions about causal distinguishability. They reflect the fact that we take the causal relation to be like the relation 'is higher than' and unlike the relation 'is to the left of' in that there seems to be a sense in which there is a *genuine* difference between 'high' and 'low' but merely a *stipulated* difference between 'left' and 'right'. In considering (1a) we assume that there is an objective difference between cause and effect. In the case of (1b), we allow for a subjective explanation of causal priority, i.e., the difference between cause and effect is tied to our perspective as agents.

I borrow this division from Horwich (1987) and Price (1992). Horwich distinguishes between a metaphysical question about causal priority and an epistemological one. The metaphysical question asks: "Why do causes typically precede their effects?" Whereas, the epistemological question asks: "Why are we right to hold that causes typically precede their effects?" (Horwich 1987, p. 141) (By "right", I think he means "justified".) Horwich answers each question in a different manner:

> "...I think the direction of causation is explained by ι oting that time order is a 'constituent' of the causal relation (in a sense to be described later), but this fact is a posteriori." (Horwich 1987, p. 131)[2]

[2]Horwich explains further, "I have characterized the structure of causation by starting from a form of nomological connection and adding requirements of causal continuity and time order. Given this theory of what causation is, the explanation of why causes typically precede their effects is simply that this feature is

His project is to provide a Humean answer to the metaphysical question and give an *a posteriori* explanation for the epistemological one.

Price, on the other hand, admits both a *de re* and *de dicto* reading of two questions: Why is the arrow of causation so well aligned with the arrow of time? and What is the difference between cause and effect? (Price 1992, p. 253) He claims that (2a) contains the implicit assumption that the direction of causation is "fixed independently of our disposition to treat the direction in question as that of the future rather than the past." (Price, 1992a, p. 253) In other words, according to (2a) an adequate theory of causal priority suggests that if there were no humans there would still be a predominant direction of causation. I call such accounts *objective*. Thus, (2a) does not allow for accounts that depend on our perspective as agents (i.e., *subjective* accounts). On the other hand, (2b) "allows that in a world in which we ourselves had the opposite temporal orientation, we might take both arrows [temporal and causal] to point in the opposite direction." (Price, 1992a, p. 253) Thereby, (2b) allows for agency-based accounts of causal priority.

(1a) and (1b) are about causal distinguishability, whereas, questions (2a) and (2b) are about causal direction. Furthermore, (1b) and (2b) allow for subjective accounts of causal priority, while (1a) and (2a) accommodate objective

---

constitutive of the causal relation -- which is tantamount to saying that the fact has *no* explanation." (Horwich 1987, p. 138) accounts only. Since Price defends an agency-based account of causal priority, he focuses on questions (1b) and (2b) which allow for such accounts.

In sum, giving a theory of causal priority is tantamount to answering questions (1a-2b) above. Our concern is, in part, identifying which questions a particular account answers and whether there is a significant difference in the answer to each. Horwich addresses (1b) and (2a) in providing his Neo-Humean account. Price answers (1b) and (2b) and advocates an agency-based account. Daniel Hausman seems to focus on (1a) taking causal connection as primitive and asserting that the difference between cause and effect is that effects are more "connected" than causes. (Hausman 1984, p. 261, 265) Thus, the difference between distinct theories of causal priority depends, in part, on which question(s) the theory answers, as well as how it answers them.

1:1:3: Our Causal Assumptions

I contend that our assumptions about the causal relation divide naturally into two classes: our *logical* assumptions and our *practical* ones. Although an adequate account of *causal priority* should be expected to reflect most, if not all, of the standard assumptions about the causal relation, some of these assumptions are better grounded in our intuitions about causation than others.

Our purely logical assumptions about "is a cause of" include the following. Given certain phenomena x, y and z:

1. Causation is transitive; if x causes y and y causes z, then x causes z.
2. Causation is antisymmetric; if x causes y, then y does not cause x.
3. Causation is irreflexive; it is never the case that x causes x.

Although these are standard logical assumptions about the causal relation, none is immune from doubt. For example, here is a putative counterexample to causal transitivity borrowed from Hartry Field. Suppose that unbeknownst to me a mad bomber breaks into my apartment while I am away and plants a remote controlled bomb. A detective in pursuit of the bomber finds her as she sets the timer on the remote and destroys the remote before the bomb explodes. Certainly, the bomber setting the timer (T) causes the detective to destroy the remote (D). Presumably, also, the detective's destroying the remote causes me to survive (S). But, it is far from clear that setting the timer on the remote causes me to survive. It seems (at least initially) that setting the timer does *not* cause me to survive.[3] This case and others like it trade on the fact that there is a phenomenon whose nonoccurrence causes some other phenomenon. It is not clear what to say about such cases. In fact, Ellery Eells insists that causation is not transitive. Rather than describing causation using chains of direct causation, he describes it in terms of causes and "causal ancestors". According to Eells, there is no such thing as indirect causation. Consider three causal phenomena types, C, X, and E. Suppose C causes X and X causes E. According to Eells, C does not cause E. Rather, C is a "causal ancestor" of E.

In addition to our logical assumptions we have certain practical assumptions. These include:

1. Causes precede their effects. (directionality) (Horwich 1987, p. 144)
2. With respect to phenomena occurring in the present, we can have records of their causes but not their effects. (fixed history) (Horwich 1987, p. 144)

---

[3] Field's original example involves causation under indeterminism. He then argues that although the probability of (T) raises the probability of (D) and the probability of (D) raises the probability of (S), the probability of (T) lowers the probability of (S). I ignore here the question of whether we are dealing with determinism or indeterminism. For reasons I discuss in this chapter I concern myself primarily with causation under indeterminism and treat determinism as a special case of indeterminism.

3. If the cause had not occurred, then the effect would not have occurred. (counterfactual dependence) (Horwich 1987, p. 10)
4. Choices are made for the sake of what they might cause. (manipulability) (Horwich 1987, p. 144)
5. Instances of correlated event-types are causally connected. (causal connection) (Horwich 1987, p. 144)

These practical assumptions (with the exception of directionality) are the reason why we want a more substantive account of causal priority than Hume's. Defining causal priority via temporal priority does not adequately account for them.

1:1:4: Other Issues that Fix the Nature of Causation

Because causal priority is intimately tied to other issues that impact on causation, I believe it necessary at least to address these topics, though I do not examine each in detail. In most cases I merely state the assumptions I make with respect to the issue in question. With respect to a few, e.g., the relationship between singular and general causal claims and the distinction between objective and subjective features of the causal relation, the issues are so firmly tied to the notion of causal priority that they require closer inspection. I discuss these issues in depth subsequently.

Causal Relata: When examining particular accounts of causal priority I employ the causal relata the author describes. A viable account of causal priority, nonetheless, must be able to account for facts, events, conditions, etc. Therefore, I look for an account which can explain causal priority in a manner independent of what entities count as causes and effects. Whether causal relata are facts or events or conditions, etc., accordingly, should have no impact on the following discussion. In any place where the reader thinks the discussion hinges on what the basic units are she should take the object of causation to be Bennett-type events. A 'Bennett-event' is the assignment of a very detailed property to a particular region of space-time. The level of detail Bennett requires is unclear, but one might require that the detailed property is the conjunction of *all* the intrinsic properties of the region.

Indeterministic versus deterministic frameworks: A theory of causation must account for the following two cases:

1) when the cause is present, the effect will definitely occur, and

2) when the cause is present the effect will probably occur.

I think it natural to assume we can construct an indeterministic causal framework to explain determinism as a special case of indeterminism. Namely, a cause determines its effect when the probability that it will bring about the effect is 1. If, on the other hand, we start off with the contention that causes determine their effects, we cannot explain situations in which the cause is merely 'likely' to produce its effect as cases of causation. For this reason I assume causation under indeterminism and consider accounts of causal priority that can handle the assumption that causes do not necessarily determine their effects.

Probability: Since I desire a formulation of the fork asymmetry principle I can use to give a non-Humean explanation of causal priority[4], I must determine whether the notion of probability employed in expressing the principle contains a built-in temporal ordering. Theories employing such temporally asymmetric interpretations of probability must explain how they can use the principle to give a non-Humean account of causal priority. Otherwise, they must base the fork asymmetry on another interpretation. I prefer a relative frequency interpretation of conditional probability. Despite problems ascertaining the value of the limiting frequency, as well as the question of whether any limit of the relative frequency even exists, unless otherwise specified I assume a relative frequency interpretation of probability. Furthermore, I advocate a frequency interpretation that does not describe P (A/B) as *the frequency with which A succeeds B.* As Huw Price points out, this interpretation is problematic for fork asymmetry theorists because it introduces "an asymmetry unaccounted for, an arrow which appears to be oriented only by convention". (Price 1992, p. 255) That is, it appears to stipulate that the orientation of probability coincides with the orientation of time.

---

[4] The account will be non-Humean in the sense that it does more than merely stipulate causal order using temporal order.

Because fork asymmetry-based accounts of causal priority are objective in the sense that they insist on a predominant orientation of causation even if there are no agents, I avoid agency-based interpretations of probability. Specifically, I avoid interpretations based on a on a person's degree of conviction in believing a given proposition even when the person whose degree of conviction in believing a given proposition is "ideal". That is, her convictions do not violate the rules of probability calculus.

Subjective versus objective basis for causal priority: In section ***1:1:3*** I discussed the division between subjective (dependent on our perspective as agents) and objective (independent of our perspective as agents) accounts of causal priority. My discussion centers around objective accounts, specifically, ones based one the fork asymmetry. Such accounts should avoid interpretations of probability that have a built-in temporal ordering, since they aim to provide an explanation that does not involve defining causal priority via temporal priority. Thus, I avoid subjective interpretations of probability in formulating the fork asymmetry principle.

Backward causation: A discussion of backward causation ties naturally into the issues surrounding causal priority. For instance, I suggested above that one requirement for an viable theory of causal priority is that it not rule out the possibility of backward causation. Whether there can be cases of backward causation is, nonetheless, controversial. We may determine upon further reflection that we cannot make sense of the notion. Since we at least seem initially to be able to make sense of backward causation, I require that any workable theory of causal priority not rule it out in principle. I do not discuss the controversy surrounding our ability to make sense of the notion.

Singular versus general causal claims: A phenomenon, P, is said to "screen-off" a correlation between two other phenomena, Q and R, if the correlation between Q and R disappears when we specify (or hold fixed) P's influence (i.e., we specify whether P is present or absent). In the case of "smoking causes lung cancer and heart disease", for instance, smoking screens of the correlation between lung cancer and heart disease. That is, once we partition the relevant population (in this case people) into smokers and non-smokers, the correlation between lung cancer and heart disease disappears among all people.

The fork asymmetry principle states, roughly, that a common cause screens-off the correlation between its joint effects, but a common effect cannot screen-off the correlation between its joint causes. Thus, the fork asymmetry principle allows us to distinguish between cause and effect by noting that causes can function as screeners-off, while effects cannot. But, talk of correlations makes sense only as a relationship between *types*. Thus, a screener-off occurs only at the type level. In the case of a singular causal claim there is no screener-off because there is no correlation. For example, suppose Cindy's smoking, her lung cancer and her heart disease all occur. Though we might count Cindy's smoking as a cause of both her lung cancer and her heart disease, is not a screener-off because the relationship described is between *instances*.

Because the fork asymmetry is cashed out in terms of screeners-off, it seems, at least initially, that the fork asymmetry principle can potentially explain causal priority only at the general (or type) level. But, a satisfactory account should explain causal priority for singular *and* general causation. Thus, the fork asymmetry strategist must describe how causal priority reaches the singular level. In other words, she has some more explaining to do.

A further complication for the fork asymmetry strategy arises from my contention that general causal claims are best described as generalizations over causal instances. I call this the *generalization strategy*. Critics have been too quick to dismiss this approach. I argue in **Part II** that it is a viable option and show why it is preferable to analyzing singular causal claims as instantiations of general ones. (see section ***1:2:*** below)

Microphysics versus macrophysics: Another means of distinguishing between causal claims is by separating them into micro and macro-level claims. Micro-level processes occur between particles. Macro-level processes occur among combinations of particles large enough to observe. Thus, upon introducing this distinction we have at least four types of causal claims: micro-level singular causal claims, macro-level singular causal claims, micro-level general causal claims and macro-level general causal claims. The scientific laws governing occurrences on the micro level appear to be symmetric. Those governing occurrences at the macro level appear to be antisymmetric. This could create a problem for the fork asymmetry theorist if she wants to explain causal priority at the micro level.

I divide the remainder of my discussion into three parts based on the following questions:

1. What is the strongest formulation of the fork asymmetry?
2. What is the relationship between singular and general causal claims?
3. How does causal priority get to the singular level?

1:2: The Strongest Formulation Of The Fork Asymmetry **(Part I)**

In **Chapter Two** I examine various formulations of the fork asymmetry principle to determine which will best aid us in giving a fork asymmetry-based account of causal priority. I use what I take to be the strongest formulation of the principle in **Chapter Three** to clarify and strengthen Paul Horwich's account of causal priority from *Asymmetries in Time.* I then defend a modified version of Horwich's account of causal priority against objections put forth by Arntzenius and Price.

I argue that a relaxed version of Frank Arntzenius' formulation of the principle is most successful in meeting certain necessary criteria for the principle to drive causal priority. Among the criteria I put forth are the following. First, the formulation must provide a means of determining *when* there is a 'common cause' or when is it reasonable to postulate a common cause. In the literature authors tend to move back and forth between the claim that "correlated event types *typically* have common causes" and the claim that "correlated event types *always* have common causes." The more desirable formulation is the weaker alternative, because it allows, for example, for cases of accidental correlations among three, non-simultaneous phenomena. Second, a proper formulation must contain a component that explains why causes precede their effects and how we distinguish between cause and effect that does not invoke temporal direction. Finally, formulations of the principle must not incorporate type-token ambiguities or prefer that a certain kind of entity counts as causes.

To illustrate, consider two of Horwich's attempts to formulate the principle. In giving his account of causal priority he presents several, sometimes incompatible formulations of the fork asymmetry. One formulation states "Genuine coincidences rarely occur. That is, given a strong correlation between events, A and B, there is generally some explanation -- some earlier event, C -- that causes them both." (Horwich, 1987, p. 72-74) This first formulation suffers from a type-token ambiguity. The fork asymmetry is allegedly a fact about the world observed through an examination of correlated events. But, talk of correlation makes sense only with respect to event *types*. The principle as stated above asserts that an earlier event causes two (or more) later events. Talk of one event causing two (or more) other events makes sense only with respect to event *tokens*. So, given the aforementioned criteria this formulation will not do as stated above.

Horwich also formulates the fork asymmetry as follows:

> "Whenever two event types are correlated with one another, they are embedded in a V-shaped chain of nomological determination, but need not be embedded in an Λ-shaped pattern; thus there is always a characteristic antecedent event, but there need be not characteristic subsequent event."

A V-shaped fork points towards the future. Thus, the common cause occurs prior to the joint effects. Whereas an Λ-shaped fork points towards the past. In this case the common cause occurs later than the joint effects. This is Horwich's strongest formulation. It suggests that whenever there are correlated event types they are embedded in a V-shaped, rather than an Λ-shaped, fork. The event at the tip of the fork is the common cause. So, the formulation clearly provides a means of distinguishing between 'common cause' and 'common effect.' Further, it avoids type-token ambiguity. Using this formulation, Horwich can also account for causal direction in a more substantive manner than Hume. The distinction between V and Λ does depend on time order. It arises from a contingent fact about the world; there are no Λ-shaped forks. Though this is Horwich's best formulation of the fork asymmetry principle, it is too strict. Horwich's use of "Whenever" and "is always" suggest that correlated event-types *always* have a common cause. So, this formulation is also problematic; it cannot account for spurious correlations.

The Spirtes, Glymour, Scheines formulation of the fork asymmetry results from their use of two axioms that show how probabilities relate to causal graphs: 1) Causal Markov Condition and 2) the Faithfulness Condition. They do not explain why causes generally precede their effects. Rather, they characterize various assumptions we make about the asymmetry of causation through a formalization of the fork asymmetry.

Though their formulation does have a built-in asymmetry, this asymmetry is not a temporal one. On their view, we can distinguish between cause and effect because they relate differently to the notion of probability (roughly, a frequency notion of probability). That is, we can explain the relation 'having a common cause' probabilistically. But we cannot give the same probabilistic explanation for 'having a common effect' (though we can give some probabilistic explanation). Thus, they indirectly show how we *distinguish* between cause. But, they do not build temporal asymmetry into their formulation of fork asymmetry. Thus, the Spirtes, Glymour, Scheines formulation of the fork asymmetry can generate the necessary asymmetry for causal priority without appeal to temporal priority, making it very attractive for a fork asymmetry account of causal priority.

I also consider why some other attempts to formulate fork asymmetry fail (i.e., Reichenbach and Lewis). I think Horwich's account makes some interesting moves though there are gaps and ambiguities in his account we must eliminate before we can appreciate the full strength of his theory. So, I use a modified version of Arntzenius' formulation of the fork asymmetry to give a clearer and stronger restatement of his explanation of causal priority.

1:3: The Relationship Between Singular And General Causal Claims (Part II)

Here I investigate the nature of the relationship between singular and general causation. Humean theorists describe singular causal claims as instantiations of general causal claims. Thus, they take general causal claims as basic and describe singular causal claims via general ones. The motivation for this approach is based on assumptions about whether the cause and effect mentioned in a causal claim must occur for the claim to be counted as true. I refer to the claim that the cause and effect in causal claims (either general or singular or both) must occur as the *occurrence requirement.* Eells in particular insists that singular causal claims must adhere to the occurrence requirement though general claims need not. For this reason, he argues, attempts to analyze general causal claims via singular ones will fail. For example, if no smoker develops lung cancer, there are no singular causal claims over which to generalize. Thus, we cannot generate a general causal claim in this case. According to Eells, this conflicts with our intuitions about true general causal claims.

If we begin at the general level, however, we can explain general causal claims directly in statistical terms. Singular causal claims are then analyzable as instantiations of general ones. The Humean strategist analyzes singular causal claims by tacking the occurrence requirement onto some description of general causal claims. For example, she might claim that c causes e if and only if c occurs, e occurs and C causes E. Thus, initially, it appears easier to analyze singular causal claims using general causal claims than to analyze general causal claims using singular causal claims. Generalization strategists, on the other hand, describe general causal claims as generalizations over singular causal claims. One option open to the generalization strategist is to accept that for any general causal claim in which either the cause or effect does not occur, the claim is false. Nancy Cartwright and John Carroll take this approach.

Critics like Hitchcock and Eells have been quick to dismiss attempts to describe general causal claims in terms of singular causal claims. As a result, not much attention has been given in the literature to formulating the generalization strategy. In **Chapter Four** I maintain that the strongest Humean account ('Early' Eells) cannot avoid appeal to individual cases as it stands. So, the generalization strategy merits more attention. There do not seem to be any serious (fully spelled out) attempts to analyze general causal claims in terms of singular causal claims. As a result, in **Chapter Five** I catalog the desirable features a viable version of the generalization strategy should possess. I consider of the accounts of Cartwright and Papineau and some suggestions from Carroll to illuminate these features. I conclude by advocating a modified version of Carroll's generalization strategy.

There are alternative analyses of the relationship that introduce the idea of *varying degrees of dependency* between singular causal claims and general causal claims. Ellery Eells ('Late', 1991) attempts to show that singular causal claims and general causal claims must be analyzed *independently* of one another. Christopher Hitchcock (1995) uses the notions of 'big' and 'little' probability spaces to suggest that there is a common structure that both sentence types satisfy. While these suggestions are useful if we cannot explain one sentence type wholly in terms of the other, the Unanimity theory and Generalization strategy are simpler, more natural, and, thus, more attractive options. Since I defend a version of the generalization strategy in this chapter, I do not discuss these alternative approaches.[5]

[5] Even if Eells' independence thesis or Hitchcock's partial reduction theory were to turn out to be more attractive than any version of the unanimity theory or generalization strategy, we still must, if we adopt the fork asymmetry account of causal priority, explain how the asymmetry gets from one level to the other.

1:4: Causal Priority at the Singular Level (Part III)

In **Chapter Six** I motivate a preliminary explanation of the mechanism by which causal priority gets from the general to the singular level. I use elements of Daniel Hausman's theory of causal priority (taking the notion of causal connection as primitive) to describe and assess one way in which causal priority might be said to come in at the general level and 'filter down' to the singular level.

PART I:
THE STRONGEST FORMULATION OF THE FORK ASYMMETRY

Chapter Two: *The Fork Asymmetry Principle*[6]

I previously argued that the fork asymmetry theorist needs to explain how causal priority reaches the singular level. But, I have not yet precisely described the fork asymmetry principle (a.k.a., the principle of common cause) itself. A common cause explanation states, roughly, that when two types of phenomena, A and B, are correlated more than mere coincidence allows, and one does not cause the other, there is another type of phenomena, C, that causes them both. We believe such explanations to be future-directed. Consider the following examples. Suppose smoking causes lung cancer and smoking causes heart disease. We expect a correlation between lung cancer and heart disease among smokers because we take smoking to be a good explanation of this correlation. There are, however, no analogous "common effect" explanations. Suppose smoking causes lung cancer and so does coal mining. Despite the fact that lung cancer is a

[6]I argued previously that our intuitions about the causal relation indicate the causal relation is *anti*symmetric, rather than *a*symmetric (i.e., not symmetric). That is, whenever x causes y, y does not cause x. It is not clear at this point in my investigation whether the fork asymmetry principle entails that an inverse fork never occurs, or rarely occurs. If an inverse fork never occurs, the principle would be accurately described as a fork *anti*symmetry. If, on the other hand, we find that an inverse fork sometimes, though rarely, occurs, the principle describes a weaker relation: one that is *not* symmetric. Thus, we may find that describing the relation as *a*symmetric misrepresents the principle. Also, there is no need, given our assumptions governing the causal relation, to allow for cases in which a phenomenon, c, is a common cause of itself and some other phenomenon.

common effect of both types of phenomena, we do not expect a correlation between coal mining and smoking, because lung cancer is not a good explanation of any correlation between them. Thus, the fork asymmetry principle works only in the future direction.

There are several objections to fork asymmetry-based accounts of causal priority. Several versions of the principle exist in the literature. So, it is essential to determine which formulation of the principle can best facilitate an explanation of causal priority before assessing whether any fork asymmetry-based account can withstand objections. In this chapter I sort through various formulations of the fork asymmetry principle and explain why some formulations are better disposed than others to drive causal priority. In the process I set out certain criteria a suitable version of the principle must meet. I contend that a relaxed version of Arntzenius' fork asymmetry principle is best able to meet the necessary criteria.

2:1: Some Versions of the Principle

Reichenbach first introduces the fork asymmetry principle (a.k.a., the principle of common cause). He uses the notion of *conjunctive fork* to show that in a common cause explanation the common cause, C, raises the probability of each of the correlated events, A and B, whereas, in the absence of C, A and B are statistically independent. In other words, C *screens-off* the correlation between A and B. According to Reichenbach, a conjunctive fork is the statistical relation between three events, A, B, and C, such that one explains the fact that the conjunction of the other two is more frequent than it would be for independent events. He defines 'conjunctive fork via the following four relations (Reichenbach, 1956, p. 159):

(1) Prob(A/C)>Prob(A/not-C)

(2) Prob(B/C)>Prob(B/not-C)

(3) Prob(A&B/C)=Prob(A/C).Prob(B/C)

(4) Prob(A&B/not-C)=Prob(A/not-C).Prob(B/not-C)

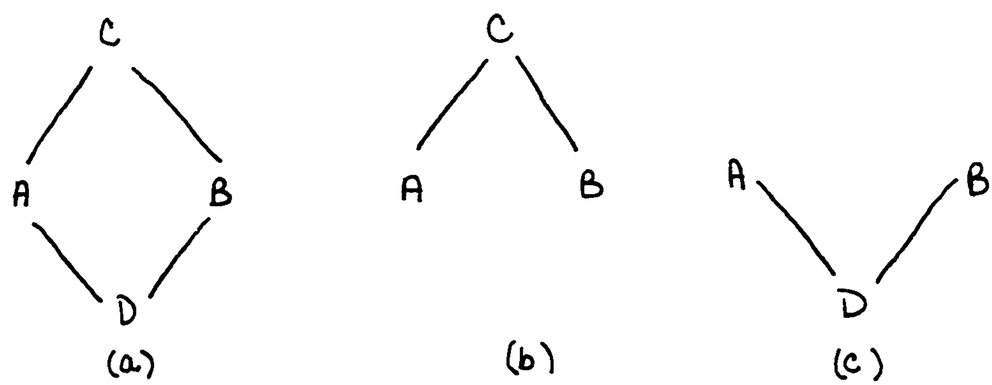

figure 1

According to Horwich and Papineau, a conjunctive fork comes in three varieties. A "normal fork" is a conjunctive fork formed by three causal phenomena[7], A, B, and C, such that the screener-off, or "vertex" of the fork, occurs prior to the pair of correlated phenomena, or "endpoints", it explains. Figure 1(b) shows a normal fork in which C screens-off the seeming probabilistic

[7]Whatever they turn out to be (e.g., events, facts, conditions, etc.).

dependence of A and B on each other and explains their frequent co-occurrence. Figure 1(c) represents an "inverse fork". An inverse fork is a conjunctive fork in which the vertex follows its endpoints. In this case D explains the frequent co-occurrence of A and B. According to the fork asymmetry, however, there are few (if any) inverse forks. That is, we rarely (if ever) find a pair of correlated causal phenomena such that some causal phenomena acts as a subsequent screener-off of their correlation, unless there is a corresponding prior screener-off. Figure 1(a) represents a conjunctive fork such that there is both a prior and subsequent screener-off. Horwich and Papineau refer to such forks as "closed".

Reichenbach, in contrast with many others who use the fork asymmetry to explain causal priority, establishes causal direction by defining *earlier than* in terms of macrostatistics: *In a conjunctive fork ACB which is open on one side, C is earlier than A or B.* (Reichenbach, 1956, p. 163) He does not take *earlier than* as primitive. Instead, Reichenbach defines the relation in terms of the conjunctive fork. According to Reichenbach, all open forks are open in the same direction. This direction, he stipulates, as "earlier than". His strategy involves answering the following two questions:

1) Why do all forks line up in the same way?

2) Given that all forks do line up in the same way, why is the vertex always earlier than the endpoints?

The reason all forks line up in the same way on Reichenbach's account is a matter of a contingent fact about the world. It just happens to be the case that all open forks point the same way. That open forks point towards the future (i.e., the vertex precedes the endpoints) is a matter of how we define "cause". That is, it arises from one of our practical assumptions about the causal relation which states that *causes precede their effects*. Since the vertex of the fork is a common cause that explains the correlation between the two endpoints or joint effects, we take the vertex to occur prior to the endpoints.

By defining the "earlier than" relation using conjunctive forks, Reichenbach explains temporal priority via causal priority. Further, he does not take the notion of causality as primitive. Instead he attempts to define "causality" probabilistically without using any temporal notions. To this end he argues that a causal connection exists between two event-types when they are highly correlated. Thus, a high correlation between event-types A and B suggests that either A causes B, B causes A, or A and B are effects of a common cause. In order to tell which of these relations holds between A and B, Reichenbach employs "the method of the mark". On this method if one event-type causes another then by intervening slightly to vary the frequency of the former the frequency of the latter occurrence will vary as well.

There are, nevertheless, problems with this method. Two standard objections are the following. First, in order to judge the specific causal relation between A and B we must interact with the causal process. But, this action itself presupposes causal (thus, temporal) order, although Reichenbach attempts to use the *de facto* irreversibility of certain causal processes to defend against this objection. Further, we have no means to determine the causal relationship between the marks of A and B. The marks themselves could be effects of a common cause. So, it is not obvious that defining "earlier than" in terms of the conjunctive fork is an advantage for Reichenbach.

It is also unclear just what the notion of "backward causation" amounts to if temporal order is defined using causal order. According to Reichenbach, given event-types A and B, if A causally influences B, then A precedes B. If, on the other hand, A and B causally influence each other, they are simultaneous. But, on this account it does not make sense to talk about *backward* causation. It cannot, by definition, occur. For this reason, Reichenbach's theory of causal priority does not meet the criteria above.

Even if Reichenbach's scheme were unproblematic, it is far from clear whether his version of the fork asymmetry principle will facilitate an explanation of causal priority that meets the criteria set out in Chapter One. Using Reichenbach's formulation of the principle can we explain causal priority in a way that a) does not rule out, in principle, backward causation or simultaneous causation and b) has explanatory power in the sense that it tells us how to distinguish between cause and effect independent of temporal order? It seems, at least initially, that we must rule out (a) (because backward causation is untenable). Nevertheless, there is one very attractive feature of his theory; he attempts to give a reduction of causation to probabilities without appeal to any temporal notion, thus, fulfilling (b) if successful.

This attractive feature of Reichenbach's theory uncovers the first criterion a suitable formulation of the fork asymmetry principle must meet. Any formulation apt to explain causal priority must have a component that accounts for causal direction in a way that does more than merely stipulate that the common cause occurs before the joint effect. Perhaps this is a matter of physics, but whatever the reason, it should be reflected in the statement of the principle.[8]

Elliot Sober uses Reichenbach's version of the fork asymmetry principle in examining the connection between the *principle of common cause* and

---

[8]Later I criticize David Lewis principle of the *asymmetry of overdetermination* on the grounds that the asymmetry it appeals to is no less dubious than causal priority. In claiming that one of the criteria for a satisfactory formulation of the fork asymmetry principle is that it addresses the issue of causal direction, I do not mean to take for granted that there can be a statement of the principle that meets this criterion. Huw Price, for example, argues that any attempt to use the fork asymmetry principle to explain causal priority either appeals to some asymmetry that is at least as confounding as causal priority or tacitly stipulates that the common cause occur before the joint effects. In Chapter 4 I defend what I take to be the strongest version of the fork asymmetry against these and other objections.

phylogenetic inference. According to Sober, inferring common ancestry is a case of postulating a common cause. He argues that *correlations always support a common cause hypothesis* no more than *positive instances always confirm a generalization.* That is, in some cases there is a high correlation between phenomena-types yet it makes no sense to postulate a com non cause. According to Sober, there is no way to assess the plausibility of the fork asymmetry except in the light of background assumptions. He argues, "We must say *when* a common cause explanation is preferable to an explanation in terms of separate causes."(Sober, 1988, p. 216)

I do not consider whether Sober's formulation of the fork asymmetry principle is useful in giving a non-Humean account of causal priority. I mention it here, nonetheless, because he observes the importance of determining *when* two correlated phenomena-types have a common cause and notes that sometimes they do not have a common cause. That is, he grasps the preferability of holding that *correlated phenomena-types usually have common causes* to holding that *correlated phenomena-types always have common causes*. His observation marks the second criterion I require for a fitting formulation of the fork asymmetry principle. Namely, an adequate formulation must provide a means of determining *when* there is a common cause or when it is reasonable to postulate one. In the literature authors tend to move back and forth between the claim that "correlated event types *typically* have common causes" and the claim that "correlated event types *always* have common causes." This ambiguity must be resolved before we can explain causal priority using the fork asymmetry.

In disambiguating the statement "correlated events have common causes" we must first identify three threads of its underlying idea that are significantly affected by whether we use "typically" or "must" in their wording. To be more precise, there are three ways a suitable formulation of the fork asymmetry principle should opt for "typically" instead of "must". I list them below, giving the weaker ('typically') interpretation in the statement with the stronger ('must') wording in parentheses. Given two highly correlated phenomena types, A and B,

1. A and B probably (must) have a common cause;
2. Given that A and B have a common cause, C; A and B rarely (never) co-occur without C, and;
3. C usually (always) occurs prior to A and B.

Consider (1). A high correlation between two phenomena types does not always indicate a common cause (or screener off) that explains their correlation. Nor do we always expect one. For example, bread prices in Britain have been increasing steadily over the last few centuries and so has the sea level in Venice. But, we concede no intuitive causal connection between bread prices in Britain and the sea level in Venice. (Sober, 1988, p216) As a result, we do not expect, nor are we disposed to accept, that the correlated phenomena types have a screener off that counts as a common cause; that is, there is no phenomena type that explains their correlation. I refer to such correlations as *accidental* since the correlation is a result of chance, rather than a result of causal connection.

A proper formulation of the fork asymmetry principle must allow for such accidental correlations. The fork asymmetry theorist will, admittedly, have to come up with some story about how we distinguish between 'genuine' and 'accidental' correlations. This is a separate issue. My present concern is the fact that any formulation of the rule she can hope to use in an explanation of causal priority must allow for correlated phenomena types that have no common cause. And, to make this allowance she must formulate the principle such that a pair of correlated phenomena types *probably* has a common cause.

Another reason a suitable formulation of the fork asymmetry principle should adopt the weaker statement of (1) is that for some extant correlations we expect a common cause explanation, though the fork asymmetry principle permits none. Consider, for instance, the following Bell experiment. We shoot two particles simultaneously from the same point in opposite directions. When we measure the spin of one particle we always find the other to have the opposite spin orientation. Thus, if particle 1 has a ↑ spin, our measurement of particle 2 will invariably show a ↓ spin. So, we find a perfect anti-correlation between the results of the spin measurements on the pair of particles. Further, if we interfere with the spin direction of one of the particles, we still find an opposite spin orientation for the other. Thus, despite our intuitive pull towards a common cause explanation for the correlation, there is not one. And, unless we allow instantaneous action at a distance, we can acknowledge no causal connection at all.

We may be able to deal with such cases by eliminating our standard requirement that causation is local (i.e., causes and their effects must occur close together in space and time). We have, nevertheless, an instance in which two highly correlated phenomena types have no common cause to explain their correlation. Unless we use the relaxed form of (1) above, the fork asymmetry principle will demand a common cause explanation for such cases.

In the case of (2) we need the weaker interpretation to permit probabilities other than 0 and 1. But we need not give up determinism. Papineau suggests that the Spirtes, Glymour, Scheines version of the fork asymmetry principle is best understood (given one added but reasonable assumption) as addressing causation under determinism in the following sense. A system is deterministic if we admit probabilities other than 1 and 0 only when we lack complete information. Once we acquire complete information, all probabilities are either 1 or 0.

Finally, we must ask whether a proper formulation of the fork asymmetry should state that C usually or always occurs *prior* to A and B. To allow for the possibility of backward and simultaneous causation, we should assert that causes usually occur prior to A and B. Whether there are any actual cases of backward or simultaneous causation is a contingent matter.

2:2: Paul Horwich s Formulation of the Fork Asymmetry

Paul Horwich also makes use of Reichenbach's principle of common cause to explain causal priority. In *Asymmetries in Time* he presents several non-equivalent formulations of the fork asymmetry principle. Some are stronger than others, but I charge that none of them are appropriate to drive the antisymmetry of causation.[9] Since Horwich gives several statements of the fork asymmetry, I find it useful to list his various (and sometimes nonequivalent) statements of the principle in an effort to clarify precisely what he takes the fork asymmetry to be. Consider Horwich's statements of the fork asymmetry:

> FA1) *Genuine coincidences rarely occur*; that is, given a strong correlation between events A and B, there is always some explanation -- some earlier event, C -- that causes them both...it is frequently not the case that correlated events A and B have a characteristic joint effect E. (Horwich 1987, p. 72-3)

> FA2) Highly correlated event types are invariably preceded by some unified common cause, but they need not have a joint effect. We never find a pattern...in which correlated events are linked only by having a joint effect. (Horwich 1987, p. 73)

[9]In Chapter Four I replace Horwich's version of the fork asymmetry with a what I take to be the strongest formulation and evaluate his theory in light of this substitution.

FA3) Given event types A, B and C: whenever A and B are correlated more than mere coincidence allows, there is a common cause, C, that screens-off their correlation. And, C almost always occurs when A and B co-occur. (Horwich 1987, p. 73-4).[10]

FA4) Whenever two event types are correlated with one another, they are embedded in a V-shaped chain of nomological determination, but need not be embedded in a Λ-shaped pattern; thus, there is always a characteristic antecedent event, but there need be no characteristic subsequent event. (Horwich 1987, p. 199)

FA5) The fork asymmetry consists in the truth of V-correlation in conjunction with the falsity of its time reverse... "V-correlation": that correlated event types are invariably associated with some characteristic *antecedent* event within a V-shaped pattern of nomological determination. (Horwich 1987, p. 201)

FA6) There are many cases of "normal forks" in which separated, simultaneous correlated events are associated with an earlier central event but not with any characteristic later event, but the time reverse of this pattern -- an inverse fork -- is never exemplified. (Horwich 1987, p. 142-3)

One confusing feature of FA1-FA6 is Horwich's tendency to switch back and forth between talk of event instances and event types. That is, strictly speaking, some of FA1-FA6 suffer from a type/token ambiguity. Since talk of correlation makes sense only with respect to types, we must express the fork asymmetry using types rather than tokens. Thus, the third criterion for a suitable version of the fork asymmetry requires that the formulation be free from any type-token ambiguity.

[10]In addition, Horwich requires in FA3 that C be a *unified* cause of A and of B, that A and B rarely co-occur without C, and that C precede A and B. (Horwich 1987, p. 73-4)

Most of Horwich's versions of the principle can be read as intending to refer to event types. He does, however, express the common cause relationship two ways: one using event-tokens and one using event-types. In terms of the common cause relationship between event-types Horwich claims, roughly, that given a pair of highly correlated event-types A and B, there is always an earlier event-type, C, (but need not be a later one, E,) that explains their correlation. On the level of event tokens, Horwich states that for most instances of A and B's co-occurrence there is an instance of C that causes them both and it is not the case that for most instances of A and B's co-occurrence there is an instance of E that causes them both. For clarity, I stick to expressing the fork asymmetry in terms of event-types. I read FA1-FA6 as being about event-types and throughout my discussion I replace "event(s)" with "event-type(s)".

Horwich's expressions of the fork asymmetry represent the following claim (and also two others) about the relationship between two correlated event-types and the event-type that explains them:

1. Given a pair of highly correlated event-types there is always some event-type that explains their correlation.

(1) appears consistently throughout FA1-FA6. The two other important components of Horwich's version of the fork asymmetry are given in a strong and weak form in different of his versions of the principle.

2a. There is *always* an *earlier* explaining event-type. (strong)

2b. There is *usually* an *earlier* explaining event-type. (weak)

3a. There *frequently is not* some later explaining event-type. (strong)

3b. There is *not always* some later explaining event-type. (weak)

FA1-5 contain 2a, while FA6 contains the weaker 2b. FA1, 3 and 6 include the stronger claim 3a, though FA2, 4 and 5 include 3b. According to Horwich, his most precise version is given in FA3 which I abstracted from the following argument he puts forth. Given event types A, B, and C:

Assumption 1: P(A) and P(B) are small.

Assumption 2: $P(A/B) >> P(A) \bullet P(B)$

Assumption 3: $P(A/C) >> P(A)$, $P(B/C) >> P(B)$

Assumption 4: $P(A \wedge B/\text{-}C) = P(A/\text{-}C) \bullet P(B/\text{-}C)$

Conclusion: $P(C/A \wedge B) \cong 1$

According to Horwich, if A and B meet assumptions (1) and (2), there is a common cause, C, such that the probability that C will occur given that A and B occur is nearly 1. Further, Horwich holds that C is a *unified* cause of A and of B, A and B rarely co-occur without C, and C precedes A and B. (Horwich 1987, p.

73-4) In order to require that C precede A and B, Horwich must appeal to 2a ("there is *always* an *earlier* explaining event-type") and a version of 3 that is at least as strong as 3a ("there *frequently is not* some later explaining event-type"). Since he endorses 3a in other of his formulations, I take him to be appealing to 3a in FA3 as well (rather than the even stronger claim that there never is some later explaining event type).

The driving force of the fork asymmetry, according to Horwich, is depicted in FA3, Assumption 4: $P(A \wedge B/\text{-}C) = P(A/\text{-}C) \bullet P(B/\text{-}C)$. Horwich claims that in those cases where A and B have a joint effect, E, we can substitute E for C in every equation of FA3 except Assumption 4. And, the reason we cannot make this substitution in Assumption 4 arises from the combination of 2a and 3a which established a predominant direction for causation that is towards the future.

FA3, though Horwich's strongest formulation, is unfit for use in an explanation of causal priority because it is excessively stringent. Horwich demands a common cause for every pair of highly correlated event-types. Many pairs of correlated event-types, nevertheless, have no earlier or later explaining event-type. That is, there is no third event-type correlated with both members of the pair that screens-off their seeming probabilistic dependence on each other *and* one member of the pair does not cause the other. Recall Sober's example -- the bread prices in Britain and water levels in Venice have both been steadily increasing over the last few centuries. But, according to Horwich, two highly correlated event-types *must* have a common cause. Thus, his most accurate formulation (FA3) cannot account for cases where there is merely an accidental (in the sense described earlier) correlation, albeit a high one, between a pair event-types.

Further, according to FA3, there is always an earlier explaining event and frequently not a later one. Thus, the earlier (more reliable) screener off has an explanatory power that the later (less reliable) screener off does not. Horwich argues that there is never a later screener off that occurs without a corresponding prior screener off (i.e., there are no free-standing inverse forks). Though this requirement (as Horwich freely admits) rules out backward causation, he justifies the move by arguing that the impossibility of backward causation is a contingent matter. Thus, backward causation is ruled out but not ruled out *in principle*.

2:3: Spirtes, Glymour and Scheines

Spirtes, Glymour and Scheines formalize the idea behind fork asymmetry-based accounts of causal priority. Their formulation arises from two axioms that show how probabilities relate to causal graphs: 1) the Causal Markov Condition and 2) the Faithfulness Condition. Rather than explaining why causes generally precede their effects, they characterize various assumptions we make about the antisymmetry of causation by formalizing the fork asymmetry idea. Though their formulation has a built-in asymmetry, it is not a temporal one. On their view we distinguish between cause and effect because they relate differently to the notion of probability (roughly, a frequency notion of probability). That is, there is a probabilistic relationship between the three highly correlated phenomena-types of a normal fork that does not exist between those of an inverse fork. Specifically, the vertex of a normal fork acts as a screener-off of the probabilistic relationship between the end points, whereas the vertex of an inverse fork does not. Thus, they indirectly demonstrate how we *distinguish* between cause and effect without building temporal direction into their formulation of fork as·'mmetry.

To see whether this strategy serves in an explanation of causal priority I now take a closer look at the connection between graphs and causal notions. First, Spirtes, Glymour and Scheines introduce the notion of "direct cause" (i.e., a cause which screens-off other, indirect, causes from the effect). Conditioning on the direct cause renders any indirect cause probabilistically independent of its effect. On a causal graph (i.e., a directed acyclic graph that represents a causal structure) a directed edge represents the relationship of direct causation between the two given vertices.[11] The authors then state that given a set of vertices {X,Y,Z} on a causal graph, G, Z is a common cause of X and Y if there is a directed path from Z to X

[11]It follows from transitivity and irreflexivity that these graphs are also acyclic.

in G that does not contain Y and a directed path from Z to Y in G that does not contain X. [12]

Spirtes, Glymour and Scheines indirectly explain causal distinguishability when they delineate the connection of these graphs to probabilities. The Causal Markov and the Faithfulness conditions match probability distributions and causal graphs that are compatible. The Causal Markov Condition sums up our intuitions about the connection by generalizing the fork asymmetry via conditional independence. A graph and probability distribution together satisfy the condition if for every vertex, V, in the graph, G, V is probabilistically independent of the vertices remaining in G other than its parents and descendants, once we hold fixed its parents. (SGS, p. 54) A probability distribution on G may satisfy the Causal Markov Condition yet include other conditional independence relations besides those entailed by the condition. But if, in fact, the only conditional independence relations of the probability distribution on G are entailed by the Causal Markov Condition, the graph and the probability distribution are said to be *faithful* to each other. (SGS, p. 56)

To identify all and only those conditional independence relations satisfying the Causal Markov Condition for a given causal graph and probability distribution that are faithful Spirtes, Glymour and Scheines introduce the notion of *d-*

[12]Note that these graphs cannot represent cases of interaction. I ignore this aspect for now.

separation. d-separation enables us to determine for *any* pair of vertices, X and Y, on a causal graph and *any* set of vertices on the graph not containing X and Y, whether or not X and Y are independent conditional on that set of vertices. (SGS, p. 71) Whereas, the Causal Markov Condition cannot serve this function alone since it provides only a sufficient condition of probabilistic independence (not a necessary one):

> "For a directed acyclic graph G, if X and Y are vertices in G, X ≠ Y, and W is a set of vertices in G not containing X or Y, then X and Y are d-separated given W in G if and only if there exists no undirected path U between X and Y, such that (i) every collider on U has a descendent in W and (ii) no other vertex on U is in W." (SGS, p. 74)

The faithfulness of a causal graph and probability distribution to each other gives rise to the connection between statistical and causal dependency via the notion of d-separation. Two vertices, x and y, on causal graph, g, are causally dependent on each other when and only when there is a probabilistic dependence relation between them. Thus, a particular kind of statistical dependency, conditional probabilistic dependency, identifies causal connection. A source on the path between two vertices that are causally connected picks out a common cause relation.

Spirtes, Glymour and Scheines' formalization of the idea behind the fork asymmetry principle (hereafter, SGS) is attractive to a fork asymmetry-based explanation of causal priority for a couple of reasons. Though they consider causes to be events, the connection between statistical dependency and causal connection demonstrated by causal graphs does not in any way depend on what things count as causes as long as the appropriate statistical relationship holds between them. This marks the fourth criterion any suitable formulation of the fork asymmetry principle must satisfy; that is, the formulation should not depend on what kinds of things count as causes.

Another advantage of SGS lies in its ability to give a more complete picture of causal relations than other versions in that it allows us to identify a unique causal structure for most sets of probabilistically related variables. In cases where we cannot identify a unique causal structure for a given set of probabilistically related variables, SGS suggests that there is always some possible wider set of probabilistically related variables for which we *can* identify a unique structure. And, SGS is attractive because it clearly does not characterize causal order using temporal order. Causal order is determined, instead, by the different ways in which cause and effect are related to the notion of probability.

SGS, despite its attractive features, violates the most important criterion set out above. Recall that SGS's strengths arise from combining the Causal Markov and Faithfulness conditions. But, the Causal Markov condition requires that every pair of correlated events has a screener-off, i.e., a vertex, V, is conditionally probabilistically independent of all vertices other than its parents, descendants and its parents' relatives. We cannot, in the case of SGS, relax the "must" to "typically" without giving up the Causal Markov condition. So, SGS is unfit for an explanation of causal priority.

2:4: Arntzenius

Frank Arntzenius (1992) divides the principle into two parts giving the following formulation[13]:

Fork Asymmetry: Version A (Arntzenius, 1990, p. 77) --

1. Correlated, simultaneous, non-directly causally connected A and B-type events imply C-type events such that the three types of events form a conjunctive fork.
2. Where a conjunctive fork exists the C-type event occurs always before the A and B-type events. Or, open conjunctive forks are always open to the future.

I contend that Version A is too weak because it restricts the fork relation to those cases in which the A-type events and the B-type events occur simultaneously, and the C-type events occur either before or after the A-type and B-type events. An

[13]Arntzenius formulates the fork asymmetry principle with the aim of showing that it will not work in an explanation of causal priority. I examine Arntzenius' objections to Version A and B in greater detail in section **4:1**. Presently I merely assess the viability of his formulation of the fork asymmetry principle.

account of causal priority should be able to handle cases where the effects of a common cause are not simultaneous. Arntzenius' next version is better.

Fork Asymmetry: Version B (Arntzenius, 1992, p. 227) –

1. If two events (or types of events, or facts, or conditions, or, ...) are correlated, and one does not cause the other, then there is a third event (type of event,...) such that the two events are probabilistically independent given the presence or absence of the third event. That is to say: for every pair of correlated events that do not have direct causal links there is a *screener-off* of that correlation.
2. This *screener off* occurs before the correlated events.

Version B gives us a rough idea of what the principle entails and helpfully separates its distinct aspects. This separation facilitates our judgment of whether the formulation meets my proposed criteria. It also aids in assessing how well the formulation explains causal distinguishability independently of how it works to explain causal direction. (1) describes how we distinguish between common cause and joint effect, while (2) states that causes usually occur before their effects. Some formulations may adequately account for one aspect of the principle, yet be unable to justify the other. SGS, for instance, does not address the issue of why the causal relation works predominately in one particular temporal direction.

Another beneficial element of Version B is that is does not depend essentially on what kinds of entities are causes. In examining a particular account of causal priority I employ the basic units the author describes. A viable account of causal priority, however, must be able to account for causation among facts, events, conditions, etc. Since I want an explanation of causal priority that is independent of what the causal relata are, I need a version of the fork asymmetry that does not depend on what kinds of entities are causes. Although Version A restricts to those cases where the correlated phenomena in question are simultaneous, Version B does not express this restriction. Thus, Version B is the stronger of Arntzenius' two formulations.

Notice, however, that Arntzenius' formulation of the fork asymmetry principle is unduly restrictive in two ways. (Since Arntzenius goes on to present several objections to the fork asymmetry principle, it is important to note whether his objections hinge on an inordinately narrow formulation of the principle.) On Version B (and, in fact, on both versions) correlated event-types, A and B, *imply* a common cause event-type, C. That is, for every pair of correlated phenomena types, A and B, there *must* be a C-type phenomena that explains their frequent co-occurrence. Thus, Version B is inordinately strong in that it violates criteria 3 above. For many pairs of highly correlated phenomena types there is no common cause as in Sober's example -- both the bread prices in Britain and the water levels in Venice have been steadily increasing over the last few centuries -- though there is presumably no causal connection between them. In such cases the co-occurrence of A and B, though frequent, is merely coincidental. Taking into consideration accidental correlations, it is not the case that every pair of highly correlated phenomena types has a common cause. Thus, a more workable formulation of the fork asymmetry principle should incorporate instead the weaker claim that given a correlation between A-type phenomena and B-type phenomena there is *typically* a C-type phenomena that explains the correlation.

Version B is also too strong in that B(2) suggests that the screener-off always occurs prior to the pair of correlated causal phenomena. To allow for the possibility of backward and simultaneous causation we should relax the 'always' in B(2) to a 'typically', as well. Additionally, Version B brings in the notion of *causal connection*, but does not give any criteria for determining when two correlated events are causally connected. So, despite its appealing qualities, Version B is problematic.

Though Version B is troublesome, I endorse a slightly diluted variant of this formulation as the one that best satisfies the criteria proposed above.

Fork Asymmetry: Version C –

1. * If two causal phenomena-types (e.g., events or types of events, or facts, or conditions, or, ...), A and B, are highly correlated and one does not

cause the other, there is *typically* a third phenomena-type, C, that screens-off their correlation and A and B rarely (though it is possible) co-occur without C.

2. * This screener off *typically* occurs prior to the correlated events.

We cannot weaken SGS to allow for the fact that there are some correlations with no screener off. Its dependence on the Causal Markov condition prevents us from making this move. Version B, on the other hand, can make use of this strategy more easily. We will still have to tell some story about how to distinguish between those correlations that do not have screeners off and those that do.[14] But, we need not give up an elemental power behind the fork asymmetry principle in the process. Thus, we arrive at B(1*). B(2*) is relaxed so as to allow for the possibility of backward and simultaneous causation, even if, as a matter of contingent fact, neither is ever instantiated. As a result of moving from 'always' to 'typically', Version C meets all criteria.

In sum, any formulation of the fork asymmetry principle that will be used to explain causal priority *must*:

1. have a component to account for causal direction in a way that does more than merely stipulate that the common cause occurs before the joint effects,

---

[14]I think we can draw this distinction by invoking certain pragmatic considerations.

2. provide a means of determining *when* there is a 'common cause' or when is it reasonable to postulate one (and in doing so it must suggest that correlated phenomena *typically*, rather than *must*, have a common cause),
3. avoid type-token ambiguities,
4. not depend on what kinds of entities count as causes, and it is made more appealing if it,
5. separates claims about causal distinguishability from claims about causal direction.

The weakened form of Arntzenius' principle satisfies these criteria. Thus, it is the best candidate for giving a fork-asymmetry based account of causal priority. It is this formulation I use to strengthen Horwich's theory of causal priority.

Chapter Three: Two Fork Asymmetry-Based Accounts of Causal Priority

Before considering objections to fork asymmetry-based accounts of causal priority and attempting to clarify and strengthen Horwich's theory of causal priority, I explain briefly why I think David Lewis' account will not work. He bases his theory on what he calls the "miracle asymmetry" which is, in turn, explained by the *asymmetry of overdetermination* (hereafter, AOD). Although AOD is not equivalent to the fork asymmetry, they are closely related. So, I find it worth discussing why his pseudo-fork asymmetry principle will not be useful in an explanation of causal priority. Assuming determinism for the actual world, @, Lewis formulates AOD:[15]

The past determines the future, but the future overdetermines the past.

Horwich restates it as follows:

> "for any event C there is, at every later time, events E1, E2,..., that each independently determine (given our laws of nature) that C occurred, but that the time reverse of this general fact does not obtain. That is, C may indeed be determined by earlier events, but not grossly overdetermined (there may be, at each earlier time, some event B1 that determines C's future occurrence, but no set of other events B2, B3,..., that also require C)."(Horwich 1987, p. 81)

Horwich contends that AOD contains the fork asymmetry as one of its components. (Horwich 1987, p. 168) The fork asymmetry, according to Horwich,

[15]Lewis does not commit himself to causation under determinism only.

states that correlated events have characteristic common causes but not always a characteristic common effect. (Horwich 1987, p. 174) He argues that AOD consists of the fork asymmetry principle plus the claim that every event is determined by, and may be inferred from, each of the effects on its own and the claim that causes (along with some other simultaneous facts) *determine* their effects. (Horwich 1987, p. 168)

Horwich's analysis of AOD is unclear, because AOD makes no reference to correlated event types at all. The fork asymmetry principle is clearly an assertion about correlations between event types and the relationship between AOD and event types is unclear. Though Horwich argues that the fork asymmetry is one component of AOD, I'm not sure just what else AOD contains. The methods of Lewis and Horwich, nevertheless, are similar in that they each attempt to establish causal asymmetry by appeal to an empirical fact about the world. So, I think it important to examine why AOD cannot generate the needed asymmetry to explain causal priority.

3:1: Lewis on the Asymmetry of Overdetermination

In *"Counterfactual Dependence and Time's Arrow,"* Lewis argues that a counterfactual theory of causation (hereafter, CTC) has three advantages over Law-based (Regularity) Accounts of Causation. First, it can distinguish between the relations *"is a cause of,"* and *"is an effect of."* Second, CTC can distinguish between the relations *"is a cause of,"* and *"are effects of a common cause."* Third, it can handle cases of pre-emption. Lewis' CTC accomplishes these three tasks by invoking AOD to establish counterfactual asymmetry and counterfactual asymmetry to establish causal asymmetry. He claims that the use of AOD and counterfactual dependence shows:

> I. The time asymmetry in causation can be explained by a time asymmetry in counterfactuals.
>
> II. The time asymmetry in counterfactuals can be explained independently of any time asymmetry of causation.

If II cannot be established I will be open to a suspicion of circularity.

His CTC has four components. First, an event c is a cause of an event e if and only if there is a chain of events from c to e, each event in this chain being causally dependent on its predecessor. Next Lewis argues that e causally depends on c if and only if, if c had not occurred e would not have occurred. Let A be read as "c occurred," and let B be read as "e occurred." Let @ be the actual world (where e causally depends on c). Then, (-A []-> -B) is true at @. According to Lewis, a counterfactual is true if and only if some accessible world in which the antecedent and consequent are both true is closer (*i.e., more similar*) to the actual world than any world in which the antecedent is true and the consequent is false.

Lewis then provides this prioritized list of the standards by which similarity of worlds is assessed [from most important to least important]:

1. Avoid big violations of law.

2. Maximize the spatio-temporal region throughout which perfect match of particular fact prevails.

3. Avoid small localized violations of law.

4. It is of little or no importance to secure approximate similarity of particular fact.

Finally, Lewis explains counterfactual asymmetry using AOD.

Let the following four lines represent four possible worlds whose similarity can be judged by Lewis' standards. The horizontal component represents time and the vertical component represents match of spatio-temporal region between the actual world @, and the other three worlds. Consider the following counterfactual: *If Nixon had pressed the button there would have been a nuclear holocaust.* Let B be a world which is like @ until t_n at which time Nixon presses the button. Let C be a world like @ until t_n at which time Nixon presses the button. At t_{n+1} a *reconvergence miracle* occurs which makes C like @ after t_{n+1} . Let D be a world which differs substantially from @ in particular fact until t_n when a *convergence miracle* occurs to make D like @ after t_n .

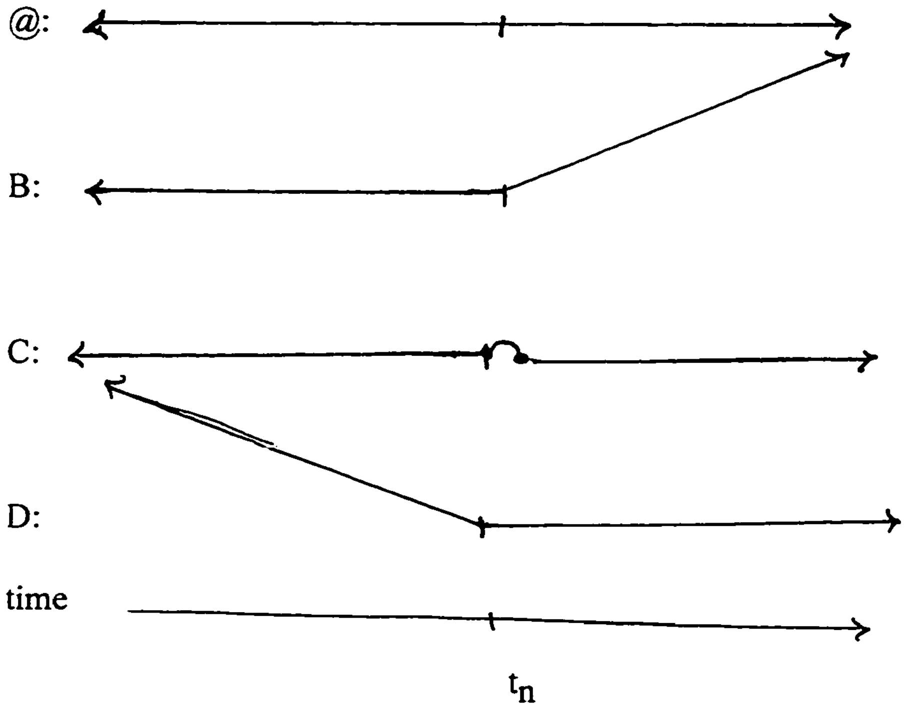

According to Lewis, counterfactuals are future-directed because the most similar possible worlds are those whose pasts are similar and whose futures are not. In the above case, Lewis insists, the most similar world to @ is B. But what makes B more similar to @ than is C or D?

C differs from B and D in the number of miracles occurring in it. Both a *divergence miracle* and a *reconvergence miracle* occur in C. Only one miracle occurs in D: a *convergence miracle.* And, only one miracle occurs in B: a *divergence miracle.* This discussion will focus on the similarity between worlds @ and D, and between @ and B. I will not discuss the similarity between @ and C. The crucial question then becomes, "what makes B closer to @ than D?"

Suppose the extent of the spatio-temporal region in which @ and B match in particular fact is the same as the extent of the spatio-temporal region in which @ and D match in particular fact. According to Lewis, B is more similar to @ because there would only be a small violation of law necessary in @ to make the course of @ diverge in such a way that @ would match B in particular fact throughout their entire spatio-temporal region (not including the small transition period required for the miracle to take place). In other words, divergence miracles can be small. On the other hand, for @'s history to differ so that it would match world D in particular fact many small violations of law would have to take place: one for each effect of each event occurring before t_n . These small violations of law add up to one huge miracle. That is, convergence miracles must be huge.

What reason do we have to believe that divergence miracles can be small, while convergence miracles must be huge? According to Lewis, the asymmetry of overdetermination explains this asymmetry of miracles. (Lewis 1986, p. 50) Lewis describes a determinant:

> "a minimal set of conditions jointly sufficient, given the laws of nature, for the fact in question. (Members of such a set may be causes of the fact, or traces of it, or neither.) The fact may have only one determinant at a given time...Or it may have two or more essentially different determinants at a given time, each sufficient by itself. If so, it is overdetermined at that time."

If we restate AOD substituting the definition of *determines* for the actual term, we get the following statement of AOD[16]:

i) "The past determines the future...": Given some event, A, there is at least one, and usually only one, minimal set of conditions (at a given earlier time) which is sufficient given the laws of nature for the fact that A occurred. And,

ii) "but the future overdetermines the past.": Given some event A, and some time, t, later than the time of A, there are at least two minimal sets of conditions occurring at t that are sufficient given the laws of nature for the fact that A occurred.

According to Lewis, it takes one small miracle to break the link between any determinant and that which it determines. (Lewis 1986, p. 50) Since the past *determines* the future, only one link need be broken between an event in @ and the event it determines for @ to diverge from its present course. Thus, *only* one small miracle is needed to divert @'s future from its present course.

[16]Notice that Lewis (*Asymmetry of Overdetermination* p49 in *Counterfactual Dependence and Time's Arrow)* talks of the determination/overdetermination of *facts*. Also, he states "Whatever goes on leaves widespread and varied traces at future times," which seems to suggest that he is talking about events. Lewis is obviously using 'fact' and 'event' interchangeably in this section. I'm not sure what the relationship between event-causation and fact-causation is for Lewis, but I take Lewis to mean by "A fact, A, was determined by a fact, B," that "The fact that A occurred determined the fact that B occurred," where A and B are events.

However, in order to change @'s past we must break the links between many events at t_n and the events they determine before t_n, because the future *overdetermines* the past. Thus, several small miracles must take place which together make up one huge miracle required to change @'s past.

Lewis' entire picture looks like this: AOD shows that divergence miracles may be small, while convergence miracles must be huge. So, with respect to -A [] -> -B the nearest (i.e., most similar) possible world, w_1, where -A is true at time, t, will be a world whose past is like @ but whose future is not. In other words, if A had not happened, something in the future (since @ and w_1 are identical until just before t), namely B, would not have happened either. So, A will typically occur earlier than B. Thus counterfactuals are future-directed. Because a causal statement is a type of counterfactual statement, causal statements are also future-directed.

I charge that there is no compelling reason to accept Lewis' miracle asymmetry. Whether any kind of miracle (i.e., divergence or convergence) is big or small in either the past or future direction depends on whether we want to:

1) change @'s course (past or future) at all, or

2) change @'s past (future) to match the past (future) of a particular world.

(1) should be small in either direction. That is, if we choose the right deterministic link we may get away with breaking only one link to shift @'s past or future from

its current path. Thus, both convergence and divergence niracles may be small. (2), on the other hand, is likely to be big in either direction. To change @'s past or future to match that of a particular other possible world we must break many deterministic links. Thus, both divergence and convergence miracles may be huge.

Lewis is correct in arguing that it takes only one small miracle to break the link between any determinant and that which it determines. He is also correct in claiming that for @ to diverge from its present course (or, to change the future of @) we can, in some cases, get away with breaking only one link between some event and another event it determines. But, his mistake lies in a hidden requirement for changing @'s past. To be more precise, changing @'s past requires forcing a match with the past of a particular other world. Whereas, in the future direction he requires merely that we divert @'s course so that its future is other than what it would have been with no interference from us.

In the future direction, breaking one deterministic link will divert @ from its present course but will not necessarily put it onto B's course. But, if we want to change's @'s future so that it matches *B's* future in particular, rather than just any world, we may need to break many links (just as when we try to make @'s past match D's).[17] Analogously, if we want to change @'s past we need only break

[17]I assume the laws of physics are symmetric.

one link to send it on a different course, but may need to break many links to make its past match the past of some particular world, e.g., *D's*.

Without the miracle asymmetry Lewis' standards of similarity are not useful in determining the closeness of possible worlds. We no longer have reason to think that avoiding big violations of law is top priority. We must, according to these standards, say that B and D are equally close to @. But, this leaves Lewis' account open to the future similarity objection; all talk of similarity between possible worlds becomes nonsensical. Since we have no reason to think violations of law are larger in one temporal direction than the other, we can no longer judge the similarity of possible worlds by the size of the miracles needed to bring them into sync. Further, the extent to which two worlds match in particular fact no longer tells us anything about the similarity of worlds. In fact, any world like @ at one instant is like @ at all instants. So, the only world similar to @ is @.

As for AOD -- I see no compelling reason to accept an asymmetry in this case either. But, I leave this for now.

3:2: Horwich's Account of Causal Priority

Horwich asserts that the problem of explaining causal priority stems from a tension between the following two observations:

1) Causation appears to be a type of determination,

2) And, all types of determination we might plausibly identify with causation seem to be time symmetric.

To resolve this tension, he argues, causation should be identified with determination plus some further factor that explains why causes precede their effects. Thus, causation = determination + _____. We can follow Hume and fill in the blank with the stipulation that the cause occurs first. Or, we can fill in the blank with a contingent fact about the world that restricts the order of cause and effect. In the case of the latter, more "substantive", approach our concept of causation explains causal priority only in conjunction with the specified contingent fact(s) about the world.

Accordingly, we must answer two distinct questions in giving an explanation of causal priority:

1) Why do causes typically precede their effects?

2) On what basis do we maintain that causes typically precede their effects?

Hume answers (1) with "this future orientation is a constituent of the causal relation" and (2) with "we know this *a priori*". The more substantive strategy answers (1) with "there is a contingent fact about the world that restricts the direction of the causal relation" and (2) with "we know this *a posteriori* from the aforementioned contingent fact about the world".

Horwich combines the two strategies, answering question (1) with "this future orientation is a constituent of the causal relation" and (2) with "we know this *a posteriori* from a cluster of contingent facts about the world". He describes his "Neo-Humean" approach as analogous in status to giving the chemical analysis of water. In other words, we describe the underlying structure of a phenomenon whose familiar symptoms merely provide a moderately reliable guide to when the stuff is present. (Horwich, 1987, p.) His theory is based in part, though not solely, on the fork asymmetry principle. Horwich adopts a Quinean picture of explanation in arguing that "causes typically precede their effects" is explained by the following non-harmonious contingent facts about the world:

1. Correlated events are causally connected.[18]
2. Choices are made for the sake of what they might cause.
3. The causes of a present event are knowable, but its effects are not.

He explains:

> "A better way to understand the role of manipulability [for example], in my opinion, is to regard it as providing one of the cluster of crude, fallible reference-fixing principles that help us to zero in on the causal relation. Thus we might roughly identify causation as "that relation believed to hold between rational choices and the desired events for the sake of which those choices are made". This maxim picks out a *future*-oriented relation." (Horwich, 1987, p. 140)

---

[18]A and B are causally connected iff either A causes B or B causes A or A and B are effects of a common cause.

So, these contingent facts about the world are merely approximate truths, though they play a significant role in establishing our belief that causes precede their effects.[19]

Horwich's strategy is not entirely clear. The key to shedding light on his theory involves taking a closer look at his water analogy. Horwich assumes the following, Neo-Humean, account of causation which he refers to as the *predetermination chain theory:*

> "A direct cause of some effect is an essential part of an antecedent condition whose intrinsic description entails, via basic laws of nature, that the effect will occur. And causation in general involves a chain of direct causation. That is to say, C causes E if and only if there are some events, e1, e2,...,eN, such that C directly causes e1, e1 directly causes e2,..., and eN directly causes E." (Horwich, 1987, p. 133)

He argues that water, like causation:

> "...is not constituted and identified by means of a set of analytically necessary and sufficient observable properties. Rather, it is recognized by fallible criteria, such as 'colorless, tasteless liquid' and 'constituent of rain'. Its nature as H_2O is discovered afterward, and is then used to correct the original principles of identification. Similarly in the case of causation we should allow that its structure, as given by the predetermination-chain theory, may conflict with the symptoms by which we recognize instances of causation...[T]hese symptoms need be neither essential, universal, nor permanent." (Horwich, 1987, p. 140-1)

---

[19]In fact, as I discussed earlier in considering Horwich's versions of the fork asymmetry principle, he maintains that causes *always* precede their effects. This, I warrant, is a major weakness in his theory. But, as we shall see shortly when I deal with Arntzenius' and Price's objections, substituting my weakened form of Arntzenius' version of the principle eliminates this flaw.

Horwich draws a distinction between the *means by which we identify instances* of a thing and a *description* of that thing's nature. H_2O represents just what water is (its nature). But, our practical identifiers of water (i.e., colorless, tasteless, liquid) do not always pick out a substance with the structure H_2O. Also, a body of H_2O is not always colorless, tasteless, and liquid. In this way, we give a metaphysical reduction of water (water > H_2O) and we identify water using certain facts known *a posteriori*. So, Horwich (like Papineau) advocates a metaphysical, rather than a conceptual, reduction of the causal relation. As a result, he need only provide *a posteriori* evidence that these contingent facts pick out instances of causation. Also, because he gives a metaphysical reduction instead of a conceptual one we can give a definition that provides the essence of causation even though it may not capture all of its features. Just as we use certain contingent facts about the world to pick out instances of H_2O, we use others to pick out instances of causation. According to Horwich, these contingent and fallible identifiers of causation amount to:

Causation is the relation believed to hold between

a. rational choices and the events for which those choices are made,

b. correlated phenomena-types, and

c. knowable and unknowable phenomena.

We do not necessarily have an instance of causation whenever one (or all) of a-c holds. Conversely, the occurrence of two causally connected phenomena- types does not entail one (or all) of a-c. According to Horwich, causation reduces to chains of direct determination. The time order of causation, he contends, is a feature of the relation that is built-in via specific restrictions placed on the method by which we map causal chains:

> "We should not impose the time-order requirement on each basic link in a causal chain. Instead, we should consider *entire* chains of determination (stretching between the distant past and the distant future) and then impose on the elements of any such chain a causal interpretation subject to the following pair of global constraints: (1) require that causes precede their non-simultaneous effects, and (2) maximize causal continuity (so that causal priority in one part of the chain may be 'smoothly' extended to adjacent parts)..." (Horwich, 1987, p. 136-7)

Thus, he stipulates time order to determine the usual causal order. But, he does so in a way that does not rule out cases of simultaneous causation or epiphenomena.

He gives up backward causation, however, on the grounds that:

> "...as far as we know, there is no such thing. One might well have objected to our theory, were it put forward as an a priori definition. For no doubt backward causation is conceivable. But since the theory purports merely to describe a posteriori the actual nature of causation, it is not threatened by the fact that under different evidential conditions we would be inclined to take a different view of the matter." (Horwich, 1987, p. 137)

Since Horwich gives a metaphysical, rather than conceptual, definition of causation, the definition need not be an *a priori* one. Instead, an *a posteriori* description is acceptable and preferred. And, using this description, he can deny the actuality of backward causation without denying its conceivability.

3:3: Objections to the Fork Asymmetry Strategy

Frank Arntzenius presents several objections to the fork asymmetry principle (a.k.a., the principle of common cause). Specifically, he objects to the principle as used by Horwich (1987) (and Papineau (1992)) to explain causal priority. He contends that the fork asymmetry does not hold in either deterministic or indeterministic systems. In deterministic systems, Arntzenius argues, the fork asymmetry principle is symmetric in the absence of any "naturalness restriction" to generate asymmetry. Thus, it cannot be used to explain causal priority.

Arntzenius (1990) raises some problems with the idea behind accounts of causal priority like Horwich's, i.e., that causal priority is generated in a deterministic system from the fact that there are no free-standing inverse forks. He contends that because such theories assume determinism and place no naturalness restriction on the fork relation, whenever two phenomena-types are correlated there *must* exist a prior screener-off of the correlation.[20] But, also, there *must* exist a subsequent screener-off. As a result, assuming determinism in the absence of any naturalness restriction on the fork relation implies that all conjunctive forks

[20]Horwich does require that the common cause of two highly correlated causal phenomena be "unified". Though this might be taken as a candidate for some naturalness restriction, I argue shortly that it cannot restore asymmetry to the fork relation in light of Arntzenius' objection.

are closed. Under these circumstances, there are no free-standing forks, inverse or normal. So, such theories cannot appeal to the fork asymmetry principle as presented by Horwich (or Papineau) to generate the asymmetry needed to explain causal priority:

> "Events that determine each other will have conditional probability 1 upon each other, and hence have the same probabilistic relations with all other events. It therefore follows from determinism that for any screener off that occurs before some pair of correlated events, there is a screener off that occurs after the correlated events, namely the event that occurs at that later time, iff the screener off occurs at the earlier times." (Arntzenius 1992, p. 229)

Though Arntzenius appears to phrase his argument in terms of event *tokens*, by "event(s)" he means "events or types of event, or facts, or conditions, etc." (Arntzenius 1992, p. 228) Whatever causal phenomena turn out to be, Arntzenius claims that the relation between causal phenomena is such that if one determines another the two phenomena each have a conditional probability of 1 upon the other. Since the notion of a screener-off makes sense only with respect to phenomena types, we should read the remainder of the quote as referring to causal phenomena types. So, at the type level, according to Arntzenius, there exists both a prior and subsequent screener-off for every pair of correlated phenomena types.

According to Arntzenius, Horwich (and also Papineau) presents a paradigm theory utilizing Version B of the fork asymmetry principle. He attempts to support the fork asymmetry principle by assuming *background determinism* and "some randomness condition on the distribution of initial conditions". (Arntzenius 1990, p. 80) Consider the example above in which smoking raises the probability that the smoker will develop both lung cancer and heart disease:

> "...to assume background determinism is to assume that for any particular person, whether smoker or not, there is a factor c, the presence or absence of which determines whether or not he will later develop cancer, and that for any person there is a factor y, the presence or absence of which determines whether the person will develop [heart disease] or not. The claim that persons who smoke are more likely to develop cancer then translates into the claim that for persons who smoke factor c is more likely to be present than for persons who do not smoke." (Arntzenius 1990, p. 80)

This arises from his characterization of a deterministic system as one for which the complete state of the system at any given time determines the complete state of the system at any other time. (Arntzenius 1992, p.228) Each of the factors y and c may represent some highly non-intrinsic feature of the person. For instance, each factor may depict a time slice of the person that includes a collection of highly detailed information unrelated to the person as such, like information about the motion of particles in a distant galaxy.

Horwich's added assumption that the initial conditions of a system have a random distribution suggests that y and c are statistically independent among smokers and among non-smokers. But, the randomness of initial conditions in conjunction with background determinism implies the later correlation between heart disease and lung cancer among all persons. That is, the presence or absence of y and c among all persons is random. But, the presence of c (y) determines that one will later develop cancer (heart disease). As a result, though c and y may each represent some highly non-intrinsic feature of the person, a correlation between c and y indicates a later correlation between cancer and heart disease, though their initial distribution among persons is random. Thus, Horwich justifies Version B(1).

In order to justify Version B(2), according to Arntzenius, he must, in addition, show that there are no inverse conjunctive forks. To eliminate inverse forks Horwich appeals to statistical mechanics. He argues that while the distribution of the initial conditions of a system are random, the distribution of final conditions are not. This suggests, in line with the Spirtes, Glymour, Scheines formalization of the fork asymmetry principle, that a common effect cannot have the same probabilistic relationship to its joint causes that a common cause has to its joint effects. While a common cause screens-off the correlation between its joint effects, a common effect does not screen-off the correlation between its joint causes. Thus, there are prior screeners but no subsequent ones.

Consider the following to illustrate. The smoking example above describes a normal fork. Factors y and c (that determine the person will get heart disease and cancer, respectively) are statistically independent among smokers and among non-smokers. There is no analogous statistical independence between such factors in the case of inverse forks. Suppose that lung cancer and heart disease each separately (i.e., they do not interact) raises the probability that a person will die. Assume as above factors y and c for any particular person, dead or alive. Factors y and c are not statistically independent among dead persons and among living persons. That is, death does not screen-off the correlation between heart disease and cancer.

In contrast, I think Arntzenius correct to maintain that without some naturalness restriction on the fork relation there will certainly be some later occurring phenomena-type that screens-off the correlation between lung cancer and heart disease:

> "Of course the event that occurs if and only if the screening off event occurs at the earlier time will in general be a very complex and unnatural event, it will be some weird condition on the distribution of elementary particles throughout some large part of the universe. But, by determinism, we know that there is some such condition which occurs after the correlated events and screens them off, if there is a prior screener off." (Arntzenius 1992, p229)

Even though this subsequent screener-off is complicated and unnatural, e.g., it may include conditions on the motion of particles in a distant galaxy, it will exist, nonetheless, for every pair of correlated phenomena-types. Thus, on Horwich's account every pair of correlated phenomena-types must have both a prior *and* subsequent screener of the correlation.

Horwich does require that if C is a common cause of A and B, C must be a *unified* cause of A and B. (Horwich 1987, p. 73) This notion of "unifiedness" initially appears to be a naturalness restriction. Horwich contends, C must be "unified in the sense that it cannot be split into two parts such that one causes A and the other B." (Horwich 1987, p. 73) But, this restriction does nothing to alleviate Arntzenius' objection. Arntzenius could easily respond that the subsequent screener-off accompanying every prior screener in a deterministic system is complicated and unnatural, yet unified in Horwich's sense.

Further, Arntzenius contends, rightly, that in a deterministic system it makes no sense to assign a dissimilar status to initial and final conditions with respect to the uniformity of their distribution. They should both have the same status. In other words, if the initial conditions of a system exhibit a random distribution, so will the final ones. According to Arntzenius, "At any time there are exactly as many correlated properties as at any other time, in the sense that for any two times one has a 1-1 mapping between properties which preserves all correlations." (Arntzenius 1990, p. 83) Thus, one cannot claim the distribution of initial conditions is random, while the distribution of some final conditions is not. So, in a deterministic system, whenever there is a prior screener-off there is also a later one.

There are two strands the fork asymmetry theorist needs to tease out of the debate in order to address these problems:

1. What version of the principle is criticized?

2. Can we draw an effective distinction between natural and non-natural screeners-off? (Such a restriction on "naturalness" is effective insofar as it restores asymmetry to the fork relation.)

Arntzenius objects to Version B of the fork asymmetry principle. Recall Version B:

1. If two events (or types of events, or facts, or conditions, or, ...) are correlated, and one does not cause the other, then there is a third event (type of event,...) such that the two events are probabilistically independent given the presence or absence of the third event. That is to say: for every pair of correlated events that do not have direct causal links there is a *screener-off* of that correlation, and,
2. This *screener off* occurs before the correlated events.

(Arntzenius, 1992, p. 227)

But, I argued earlier that a relaxed Version B (my Version C) is better able to generate the asymmetry needed to explain causal priority. Recall Version C,

1. * If two causal phenomena-types (e.g., events or types of events, or facts, or conditions, or, ...), A and B, are highly correlated and one does not cause the other, there is *typically* a third phenomena-type, C, that screens-off their correlation and A and B rarely (though it is possible) co-occur without C.

2. * This screener off *typically* occurs prior to the correlated events.

I consider how Arntzenius' objection fares against Version C.

It is not clear whether we can find a viable means of distinguishing between natural and non-natural screeners-off or whether this strategy is effective in restoring asymmetry to the fork relation. Since the common cause relation occurs across different classes of phenomena (e.g., there are microscopic screeners-off of macroscopic phenomena and vice versa) we must be careful. According to Arntzenius, the fork asymmetry principle cannot hold for a class of causal phenomena that has causes outside that class. (Arntzenius 1992, p. 230) He asserts that because microscopic phenomena have macroscopic consequences we cannot restrict what counts as "natural" common causes to macroscopic phenomena.

Consider his example. Cleopatra has a poison so strong that a person will die upon ingesting one molecule of it. She puts one molecule of the poison in each of 100 glasses of wine and orders 100 slaves to drink ½ a glass of wine. If the molecule of poison is consumed, death is preceded by a reddening of the left hand and of the right hand. Thus, consuming the molecule is a prior screener-off to the reddening of the left and right hands. Death is a subsequent screener-off (assuming death occurs in exactly those cases that the poison is ingested). If we restrict natural screeners-off to macroscopic events, the correlated phenomena-types (the reddening of the left and right hands) will have a posterior (but no prior) screener-off. Thus, we might restrict natural screeners-off to the class of microscopic phenomena (since macroscopic phenomena are reducible to microscopic phenomena). But, we cannot restrict them to macroscopic phenomena. And, if the laws of microscopic physics are deterministic, according to Arntzenius, restricting common causes to microscopic phenomena will not generate the asymmetry needed for the fork relation.[21] (Arntzenius 1992, 230)

Perhaps we should require instead that natural screeners-off be both simple and localized. What counts as simple and what counts as localized would each depend on certain pragmatic considerations. Also, what counts as simple or localized might vary with respect to context. Using this criteria we can rule out screeners-off that depend on information about things like the motion of particles in a distant galaxy (for this violates the requirement that the common cause be localized). We can also rule out screeners-off described by an entire time slice of the light cone generated by the correlated phenomena types.

Arntzenius objects to Version B of the fork asymmetry principle. Recall that Version B is inadequate because it implies that *for any* pair of highly

[21]Though drawing a distinction between macro and microscopic phenomena will not help us get a handle on the difference between natural and nonnatural screeners-off, it is interesting in itself because causation does occur across these different classes of phenomena. Thus, we should consider the relationship between the two classes as well as the relationship between this distinction and the distinction between singular and general causal claims.

correlated phenomena-types, A and B, there is *always* a phenomena-type, C, that explains their frequent co-occurrence. This claim is inordinately strong. The version of the principle I accept, Version C, does not suffer from this complaint. By replacing Version B with Version C, we may be able to give up determination for something weaker, in which case we need not require every pair of correlated phenomena-types to have a screener-off, much less for every pair to have both a prior and subsequent screener-off.

If Horwich accepts my weaker construal of Arntzenius' fork asymmetry principle, he can allow that correlated phenomena types *typically*, rather than *must*, have a common cause. In this case, the correlation of initial properties, factors y and c, does not necessarily mean the correlation of the final properties, heart disease and lung cancer. Nor does the correlation of factors s and cm imply the correlation of smoking and coal mining. Thus, he need not concede that all conjunctive forks are closed. That is, a future directed conjunctive fork need not have a corresponding inverse fork.

Huw Price (1992) argues that fork asymmetry-based accounts of causal priority invariably fail because there are *too few forks*. That is, there is not enough actual fork-type asymmetry in the world to let us discriminate between cause and effect. The problem, according to Price, appears at both the microscopic and macroscopic physical levels. His argument is, roughly, as follows.

1. Microphysical processes are symmetric in time (in the sense that any process occurring with one temporal orientation also occurs with the opposite orientation).

2. The fork asymmetry depends on the coordinated behavior of huge numbers of microscopic processes; it is a macroscopic asymmetry.

3. At the macroscopic level common causes are too infrequent, insignificant, or both to give rise to *actual* correlations between their joint effects.

4. Thus, the fork asymmetry cannot provide enough asymmetry to allow us to distinguish between cause and effect.

I accept (1) and (2), though I do not think they create problems for the fork asymmetry strategy as Price claims. According to Price, we have conflicting intuitions governing microscopic physical processes. On the one hand, we want to distinguish between cause and effect at the microscopic level in a substantive manner. That is, we hold (contra Hume) that microscopic causal phenomena depend on certain earlier microscopic causal phenomenal in a way that they do not depend on later ones. On the other hand, we charge that all microphysical processes occur with both temporal orientations, i.e., microscopic processes are symmetric in time. In the case of the temporal asymmetry of radiation, for example, Price argues that the asymmetry arises from a macroscopic asymmetry between "sources" and "sinks". But, this asymmetry disappears at the microscopic level:

> "It makes no difference at all whether we think of the transmitters as tiny sources of outgoing waves or the absorbers as tiny sinks of incoming waves. The two mathematical representations are equivalent."

Thus, Price contends, the fork asymmetry cannot explain causal priority at the microphysical level. The fork asymmetry disappears at the microphysical level as a signal fades when the power of the wave transmitter is steadily reduced. That is, it becomes difficult to distinguish the signal from the background noise. (Price 1992, p. 258)

I agree that the fork asymmetry disappears at the microphysical level. But, I do not think this is a problem for the fork asymmetry strategist. I contend, however, that Price is wrong in thinking that our intuitions demand an asymmetry of causation at the microphysical level. Macroscopic causation is uncontroversially reducible to microscopic causation. Most macroscopic causal phenomena can be described by a set of microscopic states. The relevant feature of the relationship is that macroscopic phenomena, because they are directly observable and reduce to microscopic states, are our only access to microscopic phenomena. That is, we can know about microscopic phenomena only because sometimes microscopic phenomena have effects on macroscopic phenomena as in the example above. Since, our main contact with microphysical process is via directly observable macrophysical processes, it is only when these microphysical processes occur in large enough numbers to produce macroscopic consequences that they concern us.

Price appears to confuse the distinction between macro and microphysical processes with that of singular and general causal claims:

> "[The fork asymmetry] depends on the ordered alignment of vast numbers of microscopic events, and hence simply isn't the sort of feature of the world which can be manifest when the numbers involved are too small. Again, sources need to stand out against the [background] noise, and this requires the cooperation of many individual events." (Price 1992, p. 258)

It is misleading to represent macrophysical processes as requiring the cooperation or alignment of many individual events. It is the fact that macro processes arise from *microphysical* processes (whether they be types or tokens) that is key.

We might think of the relationship between singular and general causal claims as being such that a general causal claim is a generalization over causal instances.[22] This is similar to the relationship between macro and microphysical processes in that we think of macro processes as reducible to a set of microscopic states. It is important to keep in mind that, though both distinctions affect

[22]I argue in Part II that we have good reason to prefer this analysis of the relationship to other strategies.

explanations of causal priority, they are separate. Thus, there are at least four types of causation depicted in the table below.

	General Causation	Singular Causation
Macroscopic Causation	*Macroscopic General Causation* e.g., Smoking causes lung cancer. **A**	*Macroscopic Singular Causation* e.g., Cindy's smoking causes her lung cancer. **B**
Microscopic Causation	*Microscopic General Causation* e.g., When two particles collide the resulting momentum of each is determined by the momentum of the other. **C**	*Microscopic Singular Causation* e.g., Particle a's collision with particle b causes particle b to change direction. **D**

Smoking (Cindy's smoking) and lung cancer (Cindy's lung cancer) are macroscopic phenomena. But, particles are microscopic phenomena that are not directly observable. In making claims about causation we are primarily concerned with macroscopic connections. That is, we primarily want to be able to answer questions like "Does smoking cause lung cancer?" or "Did Cindy's smoking cause her lung cancer?" One might argue that we are concerned with the motion of particles, microphysical or not. But, we are concerned with the macroscopic results of such motions. That is, we are concerned with microphysical processes when they produce macroscopic consequences. Otherwise, micro phenomena need not concern us. So, the fork asymmetry theorist need not worry about the causal symmetry of microphysics.

Although I accept (1) and (2), I take issue with Price's claim that macroscopic common causes are too infrequent or insignificant to give rise to actual correlations between their joint effects. (3), if true, is a problem for the fork asymmetry theorist. If there is not enough macroscopic asymmetry in the world to establish the fork asymmetry, the fork asymmetry cannot drive causal priority. Price calls this the problem of "too few forks". I argue that upon clarifying Price's objection, the problem of "too few forks" amounts to the claim that the fork asymmetry strategist has trouble explaining causal priority for causal instances. I agree, but rather than reject the fork asymmetry strategy on these grounds, I explore how it might handle this problem in **Chapter Six**.

Price illustrates the problem by example. Consider that fire produces both heat and smoke:

> "It is no doubt true that there is actually a significant correlation between heat and smoke in the world, but the fact that fire causes heat and smoke surely does not depend on this being so. Had there only ever been one fire in the history of the universe (and lots of uncorrelated heat and smoke due to other causes) it would still have been true that fire cause heat and smoke. Moreover, once we see this we see that there must be many common causes which are simply too infrequent to give rise to *actual* correlations between their joint effects." (Price 1992, p. 257)

In other words, a common cause may be too infrequent to produce a correlation between its joint effects. As I argued earlier, the fork asymmetry relies on these correlations to distinguish between cause and effect and to explain why causes typically precede their effects. It suggests that common causes are usually screeners-off, while common effects rarely are. Because the fork asymmetry is cashed out in terms of "screeners-off" and "correlations" it makes sense only at the type level. Thus, if many common causes are too infrequent to produce correlations between their joint effects, it is hard to see how the fork asymmetry can be successful in providing an account of causal priority.

Price acknowledges that one obvious way to deal with the problem is to maintain that common causes *tend to* (rather than *actually*) produce correlations between their effects. This move, he insists, is problematic because it appeals to modal claims. Modal claims contain asymmetries of their own. Thus, this move creates two possible pitfalls for the fork asymmetry theorist. If the asymmetry of modal claims is analytic, she appeals to *disguised conventionalism.*[23] Since she claims to provide an alternative to this Humean strategy, she should avoid stipulating causal direction as temporal direction, either directly or indirectly. A more substantive explanation of the modal asymmetry, however, is likely to involve temporal asymmetries that are at least as problematic as the asymmetry of causation. So, appealing to modal asymmetry merely pushes the problem back a step. Instead of explaining the antisymmetry of causation, the asymmetry of

[23]Price points out that we should be careful that our notion of probability does not give rise to disguised conventionalism. Since I appeal to a frequency notion of probability I ignore this (though he points out that a frequency notion of probability can be problematic if we describe P(B/A) as the frequency with which B succeeds A.

modal claims is then in question. Price call this *passing the buck.* If the fork asymmetry strategist cannot solve the problem of too few forks in a manner that does not encounter these pitfalls, she is left with too little fork-type asymmetry to explain causal priority.

Before evaluating Price's problem of too few forks, I think it necessary to examine the way he formulates the fork asymmetry principle. He argues that there are two distinct principles at play that each demand different kinds of explanations. He call the first principle *correlation productivity.* It states that "joint causes are probabilistically independent of one another, whereas joint effects tend to bc correlated." (Price, 1992 p. 256) The second principle, called *correlation explicativity,* states,

> "A remote non-causal correlation between a pair of events is typically associated with a joint correlation with a third event which (a) is earlier than the two events in question, (b) is their common cause, (c) screens off the original correlation. The corresponding kind of correlation with a later event and/or a common effect is much more rare." (Price 1992, p. 257)

According to Price, the former falls out of the asymmetry of thermodynamics. So, it is contingent and *a posteriori.* The latter, however, is a priori because it results, at least in part, from the connections between causation and agent perspective. (Price 1992, p. 257)

The fork asymmetry strategist agrees that the fork asymmetry is a contingent fact about the world known a posteriori. So, ıo conflict arises from correlation productivity.[24] Correlation explicativity is correlation productivity plus statements about (i) what it means to be a common cause and (ii) the temporal orientation of the common cause relationship. It tells us how to distinguish between cause and effect in the common cause relation and addresses the fact that causation seems to have a future-directed temporal orientation. Thus, it is correlation explicativity that would do the work in explaining causal priority.

The principle of correlation explicativity meets all my necessary criteria except one. Correlation explicativity suggests that given a pair of correlated phenomena types, if the pair is accompanied by a third phenomena type that is a screener-off (or common cause), this screener-off *always* occurs before the original pair. It, thus, fails to allow for the possibility of backward causation. My modified version of Arntzenius' fork asymmetry principle states only that the common cause *typically* precedes its joint effects.

This flaw reflects Price's views about the connection between agent-perspective and causation. Price argues that the future-directed temporal orientation of the causal relation is in part a result of how human agents perceive the world. But, if causal priority is tied to agent perspective, it is hard to see how allowances could be made for backward causation. We do not think of agent-

[24]Notice, however, that *correlation productivity* gives a modal description of common cause "...joint effects tend to be correlated". Such a characterization of common cause is unnecessary and, according to Price, problematic.

perspective as being *typically* oriented in the future direction, but as *always* future directed. Price's agency-based account, therefore, can explain the causal priority only if the causal relation always has the same temporal orientation. One reason to attempt a more substantive account of causal priority than Hume's is to allow for the possibility of backward causation. Thus, giving an explanation of causal priority that restricts the causal relation to one particular temporal orientation, as agency-based accounts do, is undesirable.

Recall Price's common cause example: fire produces both heat and smoke. He argues that though there is a high correlation between smoke and heat in the world, the fact that fire causes both smoke and heat does not depend on the correlation actually occurring. If there had only ever been one fire, there would have been no correlation between the resulting smoke and heat, yet we would not hesitate say that this fire produced both. "Fire produces heat and smoke" is a general causal claim. If there had only ever been one fire in the history of the world, we would endorse a singular causal claim instead; "This fire caused this smoke and this heat." Price objects that the fork asymmetry can explain the former, but not the latter. So, his objection amounts to -- the fork asymmetry strategist has trouble explaining causal priority for causal instances. Fair enough. The fork asymmetry relies on screeners-off to draw the distinction between cause and effect. Screeners-off make sense only in the context of general causation. So, how does the fork asymmetry strategist explain causal priority at the singular level?

This is a problem for the fork asymmetry strategist. It is not clearly insurmountable, however, as Price assumes. In **Part III** I consider a possible strategy for explaining causal priority at the singular level open to the fork asymmetry theorist using the notion of *causal connection*. There is, nevertheless, one more question that must be addressed before trying to uncover the mechanism by which the fork asymmetry strategist gets causal priority to the singular level. We must describe the relationship between singular and general causal claims.

Humean theorists hold that a singular causal claim is best characterized as an instantiation of some general causal claim. If we analyze singular causal claims in terms of general causal claims we would expect the asymmetry to fall naturally from the general to the singular level. If, on the other hand, we analyze general in terms of singular (as I think we have good reason to do), we need some explanation of how the fork asymmetry principle is relevant to singular level causal priority. In **Part II** examine possible ways to analyze the relationship between singular and general causation directly. I then explore the idea that causal priority "filters down" from the general to the singular level.

PART II:
THE RELATIONSHIP BETWEEN
SINGULAR AND GENERAL CAUSAL CLAIMS

Chapter Four: The Eells-Cartwright Debate

4:1: Preliminaries

Two kinds of causal sentences can be distinguished. There are those which assert a relation between *types* of occurrences. For example,

1. Smoking causes lung cancer. (general causal claim)

And, those which assert a relation between *particular* occurrences. For instance,

2. Cindy's smoking caused her lung cancer. (singular causal claim)

(1) and (2) are similar to each other in that they are about the same causal phenomena. They differ in the sense that (2) is a particular instance of the more general relation described in (1). In analyzing the relationship between general and singular causal claims we attempt to describe the relationship between causal claims that are relevant to each other in the way that (1) and (2) are. That is, we attempt to describe the relationship between a claim about the causal relation between two types of phenomena and claims about the causal relation between instances of those same phenomena.

The fork asymmetry strategist must analyze the relationship between singular and general causal claims before she can explain how causal priority gets from the general to the singular level. The fork asymmetry principle, because it relies on screeners-off to draw the distinction between cause and effect, explains causal priority for general causation only. To explain causal priority for singular causation the fork asymmetry strategist must describe the mechanism by which causal priority gets from the general to the singular level of causation. Before describing this mechanism, however, she must have some account of the relationship between singular and general causal claims.

There are several strategies for describing this relationship. Those who follow Hume, e.g., 'early' Eells (1988), analyze singular causal claims as instantiations of general causal claims, thus taking general causal claims as basic. I call this the Humean strategy[25] or Regularity Account. Others, e.g., Cartwright and Carroll, take singular causal claims as basic and describe general causal claims as generalizations over singular ones. I call this the generalization strategy. Eells (1991) later proposes that singular and general causal claims must be analyzed independently of one another. Finally, Hitchcock (1995) suggests that, although neither type of sentence can be analyzed via the other, there is a common structure that both satisfy. Because the first two strategies offer simpler and more intuitive explanations of the relationship, I concentrate on them.

There are several important features of the relationship between singular and general causation to keep in mind when giving an analysis of the relationship. One striking feature is that we cannot infer the truth of one from the truth of the

---

[25]Not to be confused with the Humean strategy of stipulating causal order using temporal order.

other. I call this the principle of *non-inference*. For instance, suppose smoking causes lung cancer and Cindy is a smoker. Cindy's lung cancer may be caused by her exposure to asbestos rather than her smoking. In this case, "smoking causes lung cancer" is true, yet "Cindy's smoking caused her lung cancer" is false. So, we cannot infer from the fact that smoking causes lung cancer that it causes *Cindy's* lung cancer, even if she smokes. Now suppose Cindy's smoking *does* cause her lung cancer. Further suppose that it does so because her habit causes her to go to the nearest grocer everyday to buy cigarettes. Each time she passes through the store entrance she is exposed to a burst of radiation that causes her to develop lung cancer. In this case "smoking causes lung cancer" could be false even though "Cindy's smoking causes her lung cancer" is true. This example shows that we cannot infer from Cindy's case that smoking causes lung cancer generally.

Some argue that singular and general causal claims also differ with respect to whether the causal phenomena they mention must occur for the claim to be true. In order for (2) to be true two things *must* occur: Cindy's smoking and Cindy's lung cancer. In the case of singular causal claims it is uncontroversial that both phenomena must occur in order for the claim to be true. But, in the case of general causal claims the situation is less clear. Eells maintains that (1) can be true a) if there are no smokers or b) if there are smokers (or even if everyone smokes) none of whom develop lung cancer.[26] But, he argues, any attempt to analyze general in terms of singular causal claims must impose the occurrence requirement on general causal claims. As a result, the generalization strategist must insist that (1) can be true only if there are smokers some of whom develop lung cancer. I hereafter use *occurrence requirement* to refer to the assertion that all causal phenomena mentioned in a causal claim (whether singular or general) must occur for the claim to be true.

If we attempt to analyze general causal claims as generalizations over the appropriate singular ones, there are no claims over which to generalize unless there exist smokers who have lung cancer. Thus, initially, it appears easier to analyze singular causal claims using general ones than to analyze general causal claims using singular ones. If, on the other hand, we begin at the general level we might explain general causal claims directly in statistical terms, leaving us free to explain singular causal claims as instantiations of *relevant* general ones. The Humean strategist might then analyze singular causal claims by tacking the occurrence requirement onto some description of general causal claims.

Cases of general causal claims in which either the antecedent or consequent is not instantiated are, nevertheless, peculiar. Though we might arguably count such general causal claims as true, it is far from clear that our intuitions demand

---

[26] I take issue with this shortly.

we do so. In fact, I contend that our intuitions governing such claims are too nebulous to demand that we accept general claims like (a) or (b) as true. For instance, in cases where the antecedent is never (or rarely) instantiated (and is not likely to be instantiated), such as "Eating 1kg of uranium causes death", we might easily argue that eating uranium *would* cause death (if anyone ever ate it), though it *actually* does not. Similarly, if everyone smokes and no one develops lung cancer (though highly improbable, it is conceivable) we take this (generally speaking) to be good evidence that "Smoking causes lung cancer" is false, even if our knowledge of human physiology dictates that smoking *should* cause lung cancer. Though it is possible, as Eells points out, to have a true *general* claim when either the antecedent or consequent is not instantiated, it is far from clear whether we want to allow such modal claims to count as *causal* claims. In the next chapter I discuss a possible way for the generalization strategist to dispense with such claims (i.e., by sticking to a frequency interpretation of probability as Carroll does).

In addition to the occurrence requirement and non-inference, there are three features of causation that affect the debate between Humean and Generalization strategists. I use these issues to clarify the debate and its resolution. The first pertains to the transitivity of causation; 1) Is causation transitive? Eells attempts to undermine Cartwright's approach by relying on his contention that causation is not transitive. I argue that in this case the question of whether or not causation is transitive, though the source of conflict between Cartwright and Eells, is, at least in part, merely a verbal problem.

The complexity of general claims also impacts on the relationship between singular and general causation, namely, there are two kinds of general causal claims. Some sentences, like "Smoking causes lung cancer", take the form "C-ing causes E-ing" displaying the causal relation as holding between the types 'Smoking' and 'lung cancer' themselves, rather than between individuals. Sentences like "Sunspots cause electrical disturbances" are of the form "c's cause e's" portraying the relation as occurring between individuals that fall under the types in question. The latter lends itself more easily to some version of the generalization strategy than the former. So, we must sort out this ambiguity in order to accurately describe the relationship between singular and general causation. John Carroll maintains that the correct approach is to analyze general causal claims of the form "c's causes e's" and then extend this analysis to sentences of the form "C-ing causes E-ing". He argues that a sentence of the form "c's cause e's" is true if and only if the frequency of causers of e's in c's is high (where what counts as high is relative to context).

To illustrate, presume the frequency of lung cancer caused by smoking among people who smoke one cigarette per day is .01 and among those who smoke 1 pack per day is .5. In a context such that all smokers smoke only 1 cigarette per day, the frequency of lung cancer caused by smoking among smokers will be high though it is only .01. Whereas, in a context such that all smokers smoke 1 pack per day, the frequency of lung cancer caused by smoking among smokers is high only if it reaches .5. Carroll uses relative frequency (rather than conditional probability) to argue that in cases where (a) Some people smoke but due to some coincidence no smoker gets lung cancer or (b) Everyone smokes but due to some great coincidence no one gets lung cancer, it is not true that smoking causes lung cancer.

Carroll's account points to another important factor in describing the relationship between singular and general causation: *context relativity*. Relativization to context impacts on the relationship in three ways. First, whether we appeal to relative frequency or conditional probability, what counts as 'high' is relative to context. Carroll takes this into account and as the example above shows his theory can explain true causal claims for which the frequency of causers of e's in c's is much closer to 0 than 1. Second, a causal claim can be true in some contexts, but not others. For instance, in most contexts "ingesting acid poisoning causes death" is true. Suppose, however, that Shannon has previously ingested alkaline poison in just the right amount to neutralize the acid poison. In this very unlikely context acid poisoning *prevents* death. So, by relativizing to context we can explain how a causal claim can be both true (relative to one context) and false (relative to another context).

Also, singular causal claims and general causal claims are clearest when we take into account the significance of the causal role one phenomenon plays in bringing about another. And, the significance of the causal role one phenomenon plays in bringing about another varies with respect to context. For example, suppose (context$_1$) my stove is in a building with no gas source, but a small pocket of gas is present above the front left burner.[27] In context$_1$ the presence of gas *causes* the burner to ignite when I light a match. The prominence of gas' role in the causal relation, however, is much smaller when (context$_2$) my stove is connected to a gas source and the source is activated. In context$_2$ my lighting the match plays a greater causal role in igniting the burner than does the gas. Thus, in context$_2$ my lighting the match *causes* the burner to ignite relative to the presence of gas.

Relative to context$_1$ the pocket of gas causes the burner to ignite. But, relative to context$_2$, my lighting the match is the cause. What such examples show is that whether we count some phenomenon as a cause of another depends in part

[27] Christopher Hitchcock (Hitchcock, 1995, pp. 12-13) presents a similar example, though he uses it to argue for giving independent analyses for singular and general causal claims.

on the significance of the causal role the former plays in bringing about the latter. Hartry Field suggests we modify Carroll's view such that 'A-ing causes B-ing' means something like "In the contextually relevant contexts, the conditional probability of an A-ing playing a significant role in causing a B-ing, given that an A-ing occurs, is significant." (discussion Spring 1996) He claims that whether the given probability is "significant" depends on pragmatic considerations that in turn depend on the A and B in question.

Because causal priority occurs at the general level on fork asymmetry accounts, describing the relationship between singular and general causal claims is a necessary first step in determining the success of, as well as in giving a full analysis of, the strategy. If (as I think we have good reason to do) we analyze general causal claims in terms of singular ones, fork asymmetry theorists must explain how causal priority reaches the singular level. The debate between Humean and generalization strategists influences my scrutiny of fork asymmetry-based accounts in the following manner; after motivating some version of the generalization strategy, I explore possible methods for fork asymmetry theorists to explain causal priority at the singular level.

Having clarified some of the issues surrounding the problem of determining the nature of the relationship between singular and general causal claims and stating its relevance to causal priority, I examine the debate between the Humean and generalization strategies.

4:2: The Humean Strategy

Hume's predecessors analyzed causal relations in terms of things like an item's inherent power or efficacy or agency. They took causal *instances* as basic for understanding causal relations. Hume diverged from this strategy. He took causal laws, rather than causal instances as basic. He reduced causal laws to regularities (or constant conjunctions) and then analyzed causal instances by appeal to the appropriate causal law. Thus, a Humean strategist will accept one or both of the following claims:

1) General causal claims are reducible to regularities.

2) Singular causal claims are analyzable using general causal claims.

The Unanimity theory is a slightly weaker version of the Humean strategy that accepts (2) and a restricted version of (1). The original Humean strategist requires a cause to raise the probability of its effect(s) in *all* test situations. But, there are cases where a cause a) does not raise the probability of its effect or b) raises the probability of its effect in most but not all test situations. For example, consider "ingesting acid poison causes death". If I ingest the acid, but have previously ingested the right amount of an alkaline to neutralize the acid, I will not die. Thus, there is at least one test situation (the one where I also ingest an alkaline solution) in which ingesting acid does not raise the probability that death will occur. As a result, the original Humean strategy is too strict. In response, Unanimity theorists, or "weak-Humeans", e.g., Sober and Eells, merely require a cause to raise the probability of its effect(s) in *specified* test situations. Accordingly, a cause is "unanimous" with respect to its effect when the cause raises the probability of the effect in specified test situations. I center my discussion on the debate between Cartwright (a generalization strategist) and "early" Eells (a unanimity theorist or "weak-Humean").[28]

There are two main problems facing Eells' weak-Humean strategy (and also the original Humean strategy) involving the notions of homogeneity and unanimity.

4:2:1: Homogeneity

Factors correlated with the effect other than the cause in question can generate misleading results about the probabilistic relation between cause and effect. Suppose three causal phenomena, C, X, and E, are correlated. That is, C, X, and E frequently co-occur. The conditional probability of one given either of the others is greater than their probability simpliciter. How do we distinguish between the following possibilities?:

---

[28] I discuss problems arising from homogeneity and unanimity with respect to Eell's weak-Humean theory.

a) One of the three phenomena is a common cause of the other two.

b) Two are effects of a common cause (some fourth phenomenon, Z) and the third is accidentally correlated with them.

c) All three are effects of a common cause (Z).

d) The three phenomena form a causal chain.

e) One phenomenon causes another and the third is accidentally correlated with them.

f) The three phenomena are accidentally correlated; there are no causal connections between them.

Let C be the cause, E be the effect and X be an intermediate link between C and E. The weak-Humean wants to determine the causal influence of C on E using information about the regularity with which they occur together. So, she must not fix X. Otherwise she is unable to distinguish between (a) - (f). I call this the problem of non-homogeneity.

The inability to distinguish between (a) - (f) would have a devastating effect on the first tenet of the weak-Humean account (i.e., general causal claims are reducible to regularities). Thus, she must eliminate any additional factor that has a probabilistic influence on E. This is why the weak-Humean theorist demands that C is a cause of E if and only if C increases the probability of E in populations where E's other causes, with the exception of X, are held fixed. A causal phenomenon is held fixed when its influence (presence or absence) is known or assigned.

The fact that "cause" appears on both the left and right sides of the biconditional, however, threatens circularity. To illustrate, suppose human physiology is such that smoking, X_1, and eating deep fried foods, X_2, each cause heart disease, E. If we do not hold fixed those instances where X_2 causes E and there are many people who smoke and eat deep fried foods, $P(E/X_1)$ and $P(E/-X_1)$ will be much closer in value than we would expect if it is true that smoking causes heart disease. Even worse we might find $P(E/X_1) = P(E/-X_1)$ in which case E is probabilistically independent of X_1. Under these circumstances the weak-Humean strategy yields the undesirable result that smoking (X_1) does not cause heart disease. To accurately determine the causal influence of X_1 on E, then, the weak-Humean must fix those cases where heart disease is caused by eating too many fried foods. Without this restriction, X_1 may, in contexts where X_2 also occurs, appear to have a less significant causal influence on E than it actually has.

Though fixing E's other causes is an obvious strategy for dealing with X_2, it creates a problem in dealing with causal factors intermediate between the cause and effect. We must *not* fix any intermediate factors between C and E or we *screen-off* C's influence over E. That is, E will become probabilistically independent of C. Suppose human physiology is such that there is a specific genetic factor, G, that inclines one to smoke, X_1, and that smoking tends to bring about lung cancer, L, in those who smoke. If we restrict our consideration of G to contexts where those who have the genetic factor do not smoke, G will appear to have no causal influence on LC. Thus, it seems we should not fix X_1. But, there may be some smokers who develop lung cancer and do not have this genetic factor. In these cases smoking is not an intermediate cause between G and LC, but a distinct cause of LC. So, we should fix X_1 in these cases for the reasons described above. Thus, we must fix all and only those *instances* of X_1 that G does not cause.

This is where the circularity problem arises. As Cartwright correctly observes,

> "a regularity account of any particular generic causal truth -- such as "Aspirins relieve headaches" -- must refer to other generic causal claims if we are to pick out the *right* regularities. Hence no reduction of generic causation to regularities is possible." (Cartwright 1988, p. 79)

To reduce a general causal claim to a regularity we must refer to other general causal claims to pick out the appropriate regularities. If we cannot make the reduction without appeal to general causal claims -- the reduction of general causal claims to regularities is circular. Cartwright adds,

> "to pick out the right regularities at the generic level requires reference not only to other generic causal facts but to singular facts as well. So singular causal facts are not reducible to generic ones."(Cartwright 1988, p. 79)

If this is correct, the weak-Humean must appeal to singular causal claims to deal with cases of non-homogeneity and non-unanimity. Thus, singular causal claims cannot be analyzed in terms of general causal claims in any non-circular way.

Cartwright charges that since any attempt to reduce general causal claims to regularities involves an appeal to singular claims, no non-circular reduction from general causal claims to regularities is possible. We cannot reduce singular to general causal claims, because when we try to carry out the reduction of general causal claims to regularities we must appeal to singular causal claims. So, Cartwright threatens (2) via (1). That is, (2) is of no use when we consider why (1) fails. This line, if correct, is disastrous to Unanimity (i.e., weak-Humean) theorists like Sober and Eells. It threatens to undermine attempts to analyze singular in terms of general causal claims, as well as attempts to reduce general causal claims to regularities.

There are at least two ways homogeneity can be violated. As I described earlier an additional factor, K, may be present that influences E in the same manner as C does (i.e., positively, negatively, or not at all). Or, K might *interact* with C to influence E. Recall the example from above to illustrate this second way that homogeneity can be violated. Suppose one ingests both acid poison and alkaline poison. Ingesting acid poison causes death and ingesting alkaline poison causes death. But, when one ingests both in just the right amounts, they interact in such a way as to neutralize each poison. So, acid poisoning *lowers* the probability of death when it is combined with alkaline poisoning. Conversely, alkaline poisoning *lowers* the probability of death when it is combined with acid poisoning.

Fixing E's causes (other than C and intermediate factors, K) is an obvious solution for each kind of violation. But, it causes a problem for handling intermediate causal links, K, between C and E. In these cases we must fix some occurrences of K (those that C does not cause) and not others (those that C causes). Otherwise K's influence on E may generate misleading results about C's influence on E.

4:2:2: Unanimity

Non-homogeneity, when it involves a factor intermediate between C and E is a significant problem for the Unanimity theorist. According to Cartwright, however, a more pressing problem arises when *unanimity* fails. (She considers this problem more pressing in the sense that she thinks the Unanimity theorist can handle cases of non-homogeneity, but not cases of non-unanimity. I describe how she thinks the Unanimity theorist handles cases of non-homogeneity shortly.) Cartwright states "C is unanimous with respect to E if for the factors K that C causes and that themselves are relevant to E, either all K's are positively relevant or all are negatively relevant." (Cartwright 1988, p. 85-86) In other words, C is "unanimous" with respect to E of all intermediate factors between C and E influence E in the same manner (positively or negatively). But, there may be cases where some intermediate factors between C and E have a positive causal influence on E while others have a negative causal influence on E. Thus, C has a *mixed* causal influence over E; sometimes C causes E and sometimes C prevents E. Or, C may produce one cause of E, yet prevent another. Under these circumstances anything can happen; the probability of E may decrease, increase, or remain the same given C.

Hesslow's example illustrates the problem of non-unanimity. Suppose oral contraceptives, O, both prevent pregnancy, P, and produce a chemical, C'. By preventing pregnancy O lowers the risk of thrombosis, T, but increases the risk of thrombosis by producing C' (see Figure 2). (Cartwright 1988, p. 86)

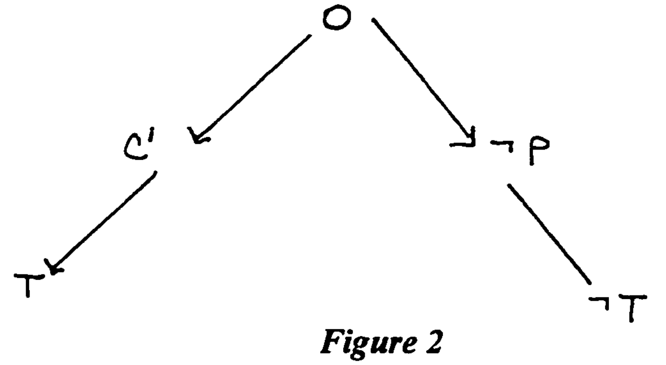

Figure 2

So, O both causes and prevents T. But the Unanimity theorist must say that O is causally neutral for T, because the Unanimity theorist holds that a cause has only one kind of (rather than a mixed) influence on its effect. According to Cartwright, this results from averaging over the relevant intermediate factors to yield the result that C is causally neutral for E. (Cartwright 1988, p. 89) She argues that this hides the *true* causal relevance of C for E; C has a mixed causal influence over E.

In sum, the Unanimity theorist, then, face two problems in formulating her strategy. I use the following terminology to refer to them.

1. ***non-homogeneity:*** She must find some way to handle cases where homogeneity is violated by some factor intermediate between C and E. This strategy must not appeal to singular causal claims.

2. ***non-unanimity:*** She must deal with cases where C has a mixed causal influence on E in a way that does not mask the true (mixed) causal influence of C on E.

4:3: Sober and Eells' Approach

Sober and Eells propose to solve both the aforementioned problems by fixing all K's causes, except C, that occur prior to or simultaneous with C (where K is some factor intermediate between C and E:

> "The basic idea...is just to hold fixed the *causes* of K and -K that are simultaneous with or prior to the time of C: by backing up like this, we claim to be paying such a K its due."(Eells and Sober 1988, p. 99)

K and -K indicate the presence or absence (respectively) of individual factors causally relevant to E, other than C, that occur (or fail to occur) at some time intermediate between the time of C and the time of E.[29] In this way we avoid appeal to causal instances.

Cartwright further explains the Sober-Eells strategy,

> "...[Non-homogeneity] can arise *only when* there are correlations between C and the other relevant factors, $F_1,...,F_n$. Sober and Eells assume... that in a population where all causes and preventatives of a factor are held fixed, the factor will exhibit no correlations with any factor that it does not cause. Thus by holding fixed all causes of our problematic K's which occur prior to or simultaneous with C, we seem to be breaking the correlations of these factors with C, *exactly when the factors are not caused by C itself*." (Cartwright, 1988, p. 84)

Sober and Eells argue that by holding fixed the causes of factors, K, intermediate between C and E (other than C) we avoid problems arising from non-homogeneity. Cartwright agrees that the Sober-Eells strategy provides a solution to non-homogeneity without appeal to singular causal claims. But, neither Cartwright nor Eells explain how holding fixed the causes of K other than C avoids appeal to singular causal claims. I take them to mean the following.

Consider my variation on an example from Eells'. (Eells 1983, p. 43) Consider the following two causal claims: "Your phoning me causes my phone to ring" and "my phone's ringing causes me to lift the receiver." On Eells' strategy we do not fix the intermediate factor ('my phone ringing') because some *instances* of my phone ringing are caused by a call from you while other instances might be

[29] I ignore the possibility of simultaneous causation in this discussion. That is, I assume here that there are no K's that occur simultaneous with their E's.

caused by a call from someone else. So, we must fix *some instances* of K and not others. Let C be your calling me. Let C' and C'' be someone *other than you* calling me. Let E be my picking up the receiver. For the sake of simplicity suppose C' and C'' are the only factors causally relevant to E other than C and K. C = 'You call me,' C' = 'Dave calls me,' and C'' = 'Shannon calls me.' K = 'my phone ringing.' (Recall that capital letter indicate types and lowercase letters indicate tokens.) The fact that we must fix *some instances* of K and not others is then illustrated by Figure 3.

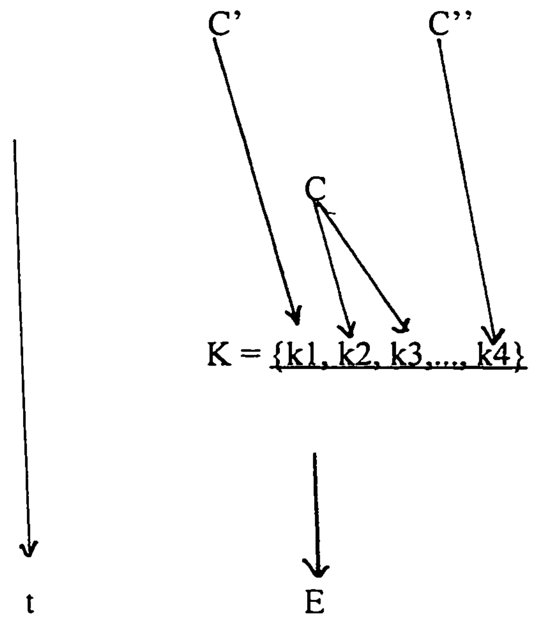

Figure 3

k_2 and k_3 are caused by C, but k_1 and k_4 are caused by events other than C. C will be correlated with C' and C'', because C is correlated with K and K is correlated with C' and C". This potentially yields misleading results concerning the causal relationship between C and E. To prevent such misleading correlations we must find some way to fix k_1 and k_4 but not k_2 and k_3. Further, we must do this without appeal to singular causal claims if we want to give a non-circular analysis singular causal claims in terms of general ones.

By distinguishing between instances of phone ringings we appeal to singular causal claims, e.g., "In the instance when I picked up the receiver, Shannon was on the line." To avoid appeal to singular causal claims Sober and Eells require that we hold fixed all causes of K (occurring prior to or simultaneous with C) except C itself. This avoids appeal to singular causal claims. We do not directly fix particular instances of K. Rather, we fix them indirectly by holding fixed the relevant types of causal phenomena, exemplified by general causal claims, e.g., "When Shannon's calls me I pick up the receiver." We fix every type of cause that either directly causes K, or initiates a chain ending in K (except C). So, we hold fixed causal types, not instances.

We hold fixed all general causes that initiate chains ending in E, except the one that begins with C. We can, therefore, make sense of Eells' claim that holding the causes of K fixed give such problematic K's their due. By fixing all causes of

K, except C, we fix just those problematic instances of K (the ones caused by something other than C), leaving the others alone. Holding the causes of K fixed avoids appeal to singular causal claims by holding the relevant general ones fixed instead. So far, so good.

There is, however, a problem with the Sober-Eells approach to cases of non-homogeneity. It must appeal to causal instances in cases where the factor that generates the misleading correlations is a *cause of* C. In such cases the problem occurs because the factor that misleads us occurs *prior to C* (assuming there is no simultaneous causation), as well as prior to the intermediate factor, K. Sober and Eells' solution is to "hold fixed all causes of our problematic K's which occur prior to or simultaneous with C."(Cartwright 1988, p. 85) And, in fact, their account demands that we hold fixed all such causes of K, except C.

To illustrate, I present the following example. Assume there exists a genetic factor, G, that predisposes one to develop lung cancer, LC. Also, assume in some instances G causes LC directly and in other instances G causes LC only by causing one to smoke, S. S causes LC and LC results in death, E. Finally, there are two (and only two) paths between G and LC. Suppose we wish to determine S's causal influence on E.

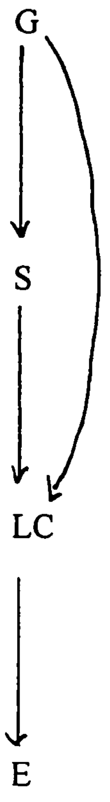

Figure 4.a

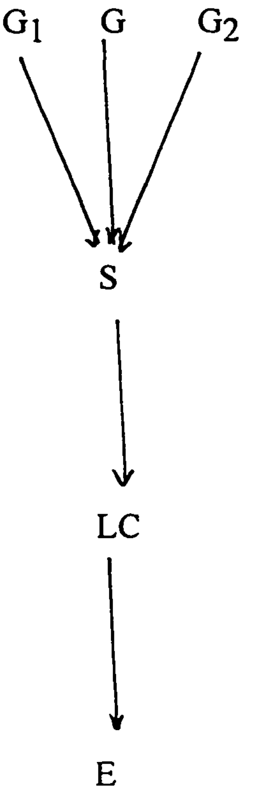

Figure 4.b

In Figure 4.a one path starts with G and moves directly to LC, without passing through S. The second path begins with G, moves from G to S, and then from S to LC. We need to fix all of LC's causes other than S. So, we want to fix only those instances where G directly causes LC, because in these cases S does not cause LC. But, to do this we must appeal to causal instances. S only screens-off those instances of G in which G causes S. So, instances in which G causes LC directly may still produce misleading probabilistic results on the relationship between S and E.

The most obvious way for Eells to answer this is to argue that the Sober-Eells strategy requires us to fix all causes of LC other than S, including *the causes of S*, i.e., including G. Initially this may not seem to be a problem. We don't mind if S screens off G from LC. In fact, that is good; no misleading correlations will result between G and E. But, when we fix *all causes of S*, we fix S. Thus, S becomes probabilistically independent of both LC and E. As a result, when we follow the Sober-Eells strategy in the above example we fix all causes of LC *including* S, thereby fixing LC as well. Previously we saw that if we fix a factor intermediate between the cause and effect, (in this case LC), then LC *screens-off* any probabilistic relation between S and E. This is exactly what Sober and Eells want to avoid. Again we are faced with the problem of finding a way to hold some instances of K fixed and not others. If we do not fix any instances of G we leave ourselves open to situations where some cause of S (in this case G) generates misleading information about the correlation between S and E by bypassing S and directly causing LC.

The solution, then, is to fix some *instances* of G and not others.[30] This approach, though, is not open to Sober and Eells because it appeals to singular causal claims. What the Sober-Eells strategy has done is to postpone the problem, rather than solve it. Thus, although the Sober-Eells strategy provides a solution to non-homogeneity without appeal to singular causal claims, their solution is illusory.[31]

---

[30]Of course, we could opt not to fix any causes of S. If we do this, however, we are left with the problem illustrated in Figure 1.

[31]Hartry Field argues that the Unanimity theorist could respond by insisting that there will always be an intermediate node between G and LC.

Chapter Five: The Generalization Strategy

5:1: Cartwright's Project

There are three features of the Sober-Eells strategy hat Cartwright rejects: their claim that causation is intransitive, their use of path analysis, and their refusal to take into consideration anything that occurs after the cause. These three features point to the major differences between Cartwright's approach and that of Sober and Eells. I argue in this chapter that a deeper consideration of these differences shows that the debate is, at least in part, a verbal one.

The first difference is that while Eells assumes that causation is not transitive, Cartwright holds that causation is *transitive*. Eells argues that in cases like Hesslow's, oral contraceptives do not cause or prevent thrombosis. Rather, oral contraceptives cause the chemical C' and prevent pregnancy. In turn, C' causes thrombosis and -P prevents thrombosis. Oral contraceptives play the role of *causal ancestor* to thrombosis via O's positive causal influence on C' and its negative influence on P.

Eells argues that O does not cause T because for an instance of probabilistic causation to be transitive one of the following must occur: 1) the causal influence of O on T is neutral when either or both of O's influence on C' and (C')'s influence on T is neutral or 2) O's causal influence on T is negative when exactly one of O's influence on C' or (C')'s influence on T is negative or 3) O's causal influence on T is positive if either a) both the influence of O on C' and C' on T are positive or b) both are negative. (Eells 1983, p. 47) In Hesslow's example O has a positive causal influence on C' and C' has a positive causal influence on T. Yet, Eells declares O has no causal influence on T. By considering only T's immediate causal predecessor as a true (or direct) cause of T, Eells attempts to eliminate any intermediate factors, K, occurring between the cause and effect. Since non-homogeneity and non-unanimity are troublesome for the Humean theorist only when the additional factors in question are intermediate between the cause and effect, the notion of direct cause eliminates the circumstances under which these problems could occur.

Eells uses the notion of direct cause to chop time into discrete chunks so small that no additional factors can occur between cause and effect. According to Cartwright the problem with this strategy is that time is not discrete. She argues that even if we allow that time can be broken into small enough chunks to make the proposal work, the only way to develop a statistical model of causation on the basis of discrete time inevitably involves reference to singular causal claims via C's capacities because,

> "If the statistics in this kind of model are to double for causation, as the Hume programme suppose, some instructions must be given on how to construct the model from the kind of information about the situation that the programme can allow. But that cannot be done, and even to get a start in doing it will require the input of facts about individual causal histories." (Cartwright, *Nature's Capacities*, p. 129)

So, to determine how small the chunks of time must be for the proposal to work, we must look at single cases.

Cartwright's focus is on how to fully describe and pick out *probabilistic causes*. A probabilistic cause, C, is such that when C occurs sometimes its effect does not follow, though all E's necessary conditions are present and operating. When the cause is present, however, (as well as all its necessary conditions) there is a fixed probability as to whether the effect will occur. According to Cartwright, in order to fully describe and pick out probabilistic causes we must draw two distinctions. First, she distinguishes between when C *occurs* and when C *operates*. Sometimes C occurs without exerting its influence on E, but C operates when it does exerts it influence on E. Cartwright assumes that whenever C operates it contributes its full influence to E and she suggests that C cannot operate unless it occurs with a set of background factors, B, that (in conjunction with C) are sufficient to produce E. So, for C to operate (or, exert its influence over E):

1. C must occur.
2. B must occur.
3. The elements in B and C must not interact.

If (1) - (3), then E is very likely to occur (though not determined).

Second, Cartwright distinguishes between the form of C's influence on E and whether C influences E. Cartwright argues that standard statistical causal models do not identify the *form of influence C contributes to E* but, rather, they decide *whether C contributes at all.* Such models use additive variables. Since it is the influences, or operations, of causes that are additive rather than the causes themselves, we cannot employ these models until we determine the *form* of the influence a cause contributes. (Cartwright, *Nature's Capacities*, p. 106) For instance, an object's mass, m, and its acceleration, a, are partial causes of its force, F. Together they contribute an influence of the form 'ma' on F. According to Cartwright, we must determine that m is related to F via F = ma before we can use any statistical model to represent the *causal relation between m and F.* But, she asserts, in determining the form of a cause's influence on its effect we must appeal to singular causal claims, which is problematic for the Humean strategy. According to Cartwright, the Humean theorist:

> "proposes to test a claim that *C* causes *E* by holding fixed a *complete set* of *E's* other causes...[I]n a case where causes are purely probabilistic...A complete set of causes for an effect *E* is a set of causes of *E* such that (i) if *E* occurs some member of that set is bound to have occurred and to have operated, and (ii) if any member of that set occurs and operates, and no preventatives of *E*...occur at the same time or during the relevant period after, then *E* occurs." (Cartwright, *Nature's Capacities,* p.111-112)

In sum, Cartwright distinguishes between a) the level of influence a cause has on its effect and b) the form of C's influence over E. Also, she distinguishes between c) whether a factor is present and whether it operates in a given population and d) whether the cause's failure to operate is due to the absence of its necessary conjoining conditions or because the cause lacks the capacity to bring about the effect. Cartwright argues that the Humean theory cannot be formulated in a way that accounts for (a) - (d) without appeal to causal instances, because in order to use any probabilistic causal model we must first determine the form of the influence C contributes to E which involves appeal to singular causal claims.

Sober and Eells attempt to solve the problem of non-unanimity by analyzing causal paths. They fix the problematic instances of K by fixing some element on every causal path to E, except the one that begins with C. The resulting population, according to Eells, shows how C affects the probability of E. C has a positive influence on E if P(E) increases, negative if P(E) decreases and neutral if P(E) is unchanged. According to Cartwright, path analysis cannot solve problems of non-unanimity because to identify unique paths we must 1) determine the form of the influence C has on its effects and 2) identify both partial and complete causes. Cartwright contends that (1) requires appeal to causal instances and given the form of C's influence on E it is possible to pick out unique causal paths only if we consider complete causes. Once we consider partial causes it is impossible to pick out unique causal paths.

Further, Sober and Eells' use of path analysis involves an added restriction: fix only elements on a given path that occur prior to or simultaneous with C. The motivation for employing this restriction arises from the significance of the causal role C plays in producing E, relative to E's other causes. Since on Eells' account C will occur nearer E on any path than the fixed element, there is no danger we will fix an element with greater, i.e., more direct, causal influence over E than C. As discussed previously, the restriction has an added benefit of allowing Sober and Eells to handle most cases of non-homogeneity.

There is a problem with using this restriction to handle cases of non-unanimity. It arises in cases like Hesslow's where the cause has dual capacities. Consider the possible states one might be in when taking the pill in Hesslow's example: C'-P; -C'-P; C'P; -C'P. When the original state is C'-P, birth control pills make no difference to the probability of thrombosis. This is because birth control pills both produce C' and inhibit P, which is the situation that already obtains. When the original state is -C'-P birth control pills produce C' thus increasing the probability of thrombosis. In the case of C'P birth control pills decrease the probability of thrombosis by preventing pregnancy. Finally, in the case of -C'P the pills both increase and decrease the probability of thrombosis by both producing C' and preventing P.

The question is then, 'in cases where unanimity fails, how do we fix C' and -P?' There are two ways. The first way (Case 1) is to maintain that C' and P occur *before* C. Since C' and P are already either present or absent at the time of C, the pill- taker is in one (and only one) of the following four conditions prior to taking the pill: C'-P; -C'-P; C'P; -C'P. Cartwright observes that in Case 1 it is uncontroversially inappropriate to average over the four possible original conditions of pill-taker because it conceals the true, mixed, relevance of C on E. In Case 2, C' and P must occur after C, because their efficacy is short-lived. While Sober and Eells admit that it is uncontroversially inappropriate to average in Case 1, they argue that it is appropriate in Case 2. Cartwright, on the other hand, contends that there is no significant difference between Case 1 and 2 to justify averaging in the latter.

According to Cartwright,

> "We see in this [Hesslow] example that birth control pills make a difference in *every single* ...[background] situation, with the exception of the one where all of their effects are already present. So they are certainly of causal relevance to thrombosis, though the relevance is mixed -- sometimes they produce it and sometimes they inhibit it. But we can make this relevance completely disappear by averaging..." (Cartwright 1988, p. 89)[32]

Assume that the different possible original conditions of pill-takers occur with the following frequencies (as assigned by Cartwright): C'P: 0.3; -C'P: 0.3; -C'-P: 0; C'-P: 0.4 (Case 1). When we calculate the partial conditional probability of E (Thrombosis) in the presence and absence of C (without averaging over the four possibilities) we get:

---

[32]Cartwright's terminology, i.e., "background situations", is vague. But, I take her to be referring to the four possible states of C' and P: C'P; -C'P; -C'-P; C-P prior to taking oral contraceptives.

For C'P: P(E/-C) > P(E/C);

For -C'P: P(E/-C) < P(E/C);

For -C'-P: P(E/-C) < P(E/C);

And for C'-P: P(E/-C) = P(E/C).

These results reveal the true causal relevance of C for E; they reveal a *mixed* influence of C on E. When we average over the four possible effects of C on E, our determination of C's causal influence on E becomes P(E/C) = P(E/-C) = .52.

Sober and Eells argue that in Case 2 it is acceptable to average over the four cases. (Eells 1988, p. 102) Cartwright, by contrast, admits of no significant difference between Case 1 and Case 2. The only difference she does acknowledge is the difference in the time that the factors occur. In Case 1 the factors occur before C and in Case 2 they occur after. This difference, she argues, does not justify averaging over C's different products in one case (Case 1) but not the other (Case 2):

> "What difference should it make whether C' and P are already in place before the decision about birth control pills is taken, or whether they befall you afterwards, so long as their doing so is completely independent of the pills? Yet in the second instance Sober and Eells average -- and conclude that the pills are irrelevant to thrombosis -- whereas they never average in the first...In sum, we agree that averaging is bad for independent factors that occur before C. But if these factors operate just after C rather than just before -- and everything else stays the same -- the Sober and Eells rule averages in exactly this bad way." (Cartwright 1988, p.89-90)

Although there are other relevant issues in the debate, it culminates in this question of whether it is acceptable to average across C's different outputs in Case 2. Since the factors are intermediate between C and E in Case 2, it is not immediately obvious how to hold fixed C' and P without appeal to singular causal claims. Eells and Cartwright disagree on how to treat Case 2; Cartwright argues that the true causal relevance is mixed, while Eells argues that C has no causal influence on E. Eells argues that the set of all relevant factors occurring prior to C plus C or not C, are all that is causally relevant, at the time of C or -C, to the probability of E:

> "It must be kept in mind that the probability values [for P(C'P), P(-C'P), P(C'-P), and P(-C'-P)] pertain to a particular background context *B* in which all factors, other than *C*, that are causally relevant to *C'* or to *P* are held fixed in one way...[I]t is obviously an indeterministic matter in this example which (if any) individuals in *B* will get *C'* and which (if any) will have *P*...So, given that all chains are Markovian, as Cartwright assumes, *B*, together with whether or not *C*, are all that is causally relevant, at the time of *C* or *-C*, to whether or not one will get *E*. *BC* and *B-C* specify (or hold fixed in one way) everything that, at the time of *C* or *-C*, is causally relevant to whether or not *E* will occur. Thus, since *BC* and *B-C* confer exactly the same probability on *E*, it is evident that *C* and *-C* make no causal difference, in background context *B*, as to whether or not one gets *E*." (Eells 1988, p.101-102)

To assume that all chains are Markovian, according to Eells, is to assume that whether or not C is present does not affect the probability of E once the presence or absence of C' and P is specified.

It is unclear from Eells' discussion exactly what he thinks the difference between the two cases is. I can think of only two differences between Case 1 and

Case 2 that might warrant Eells' averaging over the four possible causal influences of C on E in Case 2 and not in Case 1. I contend that neither is sufficient to justify averaging in one case and not the other. First, as I noted above, it is already determined which of the four situations is relevant in Case 1, whereas in Case 2 it is yet to be determined whether C' or P will occur. Second, in Case 1, C' and P are not intermediate between O and T. C' and P are partial causes/preventatives of T that are *not as significant to the production of T or -T as O.* Yet, in Case 2 C' and P occur *between* O and T. That is, on the causal path from O to T, O is an *indirect* cause of both T and -T, while C' is a *direct* cause of T and -P is a *direct* cause of -T. Because it more immediately affects T than O, C' plays a more significant role in bringing about T than does O. Analogously, -P plays a more significant role in bringing about -T than does O.

Part of the mystery surrounding this debate involves conflating these two differences. Cartwright and Eells enter and continue the debate as though they each invoke the same one. In fact, Cartwright concerns herself with the former and Eells addresses the latter. If we consider merely the time at which C' and P occur, there appears to be no justification for averaging in Case 2. Focusing instead on the significance of the causal role C' and P play in each case provides Eells with at least some initial motivation for averaging in Case 2.

By stressing the significance of C' and P's role in causing or preventing T, Eells is able to argue that in Case 2, C' has a positive causal influence on T, -P has a negative causal influence on T, and O is *causally neutral* for T. Recall Eells' contention that causation is intransitive. He argues instead that C'/-P is a direct cause/preventative of T, whiie O is merely a *causal ancestor* of T. So, according to Eells, while O affects the 'ways in which T may come about,' it does not affect 'whether one will get T in B.' Since O is not a *true* cause of T on Eells' account, when we average over the four background situations in Case 2 we get the correct result that O is causally neutral for T.

Fells offers an explanation of Cartwright's intuitions in terms of his own approach. He argues that Cartwright conflates 'ways in which E may come about' with 'whether one will get E in B.' He iliustrates Cartwright's mistake through the following example. Suppose I have an urn containing equal number of red, green, and yellow balls. First, I pull out a ball, then role a fair die. Drawing a red ball wins $1 if the die lands ace, nothing otherwise. Drawing a green ball wins $1 if the die lands with an even number up, nothing otherwise. If I draw a yellow ball, I win $1 if the die lands with a 2-6 up, nothing otherwise. Before drawing a ball, I may either paint all of the red and green balls yellow and paint one of the yellow balls green (C), or leave the balls alone (-C). R represents drawing a red ball. G represents drawing a green ball. And, Y represents drawing a yellow ball. E is my winning $1. According to Eells, "[t]he background context can be thought of as given by the original constitution of the urn." (Eells 1988, p. 100) Eells argues that in this case "C is causally positive for E, and C would surely be the right option to choose if you wanted to win $1. But Cartwright's approach implies that there is no single causal role of C for E in this case either." (Eells 1988, p. 104)

Eells' urn example unfortunately interposes an additional factor that complicates the discussion; it focuses on questions of strategy as opposed to questions of causal efficacy. Eells contends Cartwright's approach fails because it does not prescribe the correct action for winning $1. Cartwright maintains that whether choosing C is a good strategy for winning $1 is irrelevant to C's causal influence on E. She explains,

> "We would both agree that it "would surely be the right option to choose if you wanted to win $1," but I do not think that this bears on its univocality. Here is a single action with opposed causal tendencies. It is the right one to choose because the positive tendencies outweigh the negative."

Cartwright introduces a second example designed to restrict the decision factor in Eells' urn example:

> "a ruthless experiment designed to test three new drugs that may be used to treat manic-depressives. The subjects will randomly be given red, green, or yellow pills, with equal probability. Those who get the yellow pills have a 5/6 chance of becoming so depressed that they commit suicide; those given red pills have only a 1/6 chance; with green the chances are 50-50. Now you have the option to substitute green pills for the red and yellow ones. Surely this action has a significant impact on the survival or death of the subjects, even though it leaves their relative numbers unchanged." (Cartwright, article)

I agree that the decision factor makes the situation unnecessarily complicated. Cartwright's modified version of the urn example, however, fails to eliminate the question of strategy. There remains in the example an agent who must choose whether to substitute the green pills for the red and yellow ones. Also, the agent will employ a particular strategy dependent on how she wants to affect the survival of the test subjects. For example, if she believes it immoral to actively increase the chances of death for those taking the yellow pills, she will refrain from making the substitution. Due to the added complications presented by these two examples, I focus on Hesslow's example.

I mention the other examples here, however, because they play an important role in Eells' account. Questions of strategy involve epistemological considerations about what we know or can know about the causal relation. As such, it allows Eells to exclude accidental generalizations from his list of true general causal claims. So, a general causal claim like "smoking causes lung cancer" is true, according to Eells, even if there are, accidentally, no smokers. Even though Eells does not state this, he counts the claim as true because it is a better action-guiding principle even though there are no smokers. Cartwright argues, however, that the relevant question here is metaphysical. It involves what the causal relation *is*, irrespective of our knowledge of causal processes. Carroll more clearly and explicitly makes the same point. Carroll argues that just as there are many *accidentally true* generics, there are also many accidentally true causal propositions (or causal generic sentences). For example, when there are exactly 3 coins in my pocket and each is a nickel, "coins in my pocket are nickels" is an accidentally true generic sentence. If the only people who ingest alkaline poison also ingest just enough acid poison to neutralize the damaging effects of the first poison, "Acid poison prevents death" is an accidentally true causal proposition.

The generalization strategist counts accidentally true causal propositions as true. Humean theorists do not, even though they should. And the reason they should, according to Carroll (and Cartwright), is that although generic sentences have interesting properties over and above their truth, e.g., their explanatory ability and their action-guiding ability, the relevant issue here is what makes a causal proposition true. Carroll explains that the Humean theorist is confused in thinking that what we are after is an account of what makes a causal proposition true *and* a good action-guiding principle *and* a good strategy to follow. But we cannot determine this by merely considering what makes it true that, for example, smoking causes lung cancer. Carroll argues that the difference between accidentally true causal statements and more useful causal statements amounts to the difference between (1) and (2):

(1) All ravens have speeds less than 31 mps.

(2) All signals have speeds less than 300,000,001 mps. (Carroll, 1991)

To identify the difference between (1) and (2) we must determine what makes (2) a *causal law*. This issue is separate from what makes causal statements *true*. The generalization strategy says nothing about the difference between accidentally true causal claims and more useful ones, while the Humean theorist mistakenly does.

To return to the previous point, the core of the debate involves the two main differences between Case 1 and Case 2: a) whether the occurrence of C' and P are determined at the time C occurs and b) the significance of the causal role C' and P play in bringing about T. Eells needs his notion of direct cause if he is to succeed in addressing cases of non-unanimity without appeal to singular causal claims, because it is this notion that allows him to average over background factors in Case 2. Earlier we saw that Eells' notion of direct cause which relies on the notions of causal intransitivity and discrete time cannot work to solve the non-unanimity problem because time is not discrete. And any statistical causal model involving discrete time we devise for our practical purposes must appeal to singular causal claims. As Cartwright notes, "there does seem to be in many cases a distinction between generically causing a proximate cause of E, and generically causing E. But it will be a hard one to draw since most legitimate generic-level causes operate by producing more proximate causes." And, it is hard to see how we can draw the distinction without examining individual cases.

After taking these various aspects of the debate into consideration we see that it turns, in part, on a verbal issue. Specifically, it relies on the issue of whether we talk about causation in terms of causal chains or in terms of direct and indirect causes. This distinction, however, does not amount to much. Whether we talk about chains or causal ancestors we must point to some compelling, substantive difference between Case 1 and Case 2 that will allow us to average over factors in Case 2 even though it is *uncontroversially* wrong to do so in Case 1. And, a mere verbal difference in the way we describe causal paths is not enough to do the job. Cartwright's objection to the Humean strategy, in sum, is that Humean theorists must make their account contextual and probabilities cannot identify context in a fine-grained enough way. So, the Humean theorist must appeal to the single case.

Before I leave the Cartwright-Eells' debate I would like to make one more comment concerning Cartwright's distinction between C's occurrence and C's operation. Recall that in order for C to operate C must occur and E must be very likely to occur (though not determined since we are dealing with probabilistic causes). In effect, Cartwright builds the occurrence requirement into general causal claims. By holding out for C's operations, she introduces requirements about C's occurrence into an analysis of general causal claims. As a result, it is hard to see how her account will handle causal sentences like "Eating 1 gram of uranium causes death," where we take the sentence to be true even though there are no individual cases. I discuss this point in more detail in the next section.

In sum, although the debate between Cartwright and Eells collapses partially into a verbal one, Eells has not given adequate justification for averaging over background factors in Case 2. Also, he mistakenly uses this forum to draw a distinction between accidentally true causal claims and causal laws. In addition, to make any aspect of his attempts to handle cases of non-unanimity work (time chopping, path analysis, or considering only events before or simultaneous with C) he must appeal to C's capacities which involves an examination of individual cases where C occurs. He must appeal to singular causal claims to solve problems of non-homogeneity and non-unanimity. Critics like Hitchcock and Eells have been quick to dismiss attempts to describe general causal claims in terms of singular causal claims. As a result, not much attention has been given in the literature to developing the generalization strategy. Since the strongest version of the Humean strategy has difficulty avoiding appeal to individual cases, I hold that the generalization strategy merits more attention. In the next section I attempt to motivate some version of the generalization strategy that can handle its most serious obstacles.

5:2: The Generalization Strategy

I consider here several formulations of the Generalization strategy and attempt to (1) show that some version of the Generalization strategy is potentially able to handle standard objections and (2) assess the suitability of these forms of Generalization strategy as analyses of causal statements. I begin the discussion with Cartwright's proposal.

Cartwright argues, "A generic claim, such as 'Aspirins relieve headaches,' is best seen as a modalized singular claim: 'An aspirin can relieve a headache'; and the surest sign that an aspirin can do so is that sometimes one does do so." (Cartwright, *Nature's Capacities*, p. 95) She explains further,

> "To assert the causal law that aspirins relieve headaches is to claim that aspirins, by virtue of being aspirins, have the capacity to make headaches disappear." (Cartwright, *Nature's Capacities,* p. 136)

In other words, we gathered information via our empirical methods to reveal a precise set of circumstances in which an aspirin relieved a headache. Since we can then point to one case where an aspirin made a headache disappear, it is true that aspirins have the capacity to relieve headaches. So, aspirins *can* relieve headaches, making the general claim true.

Cartwright's motivation for this approach arises from what she refers to as the 'one-shot' experiments of physics. In such cases the experimenter searches for the exact set of arrangements needed to achieve the desired effect. Once she achieves this result we need not repeat the experiment: one case is enough to identify the causal relation and the method involves no statistics. As with 'one-shot' experiments, the aim in identifying a causal relation is to isolate one instance of it.

In a causal model the probabilities signify the fixed parameters of the population. The fixed parameters, in turn, represent a statistical average over the size of the influence C contributes to E each time it occurs. It is this notion, C's *operation*, that is fundamental to identifying the causal relation between C and E. Also, this notion relies on the single case. Cartwright explains, "The task is always to find a special kind of population, one where the individual occurrences of the process in question will make a predictable difference to the probabilities, and, conversely, where the probabilistic difference will show up only if some instances of the process occur." (Cartwright, *Nature's Capacities*, p. 134) Thus, we support our causal laws by establishing that the law is successfully instantiated at least once, implying that in order for a general causal claim to be true, it must be instantiated. So, the probabilities we get through our empirical methods provide support for the single case.

Cartwright's methodology is attractive given its resources for handling problems resulting from epiphenomena. We can conceive of cases where some event, c, in a certain kind of population seems to bring about some other event, e, in virtue of being a C, when c and e are really common effects of some other event, x. But since c does not cause e, we have not actually used our empirical methods to bring brought about e via c. In this case we mistakenly identify c as the cause and x as a necessary background condition. We could, however, try fixing c and then x. Fixing either phenomenon catalogs the cause and the effect, with e, of the common cause. In this case fixing x shows that the probability e will occur (given x) is independent of c; x screens-off c from e. If we fix x, we see, again, that c makes no probabilistic difference to e. So, Cartwright's account picks out x (but not c) as a cause of e. In other words, the original experiment was not the 'right' one to determine the cause of e. We did not bring about e using c, but, rather, using x.

Her account is too loose, nonetheless, in that we can conceive of cases where some event, c, by virtue of being a C and in a certain kind of population, brings about some other event, e, though this occurrence does not indicate the existence of a causal law. This occurs when we bring about e using c, but e occurs as a result of c in this context only on this one occasion; c *accidentally* causes e. It also occurs when c causes e in some unlikely context. For example, suppose that using our empirical methods we prevent Kate's death by making her ingest acid poison. The acid poison prevents death because Kate has previously ingested alkaline poison. We have given her just the right amount of acid poison to neutralize the alkaline poison. Thus, under these conditions acid poison prevents Kate's death. So, acid poison has the capacity to prevent death. That is, it can prevent death and we know this because sometimes it does so. Cartwright then commits herself to accepting 'Acid poison prevents death' as a causal law.

This "looseness" arises from the failure of Cartwright's account to handle those special features of the relationship between singular and general causal claims mentioned earlier:

1. Singular and general causal claims are not sensitive to context in the same way.
2. A true singular causal claim demands that the cause and effect occur, but a true general causal claim does not. (At least initially, general causal claims appear not to demand that the cause and effect occur.)

We need to modify the Generalization strategy so that it takes into account (1) and (2). We can do this by analyzing causal claims relative to context, relative to alternative causes and by recognizing the significance of the role various causes play in bringing about their effects. Some of these considerations are suggested by Cartwright's methodology, though not explicitly stated in the formulation of her account. Two features in particular suggest that she does recognize the importance of relativization to context. First, she states that some event, c, brings about another event, e, *by virtue of being a C*. This implies that c will probably bring about e given that *the other necessary factors to produce e are present and operating*. In identifying the other necessary factors for e, we essentially state that 'c causes e' relative to this set of other necessary factors, or relative to this certain context. Second, she bases her strategy on the one-shot experiments of physics. Such experiments involve identifying the set of necessary conditions that, together with c, produce e. Though we only instantiated the law once, we identified the right combination of necessary conditions needed in conjunction with c to produce e. So, we presume that we can instantiate the law again when we wish. Thus, we have conjured up some context relative to which 'C causes E.'

As I mention at the beginning of this chapter, there are certain features a viable version of the Generalization strategy must possess. Carroll's suggested modification of generalization strategy takes into account both the complexity of general causal claims and analyzes causal claims relative to context. He maintains that the correct approach in handling the complexity of general claims is to analyze claims of the form "c's causes e's" and then extend this analysis to sentences of the form "C-ing causes E-ing". Thus, he suggests, (GS_1) 'C causes E' is true if and only if the *frequency* of *causers* of e's in c's is high, where what counts as high is relative to context. (Carroll, 1991, p. 257)

Carroll also identifies the importance of *context relativity* explaining the relationship between singular and general causal claims. He takes into account the fact that whether we appeal to relative frequency or conditional probability, what counts as 'high' is relative to context. Thus, he can explain true causal claims for which the frequency of causers of e's in c's is much closer to 0 than 1. Recall the example from Chapter One. Suppose the frequency of lung cancer caused by smoking among people who smoke one cigarette per day is .01 and among those who smoke 1 pack per day is .5. In a context such that all smokers smoke only 1 cigarette per day, the frequency of lung cancer caused by smoking among smokers will be high though it is only .01. Whereas, in a context such that all (or most) smokers smoke 1 pack per day, the frequency of lung cancer caused by smoking among smokers is high only if it reaches .5.

Notice that we do not need a high frequency of persons with lung cancer among smokers to establish the truth of the claim. Rather, we must establish that the frequency of persons with lung cancer *caused by smoking* among smokers is relatively high. More formally, then, (GS_1) reads "S causes LC iff Fr [((y) (LCy and Cxy))/Sx] is high." There are upper and lower bounds on the context-dependence of generic conditional sentences. If the frequency of cases of lung cancer caused by smoking among smokers is 0, "Smoking causes lung cancer" is false in all contexts. If the frequency of cases of lung cancer caused by smoking among smokers is 1, "Smoking causes lung cancer" is true in all contexts. In addition, (GS_1) holds only in those cases where the frequency of LC's in S's exists. Fr (LCx/Sx) does not exist iff either: (1) there are no smokers (in which case 'Smoking causes lung cancer' is vacuously true) or (2) there are infinitely many smokers (Carroll is unsure on the treatment of this case).

Next, a causal claim can be true in some contexts, but not others. For instance, in most contexts "ingesting acid poisoning causes death" is true. Suppose, however, that Shannon has previously ingested alkaline poison in just the right amount to neutralize the acid poison. In this very unlikely context acid poisoning *prevents* death. So, by relativizing to context we can explain how a causal claim can be both true (relative to one context) and false (relative to another context).

The impact of context relativity marks the importance of (GS_1)'s reliance on relative frequency rather than conditional probability. That is, Carroll uses relative frequency to argue that in contexts where (a) Some people smoke but due to some coincidence no smoker gets lung cancer or (b) Everyone smokes but due to some great coincidence no one gets lung cancer, it is not true that smoking causes lung cancer. Thus, relative to contexts like (a) and (b) general causal claims like "Smoking causes cancer" is false. Here he diverges from those who employ conditional probability and claim that our intuitions indicate that we must accept the general causal claim in such contexts as true. By adhering to a relative frequency interpretation of probability, he can argue that when the frequency of e's among c's is zero, so is the causal influence of c's on e's. So, "Smoking causes lung cancer" is true if and only if the number of persons with lung cancer caused by smoking among smokers, is high relative to a particular context.

In other words, Carroll maintains that by employing relative frequencies we can count accidentally true generics, as well as more useful generics, true. Whereas, theories that use conditional probabilities, like Eells' above, stamp accidentally true generics as false. Carroll explains,

> "A modal probability function is sometimes used in the analysis of property-level causation because some worry that frequencies can be misleading. If there are only five sunspots and, by accident, not one of them causes an electrical disturbance, then the frequency of causers of an electrical disturbance in sunspots is zero... Were [(GS_1)] to use some more modal probability function, like a propensity function or a hypothetical frequency function, that conclusion might be avoided...[But (GS_1)] gives the correct consequence about both of the cases just described. In a case in which there are five sunspots and not one causes an electrical disturbance, it seems to me clear that it is not the case that sunspots cause electrical disturbances." (Carroll, 1991, p. 263)

According to Carroll, the frequency function belongs in (GS_1) because there are many accidentally true generics. Carroll argues that fleshing out the difference between accidentally true generics and more useful ones corresponds to showing why "Nothing travels faster than the speed of light" counts as a law of nature while "No automobile travels faster than the speed of sound" does not. Since this investigation does not depend on what determines the truth of the general causal claim, Carroll argues that we should not concern ourselves with it here.

But, one factor Carroll does not distinctly take into account is the significance of the causal role c plays in bringing about e. Consider the example from **4:1**: Suppose ($context_1$) my stove is in a building with no gas source, but a small pocket of gas is present above the front left burner.[33] In $context_1$ the presence of gas *causes* the burner to ignite when I light a match. The prominence of gas' role in the causal relation, however, is much smaller when ($context_2$) my stove is connected to a gas source and the source is activated. In $context_2$ my lighting the match plays a greater causal role in igniting the burner than does the gas. Thus, in $context_2$ my lighting the match *causes* the burner to ignite relative to the presence of gas. What such examples show is that whether we count some phenomenon as a cause of another depends in part on the significance of the causal role the former plays in bringing about the latter. Thus, there are two tactics for formulating the generalization strategy:

1. Employ relative frequency, thus extending the occurrence requirement to general causal claims and accepting that sentences like "Smoking causes lung cancer" are false whenever there are smokers and none get lung cancer (i.e., Carroll).

[33]Christopher Hitchcock (Hitchcock, 1995, pp. 12-13) presents a similar example, though he uses it to argue for giving independent analyses for singular and general causal claims.

2. Adopt conditional probability and give some counterfactual explanation about why "Smoking causes lung cancer" is true even when there are smokers and none get lung cancer (given our knowledge of human physiology).

Each kind of approach has certain obstacles with which to contend. With respect to relative frequency it is difficult, if not impossible to ascertain the values of limiting frequencies. For instance, even if at time, t, no smoker has ever developed lung cancer from smoking, the next billion smokers could get lung cancer from smoking. This would shift the relative frequency upward. So, even though in this example smoking is *not* a cause of lung cancer at t, it *becomes* a cause of lung cancer some time after t. With respect to conditional probability, on the other hand, it is difficult if not impossible to say just what probability is. These problems, nevertheless, are such that any theory of causation (whether it is of the Humean or generalization variety) employing either of the notions of relative frequency or conditional probability will have to deal with them. So, it is not so alarming that some version the generalization strategy arising from Carroll's suggestion will have to deal with them as well. Suffice it to say that there is some promising version of the generalization strategy. So, in addition to the fact that we have reason to prefer the generalization strategy to the Humean theory, it is possible to analyze general causal claims in terms of singular causal claims.

PART III: CAUSAL PRIORITY AT THE SINGULAR LEVEL

Chapter Six: Causal Priority at the Level of Singular Causal Claims

The generalization strategist attempts to analyze general causal claims using singular causal claims. She then attempts to provide an independent analysis of singular causal claims. In **Part II** I argue that this strategy is a live option for explaining the relationship between singular and general causal claims. Also in **Part II** I contend that the generalization strategist must take into account two facts about the relationship between singular and general causal claims:

1) *Context Relativity*: Singular and general causal claims are not relative to context in the same way, and

2) *Occurrence*: For a singular causal claim to be true, the cause and effect must occur. But, according to Eells, our intuitions suggest that the cause and effect need not occur for a general causal claim to be true.

To deal with context relativity she must analyze causal claims relative to context and relative to the significance of the role a cause plays in bringing about its effect. With respect to occurrence she must explain why the requirement that cause and effect occur disappears at the general level, if general causal claims are reducible to singular causal claims.

Carroll proposes that "C causes E" is true if and only if the frequency of causers of e's in c's is high, where what counts as high is relative to context. He analyzes general causal claims using a relative frequency notion of probability and maintains that general causal claims, like singular ones, must satisfy the occurrence requirement. That is, for a general causal claim to be true, the cause and effect must occur. I argue that, by altering his theory to take into account the significance of the causal role a cause plays in bringing about its effect, his version of the generalization strategy is best able to address both issues.

If singular causation were reducible to general causation, causal priority would be easily carried from the general to singular level. On such Humean accounts singular causal claims are instantiations of general causal claims. So, the causal priority of a singular causal claim is explained by the causal priority of the general causal claim that governs (instantiates) it. On the other had, if a general causal claim is best described as a generalization over some group of singular claims, causal priority at the general level does not in any immediately obvious way indicate causal priority at the singular level. For instance, there may be a correlation between two phenomena types, A and B, though most a's cause b's, and a few b's cause a's.

By adopting a version of the generalization strategy to explain the relationship between singular and general causal claims, rather than a Humean strategy, it becomes more difficult for the fork asymmetry theorist to explain causal priority for singular causal claims. The fork asymmetry strategist argues that a common cause can be a screener-off, but a common effect cannot. This is how she distinguishes between cause and effect. She maintains that it is a matter of contingent fact that screeners-off tend to occur earlier than their common effects. In this way she explains causal direction. But, screeners-off occur only at the type level. By definition, a causal phenomena token cannot be a screener-off. If some version of the generalization strategy is correct the fork asymmetry explains causal priority only at the general level. So, fork asymmetry theorists are left with the burden of explaining causal priority at the singular level.

I argued earlier that on the best available version of the generalization strategy (roughly, Carroll's version) general causal claims are true if and only if the frequency of causers of a's in s's is high, where what counts as high is relative to context. So, general causal claims are generalizations over causal instances. But, "significant" does not mean "with a frequency greater than .5". In fact, we have seen that the frequency of causers of e's among c's can be significantly less than .5. So, even though Carroll's strategy employs a frequency notion of probability we cannot maintain that what is true of a general causal claim is true of all or even most singular causal claims over which it generalizes.

Perhaps this calls for an independent account of causal priority at each level. I think, however, this move is undesirable. If general and singular causal claims are connected via the reducibility of general causal claims to singular causal claims, it makes sense that the mechanism generating the distinguishability and direction of causation at the general and singular levels are connected, as well. Perhaps causal priority somehow 'filters down' from the general to the singular level.[34] I examine the latter suggestion here.

One candidate for explaining this filtering process is Daniel Hausman's notion of *causal connection*. It is engaging because it is closely related to the fork asymmetry principle. Daniel Hausman's account of causal priority presents a fork asymmetry-like explanation of causal priority and takes singular causal claims as basic. Though his account is problematic, this feature of his theory is interesting and sheds light on how we might explain this filtering process.

6:1: Causal Connections[35]

Hausman attempts to distinguish between cause and effect (thus, also explaining why causes typically precede their effects) in a more substantive manner than Hume by taking the notion of *causal connection* as primitive:

> "I shall offer explicit truth conditions for "x is causally prior to y" in terms of a primitive relation written here "cc." "x cc y" should be read "x and y are connected or related by 'some fact of causation'"... More briefly, one can read "x cc y" as "x is causally connected to y," but this second reading may be misleading. Effects of a common cause are connected by some fact of causation, although one might hesitate in ordinary speech to say that they are causally connected. A lightning flash and a thunderclap may be connected by some fact of causation." (Hausman 1984, p. 262)

[34]This idea was suggested by Hartry Field in class discussion, Spring 1996.

[35]Hausman uses capital letters to refer to cause and effect *tokens*. Since I have been using capital letters to refer to *types* throughout my discussion, I substitute lowercase letters for capital letters in quotes from Hausman where he refers to token cause and effect, as well as in my discussion of his work.

Reading "x cc y" as "x is causally connected to y" tends to be misleading because Hausman maintains that effects of a common cause are causally connected. Others do not. Mackie, for instance, sometimes speaks of two causal phenomena, x and y, as causally connected if and only if either x causes y or y causes x. According to Mackie, effects of a common cause are not causally connected:

> "*A* and *B* are too alike for their resemblance to be coincidental; so there is *some* causal connection between them. Therefore either *A* is causally prior to *B* or *B* is causally prior to *A*." (Mackie 1974, p. 176)

Whereas, Hausman maintains that "x is causally connected to y" or "x and y are connected by 'some fact of causation'" if and only if either (1) x causes y or (2) y causes x or (3) x and y are effects of a common cause.[36] (Hausman 1984, p. 263) Hausman refers to this as the principle (CC). He maintains that the causal connection relation (which he calls, cc) differs from the causal relation. While the causal relation is antisymmetric and transitive, cc is symmetric and not transitive. (Hausman 1984, p. 262) That is, whenever x cc y, y cc x. Further, x cc y and y cc z does not imply that x cc z.

It is not clear what Hausman takes being a 'primitive relation' to involve. On his account, cc is a non-logical, relational predicate that he takes as primitive to explain causal priority. We normally consider a primitive relation to be one that

[36]According to Hausman, "x causes y" means that "x is a cause or a causal condition of y". (Hausman 1984, p. 263).

needs no further analysis. Hausman insists that (CC) is neither an analysis nor a definition of the notion of causal connection:

> "It is important to note that *(CC) is not a definition of "x cc y"*. "Causal connection" is a primitive which observations of dependencies or correlations help us to grasp. (CC) instead expresses a plausible conviction which leads to truth conditions for "x causes y."'" (Hausman 1984, p. 263)

Yet, he admits that an analysis of cc must eventually be given that is independent of the notion of causation:

> "A full theory of causation would require an analysis of causal connection (connection by "some fact of causation") as well as an explicit consideration of the ontology of events. The account of causal priority given here breaks down if no relatively independent analysis of causal connection and no compatible theory of events can be given." (Hausman 1984, p. 275)

Since cc is ultimately used to determine what it means to say that "x is a cause of y", a non-circular analysis of cc should not make reference to causation. Perhaps Hausman takes cc as primitive for the purposes of his paper, but recognizes the need, ultimately, for reduction.

Hausman attempts to explain causal priority for causal instances by using causal antisymmetry, causal transitivity and (CC) to derive truth conditions for "x causes y". He argues that causal antisymmetry, causal transitivity and (CC) jointly entail the following necessary condition of "x causes y":

> **Necessary Condition (N):** If x causes y, then x cc y and everything causally connected to x is causally connected to y. (Hausman 1984, p. 265)

To get the asymmetry he needs to explain causal priority, Hausman defends a sufficient condition for "x causes y" that does *not* follow from the conjunction of causal antisymmetry, causal transitivity and (CC):

> **Sufficient Condition or *Independence Among Causes* (I):** If x does not cause y, something is causally connected to x and not causally connected to y, or x is not causally connected to y.[37] (Hausman 1984, p. 266)

Hausman confesses that (I) "appears rather dubious and arbitrary." (Hausman 1984, p. 265) But, he argues, we have good reason to believe that it is true. Consider his example. Suppose the universe consists of a single ameba that splits at time, t, and the splitting at t (x) is the sole cause of the existence of two separate amoebae at t' (y). (Hausman 1984, p. 267) This appears to be a counter-example to (I), since everything causally connected to y is causally connected to x and *vice versa* though clearly x causes y. But, as Hausman says, a possible solution is to say that the absence of an ameba-eater at t' is causally connected to y, but not to x. And, this general approach can be used to solve other counter-

---

[37]Below are two other ways Hausman formulates of (I):

1. If x cc y and everything causally connected to x is causally connected to y, then x causes y. (Hausman 1984, p. 264)

2. Assuming that x cc y, if y causes x or x and y have a common cause only, something is causally connected to x (i.e., a cause of x) and not causally connected to y. (Hausman 1984, p. 265) [Remember that on Hausman's account causal conditions are causes.]

examples, as well. Hausman rejects this option, however, because introducing the non-occurrence of events trivializes (I).

We should instead, Hausman contends, consider the persistence of certain properties through a causal episode:

> "In the amoeba case, for example, the persistence of the polypeptide bonds within various proteins during the time interval from t to t' is a causal condition of the existence of the separate amoebae at t' (y) which is not causally connected to the splitting at t (x)." (Hausman 1984, p. 268)

When we look for the persistence of certain properties through a causal episode, rather than some non-occurrence, we find that there is something causally connected to the effect that is not causally connected to the cause:

From (N) and (I) Hausman generates the following account of causal priority:

(CP) x causes y if and only if

(1) x cc y,

(2) everything causally connected to x is causally connected to y, and

(3) something is causally connected to y, but not to x.

(Hausman 1984, p. 265)

According to Hausman, we have three good reasons for accepting (CP). First, since (CP) is deducible from causal antisymmetry, causal transitivity, (CC) and (I), it is at least as credible as they are. (Hausman 1984, p. 268) Second, if we use (CP), causal priority is observable through experimentation. (Hausman 1984, p.

268) Finally, it allows us to explain why other theories of causal priority are partially successful. (Hausman 1984, p. 268)

Unlike the causal relation, cc is nontransitive. That is, if x cc y and y cc z, it is not always the case that x cc z. To clarify the relation between cc and causation, consider Hausman's example. (Hausman 1984, p. 265) Ellen desires light. So, she flips on the nearest light switch. But, the insulation is frayed and there is a short circuit. Ellen's flipping the switch *causes*, and, therefore, *is causally connected* to the short circuit. Since Ellen's desire causes her to flip the switch, her desire is causally connected to her flipping the switch. Although cc is not transitive, the causal relation is. So, by transitivity of causation Ellen's desire causes the short circuit. Thus, Ellen's desire for light is causally connected to the short circuit.

While the frayed insulation is a cause of the short circuit on Hausman's account (thus, causally connected to the short circuit), the frayed insulation is not causally connected to Ellen's flipping the switch. That is, it is not the case that flipping the switch causes the frayed insulation or that frayed insulation causes Ellen to flip the switch or that flipping the switch and the frayed insulation have a common cause.

According to Hausman, (CP) provides a means of distinguishing between cause and effect without stipulating that the cause is the earlier one by showing that causes are partially independent. Further, (CP) distinguishes between cause and effect despite the fact that cc is symmetric. Hausman argues that asymmetry of causation is generated by the combination of (2) and (3); it follows from (2) and (3) that effects are connected, while causes are not fully connected. Thus, everything causally connected to the cause is causally connected to the effect. But, *not everything* causally connected to the effect is causally connected to the cause. In this sense the cause is *independent*, while the effect is not. This *independence of causes* provides the asymmetry needed to explain causal priority.

6:2: A Minor Adjustment to (CP)

Judith Thomson argues that "there really is no such thing as the construction of a reasoned analysis of "X causes Y" without the concurrent, or prior, construction of (a) a reasoned account of what principles are to be taken to govern "X causes Y" (where X and Y are event types) and (b) a reasoned account of what events there are." (Thomson 1984, p. 250) According to Thomson, since (a) and (b) provide adequacy conditions for an analysis of causation, Hausman is mistaken in giving an analysis of "X causes Y" in terms of "X cc Y" leaving (a) and (b) to be settled subsequently. I contend that an analysis of "A causes B" can be given without a prior or concurrent theory of events. Such analyses should be viewed, in part, as providing a framework for some theory of events; in light of cc, there are certain restrictions with which an acceptable theory of events must comply. Hausman's theory only suffers from taking this approach if a theory of events cannot be found which is compatible with cc.

David Lewis echoes this point in "Events":

> "Events are not much of a topic in their own right. They earn their keep in the discussion of other topics...In this paper I shall consider what sort of theory of events I need to go with my theses about causation. If none could be found, that would be reason to reject what I say about causation. But I think a suitable theory can be found -- or at least sketched -- and I think it is a reasonably attractive theory in its own right. What other purposes it might serve, if any, I cannot say."
> (Lewis, Philosophical Papers: Vol. II, "Events," p241-3)

While it is true that Hausman must ultimately find a compatible theory of events to make his analysis work, it is not, in principle, necessary that he determine an adequate theory of events prior to or concurrent with an analysis of "x causes y".

Although Thomson's criticisms above do not have the consequences for Hausman that she thinks they do, she raises an independent point which does suggest the need for revision of Hausman's account. She points out that the conjunction of (C), (CC), (N), and (I) with "A causes B" entails the following contradiction[38]:

1. A causes B	
2. B cc A	from (1) and (cc)
3. B cc B	from (1), (2) & (N)

[38](C) is the claim that causation is transitive and antisymmetric. Also, recall that by "events" Hausman means "individual events or states of affairs". (Hausman 1984, p. 261)

4. (Z cc B ⊃ Z cc B) tautology

5. B causes B from (3), (4) & (I)

6. - (B causes B) from (C)

(Thomson 1984, p. 250)

To repair Hausman's theory we must drop one of the following two claims:

1. a phenomenon can be causally connected to itself, or

2. a phenomenon can cause itself.

So, we must either revise (N):

(N*): If X causes Y, then everything causally connected to X, other than Y, is causally connected to Y. (Field, discussion, June 1995)

Or revise (I):

(I*): If X cc Y and everything causally connected to X is causally connected to Y, and if also X ≠Y, then X causes Y.

Thomson correctly argues that Hausman must revise (I) rather than (N). According to Thomson this is because Hausman's proof of (N) does not go through unless an event can be taken to be causally connected to itself. Hausman's proof of (N) is the following:

"Assume x causes y. Then if (case a) z causes x, z causes y by (C). If (case b) x causes z, then y and z have a common cause. Finally, if (case c) some w causes both z and x, then w causes y and y and z have a common cause. In each case z cc y." (Hausman 1984, p. 264)

In (case c) if we take z to be identical to x, we could still allow that z cc x and z cc y. On the other hand, we do not want to allow that a phenomenon can cause itself. This is inconsistent with our assumption that causation is antisymmetric. Yet, it is reasonable to suppose that "if an event B has a cause at all, then it does no harm to allow that B and itself have a common cause". (Where B is an event token.) (Thomson 1984, p. 250) Given (I*), a phenomenon cannot cause itself, though it can be causally connected to itself. Thus, we should revise (I) rather than (N).

6:3: The Observability of Causal Priority

According to Hausman, one of the advantages of taking the notion of causal connection as primitive is that cc, in conjunction with causal transitivity and causal antisymmetry, renders causal priority *objective* and *detectable through experimentation*:

> "(C) and (CC) taken together form an important fragment of a theory of causal priority. They show that causal claims have substantial testable implications concerning the existence or non-existence of causal connections among various events." (Hausman 1984, p. 263)

Our understanding of causal connections, Hausman claims, comes from observations of dependencies or correlations.

Recall (CC); "x cc y if and only if x causes y or y causes x or x and y are effects of a common cause". Hausman makes reference to "causes" in his characterization of cc. Thus, (CC) runs the risk of circularity. An analysis of cc should not refer to "cause" if it will be used to distinguish between cause and effect. Hausman admits that he is unsure "whether our understanding of causal connections is completely independent of all judgments of causal ordering". (Hausman 1984, p. 277) Since (CC) is neither an analysis nor definition of "causal connection", this may not be a serious problem for Hausman. But, he also expresses the need for a complete theory of causation to provide an analysis of cc. So, eventually Hausman must produce an analysis of cc that does not make reference to causation.

He maintains, nevertheless, that observed dependencies or correlations are reliable *indicators* of causal connections at the token level: "There will generally be some observed stochastic dependence or correlation between X and Y (which may well be spurious), if and only if x cc y." [Where x is an instance of X and y an instance of Y.][39] (Hausman 1984, p. 262) He admits this evidence is not infallible. There is a chance that observed correlations or dependencies may be spurious:

> "Observed dependencies or correlations are not conclusive proof that x cc y, since the dependency or correlation may result from chance or from the peculiarities of the particular sample. By chance or error or ignorance, one may also find no dependency or correlation between X and Y, even though x cc y." (Hausman 1984, p. 262)

Though fallible, Hausman claims that observed dependencies and correlations give us a somewhat dependable sign that two phenomena are causally connected. Such

[39]Hausman leaves unanswered the question of how such dependence at the type level implies the same dependence at the token level. I discuss this weakness in his account shortly.

indicators of causal connections do not require any knowledge of causal ordering. So, according to Hausman, we can obtain some understanding of when causal connections are present without knowledge of causal ordering:

> "One can in this way understand what a causal connections is and know when causal connections obtain without first understanding and being able to judge the causal ordering." (Hausman 1984, p. 262)

Thus, to understand what a causal connection is and when it obtains we need not make judgments about causal ordering. We must, however, make judgments about other causal connections:

> "In looking for correlations or dependencies, one has to decide which samples to study and which results to take seriously. In making these decisions, one must rely on judgments concerning other causal connections, if not concerning causal ordering as well. But the fact that judgments concerning causal connections depend on prior knowledge or presumptions concerning other causal connections does not render such judgments untestable or nonobjective." (Hausman 1984, p. 277)

So, Hausman argues that this circularity is unproblematic.

The procedure for determining whether "x causes y" is true is, roughly, as follows. First, find a type, X, of which x is a token. Find a type, Y, of which y is a token. (Let's not talk yet about how we pick them.) Suppose X and Y co-occur with a frequency greater than mere coincidence allows. According to Hausman, this indicates some (either positive or negative) causal connection between x and y. Suppose we find another phenomenon, Z, that has some causal connection with Y (i.e., Z co-occurs with Y with a greater frequency than mere coincidence allows), and Z has no causal connection with X (i.e., Z does not co-occur with X more than coincidence allows). This indicates, Hausman insists, that y does not cause x. Next, we determine whether x causes y or x and y are effects of a common cause by "controlling for possible common causes and examining the partial correlation between X and Y". (Hausman 1984, p. 269)

Though this information gathering process is helpful in determining the truth of general causal claims, it is not, on its own, useful for determining the truth of singular causal claims. Hausman jumps from the determination that 'X causes Y' is true to the determination that 'x causes y' is true without any explanation of why we should accept that this entailment holds. In arguing that there will usually be an observed dependence between X and Y if and only if x cc y, he patently assumes that the truth of a general causal claim implies the truth of *all* the singular causal claims that fall under it. That is, he assumes that if X causes Y, then for any x and y that co-occur, x causes y. But, he never says precisely what the relationship between singular and general causal claims is. Whether "X causes Y" indicates that every claim of the form "x causes y" is true depends on how we describe the relationship between singular and general causal claims.

6:4: The Filtering Process

To utilize Hausman's account as a possible means of explaining how causal priority filters from the general to the singular level in the fork asymmetry based account of causal priority presented here, the following must be compatible:

1. *fork asymmetry principle*: If A and B are highly correlated and one does not cause the other, there is typically some C that screens-off their correlation. A and B rarely co-occur. C typically occurs prior to A and B.

2. *generalization strategy*: General causal claims are true if and only if the frequency of causers of e's among c's is high. What counts as high is relative to context and to the significance of the causal role C plays in bringing about E.

3. *(CP)*: x causes y if and only if --

   a. x cc y

   b. everything causally connected to y is causally connected to x

   c. something is causally connected to y but not to x.

4. Horwich's explanation of our understanding of causal priority.

I contend that (1), (3) and (4) are compatible. Horwich argues that our understanding of causal priority emanates from certain observable evidence we encounter. Among this evidence is the claim that correlated phenomena are usually causally connected. This fact, in conjunction with other contingent facts about the world like "choices are made for the sake of what they may cause" and "the causes of a present phenomenon are knowable but its effects are not" govern our understanding of causal priority. Such facts, according to Horwich, do not constitute a description of the *nature* of causation. Still, they play a significant role in establishing our belief that causes precede their effects; they provide a reliable, though fallible, guide for *identifying instances* of causation.

The fork asymmetry principle is another contingent fact about the world that shapes our understanding of causal priority. Nevertheless, it is of a different sort than the other contingent facts Horwich mentions. The fork asymmetry principle plays a more significant role in the sense that it drives a metaphysical reduction of causation to correlations. Roughly, "cause" reduces to "screener-off" like "water" reduces to "H_2O".

Horwich claims that we come to understand "causation" in the same manner that we come to understand notions like "water". We identify water using certain contingent facts (e.g., water is colorless, odorless, liquid, etc.). Analogously, we use the contingent facts mentioned above to identify cases in which causation obtains. We give a metaphysical reduction of water o H_2O. We give a metaphysical reduction of causes to screeners-off via the fork asymmetry principle; causes can be screeners-off, effects cannot. By giving a metaphysical, rather than conceptual, reduction we need only *a posteriori* evidence that the contingent facts in question pick out instances of water (or, causation). We give a definition of water (i.e., H_2O) that characterizes the essence of water, though it does not capture all it features (e.g., water is "colorless" and "odorless"). In the same way, the fork asymmetry principle characterizes the essence of "cause" without capturing all the features of causation (e.g., we can "manipulate" causes).

Horwich's description of how we come to understand the notion of causal priority is similar to Hausman's. This stems from their mutual reliance on the fact that usually two phenomena are correlated, if and only if, either one causes the other or they are both effects of a common cause. This idea is closely related to both the fork asymmetry principle and Hausman's notion of causal connection.

The fork asymmetry principle indicates that if two phenomena are causally connected, one does not cause the other, and a third phenomenon screeners-off their correlation, that screener-off is a common cause of the two original phenomena. We observe causal connections by observing correlations among phenomena types.

The problem is to show how the combination of (1), (3) and (4) is compatible with (2). While there seems to be no elemental incompatibility between (1), (2) and (4), the framework of Hausman's theory as it stands renders (2) and (3) inconsistent. To be more precise, Hausman draws support for (CP) from observed correlations. But, to explain how we observe causal instances he postulates a principle I will call H*:

H* -- There is usually a correlation between X and Y if and only if x cc y.

The relationship between "X and Y are correlated" and "x and y are causally connected" is unclear. Perhaps the correlation between X and Y establishes the causal connection between x and y by exposing a causal connection between X and Y. Or, it does so independently of any causal connection between X and Y. It is hard to imagine how we could use the correlation to establish a causal connection between x and y except via a causal connection between X and Y. Thus, Hausman should say that "X and Y are correlated" indicates that "X and Y are causally connected". This, in turn, marks a causal connection between x and y. But, if we interpret H* in this manner, Hausman jumps from the determination that, for example, 'X causes Y' is true to the determination that 'x causes y' is true without any explanation of why we should accept this entailment.[40] Whether the truth of "X causes Y" indicates the truth of "x causes y" depends on how we describe the relationship between singular and general causal claims. Moreover, Hausman never specifies the relationship between singular and general causal claims.

---

[40]Or, he jumps from 'Y causes X' to 'y causes x' or from 'X and Y are effects of a common cause' to 'y causes x' or 'x and y are effects of a common cause', depending on how the causal connection pans out.

Perhaps Hausman endorses a regularity approach to analyzing causal claims. In that case the truth of a general causal claim implies the truth of *all* the singular causal claims that fall under it. This is clearly incompatible with the modified version of Carroll's generalization strategy given above (i.e., (2)). So, if we want (3) and (2) to be compatible we must drop H*. Hausman is correct that there is usually a correlation between X and Y if and only if x cc y. But, we can observe x cc y directly without first observing the correlation between X and Y. Suppose I observe myself sticking my hand into fire and getting burned. A Humean would say that what I observe is the co-occurrence of sticking my hand into the fire and my burnt hand. But, on my account, I observe an instance of causation. I observe that everything causally connected to my sticking my hand into the fire (e.g., my desire to see what will happen) is causally connected to my hand being burned. And, I observe at least one thing causally connected to my hand being burned that is not causally connected to my sticking my hand into the fire (e.g., who lit the fire). Once I drop H*, (1-4) are compatible and we have an explanation of causal priority at both the general (via the fork asymmetry) and singular level (via causal connection).

6:5: Conclusion

My project has been to:

A) Clarify what we mean by "causal priority",

B) Formulate a version of the fork asymmetry principle that best facilitates an explanation of causal priority,

C) Defend a slightly modified version of Horwich's fork asymmetry-based account against objections from Arntzenius and Price,

D) Endorse, roughly, Carroll's version of the generalization strategy, and

E) Show how we might explain causal priority for singular causal claims given the conjunction of (A) -(D).

I argue that by causal priority we mean causal distinguishability (i.e., causes and effects are distinguishable in a way that is independent of any temporal ordering) plus causal direction (i.e., effects rarely (if ever) precede their causes). Thus, there are two questions to be answered in giving a complete account of causal priority:

1) Why do causes precede their effects?

2) How do we distinguish between cause and effect?

That causes precede their effects is constitutive of the causal relation, according to the fork asymmetry theorist. She also insists that we distinguish between cause and effect via some contingent fact about the world (i.e., the fork asymmetry).

The best formulation of the fork asymmetry principle:

1. has a component to account for causal direction in a way that does more than merely stipulate that the common cause occurs before the joint effects,
2. provides a means of determining *when* there is a 'common cause or when is it reasonable to postulate one (and in doing so it must suggest that correlated phenomena *typically*, rather than *must*, have a common cause),
3. avoids type-token ambiguities,
4. does not depend on what kinds of entities count as causes, and,
5. it is made more appealing if it separates claims about causal distinguishability from claims about causal direction.

Thus, I arrive at the following formulation of the fork asymmetry principle: If two causal phenomena-types (e.g., events or type of events, or facts, or conditions, or, ...), A and B, are highly correlated and one does not cause the other, there is *typically* a third phenomena-type, C, that screens-off their correlation and A and B rarely (though it is possible) co-occur without C. And, this screener-off *typically* occurs prior to the correlated events.

I contend that a modified version of Horwich's account can meet objections posed by Arntzenius and Price. I replace Horwich's version of the fork asymmetry principle with what I take to be a stronger version (i.e., my Version C). I then add Carroll's version of the generalization strategy to Horwich's program. I also annex Hausman's scheme for explaining causal priority at the singular level. The resulting strategy explains causal priority at the general level, as well as explaining how causal priority filters from the general to singular level. Thus, the fork asymmetry-based approach to explaining causal priority is viable and can meet its strongest objections.

Bibliography

1. Arntzenius, Frank. (1990) "Physics and Common Causes", *Synthese* 82: 77-96.

2. ..., (1992) "The Common Cause Principle", *PSA*, Volume 2, pp. 227-237.

3. Bennett, Jonathan. (1988) *Events and Their Names*. Hackett Publishing Co.

4. Carroll, John. (1991) "Property-level Causation?", *Philosophical Studies* 63: 245-270.

5. Cartwright, Nancy. (1988) "Regular Associations and Singular Causes," in Brian Skyrms and William L. Harper (eds.), *Causation, Chance, and Credence, Vol. 1*, Kluwer Academic Publishers, 79-97.

6. ..., (1988), "Reply to Ellery Eells," in Brian Skyrms and William L. Harper (eds.), *Causation, Chance, and Credence, Vol. 1*, Kluwer Academic Publishers, 105-108.

7. Eells, Ellery. (1988) "Eliminating Singular Causes: Reply to Nancy Cartwright," in Brian Skyrms and William L. Harper (eds.), *Causation, Chance, and Credence, Vol. 1*, Kluwer Academic Publishers, 99-104.

8. Field, Hartry. (1995) Discussion.

9. Hausman, Daniel M. (1984) "Causal Priority," *NOUS* 18: 261-279.

10. Hitchcock, Christopher. (1995) "The Mishap at Reichenbach Falls", forthcoming in *Philosophical Studies*.

11. Horwich, Paul. (1987) *Asymmetries in Time*, MIT Press.

12. Lewis, David. (1986) *Philosophical Papers Volume II*, "Causation", Oxford University Press.

13. ..., (1986) *Philosophical Papers Volume II*, "Counterfactual Dependence and Time's Arrow", Oxford University Press.

Bibliography (continued)

14. Papineau, David. (1991) "Correlations and Causes," *British Journal of Philosophy of Science,* 42: 397- 412.

15. Price, Huw. (1992) "Agency and Causal Asymmetry", *Mind* 101: 501-20.

16. ..., (1992a) "The Direction of Causation: Ramsey's Ultimate Contingency", *PSA,* Vol. 2, pp. 253-267.

17. Reichenbach, Hans. (1956) *The Direction of Time,* ed. Maria Reichenbach, University of California Press.

18. Sober, Elliot. (1988) "The Principle of Common Cause," in *Probability and Causality,* James H. Fetzer (ed.). Dordrecht: Reidel, pp. 211-229.

19. Spirtes, Glymour, and Scheines. (1993) *Causation, Prediction and Search.* New York: Springer-Verlag.

20. Thomson, J.J. (1988) "Causal Priority", *British J Phil* p. 245-253.

IMAGE EVALUATION TEST TARGET (QA-3)

150mm

6"

APPLIED IMAGE, Inc
1653 East Main Street
Rochester, NY 14609 USA
Phone: 716/482-0300
Fax: 716/288-5989