

*Causality and Responsibility*¹

Glenn Shafer²

gshafer@andromeda.rutgers.edu

May 13, 2000

I. Introduction.....	2
II. Causality.....	4
1. The Structure of Possibility.....	4
2. Refinement	6
3. Action	8
4. Causal Relations Among Instantaneous Events	9
5. Merely Valid Prediction	11
6. The Meaning of Predictive Probability	13
III. Responsibility	15
7. Court, Defendant, and Nature.....	15
8. Action Defined by Intention.....	16
9. Approximate Cause	18
10. Toxic Torts	20
References.....	22

¹ Research for this article was supported by grant SES-9819116 from the National Science Foundation. The article grows out of my Sewall Wright Lectures at the Fields Institute for Research in the Mathematical Sciences, given in October and November of 1999. I am grateful to Steffen Lauritzen for arranging these lectures and to the Fields Institute for their hospitality. I have benefited from comments by Peter Gillett, by the participants in the Program in the Causal Interpretation and Identification of Conditional Independence Structures at the Fields Institute (Fall 1999), and by participants in the Symposium on Artificial Intelligence and Judicial Proof organized by Peter Tillers at the Cardozo School of Law (April 30, 2000).

² Glenn Shafer is at the Department of Accounting and Information Systems, Rutgers University, 180 University Avenue, Newark, New Jersey 07102. He is also Visiting Professor at the Department of Computer Science at Royal Holloway College, University of London.

I. Introduction

Attribution of responsibility, whether as blame or praise, always rests on a claim about causation. In order to be responsible for an event, a person must take an action that causes the event. This creates a philosophical puzzle. As philosophers, we tend to be skeptical about claims of causal knowledge. No matter how thoroughly we have studied a phenomenon, we say, we may not have gotten to the bottom of it and really understood its causal structure. But as judges and juries, we are prepared to conclude that a certain action has caused a certain result. How can we reconcile our philosophical and scientific skepticism with our practical credulity?

The most common response to this question emphasizes differences in standards of proof. Science, we are told, demand an extraordinarily high standard of proof, whereas deliberation about particulars, even in matters of life and death, never requires more than proof beyond a reasonable doubt. There is something to this, but it is not the whole story. Another aspect of the matter is the vast difference between understanding a whole causal structure and drawing a conclusion about one particular causal relation. It is eminently plausible that we might be practically certain about one aspect of a causal structure without understanding the structure down to its bottom.

This article discusses some ways we can know things about a causal structure without knowing everything about it. It is based on the predictive understanding of causality that I have advocated in *The Art of Causal Conjecture* (1996), in subsequent articles (especially Shafer 1998 and Shafer, Gillett, and Scherl 2000), and in my forthcoming book with Volodya Vovk, *Probability and Finance: It's Only a Game*. According to this predictive understanding, causal structure is the structure of the predictions that would be made by a superior intelligence, who witnesses everything that could conceivably be witnessed by a human-like witness and predicts everything that could conceivably be predicted by a human-like scientist. For brevity, I call this superior intelligence *Nature*. Like us and unlike God, Nature moves through time. Her knowledge increases as she sees what happens. Like us, Nature can predict some events in advance but cannot predict everything. In a given situation, she regards as possible some events that later turn out not to happen. She differs from us in two ways. First she is infallible; when she does make a prediction, it is verified by later experience. Second, she can make more predictions than we can; she can rule out some events we regard as possible.

Causality, in this way of looking at the matter, is an aspect of the structure of Nature's predictions. If Nature cannot predict that event B will happen before I take action A but can predict that B will happen if I do A, then we may say that my doing A causes B. The imperfection of our knowledge of causes is thus merely an aspect of the imperfection of our knowledge of Nature's predictions. Nature's knowledge and predictions represent, by my definition of *Nature*, the limiting totality of what all human-like agents could know and predict. We will never know the causal structure of a phenomenon fully, simply because our knowledge and prediction will never reach that limiting totality. But the knowledge we do gain of the world, to the extent it is valid, can be thought of as partial knowledge of Nature's predictive or causal structure.

The key point here is that Nature's superior knowledge allows her to rule out some things we regard as possible but does not, if our knowledge is valid, allow her to rule in things we know to be impossible. This means that we are on safe ground in attributing responsibility if we do so based on our knowledge of impossibilities. It is not surprising, therefore, that the classical legal concept of cause—necessary and sufficient cause—is defined in terms of impossibilities. According to this concept, an action *causes* an event if the event must happen (it is impossible for it not to happen) when the action is taken and cannot happen (it is impossible for it to happen) if the action is not taken. From a philosophical or scientific point of view, this is merely one of many concepts of cause (it is not quite the same, for example, as the concept of cause I offered in the preceding paragraph), but it is the one concept that is fully characterized by impossibilities and hence will not be invalidated as our knowledge of Nature increases. This may help explain why legal thought has clung to it so tightly.

An important contemporary challenge to legal thought is the extension of the concept of cause to situations that seem intrinsically probabilistic. This challenge appears most often in discussions of toxic torts: how can we hold a polluter responsible for impairing the health of many people when it appears to be a matter of chance whether damage will be done to any particular individual? This question is addressed by a growing literature, to which philosophers and computer scientists, as well as law professors and epidemiologists, are now contributing. In this article, I argue that it is best addressed by combining the classical concept of necessary and sufficient cause with a philosophy of probability that locates the empirical meaning of probability in the impossibility of successful gambling strategies. As we shall see, this argument supports the view that toxic torts should be class actions rather than individual actions.

II. Causality

Before looking closely at the relation between causality and responsibility, we must first explore in a moderate amount of detail how Nature's predictions can be represented mathematically. A simple representation that I find useful begins with the idea of an event tree—a tree in which the nodes are instantaneous events and possibility is represented by the links between them.

Much of what I have to say here covers ground already tilled in *The Art of Causal Conjecture* and some of my other articles. But the ideas in Sections 5 and 6, which are key for the discussion of responsibility in Part III, are likely to be novel for most readers.³

1. The Structure of Possibility

Figure 1 provides an example of an event tree, with the direction of time downward. We can imagine Nature moving down the tree as events unfold. At each node in the tree, the immediate steps downward represent all Nature's possibilities for what can happen next. (God may already know what will happen, but Nature does not.) A complete path down the tree is a *history*—one possibility for how events may evolve.




Figure 1. An event tree for what Rick may do after school. He may watch television right away, he may delay watching television, or he may even end up reading instead. We assume that the different paths down the tree represent all the ways in which the events shown can happen. At the outset, for some reason, Nature can rule out the possibility that Rick might call his mother and then pump up his bicycle tire afterwards.

³ The idea of a merely valid event tree, which is introduced in Section 5, is discussed in Section 13.6 of *The Art of Causal Conjecture*, but its importance for causal inference is not emphasized there.

A node E in an event tree has a dual meaning. On the one hand, E is an instantaneous event—something that happens at a particular instant. On the other hand, E is a situation that the world is in at a particular instant. The situation E is the situation that arises when the event E happens. The event E is the event that the situation E arises. The node F in Figure 1, for example, represents both the event that Rick calls his mother and the situation in which he does so.

Here are some additional aspects of the meaning of the situations in an event tree, which bear on their relation to the tree as a whole.

- In order to specify a situation fully, we must say how we got there. In other words, we regard situations that are similar but are preceded by different histories as different situations. This is why we assume that the situations form a tree rather than a graph in which paths can come together again after diverging. If we can get into a situation G after being in a situation E, then we cannot also get into G after being in a situation F that is not in the same history (path down the tree) as E.
- By the same token, the same situation cannot arise twice as the world evolves (moves down the tree). And hence the same instantaneous event cannot happen twice.
- On the other hand, a situation does not specify completely everything that has happened in the world. It is only as detailed in its meaning as the tree in which it is situated.

The interpretation of an event tree also involves a number of substantive assumptions, which are not necessarily made in all contexts where the words “situation” or “event” are used, and which should therefore be stated explicitly. In stating these assumptions, I speak of two nodes that are linked as “mother” and “daughter”; the daughter is below the mother.

1. A daughter cannot happen until after its mother happens (strictly after; not at the same time).
2. Two distinct daughters of a situation are mutually exclusive; they cannot both happen. In situation I in Figure 1, Rick may pump up his bicycle tire (G) or watch television right away (E_1), but he cannot do both.
3. The daughters of a situation (and even all paths down from the situation) are all possible in that situation, no matter what else may be said about what has happened. In situation G in Figure 1, where Rick has just pumped up his bicycle tire, it is possible that he will call his mother at the office and then read, no matter how much pressure he has put in the bicycle tire.

4. The daughters of a situation are exhaustive; once the world is in that situation, one of the daughters must happen. In situation I in Figure 1, if Rick does not pump up his bicycle tire (G), then he watches television right away (E_1). There is no third possibility.

This is a rather imposing list of assumptions. In many problems we would surely prefer a more flexible mathematical representation of causality, in which we could adopt some but not all of these assumptions. Shafer, Gillett and Scherl (2000) develop such a representation. The event-tree representation is adequate, however, for developing the ideas of this article.

An event tree is a succinct way of showing the predictions of a particular person who witnesses the unfolding of events. In each situation, it shows the possibilities the person envisions for the further unfolding of events. The person can be Nature, or it can be someone less knowledgeable (we will look at event trees for less knowledgeable persons in Section 5). When the person is Nature, I call the event tree *causal*.

2. Refinement

Can a picture as simple as Figure 1 be an accurate representation of the complex causal order of the world? This picture obviously tells about only a tiny part of what is going on even in the small corner of the world inhabited by Rick on the afternoon in question. Is it conceivable that what it says about this tiny part could still make sense when we look at a larger picture?

Yes, it is conceivable. A simple event tree is not necessarily falsified when we add more information. Figures 2 and 3 illustrate this point by showing event trees that result from adding and subtracting from the information in Figure 1. Figure 2 adds information about how much air Rick puts in his bicycle tire, and Figure 3 omits everything except for the point that Rick will end up either reading or watching television.

Shortly we will study event trees that disagree. But the event trees in Figures 1, 2, and 3 do not disagree. They merely provide different levels of detail. Figure 2 says more than Figure 1, and Figure 3 says less, but none of the three deny any assertion about possibility or impossibility made by one of the others. There is no inconsistency in supposing that all three of these trees are accurate descriptions of Nature.

Notice that the instantaneous event represented by a node in one event tree may be represented by a group of nodes in a different event tree. The instantaneous event that Rick calls his mother is represented by the single node F in Figure 1, but by two nodes, F_1 and F_2 , in Figure 2. The

instantaneous event that he watches television is represented by the single node E in Figure 3 but by the three nodes E_1 , E_2 , and E_3 in Figure 1.

An instantaneous event, as I will use the term, can happen at most once as the world evolves. So I will say that a set of nodes in an event tree represents an instantaneous event if and only if none of its elements precedes another on a path down the tree. I call a set of nodes satisfying this condition a *clade*. We should not think of an instantaneous event represented by a single node in a particular event tree as fundamentally different from an instantaneous event represented by larger clade. The difference lies in our choice of representation, not in the events themselves. Indeed, any instantaneous event represented in one event tree by a clade consisting of several nodes can equally well be represented in a different event tree by a single node, and vice-versa.

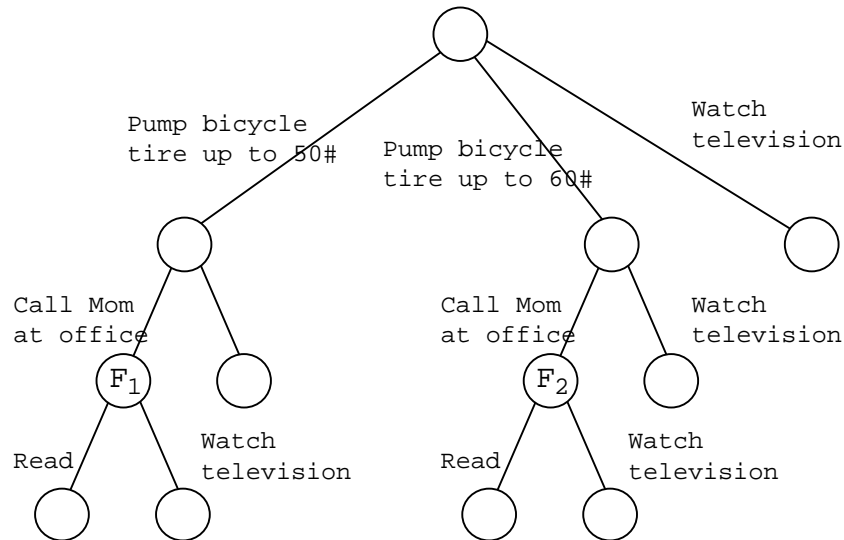


Figure 2. More or less air in the bicycle tire. This is a refinement of Figure 1. As a refinement, it respects all the statements about possibility and impossibility that Figure 1 makes. For example, Figure 1 says that in the situation where Rick pumps up his bicycle tire, it is possible that he will read later, and Figure 2 agrees; it tells us that no matter which pressure he puts in the tire, it is possible that he will read later.




Figure 3. Watching television as a single node. This event tree is refined by Figure 1.

It is entirely conceivable that the event tree Figure 2 is causal—i.e., that it is an event tree for Nature. In this case, the simpler event trees in Figures 1 and 3, because they are refined by the one in Figure 2, are also causal.

Does Nature have a single ultimately complicated event tree, which makes predictions about everything for all time, and which refines all the simpler causal event trees that we might consider? There is some didactic value in assuming that Nature does have such an ultimate tree, but this assumption is not really needed and may be unreasonable. As I have already remarked, one can make weaker forms of causal judgment than those made in an event tree, and it is not clear that we have reason to insist that Nature forgo these weaker judgments in favor of judgments that fit snugly together into an ultimate event tree. The representation developed by Shafer, Gillett, and Scherl (2000) avoids supposing that there is an ultimate tree; it allows an open system, which can be indefinitely expanded.

3. Action

An *action* is represented in an event tree as a step from one event to a later event. An action can be taken by a human being, an animal, or an inanimate agent, such as a storm or a meteor.

Figure 4 shows an example. Here Bill hits Joe, knocking out a tooth. This action is represented by the step from a situation where Joe has his teeth (shown as a square) to a situation where he is missing a tooth (shown as a circle). The circle can also be thought of as the instantaneous event that Joe loses a tooth; this instantaneous event is the *effect* resulting from the action.




Figure 4. An action that causes an effect.

The action causes the effect in the classical legal sense. The effect is a necessary result of the action, assuming the tree is correct, and the effect would not have happened had the action not been taken.

I hasten to acknowledge that blameworthiness involves more than cause and effect. People are not necessarily blamed for bad things they cause. Figure 5 uses a silly example to make the point.

Since birth is inevitably followed by death, we may say that when parents cause a birth they are also causing a death. But we do not blame them for the death. We blame someone for causing an effect only when the action they have taken transgresses certain principles or rules of behavior.




Figure 5. Causing birth and death. Having a baby causes both its birth and its death in the classical legal sense, because both birth and death are inevitable results of having the baby, and neither happen if the parents do not have the baby. Weaning the baby, although it is inevitably followed by the child’s death, does not cause the death in the classical legal sense, because the child will die in any case.

In an event tree, causes and effects are different kinds of objects. We may call them both events, but they are different kinds of events. Effects are instantaneous events—i.e., situations—in Nature’s tree. Causes are actions. An action is a transition—a change from one situation to a later situation. In *The Art of Causal Conjecture*, I called actions *Humean events*.

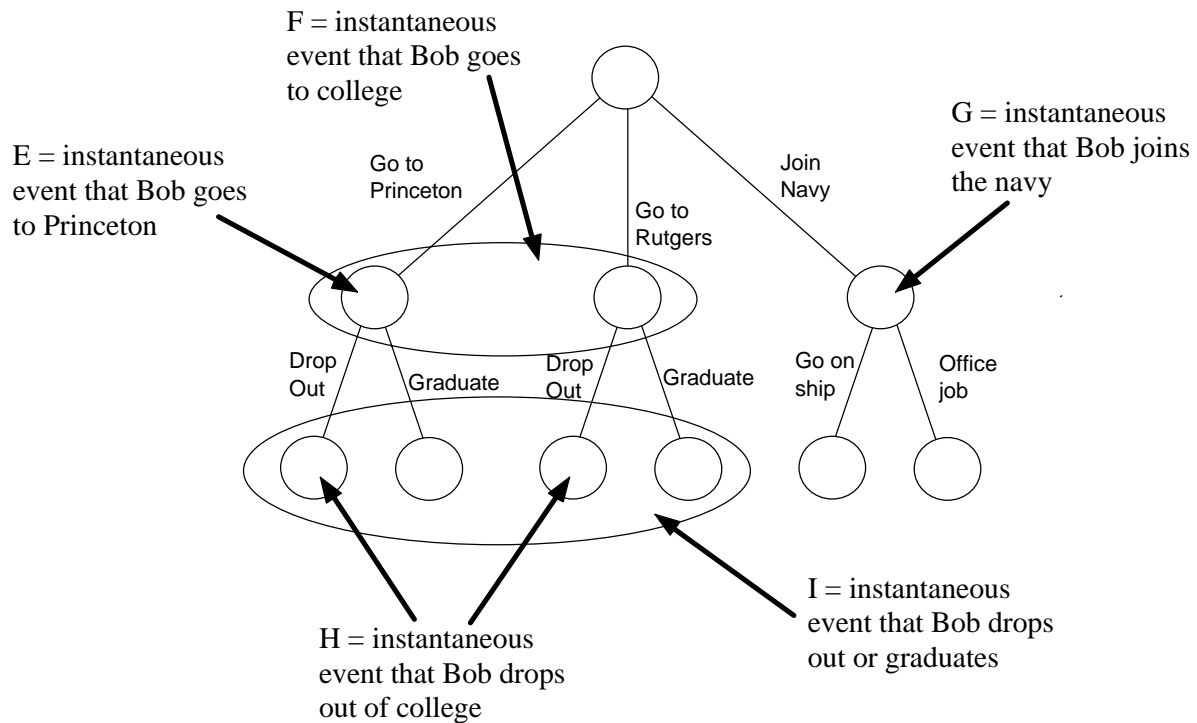
In Figures 4 and 5, the situations at the beginning and end of the action are represented by single nodes. But of course they could also be represented by clades consisting of several nodes. We will see several examples of this in Part II.

4. Causal Relations Among Instantaneous Events

The word “cause” is used in a wide variety of ways, most of them not very precise, and every scholar has their own opinion about how the word should or should not be made precise. We can safely predict that most of these opinions will conflict with the definition of cause I gave in the preceding section. So I would like to emphasize that I am not putting this definition forward as a way

of understanding causality in general. I present it only as a way of understanding the very special classical legal concept of cause: an action that is necessary and sufficient for an effect.

In my explanation of the classical legal meaning of “A causes B,” A is an action and B is an instantaneous event. Many philosophers prefer to say “A causes B” when A and B are both instantaneous events. But as I argued in Shafer (1998), nothing seems to be gained by giving the name “cause” to any particular relation between instantaneous events. It is more profitable simply to catalog the extensive variety of causal relations among instantaneous events, relations that capture various aspects of the polymorphous meaning of “cause” in ordinary language. Figure 6 lists a few of these relations.



<i>E allows H.</i>	No matter how E happens, H can happen later.
<i>H requires F.</i>	No matter how H happens, F has already happened.
<i>F precedes H.</i>	F allows H and H requires F.
<i>E foretells I.</i>	When E happens, I's later happening is inevitable.
<i>F always foretells I.</i>	F foretells I and I requires F.

Figure 6. Some causal relations among instantaneous events. These are a few of the many relations studied in Shafer (1998).

5. Merely Valid Prediction

We now turn to an issue that is crucial for causal inference: the relation between Nature’s predictions and predictions that are merely valid. As I explained in the introduction, we may be able to make valid predictions about a topic even though we know much less about it than Nature. We predict that A or B will happen, ruling out C. Nature goes farther, by ruling out B and predicting that A will happen. Then A happens. We predicted less than Nature did, but our prediction was correct.

Figure 7 illustrates how this idea plays out in terms of event trees. In this example, I imagine that I have learned from experience that my cow Bessie gives more milk, on average, when I feed her silage than when I feed her hay. On a day when I feed her silage, I get either 5 gallons or 7 gallons of milk—I cannot predict which. On a day when I feed her hay, I get either 3 gallons or 5 gallons—again, I cannot predict which. Nature, on the other hand, knows more than me. In the morning, before I decide whether to feed silage or hay, Nature observes events that tell her whether or not it is going to be a 5-gallon day. If so, Bessie will give 5 gallons of milk no matter how I feed her. Otherwise, she will give 7 gallons or 3 gallons, depending on whether I feed her silage or hay.




Figure 7. Glenn’s valid but not causal event tree. When Glenn is in situation S, he can rule out 3 gallons and predict that Bessie will give either 5 or 7 gallons of milk. Nature, meanwhile, is in S₁ or S₂, where she can predict exactly how much milk Bessie will give.

Let us call an event tree *valid* if its predictions are always correct. A causal event tree (an event tree for Nature) is valid, but as we have just seen, an event tree can be valid without being causal. So there are three distinct types of event trees, as listed in Table 1: valid, causal, and simply wrong. An

event tree that is wrong will eventually make a prediction that is falsified by experience. Event trees that are merely valid will not be falsified, but they will say some things are possible that Nature has ruled out in advance.

<i>Causal</i>	<i>Valid</i>	<i>Wrong</i>
Nature agrees when I say something is possible and also when I say it is impossible.	Nature knows more than me, but I may never find it out.	I will find out that my tree was wrong.
I say E is impossible \Rightarrow E is impossible	I say E is impossible \Rightarrow E is impossible	
I say E is possible \Rightarrow E is possible		

Table 1. Three kinds of events trees. Merely valid event trees fall short of being causal because in at least one situation they declare possible an event that Nature knows is impossible.

The distinction between causal and merely valid event trees is important in causal inference, because validity is the most we can hope for in general for event trees constructed from our experience. If our predictions are wrong, we can hope to learn this from experience. But if they are merely valid, because Nature knows something more that enables her to predict more, we may or may not find this out down the road. It follows that the only causal relations that we can hope to infer from our experience are those that can be inferred from merely valid event trees—i.e., those that depend only on impossibilities.

As I pointed out in the introduction, the classical legal concept of cause meets this condition.

An action A *causes* and event E in this sense when

1. A makes E inevitable, and
2. if A is not taken, E is impossible.

These are both statements of impossibility. The first statement says that in the situation resulting from A, the failure of E is impossible. The second statement says that in a situation resulting from any alternative to A, E is impossible. So if A appears to cause E in the classical sense in an event tree that is merely valid, then A causes E in the sense of the definition. This causal inference can be refuted only by showing that our event tree is outright wrong, not merely by showing that there is a deeper causal structure of which we were unaware.

6. The Meaning of Predictive Probability

In addition to teaching us what is possible and what is impossible, experience also teaches us what is probable and what is improbable. We often incorporate this additional information in our predictions. We might elaborate Figure 1, for example, by adding probabilities to the branches of the tree, as in Figure 8.

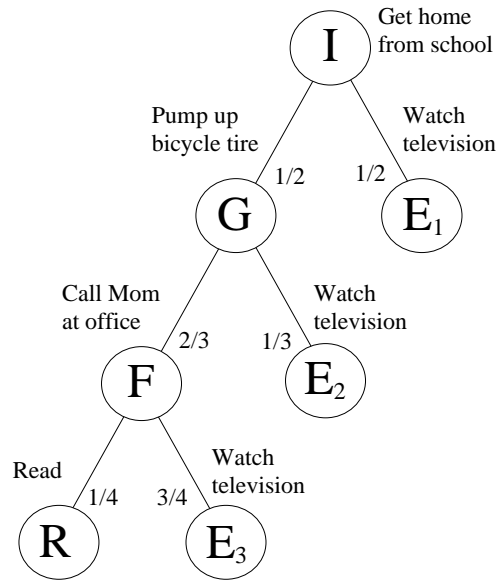


Figure 8. Probabilities for what Rick will do.

Suppose Figure 8 is merely valid rather than causal. Suppose, for example, that when I am in situation G, Nature can sometimes rule out the possibility that Rick will watch television right away (E_2) rather than calling his mother (F). Then Nature evidently disagrees with the probabilities given by the tree; for her, the probability for E_2 in G is zero, not $1/3$.

What causal meaning, then, can be given to valid probabilistic prediction? If our experience teaches us certain probabilities, and these probabilities are borne out in experience, then don't they tell us something about the true causal structure? Can they be simply wrong from the deeper viewpoint of Nature? Can they be completely uninformative about Nature?

One way of answering this question is to suppose that the probabilities in Nature's correct tree (the causal tree) are somehow arranged so that they are consistent with our probabilities for what we can observe. Figure 9 shows how this might happen in the case of Bessie. (For additional examples, see Section 13.6 of *The Art of Causal Conjecture*.) When I wrote *The Art of Causal Conjecture*, this was the only way of answering the question that I could see, but I was troubled by its inadequacies.

One inadequacy is that it seems to require that Nature have a complete probability distribution for her event tree (she cannot, for example, acknowledge my freedom of will by refusing to give probabilities for how I will feed Bessie). Another is that it does not make room for the seemingly radical differences between shallower and deeper levels of causal description—for example, the difference between classical and quantum mechanics (see the discussion of “multiple Natures” in Section 1.2 of *The Art of Causal Conjecture*).




Figure 9. Probabilities for Bessie. Glenn and Nature agree that Bessie gives 5 gallons half the time, no matter whether she is fed silage or hay, and that she gives 7 gallons half the time she is fed silage, and 3 gallons half the time she is fed hay.

In more recent years, I have come to believe that we must look deeper into the philosophy of probability in order to account for how radically probabilities can change when we move from a merely empirically valid description of the world to a deeper causal description, or even from one description that we are willing to label as causal to a deeper one. A key insight, I believe, is a point I have learned from Volodya Vovk: probabilities not close to zero or one have no meaning in isolation. They have meaning only in a system, and their meaning derives from the impossibility of a successful gambling strategy—the probability close to one that no one can make a substantial amount of money betting at the odds given by the probabilities.

We validate a system of probabilities empirically by performing statistical tests. Each such test checks whether observations have some overall property that the system says they are practically certain to have. It checks, in other words, on whether observations diverge from the probabilistic model in a way that the model says is practically (approximately) impossible. In *Probability and*

Finance: It's Only a Game, Vovk and I argue that both the applications of probability and the classical limit theorems (the law of large numbers, the central limit theorem, etc.) can be most clearly understood and most elegantly explained if we treat these asserted practical impossibilities as the basic meaning of a probabilistic or statistical model, from which all other mathematical and practical conclusions are to be derived. (See also Dawid 1985, Shafer 1990, Vovk 1993, and Dawid and Vovk 1997.) I cannot go further into the argument of the book here, but I do want to emphasize one of its consequences: because the empirical validity of a system of probabilities involves only the approximate impossibilities it implies, it is only these approximate impossibilities that we should expect to see preserved in a deeper causal structure. Other probabilities, those not close to zero or one, may not be preserved and hence cannot claim the causal status.

We will return shortly to the implications of this insight for legal reasoning.

III. Responsibility

As I acknowledged in the introduction, blame or praise is based on more than judgments about causation. In order to blame or praise someone for an action, we must also place that action in the context of norms and expectations: Do we expect and demand that the person behave as he or she did, or, on the contrary, did the person's actions violate norms or even laws? In many cases, we must also understand the person's own thinking: What did the person mean to do? These issues affect even the causal questions we ask.

7. *Court, Defendant, and Nature*

It is wise to begin by reminding ourselves that several points of view are in play when we talk about attributing responsibility. The three most prominent are:

1. The person who makes the attribution. Anyone can play this role, but in a society with laws it is played most authoritatively by courts—judges and juries.
2. The person to whom responsibility is attributed. Because we must assess the person's intention, we must ask how he or she thinks the world works.
3. Nature. We must understand what effect the person's actions really had.

To fix ideas, we may suppose that each of these three characters has their own event tree for the events under discussion, and perhaps their own subjective probabilities within this tree. The subjective

probabilities of Nature are, by definition, really objective probabilities—probabilities that are validated by empirical tests.

We need not enter into a detailed discussion of the complexities introduced by this multiplicity in viewpoints, but it may be wise to make a few distinctions.

1. When we speak of probabilistic standards of proof (e.g., more likely than not for civil cases, beyond a reasonable doubt for criminal cases), we are talking about the reasoned beliefs of the court. We are not talking about causal probabilities (Nature's probabilities).
2. Intention involves the beliefs of the accused, but we can demand that those beliefs should be reasonable, and in particular we can demand that they take account of generally understood causal relations.
3. According to the argument of Section 6, we are never on safe ground in claiming a causal status for intermediate probabilities—those not close to zero or one. This is true whether we are engaged in civil or criminal litigation.

8. Action Defined by Intention

I have already pointed out that the classical legal concept of cause—an action both necessary and sufficient for the effect—can be explained by the requirements of causal inference. As mere mortals, we cannot be sure that moderate regularities we see in the world (probabilities not close to zero or one) will carry over to a deeper understanding of the world's causal structure. So we are unwilling to infer causation from such moderate probabilities. From our experience that an action is likely to produce an effect, we cannot infer that Nature would expect the effect from the action in a particular case, because Nature may know something that rules out the effect in that case. And from our experience that the effect is unlikely without the action, we similarly cannot infer that Nature thinks it unlikely in a particular case. We can expect Nature to agree with us in the particular case only if experience teaches us that the connection between the action and the effect is certain or nearly so.

This is clear enough, but a full understanding of the roles of necessity and sufficiency in the legal definition of cause must also take into account the way intention is built into our concept of action. Even the simplest action, such as the lifting of a finger, extends through time and requires planning and the assessment of feedback to reach a goal. Thus the goal cannot be separated from the action. The event-tree framework makes this very clear, for it defines an action in terms of an initial

situation and a later situation, which is the goal. Almost by the definition of action in this framework, an action is necessary and sufficient for its goal.

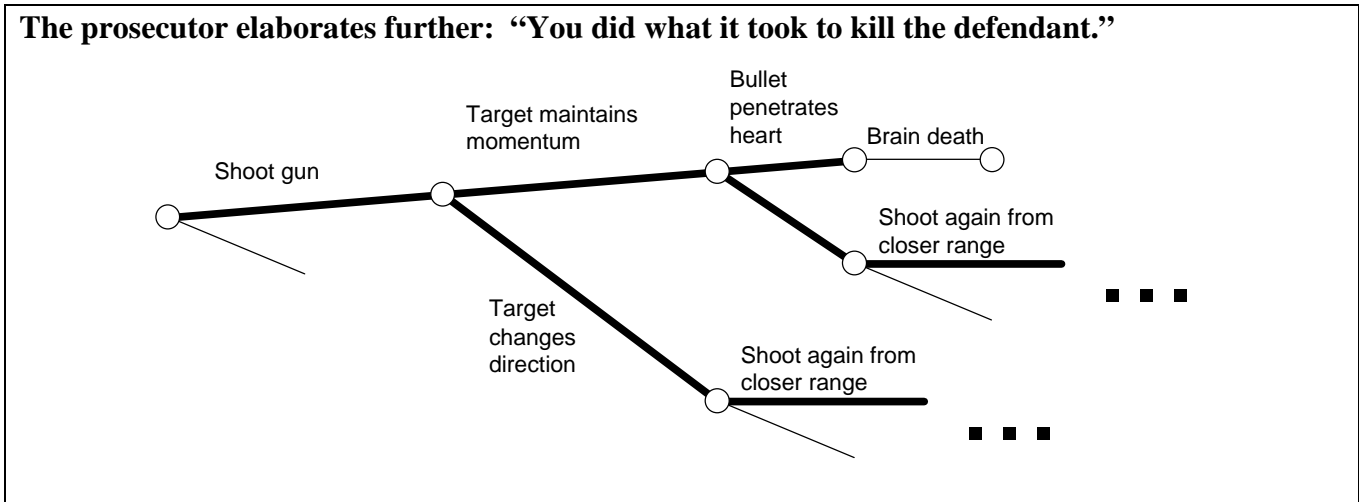




Figure 10. Intent to kill. The defendant, who protests that he merely aimed and pulled the trigger, stands accused of killing the victim.

Consider the murder case illustrated in Figure 10. A child who pokes at a gun's trigger out of curiosity will not be held culpable for resulting injury or death. But in a prosecution for murder the prosecutor will not allow us to say that the defendant merely aimed the gun and pulled the trigger. The prosecutor undertakes to prove that the person had murder as the goal, so that the damage done the victim by the gun is part of the action that the person intended and performed. As the figure illustrates, this action extends in time past the discharge of the gun to the damage that was intended, and it extends "counterfactually" across unrealized possibilities to take account of the defendant's intention to do what is needed to carry out the deed.

9. *Approximate Cause*

People are held responsible, of course, for negligence as well as for intentional actions. The endpoint of an act of negligence is defined not by intention but by obligation. This kind of action is usually only an approximate cause: it is only approximately necessary and sufficient for its result. Let me give an example.

This example is from my childhood—from June 1957 to the best of my recollection. That is when a neighbor, whom I will call Joe, came to our farmhouse to collect \$300 from my father, Dick Shafer, for the damage our cows had done to his corn. Both the clarity and emotional charge of this exchange were heightened by the fact that the obligations involved had been established very precisely three years before, when Joe had paid Dick exactly the same amount for exactly the same infraction—failing to repair a fence that was his responsibility and thus permitting his cows to get into his neighbor's corn. The two neighbors had long before agreed on which stretches of fence were whose responsibility.

A fence that is adequate to hold cows at bay when there is only an ordinary pasture on the other side may not be adequate to hold them when there is a succulent corn crop on the other side, and one might imagine that the person planting the corn would notice the inadequacy of the fence and take some responsibility, at least by calling the neighbor's attention to the problem. But in the earlier incident, this argument had cut no ice with my father. He had insisted that the person responsible for the fence was responsible, period. So he had no room for argument when the tables were turned.

There is more to tell. The truth is that my father had collected the \$300 in hard times, at a time when he had no paid work and was struggling with a continuing drought as a farmer. As it turned out, the corn crop failed for lack of summer rain in 1954, and the \$300 Dick collected from Joe was the

only money he made from that corn field. Was it really reasonable to collect \$300 for damage to corn plants that turned out to be worthless? Dick's theory had to be that Joe had paid him for the expected worth of the corn plants at the time. In any case, he had needed the money badly and he kept it. His financial condition was much better in 1957, when he had to pay the \$300 back. But I do believe that the money was less important, both to him and to Joe, than the principle. The fence in question was near Joe's house but far away and out of sight from our house, and I have always believed that Joe watched that fence as his corn grew, waiting with relish for the moment when our cows would come through it.

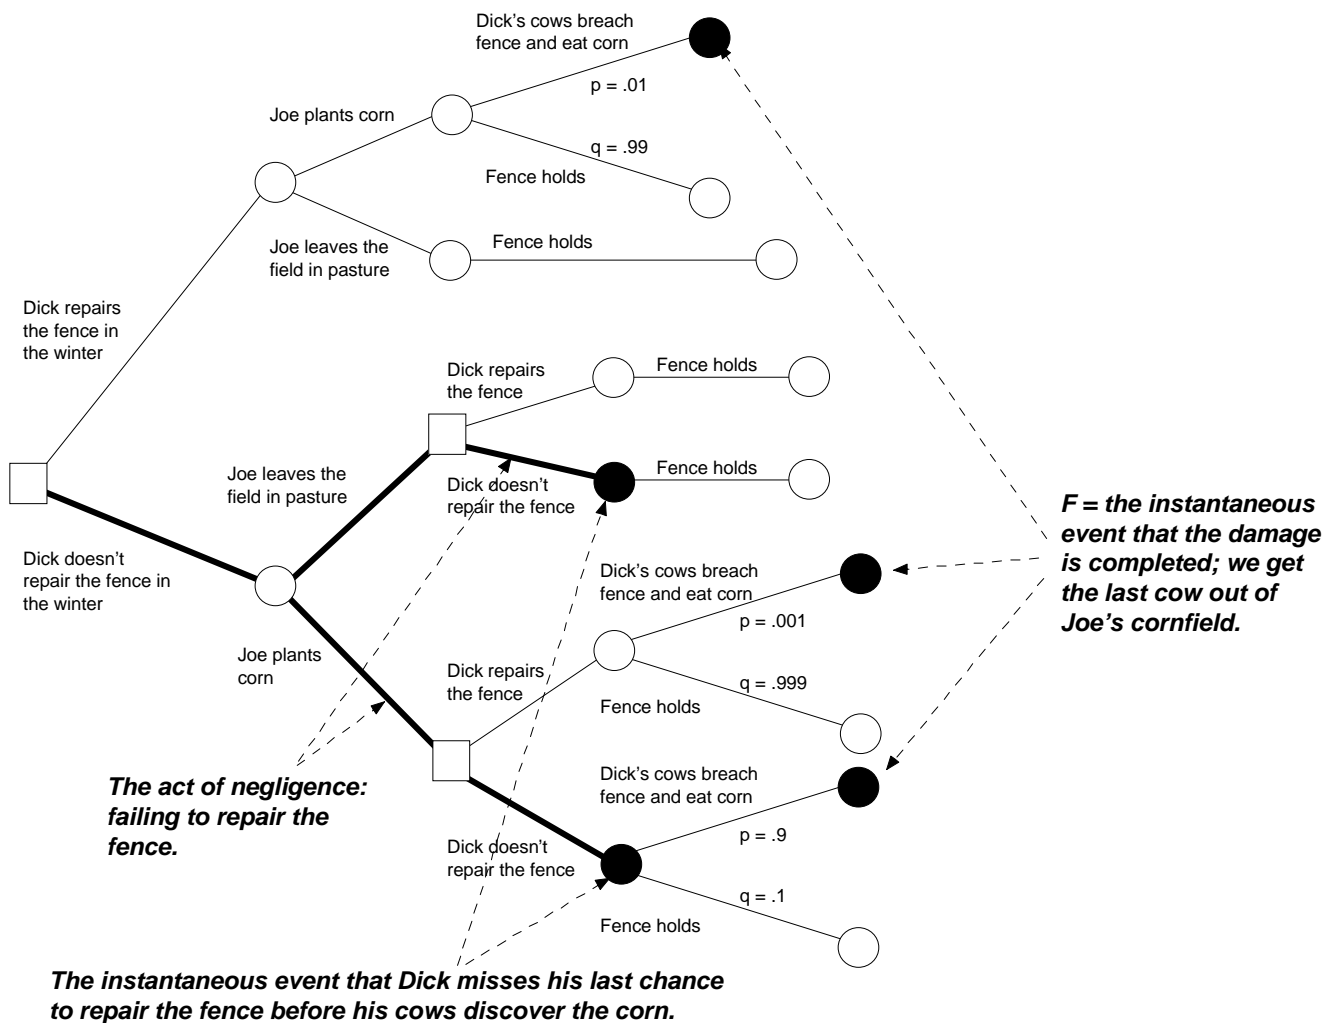


Figure 11. Cows in the corn. My father had to pay \$300 after his cows frolicked in his neighbor's succulent corn field. His act of negligence is represented in the figure by the two paths marked by heavy lines.

In any case, Figure 11 uses this story to illustrate the idea of an approximate necessary and sufficient cause. Had my father repaired the fence, the probability of his cows going through it would have been very small—perhaps only 1%. My father's negligence, which extended up to the moment when his cows went through the fence, produced a situation where that result was practically certain.

10. Toxic Torts

Let us now turn to the problem of toxic torts—the problem of assessing responsibility for actions that only raise the probability of damage to a large number of individuals.

Consider, for concreteness, the story told by Figure 12. Suppose we know, from experience, that Joe's spraying increases the probability of my catching a cold from 20% to 30%. Suppose I do catch cold, and suppose the cold costs me \$300 in pain and lost productivity. What should Joe pay?




Figure 12. Chemicals in the air. The insecticide that Joe sprays in his field, next door to my house, raises the probability of catching cold, for me and my neighbors.

For a statistician innocent of the law, the obvious answer is that Joe should pay \$30 whether I catch cold or not. If he does not spray, my expected loss to a cold is 20% of \$300, or \$60. If he sprays, then my expected loss goes up to 30% of \$300, or \$90. The difference, \$30, is the expected damage he causes me. We might also talk about punitive damages, especially if Joe's economic gain from the spraying is much greater than the damage he is causing me and my neighbors.

The law does not work this way of course. You cannot sue unless you are actually harmed. Expected harm is not enough. Moreover, for most legal authorities, the crucial issue is the probability of whether my particular cold can or cannot be attributed to Joe's spraying. One frequently adopted argument, for example, concludes that the probability that Joe caused my cold is only 1/3, because 2 out of 3 of the colds that occur after Joe sprays would have occurred anyway. Because 1/3 is less than

the standard of proof in civil suits (more likely than not), I cannot collect damages from Joe. The fraction $1/3$ is called the *attributable fraction* in epidemiology. In the legal literature, it is often called the *probability of causation*. It seems fair to say that most lawyers have a picture like Figure 13 in mind when they talk about the probability of causation. Nature can tell whether my cold is a Joe cold or a regular cold, but we cannot tell them apart. All we can do is give a subjective probability of $1/3$ that mine is a Joe cold.

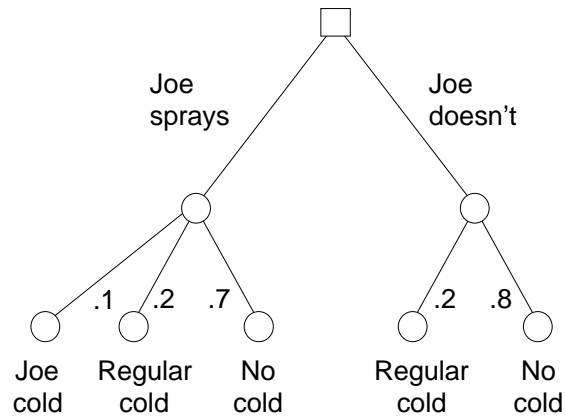


Figure 13. The theory of Joe colds. Joe’s spraying causes a special kind of cold.

There is a large literature criticizing the probability of causation, on the grounds of possible heterogeneity in susceptibility and exposure, non-additivity of effects, and so on (see, for example, Paracandola 1968, Greenland and Robins 1988). The general tendency of this literature is towards complication, and as the proposed causal models become more complicated, the hope that we learn the probabilities in them becomes more remote. In the most recent literature, there is even a tendency to introduce models that are non-identifiable in principle, because they involve so-called “counterfactual” probabilities (Freedman 1999; Pearl 2000; Pearl and Tian 2000).

Without going into the details of this literature on the probability of causation, I would like to suggest that the appropriate solution is to acknowledge our intrinsic inability to identify causal probabilities for the particular case. No matter how complicated we make our models, any probabilities that we identify empirically will be only valid in the sense of Section 5, not certifiably causal. The only probabilities that we can take to be causal are empirically valid probabilities close to zero or one.

This is not message of despair, because in the kind of situation that we have been describing we do have probabilities close to zero and one. What we know empirically is that the spraying increases the frequency of colds from 20% to 30%, and our models, in the end, will give this back to us a practical certainty; the probability is close to one that the frequency of colds in the population will increase from 20% to 30%. This suggests that the problem be treated as a class action, not as individual suit, and that the affected population be compensated for the increased cost. When we have distributed the proceeds across the population, we end up with the result I suggested initially: Joe pays me \$30, whether or not I get a cold. This solution agrees with the stance taken by Berger (1997).

References

- Berger, Margaret A. (1997). Eliminating general causation: Notes towards a new theory of justice and toxic torts. *Columbia Law Review* **97** 2117-2152.
- Dawid, A.P. (1985). Calibration-based empirical probability (with discussion). *Annals of Statistics* **13** 1251-1273.
- Dawid, A.P., and Volodya Vovk (1997). Prequential probability: Principles and properties. *Bernoulli* **31** 1-38.
- Freedman, David (1999). The swine flu vaccine and Guillain-Barré syndrome: a case study in relative risk and specific causation. Technical Report No. 546. Statistics Department, University of California, Berkeley.
<http://www.stat.berkeley.edu/tech-reports/index.html>
- Greenland, Sander, and James M. Robins (1988) Conceptual problems in the definition and interpretation of attributable fractions. *American Journal of Epidemiology* **128** 1185.
- Parascandola, Mark (1998). What is wrong with the probability of causation? *Jurimetrics* **39** 29-44.
- Pearl, Judea (2000). *Causality: Models, Reasoning, and Inference*. Cambridge University Press.
- Pearl, Judea, and Jin Tian (2000; to appear). Probabilities of causation: Bounds and identification. *Annals of Mathematics and Artificial Intelligence*.
- Shafer, Glenn (1990). The unity of probability. Pp. 95-126 of *Acting Under Uncertainty: Multidisciplinary Conceptions*, edited by George von Furstenberg. Boston: Kluwer.
- Shafer, Glenn (1996). *The Art of Causal Conjecture*. Cambridge, Massachusetts: MIT Press.
- Shafer, Glenn (1998). Mathematical foundations for probability and causality. Pp. 207-270 of *Mathematical Aspects of Artificial Intelligence*, edited by Frederick Hoffman. American Mathematical Society, Symposia in Applied Mathematics, Volume 55.
- Shafer, Glenn, Peter Gillett, and Richard Scherl (2000; to appear). The logic of events. *Annals of Mathematics and Artificial Intelligence*.
- Shafer, Glenn, and Volodya Vovk (to appear). *Probability and Finance: It's Only a Game*. Wiley.
- Vovk, Volodya (1993). A logic of probability, with applications to the foundations of statistics (with discussion). *Journal of the Royal Statistical Society, Series B* **55** 317-351.