

2008

Causation, Truth, and the Law

Richard Scheines
Carnegie Mellon University

Follow this and additional works at: <http://repository.cmu.edu/philosophy>

 Part of the [Philosophy Commons](#)

This Article is brought to you for free and open access by the Dietrich College of Humanities and Social Sciences at Research Showcase @ CMU. It has been accepted for inclusion in Department of Philosophy by an authorized administrator of Research Showcase @ CMU. For more information, please contact research-showcase@andrew.cmu.edu.

Causation, Truth, and the Law

Richard Scheines[†]

I. INTRODUCTION

Deciding matters of legal liability, in torts and other civil actions, requires deciding causation. The injury suffered by a plaintiff must be caused by an event or condition due to the defendant. The courts distinguish between cause-in-fact and proximate causation, where cause-in-fact is determined by the “but-for” test: the effect would not have happened, “but for” the cause.¹ Proximate causation is a set of legal limitations on cause-in-fact.

As this conference is entitled *A Cross-Disciplinary Look at Scientific Truth: What’s the Law to Do*, I will ignore the distinction between cause-in-fact and proximate causation, and instead focus on both the sense in which cause-in-fact claims can be considered true or false and on the challenges to establishing them.

Before a court can decide on proximate causation, and thus on liability or damages, it must decide on the truth of the cause-in-fact question: Was the injury suffered by the plaintiff caused by the action(s) of the defendant? For example, was John Smith’s liver cancer caused by his exposure to trichloroethylene (“TCE”) in a factory that employed him for ten years?

I will focus on what we must assume before cause-in-fact claims can even be said to have a truth value, that is, objectively true or false independent of whether we can know it. Philosophers distinguish individual level causal claims (cause-in-fact claims) from “general causal claims.”² I will try to

[†] Dr. Scheines is a Professor (and Head) of Philosophy at Carnegie Mellon University, with courtesy appointments in the Department of Machine Learning and the Human-Computer Interaction Institute.

¹ See Richard W. Wright, *Causation in Tort Law*, 73 CAL. L. REV. 1735, 1775 (1985).

² David Lewis, *Causation*, 70 J. PHIL. 556, 556-57 (1973); see Richard Scheines, *Causation*, in 1 THE NEW DICTIONARY OF THE HISTORY OF IDEAS 280. 280-88 (Maryanne Cline Horowitz, ed. 2005).

make the difference clear and then discuss the difficulties that arise in claiming that either have a determinate truth value. I will argue that if there is a truth to the matter about individual causal claims, it is parasitic upon the truth of general causal claims, and that therefore the relevant issues for the law involve deciding on the truth of general causal claims.

Given the limited amount of relevant empirical scientific evidence that is typically available, however, deciding in a legal setting whether such claims are actually true or false can be extremely difficult. As courts have no choice but to decide such matters, they need a rational process by which to synthesize the evidence for or against causal claims—both with respect to our best scientific guess about the truth of the claim and with respect to the scientific uncertainty about such a guess.

I will sketch the various forms of evidence that are used to prove general causal claims and then describe the strategies and problems associated with synthesizing the totality of this evidence into a single judgment, both with regard to the truth of a causal claim and with regard to the uncertainty with respect to that judgment.

II. INDIVIDUAL VERSUS GENERAL CAUSATION, INDETERMINISM, AND TRUTH

A. *Individual versus General Causal Claims*

Consider first the difference between individual and general level causal claims. In legal contexts, the goal is often to establish whether one particular event or condition was the cause of another particular event or condition. For example, if John Smith contracts liver cancer, a court might seek to establish whether or not his exposure to TCE in a factory that employed him for ten years was the “proximate cause” of his particular cancer. In such cases, by saying that exposure to TCE *caused* the disease, courts typically ask whether the cancer would not have occurred “but for” the exposure to TCE.³ This is an individual level causal claim, and one whose truth, if

³ See Wright, *supra* note 1, at 1775; see also Jane Stapleton, *Legal Cause: Cause-in-Fact and the Scope of Liability for Consequences*, 54 VAND. L. REV. 941, 958-61 (2001).

it has one, depends on a *counterfactual claim*: what *would have* happened to John Smith *had he not* been exposed to TCE.⁴

General causal claims refer to a *population* of individuals, and concern the *probability* or *average severity* of a property (for example, a disease) in that population. For example, a qualitative general causal claim about TCE and liver cancer might be: in a population of factory workers who were exposed to TCE, the probability of getting liver cancer (the risk among the exposed) is higher than it would have been in the same population had they not been exposed to TCE. A quantitative version of the same general causal claim: the risk of liver cancer among those exposed to TCE was 2%, while the risk in the same group, had they not been exposed, would have been 1%. Another example is the following: among middle class American children between the ages of five and ten, if everyone had watched one less hour of TV per day, then the average Body Mass Index (BMI)⁵ of the group would have been .5 point lower than it was.

Clearly, claiming that TCE causes liver cancer in a population of workers does not entail that every worker who was exposed to TCE will develop liver cancer, and it does not entail that every case of liver cancer among the workers would not have happened but for TCE exposure. Similarly, not all children would have lost ½ point of BMI had they decreased their TV watching by one hour per day, etc. Again, the general causal claims each make a counterfactual claim. In the TCE and liver cancer case, the claim is: had the same population lived the same life, with the exception of not being exposed to TCE, then the *probability* of liver cancer would have been lower than it was in the actual world. In the TV and obesity case, the claim is: had the same population of American children lived the same life, with the exception that they had watched an hour less of TV per day, then the average BMI in the population would have been .5 point lower than it was in the actual world.

⁴ David Lewis developed the most influential account of counterfactuals in philosophy, see DAVID LEWIS, *COUNTERFACTUALS* (2001); and Donald Rubin developed the most influential account of causation based on counterfactuals in Statistics, see Donald B. Rubin, *Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies*, 66 J. EDUC. PSYCHOL. 688 (1974).

⁵ BMI is calculated as $703 \times (\text{weight in pounds}) / (\text{height in inches})^2$.

B. General Causal Claims and Truth

Since our target is the law, which for the most part deals with events that have already occurred, I will restrict my attention to causal claims about things that have already happened. For example, was John Smith's liver cancer caused by his exposure to TCE, or was the rise in obesity in the late twentieth century caused by an increase in TV watching?

On this sort of counterfactual account of causation, consider what it might mean to assert that general causal claims have a truth value, that is, that they are either true or false in the actual world. Two objections arise immediately: the vagueness and objectivity of counterfactual possible worlds, and the meaning of probability.

That causal claims in the actual world might depend upon what would or would not happen in alternative possible worlds bothers almost everyone. The problem is that most descriptions of alternative possible worlds seem intolerably vague. For example, in the TV and obesity claim we are asked to consider a world in which the same children lived the same life but watched one hour less of TV per day. How exactly do we imagine the change in their world so that they each watch one hour less of TV per day? Do we make the TV inoperable one hour before they would have turned it off anyway? Convince their parents to intervene and select an hour of TV every day the child will no longer watch? Offer them just enough of a candy reward to get them to voluntarily shut down the TV one hour before they would have anyway? Make them replace the first hour of TV they would have watched with exercise? The counterfactual as we stated it is vague—it can't answer any of these questions even though they all obviously matter for assessing the causal claim.

One can, however, fully specify a manipulation or intervention that would change the actual world to the possible world we are considering in a way that eliminates all this vagueness. Donald Rubin famously articulated a counterfactual theory of causation based on drug trials.⁶ If we consider an experiment in which some people received a pill containing a drug (the treatment) and the other half received a pill identical in appearance, taste, etc. (the control), then the causal inference problem with respect to the drug amounts to missing

⁶ See Rubin, *supra* note 4.

data from alternative possible worlds. For the people who took the treatment pill, we are missing the data on what would have happened to them if they had taken the control pill. For the people who took the control pill, we are missing the data on what would have happened to them if they had taken the treatment pill (Table 1).

	Took treatment pill	Took control pill
Jane Doe	Cured	???
Person 2	Ill	???
...		
Person N - 1	???	Ill
Person N	???	Ill

Table 1 Missing data needed to assess causal efficacy of treatment.

Since the pills are identical in appearance, etc., surely it is not difficult to be perfectly precise about the alternative worlds under discussion. Jane Doe took the treatment pill and recovered—what would have happened if she had taken the control pill? It isn't hard to imagine the antecedent: an alternative world in which we leave everything as we found it in the real world, except for removing the drug from the pill Jane took. Knowing whether or not, in this hypothetical world, Jane would have remained ill or recovered is not so simple.

The point is this: the problem of vagueness is not insurmountable. It requires being clear about the intervention performed to transform the actual world into the counterfactual world.⁷ In the TCE and liver cancer case, we might describe a world in which the factory workers behaved identically, but the de-greasing chemical used in the factory in

⁷ The idea of making counterfactuals clear by formalizing the idea of an intervention has been developed extensively. *See generally* JUDEA PEARL, CAUSALITY: MODELS, REASONING, AND INFERENCE (2000); PETER SPIRITES ET AL., CAUSATION, PREDICTION, AND SEARCH (2d ed. 2000); James Robins, *A New Approach to Causal Inference in Mortality Studies with Sustained Exposure Periods—Application to Control of the Healthy Worker Survivor Effect*, 7 MATHEMATICAL MODELLING 1393 (1986) (errata appears in 14 COMPUTERS & MATHEMATICS WITH APPLICATIONS 917 (1987)).

which they worked was changed from TCE to a specific alternative known to be non-carcinogenic.⁸

Assuming that vagueness is not an issue, what about the probability part of general causal claims? Recall that to assess the TCE and liver cancer claim, we need to know the probability of liver cancer in the population exposed to TCE and in the same population not exposed to TCE.

What is it for a probability claim to be true? Given two claims about a coin in my pocket:

H₁: the coin is fair (probability of heads = .5)

H₂: the coin is loaded 75% toward heads (probability of heads = .75)

What is it for H₁ to be true but H₂ false? Unfortunately, neither hypothesis puts any binding constraints on any experiment we might conduct in the actual world. We might say that H₁ implies that the proportion of heads in a very long sequence of flips should converge to .5 as the sequence gets longer and longer, but that this is not the case for H₂. But the “should” in this sentence is itself probabilistic. *Any* finite sequence of coin flips is consistent with both of these hypotheses. Perhaps this is just philosophical obstructionism. Even though we don’t yet possess an entirely satisfactory account of what it means for H₁ to be true and H₂ false in the actual world, several accounts are out there.⁹ We don’t want to put the legal system on hold until the philosophers can agree on a semantics for probability.

To summarize, provided we can be sufficiently precise about the manipulations (interventions) that will transport us from the actual to a counterfactual possible world, and provided probability claims have a coherent semantics, then general causal claims have a truth value as well. At minimum, their truth depends on the probability of the effect in two populations.

C. *Individual Causal Claims and Truth*

Now consider whether *individual* causal claims have a determinate truth value. Was John Smith’s liver cancer caused by his exposure to TCE? Was Jane Doe’s illness cured by the experimental drug she took?

⁸ Presuming there is such a thing!

⁹ See *Interpretations of Probability*, in STANFORD ENCYCLOPEDIA OF PHILOSOPHY (rev. July 7, 2007), <http://plato.stanford.edu/entries/probability-interpret/>.

Again, these claims depend upon evaluating counterfactuals. What happens to John Smith in the world in which he wasn't exposed to TCE? What happens to Jane Doe in the world in which she takes the placebo pill? Having argued that these counterfactual worlds are not necessarily insurmountably vague, let us now consider whether these questions have definite answers.

Again, the answer depends on probability, but in a completely different sense than we have already discussed. It depends upon whether the world is deterministic, as we are psychologically built to expect, or indeterministic, as the physicists tell us it is.

Roughly, in a deterministic world, the past fully determines the future. What happens to Jane Doe after she takes the pill is not a lottery, but a sure thing. We might not know enough about the details to be able to predict what will happen, but that is an epistemic limitation and not a feature of the world we inhabit. In an indeterministic world (still guided by physical laws), the past does not determine the future, but at most determines the probabilities of possible futures.

Although Einstein detested the thought of an indeterministic world and said so (“[God] does not play dice”¹⁰), in modern physics God is indeed a gambler. When a single electron is shot at a screen that will record the point where it “hits,” all that can be predicted about its landing site is its probability. The *probabilities* can be predicted perfectly, and experiments using thousands of electrons have confirmed the accuracy of these probabilistic predictions, but the *exact position* on a single trial cannot be predicted perfectly, *no matter what we know* about the electron. Even if we learned everything there was to know about the electron midway through its flight, and the universe within which it traveled, it would not be enough to determine where it will hit the screen. We are simply not evolved to accept this idea fully, but this is the way the world seems to be.

Translated back to Jane Doe, if the world is truly indeterministic, then although she took the treatment pill and recovered in the actual world, there is *no fact to the matter* about what would have happened to her if she had taken the

¹⁰ Albert Einstein, Letter to Max Born, Dec. 4, 1926, in LEWIS S. FEUER, EINSTEIN AND THE GENERATIONS OF SCIENCE 80 (2d ed. 1982).

control pill. And it has nothing to do with the vagueness of the counterfactual.

If the world is truly indeterministic, then even if we could go back in time and replay the world millions of times from the exact spot we like, for example, leaving everything the same but changing the contents of Jane's pill, the outcome on each play is still truly a lottery. Analogous to the electron's landing site on the screen, the *probabilities* for Jane's outcome might be determined and perhaps knowable, but the *outcome* on each play is not. Thus, if one uses the basic legal test for causation, that is, the "but for" test, then individual level causal claims simply do not have a truth value in a world in which God actually does play dice.

If the world is deterministic, then probability statements capture only our epistemic uncertainty, not something more fundamental about the world. For example, suppose that a population exposed to TCE has a probability of getting liver cancer of .02. It might be the case that some people in the population have an unusual genetic makeup such that, if they are exposed to TCE they will definitely get liver cancer, but if they are not exposed they definitely will not get cancer, and that 2% of this population has the unlucky gene. In 1960, before we could sequence an individual's genes, we would not be able to tell which of the individuals were lucky or unlucky. Thus, for any individual J, chosen at random, J's fate is determined, but our epistemic access to it is limited to the true claim: "J was exposed to TCE and thus the probability that he will get liver cancer is 2%." The underlying situation is deterministic, but due to our limited access *appears* indeterministic. Philosophers refer to such a world as "pseudo-indeterministic."¹¹

In a pseudo-indeterministic world, individual level causal claims *do* have a truth value. Although we might not have access to his genome, individual J either has the unlucky gene or he doesn't, and whichever it is determines his cancer outcome. If individual J got liver cancer, and he had the unlucky gene, then the claim that J would not have gotten liver cancer but for the TCE exposure is true, even though we cannot know it until we can sequence his genome or find some other marker that correlates perfectly with the unlucky gene.

¹¹ See SPIRITES ET AL., *supra* note 7, at 19-29.




Figure 1 Pseudo-indeterministic population.

So which world are we in? Electrons may be truly indeterministic, but is cancer? Even if a gene exists which makes an individual vulnerable to TCE exposure, cancer requires a complicated series of genetic mutations and other developments, all of which can happen in a number of different ways, including insult from TCE, cosmic rays, a failure of the DNA repair mechanisms, etc. Perhaps quantum mechanical indeterminism does play a role in TCE exposure. Perhaps TCE interacts with some molecule in codon 61 of the *H-ras* protooncogene,¹² the result being to move the probability of mutation in this gene slightly higher but leaving us with nothing, even in principle, that we could measure or observe about an individual of whom we could say that but for TCE, they would not have gotten liver cancer. I don't know, and right now I think it is safe to say that no one else does either.

If the question of whether individual level causal claims have a truth value depends on whether the world is pseudo-indeterministic or truly indeterministic, this is not so for general causal claims.

General causal claims involve the probability of the effect in a population that was exposed to the cause and the probability of the effect in the same population not exposed to the cause. These probability claims have a truth value regardless if the world is pseudo-indeterministic or truly indeterministic. Consider again the probability of liver cancer and TCE exposure, and suppose our general claim is that the probability of liver cancer is .02 if you are exposed to TCE, and

¹² This is the suspected loci of TCE's effect on mice in tumorigenesis studies. Richard J. Bull, *Mode of Action of Liver Tumor Induction by Trichloroethylene and Its Metabolites, Trichloroacetate and Dichloroacetate*. 108 (Supp. 2) ENVTL. HEALTH PERSP. 241, 254 (2000), available at <http://www.pubmedcentral.nih.gov/picrender.fcgi?artid=1637759&blobtype=pdf>.

.01 if not. In the pseudo-indeterministic world depicted in Figure 1, 10% of the population has an unlucky gene (gene L) that produces liver cancer always, and another, separate 10% of the population has a vulnerability gene (gene V) that gives you liver cancer whenever you are exposed to TCE (and fails to prevent the first gene from giving you cancer when you are not exposed to TCE).

Likewise, we can imagine a truly indeterministic world in which everyone's propensity for getting liver cancer is moved from 10% to 20% with TCE exposure, but that there is nothing hidden to be discovered that will determine the outcome.

D. The Primacy of General Causal Claims

So individual causal claims have a truth value in a pseudo-indeterministic world, but upon what does their truth depend? Essentially, it depends on answering a counterfactual question of the form would the "effect" have failed to occur "but for" the cause? Ignoring the many difficulties with this simple account of actual causation (such as overdetermination¹³ or pre-emption¹⁴), consider again what is required to evaluate such counterfactual claims. Consider two claims: (a) John Smith would not have gotten liver cancer but for exposure to TCE, and (b) John Smith would not have gotten liver cancer but for wearing brown socks to work on Mondays. We must be able to assess whether John Smith's life would have produced liver cancer in each of the two counterfactual worlds: (a) the world in which he lives his life exactly as before but is not exposed to TCE, and (b) the world in which he lives his life exactly as before but does not wear brown socks to work on Mondays. Assessing whether he gets liver cancer in these worlds requires the general causal knowledge about how the world would have responded to such changes.

Once we decide, in our counterfactual world, exactly how to change John Smith's circumstances, then the question of whether or not he gets liver cancer in this alternative world

¹³ For example, when several soldiers in a firing squad shoot real bullets accurately, the prisoner's death is overdetermined. For any individual soldier, it is false to say that the prisoner would not have died but for the soldier. We want each soldier to come out as a cause.

¹⁴ For example, when spy 1 pokes a hole in the canteen of an enemy about to cross a desert, he pre-empts the effect of spy 2, who had previously filled the canteen with poison. It is false to say that the enemy would not have died, but for either spy. We want spy 1 to be the cause and not spy 2.

depends entirely on the causal laws we take to hold in *all* the possible worlds we consider. Given that the shift from wearing brown socks to black socks on Mondays is sufficiently minimal, then Smith still gets cancer because of the general causal claim: sock color has no causal influence on liver cancer.

Halpern and Pearl,¹⁵ Woodward,¹⁶ and others who have articulated clear accounts of individual level causal claims all require as input: (1) what happened in the real world, (2) precision about how the counterfactual world is to differ from the actual world as a result of removing or adding the “cause,” and (3) the general causal laws (usually called the “structural equations”) relevant to the events discussed. The moral is clear: we cannot assess the truth of individual level causal claims until we have the general causal laws relevant to the events at issue.

E. The Probability of Causation

If we know the general causal claims, that is, the risks of those exposed and of those not exposed, then we can turn to a weaker notion than truth for assessing our individual “but for” causal claims. We can compute what is called the probability of causation (PC), a number that roughly corresponds to the probability that someone exposed would have avoided the disease had they not been exposed.¹⁷ The PC is based on what is called the attributable fraction of risk in a population (AF):

$$AF = \frac{Risk(exposed) - Risk(unexposed)}{Risk(exposed)}$$

$$= \frac{P(disease/exposed) - P(disease/unexposed)}{P(disease/exposed)}$$

¹⁵ Joseph Y. Halpern & Judea Pearl, *Causes and Explanations: A Structural-Model Approach. Part I: Causes*, 56 BRIT. J. PHIL. SCI. 843, 843-87 (2005).

¹⁶ JAMES WOODWARD, MAKING THINGS HAPPEN: A THEORY OF CAUSAL EXPLANATION 3-24 (2003).

¹⁷ COMMITTEE ON EVALUATION OF THE PRESUMPTIVE DISABILITY DECISION-MAKING PROCESS FOR VETERANS, INSTITUTE OF MEDICINE, IMPROVING THE PRESUMPTIVE DISABILITY DECISION-MAKING PROCESS FOR VETERANS 7-5 to -6 (Jonathan M. Samet & Catherine C. Bodurrow eds., Nat'l Academies Press 2007) [hereinafter DISABILITY DECISION-MAKING]; Sander Greenland & James M. Robins, *Epidemiology, Justice, and the Probability of Causation*, 40 JURIMETRICS J. 321, 321-22 (2000).

For example, if the risk of liver cancer among TCE exposed workers is 2%, and would have been 1% had they not been exposed, then the AF = .5, so half of the liver cancers observed in TCE exposed workers can be assumed to be the result of TCE exposure.

In many cases we have data on widely defined populations, for example, factory workers in Ohio, but no data on subpopulations, for example, workers 40 to 45 years in age, or workers who have a close ancestor with liver cancer, etc. In such cases it is typical to use the AF from the most narrowly defined and most informative population possible. For example, the risks for the Ohio factory workers might be 2% for exposed and 1% for unexposed, but among 40- to 45-year-old workers with a history of liver cancer in their family, the risks might be 6% for exposed and 5% for unexposed, making the AF = $.06 - .05/.06 = 16.67\%$. So if Robert Jones is a non-descript Ohio factory worker who was exposed to TCE and got liver cancer, then the probability that his liver cancer was caused by TCE was 50%, but if Tim Lewis is a 43-year-old factory worker exposed to TCE with a father who had liver cancer, then the probability that his liver cancer was caused by TCE was much lower: 16.7%.

As to truth, the probability of causation is no help at all. Just as it is not true to say that a given coin flip came out heads because the coin was loaded 75% heads, it is not true to say that a given individual's cancer was caused by TCE exposure because the probability of causation was 75%. A high probability of causation,¹⁸ or something like it, might be what the Law must resort to in deciding torts and similar issues, but it should not be confused with assenting to the truth of a "but for" claim.

III. EVIDENCE FOR GENERAL CAUSAL CLAIMS

So the truth or probability of an individual causal claim depends upon the general causal laws. It is to the problem of deciding on the general causal laws that I now turn. As I have

¹⁸ PC is objectionable for several reasons, not just that it falls short of truth. Sander Greenland and Jamie Robins have argued that a better measure of the plaintiff's injury is an estimate of the years of life lost ("YLL"). Since YLL is also either something that might be determined by exposure, or only have its probability determined by exposure, I have no stake in this debate: all the issues regarding the truth of individual causal claims apply to YLL and PC equally. See Greenland & Robins, *supra* note 17, at 346.




Figure 2 Association versus causation.¹⁹

already stressed, assessing general causal claims requires comparing a real population (e.g., Actual Population 1, Figure 2) with a counterfactual one (e.g., Counterfactual Population 1, Figure 2). Clearly we cannot go back in time and remove TCE or change the commercials children were exposed to. Can we compare workers exposed to TCE to other workers not exposed to TCE? Yes, but comparing an exposed population (Actual Population 1, Figure 2) to another population that we *observe* not to be exposed (Actual Population 2, Figure 2) reveals association, but not necessarily causation.

This is because an actual population of workers not exposed to TCE (Actual Population 2, Figure 2) might differ from the workers exposed to TCE (Actual Population 1, Figure 2) in other ways that affect cancer, for example, diet, income, etc.

For example, it is nearly certain that children five to eight years of age who were in fact exposed to fewer than ten junk food commercials per day had a lower frequency of obesity in 2005 than children who were exposed to more than ten. But because it is also nearly certain that this group of children differs in other important ways from the population described in the claim, for example, their parents are more educated,

¹⁹ The expression $P(\text{Cancer} \mid \text{TCE})$ denotes the probability of cancer among those exposed to TCE. It also might be referred to as the conditional probability of cancer, given exposure to TCE. The expression $P(\text{Cancer} \mid \text{Set to TCE})$ denotes the probability of cancer among those upon whom we intervene to force exposure to TCE.

wealthier, etc., this is *not* the appropriate contrast class for causation. The appropriate contrast class is the *same* group of children, living the same life, but with the junk food commercials removed, or replaced with other sorts of commercials. In the first case, we are comparing children we *observed* being exposed to a lot of junk food commercials to children we *observed* being exposed to few commercials, and this comparison would undoubtedly reveal a statistical association between junk food and obesity. In the second case, we are comparing children we *observed* being exposed to a lot of junk food commercials to the *same* children after a hypothetical *intervention* in which we go back in time and change their exposure to commercials from a lot to a little. It is this comparison that reveals the causal relationship, but which is, in the deepest sense, unobservable.

As the population under a counterfactual, hypothetical intervention cannot be observed, how are scientists to gather evidence about this counterfactual population? This is the problem of causal inference.

A. *Randomized Trials*

Sir Ronald Fisher, the brilliant and prolific British statistician, provided in the 1930s what is still the gold standard today for causal inference: the randomized trial (RT).²⁰ In its simplest form, an RT randomly splits a population into two subgroups (which we can expect on average to be identical), thus creating two versions of the same population, and then exposes one subpopulation to the cause (the “treated” group) and one to the absence of the cause (the “control” group). The frequency of the effect in the two groups provides evidence of the probability of the effect in the two populations we seek: one in which the cause is present, and an identical copy in which the cause is not present. Subtleties abound, but the basic strategy is sound and taught in every introductory research methods course.

The problem, of course, is that in a number of situations performing an RT is either ethically or practically impossible. We simply cannot intentionally expose half of a population to TCE and look for liver cancer.

²⁰ See RONALD A. FISHER, *STATISTICAL METHODS FOR RESEARCH WORKERS* (4th ed. 1932).

There are essentially two recourses to an RT: (a) we can statistically adjust for naturally occurring differences in two populations, or (b) we can perform very small versions of RTs on animals we don't seem to mind harming, for example, rodents.

B. Epidemiological Studies

Epidemiological studies involve observing human populations in which we do not control exposure to a cause and thus must resort to recourse (b): that is, epidemiological studies must statically adjust for naturally occurring differences in the exposed and nonexposed populations before they can claim evidence of causation. Statistical adjustment requires that we know *all* the relevant features upon which individuals differ besides being exposed to the cause or not. For example, a subpopulation that is exposed to TCE, for example, automotive factory workers who handle paint strippers that contain TCE, and a subpopulation that is not exposed to TCE, for example, workers on a chicken farm, may differ in more ways than just TCE exposure. The chicken farm workers may be different in age, have different diets, etc.

If we measure *all* the relevant differences, that is, those that also might cause liver cancer, then we can often adjust for these differences statistically and test for differences in liver cancer rates among the groups after this adjustment. If we do not know all the relevant differences, however, then this strategy fails. For example, if, unbeknownst to us, the autoworkers' drinking water contains some other set of chemicals that cause liver cancer, while the farm workers' water does not, and we don't adjust for this, then our inference will be unsound.

A raft of other methodological issues confront epidemiologists, but the scientific evidence from such studies can in some instances be compelling, for example, cigarette smoking and lung cancer.

C. Toxicological and Animal Studies

In many cases, animals like rats or mice or rabbits or chimps share enough of human physiology to make it plausible to extrapolate from experiments with animals to what would happen in a similar experiment with humans. Biologists frequently perform controlled experiments on rodents to garner

evidence to show whether some chemical causes cancer. They expose some rodents to a “control,” and others that are genetically identical and raised in the same environment to the chemical of interest, and then compare the frequency of cancerous tumors. In some cases they can examine extremely detailed mechanisms by which the chemical might lead to a tumor by doing cell physiology on both the animals under study and human cells.

There are a number of informative and accessible discussions of the wide variety of evidence that can be used for causal inference.²¹ Although in my view the topic deserves dozens of books, it is out of my scope to say more here. In summary, scientists have long recognized that there are at least three distinct kinds of evidence that bear on the truth of general causal claims:²²

- Interventional studies on humans (e.g., RTs)
- Non-interventional (observational) studies on humans (e.g., Epidemiological studies)
- Mechanistic/toxicological evidence (e.g., animal and cell studies)

IV. COMBINING THE EVIDENCE FOR CAUSATION

To come to a reasoned position on the status of a general causal claim, especially in a legal setting, we must (1) combine all the available evidence into a single judgment on whether the claim is true, and (2) express the degree of our uncertainty about the claim.

In many cases, the evidence for a general causal claim is mixed. On some questions, there are RTs that show that a drug or treatment has a positive effect, others which show no effect, and still others which show a negative effect.²³ As they are more complicated methodologically, epidemiological studies

²¹ See, e.g., DISABILITY DECISION-MAKING, *supra* note 17, at 7-2 to -5; CENTERS FOR DISEASE CONTROL AND PREVENTION, DEP'T OF HEALTH AND HUMAN SERVS., SURGEON GENERAL'S REPORT: THE HEALTH CONSEQUENCES OF SMOKING. (2004), available at http://www.cdc.gov/tobacco/sgr/sgr_2004/index.htm; FEDERAL JUDICIAL CENTER, REFERENCE MANUAL ON SCIENTIFIC EVIDENCE (2d ed. 2000), available at <http://www.rashkind.com/researchlinks.htm>.

²² See, e.g., sources cited in *supra* note 20.

²³ See, e.g., A.R. White et al., *A Meta-Analysis of Acupuncture Techniques for Smoking Cessation*, 8 TOBACCO CONTROL 393, 393-97 (1999).

often present mixed evidence for a general causal claim. In many situations animal studies also show mixed results. Rationally combining multiple pieces of similar evidence, for example multiple RTs, is difficult, but often feasible. The real challenge is to rationally combine all sources of scientific evidence into a single judgment about the truth of a general causal claim.

A. *Meta-Analysis*

For many general causal claims, for example, hormone therapy and breast cancer, there are often several distinct RTs in the published literature. Optimally combining the evidence from such studies is a topic of its own, called “meta-analysis.” A meta-analysis involves mathematically combining the results from multiple but comparable RTs to derive a summary estimate of the effect of some cause on some effect, often involving health, that appropriately combines the results of all the individual studies.²⁴

The technique is not limited to combining RTs, but can also be used to combine the evidence from several epidemiological (observational) studies, provided the populations studied are comparable. If the populations vary, then a related technique called meta-regression²⁵ sometimes allows pooling of data in a principled way. Recent work by Eloise Kaizar²⁶ improves on meta-regression when data from both RTs and observational studies are available on a similar general causal claim.

Meta-analytic methods are in general quite useful when there are multiple studies on the same causal claim. For example, the 2006 Institute of Medicine Committee on Asbestos did a quantitative meta-analysis on studies which individually estimated the effect of asbestos exposure on any of

²⁴ See, e.g., Kay Dickersin & Jesse A. Berlin, *Meta-Analysis: State-of-the-Science*, 14 EPIDEMIOLOGIC REV. 154 (1992); P. Easterbrook & J. Berlin, *Meta-Analysis*, 341 LANCET 965 (1993); K.A. L'Abbe et al., *Meta-Analysis in Clinical Research*, 107 ANNALS INTERNAL MED. 224 (1987).

²⁵ See Sander Greenland, *Can Meta-Analysis Be Salvaged?*, 140 AM. J. EPIDEMIOLOGY 783 (1994); Sander Greenland & Keith O'Rourke, *On the Bias Produced by Quality Scores in Meta-Analysis, and a Hierarchical View of Proposed Solutions*, 2 BIostatistics 463 (2001).

²⁶ Eloise E. Kaizar, *Combining Information from Diverse Sources* (2006) (Ph.D. Thesis, Dep't of Statistics, Carnegie Mellon University); Eloise E. Kaizar, *MetaAnalyses Are Observational Studies: How Lack of Randomization Impacts Analysis*, 100 AM. J. GASTROENTEROLOGY 1233 (2005).

a number of different cancers, reporting a quantitative estimate that is a combination of the estimates from the individual studies, as well as measures expressing the uncertainty of such estimates.²⁷

No matter how sophisticated the meta-analytic technique, however, it is still limited to combining statistical evidence from different human studies into a single statistical estimate of the effect size, for example, the probability of the effect in the population exposed to the cause and the probability of the effect in the same population not exposed to the cause.

Meta-analysis of any form, however, cannot incorporate toxicological/mechanistic knowledge, nor can it easily factor in the quality of the studies being combined into a single estimate.

There is a technique, called the “Bayesian approach,”²⁸ for combining *all* available evidence, including a scientist’s background knowledge, judgment, etc., into a single judgment about the nature and uncertainty of a general causal claim.

B. The Bayesian Approach

Several forests have been sacrificed explicating and debating the pros and cons of the Bayesian approach,²⁹ so I will try to avoid piling on and provide only the briefest of sketches. There are many forms of the Bayesian approach, but the most appealing, in my view, is the most extreme. In this view, almost any statement, for example “TCE causes liver cancer,” or “your next child will be born with blond hair” or “the moon was formed by a collision between a proto-planet and the Earth” can be assigned a credence, or degree of belief between 0 and 1, and the degrees of belief can be interpreted as probabilities. In some cases the probability can be assigned objectively, for example, the objective probability of your next

²⁷ COMMITTEE ON ASBESTOS, INSTITUTE OF MEDICINE, ASBESTOS: SELECTED CANCERS 2 (Nat’l Academies Press 2006).

²⁸ Named after the Reverend Thomas Bayes, who lived during the first half of the eighteenth century.

²⁹ For a sampler, see JAMES O. BERGER, STATISTICAL DECISION THEORY AND BAYESIAN STATISTICS (2d ed. 1985); WILLIAM M. BOLSTAD, INTRODUCTION TO BAYESIAN STATISTICS (2d ed. 2007); ANDREW GELMAN ET AL., BAYESIAN DATA ANALYSIS (2d ed. 2004); COLIN HOWSON & PETER URBACH, SCIENTIFIC REASONING: THE BAYESIAN APPROACH (1993); RICHARD E. NEAPOLITAN, LEARNING BAYESIAN NETWORKS (2004); A.P. Dawid, *Probability, Causality and the Empirical World: A Bayes-De Dinetti-Popper-Borel Synthesis*, 19 STATISTICAL SCI. 44-57 (2004).

child being born with blond hair, or any of a number of heritable diseases, can be worked out by a genetics counselor with appropriate access to your family history and perhaps some of your blood. In other cases the probability corresponds to nothing more than a subjective degree of belief. For example, having little or no evidence to go on, in 2007 I may assign a probability of .2 to the statement: Blu-Ray will win the format wars over HD-DVD for the next generation of DVDs.

The approach provides a principled way to compute the probability one should assign to a hypothesis H , after you have seen a new piece of evidence E , notated as $P(H|E)$.³⁰ The fundamental theorem which drives the approach is extremely simple to state and prove:

$$P(H|E) = \frac{P(E|H)P(H)}{P(E)}$$

The numerator on the right involves $P(E|H)$, called the “likelihood” since it represents the probability of the evidence E given H is true, and $P(H)$, called the “prior,” the probability assigned to H prior to seeing the evidence. The denominator, $P(E)$, is the probability of the evidence without any consideration of the hypothesis H . The target, $P(H|E)$, is called the “posterior” as it reflects the probability of H after seeing the evidence E .

A classic use of the formula is in computing the probability of having a disease, given a diagnostic test result. For example, suppose a 20-year-old upper middle class heterosexual male Jim gets a blood test for HIV, and it comes out positive. Jim is scared, but what *is* the probability of H : that he is actually infected with HIV, given the evidence from the blood test E ? First suppose that the test is 98% reliable. That is, suppose that the probability of the test coming out positive given HIV infection, $P(E|H)$, is .98. Now Jim is truly terrified. Next suppose that $P(H)$, the prior probability of a 20-year-old upper middle class heterosexual male having HIV, is 1 in 1000 (.001). Finally, suppose that $P(E)$, the probability of a blood test coming out positive is 1 in 125 (.008). Then, to Jim’s relief, the posterior probability of HIV = .1225:

$$P(H|E) = \frac{P(E|H) = .98 * P(H) = .001}{P(E) = .008} = .1225$$

³⁰ Also referred to as the probability of H conditional on E , or H given E .

The apparent discrepancy between the 98% reliability of the HIV test and posterior probability of HIV of 12.25% is due to the low prior probability of HIV in Jim's cohort, and the frequency of positive blood tests. If one views the situation as a Bayesian, however, the evidence did make a big difference. Before taking the test, the probability of Jim having HIV was 1 in 1000. After the positive test, the probability moved to about 1 in 8, a huge jump.

The Bayesian approach has been used (and misused) in assessing forensic evidence like DNA testing in courtrooms for well over a decade.³¹

In our context, the potential utility of the approach should be apparent. Beginning with some prior belief over a general causal claim H_c , for example, that TCE causes liver cancer, then for each new piece of evidence, in the form of an RT, or epidemiological study, toxicological study, or what have you, one can use the Bayesian approach to compute the probability that H_c is true given this evidence. When all the evidence is in, we emerge with a posterior probability of H_c , that is, the probability of H_c given all the evidence available.

In practice, the Bayesian approach is *far* from a panacea. In typical scientific contexts involving a community of scientists, it is very difficult to move from inchoate and diverse sorts of background knowledge to a "prior," that is, a degree of belief in the causal claim prior to reviewing the evidence. Further, while "updating" to a posterior from certain kinds of evidence is reasonably straightforward, updating from other kinds of evidence is not. For example, consider computing the probability of liver cancer as a function of TCE exposure. After a few studies estimating the dependence of liver cancer risk on TCE concentration in work environment air, we can use the Bayesian approach to incorporate a new sample of 400 factory workers who were exposed to air with varying amounts of TCE concentrations for ten years, for example. As the probability of liver cancer as a function of TCE exposure is the hypothesis under study, the likelihood in Bayes' formula is objective. That is, the probability of seeing a particular frequency of liver cancers given that TCE does cause liver cancer, is objectively, mathematically derivable, and from that we can apply Bayes' theorem.

³¹ See, e.g., Joseph B. Kadane, *Misuse of Bayesian Statistics in Court*, CHANCE, Spring 2006, at 38.

What we cannot do, at least in any objective way, is to use the approach to update on evidence that shows rats exposed to 50 ppb TCE get tumors at three times the rate of those exposed to 1 ppb. What is the likelihood of this evidence, assuming TCE does cause liver cancer in humans? Here we are often beyond mathematics and statistics and into opinion about the comparability of rats and humans.³² Other features of studies, for example the quality of measures employed in an epidemiological survey, are also extremely hard to incorporate in any objective way into a Bayesian analysis. In general, this is referred to as the “objectivity of the likelihood” problem. Besides the sheer difficulty in performing the appropriate computations, the approach is a regulative ideal, but it is still far from the practical device we want to get us to a rational judgment on the truth of a general causal claim.

Finally, the probability of a hypothesis does not correspond in any simple way to the chances of it being true in the world. The justification for the technique is decision-theoretic, and is based more on a theory of rationality than it is on correspondence to the truth.³³

This is not to say that the technique is hopeless. It isn't. An actual example of using it for assessing a general causal claim relevant to the law comes from the National Academy of Science's (“NAS”) BEIR IV report, which sought to estimate the carcinogenicity of plutonium in humans.³⁴ By assuming that the ratio of carcinogenic potencies of plutonium to various other radionuclides like radium would be roughly constant across species, the Committee managed to combine very limited human data involving plutonium, with extensive animal data on plutonium and radium, and more extensive human data on radium, to emerge with a posterior over the hypothesis concerning the carcinogenicity of plutonium in humans.

³² This is not always the case. See William H. DuMouchel & Jeffrey E. Harris, *Bayes Methods for Combining the Results of Cancer Studies in Humans and Other Species*, 78 J. AM. STATISTICAL ASS'N 293 (1983).

³³ For a recent philosophical discussion of the epistemological view of the Bayesian approach, see LUC BOVENS & STEPHAN HARTMANN, *BAYESIAN EPISTEMOLOGY* (2003).

³⁴ COMMITTEE ON THE BIOLOGICAL EFFECTS OF IONIZING RADIATIONS, NATIONAL RESEARCH COUNCIL, *HEALTH RISKS OF RADON AND OTHER INTERNALLY DEPOSITED ALPHA-EMITTERS: BEIR IV* (Nat'l Academy Press 1988).




Figure 3 Posteriors over inches of rain on Labor Day 2010.

Although I have been focusing on the problems of combining different sorts of evidence and ignoring the issue of expressing the uncertainty over causal claims, this comes for free with the Bayesian approach. For simple propositions, like R: “it will rain on Labor Day 2010 in Brooklyn,” the posterior will be a probability between 0 and 1. If, in scientist A’s posterior R has a probability of .93, and in scientist B’s posterior R has a probability of .55, then A is in some sense more certain about R than B. Both scientists are more confident that R is true than that R is false—so they both in some sense believe R is true—but their degree of uncertainty is not the same. They should both take an even bet on R, but A would give much longer odds than B.

For more complicated hypotheses (e.g., “it will rain x inches on Labor Day 2010”), where we are asked to put a probability over each possible value of x , then the posterior is not a number but a probability *distribution*. For example, scientist A might have a posterior like the left side of Figure 3, while scientist B has a posterior like the right side of Figure 3. Roughly, the height of the graph corresponds to how much probability the scientist distributes over that number of inches of rain. The rectangle over 0 represents the probability of no rain (0 inches), which is .07 for Scientist A and .45 for Scientist B, while the rest of the posterior is distributed over rain from 0-1.2 inches on that day. If it does rain, then Scientist A seems to put the most probability over around .3 inches, while Scientist B, although he is less confident that it *will* rain, if it does, he deems it most probable to rain a little over .4 inches.

We can do the same with causal claims regarding the relative risk of liver cancer after 10 years of exposure to 10 ppb TCE exposure compared to no TCE exposure. A relative risk of 1.0 means that TCE has no effect on liver cancer. A relative




Figure 4 Posteriors over relative risk of liver cancer.

risk of 2.0 means that a person exposed to 10 ppb for 10 years has twice the probability of getting liver cancer as someone not exposed. In Figure 4, Scientist A believes that overall there is an effect of TCE, but it is relatively small, that is, a relative risk of around 1.3. Although her assessment of the size of the effect is small, she is quite confident that the relative risk is close to 1.3 as her posterior is narrowly distributed around 1.3. Scientist B seems to put the most probability around a larger effect, a relative risk of 2.0, but as her posterior is much wider and more diffuse, she is more uncertain about the size of the causal effect than is Scientist A.

So the Bayesian approach, although imperfect in many ways and practically always a challenge to apply, provides one way to synthesize the evidence and to express uncertainty about general causal claims that might be appealing to the Law.

V. WHAT'S THE LAW TO DO?

The final question to tackle is the hardest: what's the Law to do? Currently, the courts deal with complicated matters of causation in something like the following way. The judge must act as gatekeeper and decide which experts will be allowed to testify as to the scientific case for or against the general causal claim.³⁵ The plaintiffs then mount a case by summoning the experts (whom the judge allowed) to argue to the jury that the scientific evidence for the general and the individual causal claim is compelling. The defense then

³⁵ This is true since *Daubert v. Merrell Dow Pharmaceuticals, Inc.*, 509 U.S. 579, 592-93 (1993).

summons their own experts, who argue that the scientific evidence is not compelling, sometimes by impeaching the credibility of the plaintiff's experts, sometimes by emphasizing alternative evidence, etc.

Neither the judge nor the jury is trained to synthesize a diverse body of complicated scientific evidence, especially evidence presented by highly partial, highly sophisticated experts. A better system would have the community of scientists—who are presumably both less invested in the outcome and more qualified to rationally assess a wide body of complicated evidence—come to consensus as to the truth of the general causal claims at issue as well as the scientific uncertainty around these claims. This output from the scientific community could then be used as input to the legal system. This would not preclude plaintiffs or defendants from mounting their own experts and cases, but it would give the judge and jury a perspective on the science to fall back to when they are overwhelmed by the briefs submitted or the pyrotechnics in the courtroom.

Such a system is in fact already in place and used widely. For example, in decisions as to whether to compensate Vietnam veterans who were exposed to Agent Orange and now have some illness (like liver cancer), the Veteran's Administration does not restrict itself to hearings involving experts from both sides; rather, it consults a bi-annual report on the general causal claims true of Agent Orange exposure produced by a distinguished panel of independent scientists retained by the Institute of Medicine, a branch of the National Academies of Science.³⁶

The International Agency for Research on Cancer ("IARC"),³⁷ the U.S. Environmental Protection Agency,³⁸ the Centers for Disease Control and Prevention,³⁹ the National Institutes of Health,⁴⁰ and the National Toxicology Program⁴¹

³⁶ See COMMITTEE TO REVIEW THE HEALTH EFFECTS IN VIETNAM VETERANS OF EXPOSURE TO HERBICIDE, INSTITUTE OF MEDICINE, VETERANS AND AGENT ORANGE: UPDATE 2004 (Nat'l Academies Press 2005).

³⁷ INTERNATIONAL AGENCY FOR RESEARCH ON CANCER, *Preamble*, IARC MONOGRAPHS ON THE EVALUATION OF CARCINOGENIC RISKS TO HUMANS (W.H.O. 2006), available at <http://monographs.iarc.fr/ENG/Preamble/CurrentPreamble.pdf>.

³⁸ RISK ASSESSMENT FORUM, U.S. EPA, EPA/630/P-03/001F, GUIDELINES FOR CARCINOGEN RISK ASSESSMENT (2005), available at <http://www.epa.gov/iris/cancero32505.php>.

³⁹ CENTERS FOR DISEASE CONTROL AND PREVENTION, *supra* note 21.

⁴⁰ National Institutes of Health, News Release, *Fact Sheet: The "Report on Carcinogens" 9th Edition*, <http://www.nih.gov/news/pr/may2000/niehs-15.htm>.

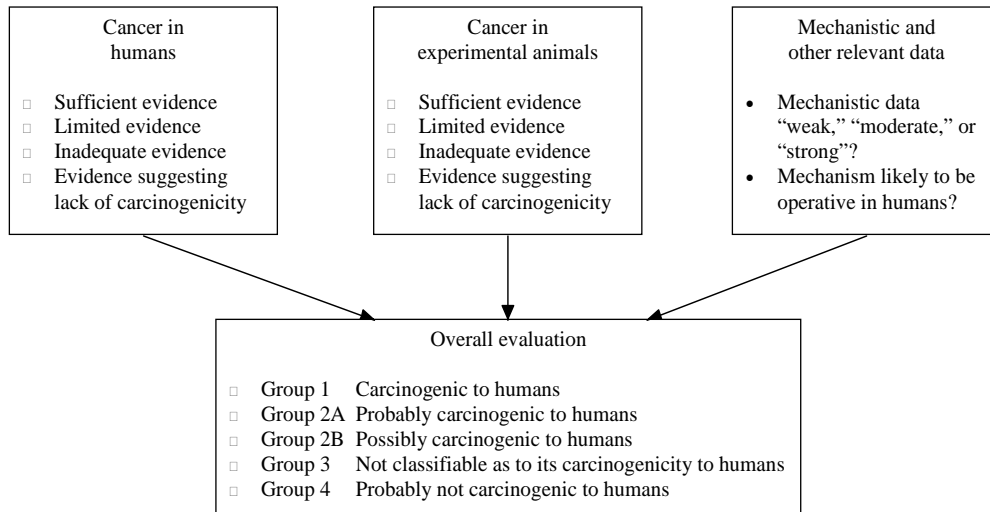


Figure 5 IARC categorization scheme.

have all developed systems for retaining panels of scientists to assess the status of general causal claims and to classify the level of evidence in support of such claims, particularly where such claims regard chemicals and cancer.

These panels rarely do a full Bayesian analysis to synthesize all the evidence, although as I illustrated with the NAS’s BEIR IV report, it is not unknown. Almost universally, however, they assign the causal claim to one of four or five discrete categories which differ on the truth of the claim and on the scientific uncertainty surrounding the claim.⁴²

For example, IARC forms committees consisting of biologists, epidemiologists, and toxicologists. They instruct these scientists to first categorize the level of evidence within three subcategories—human, animal, and mechanistic—and then to synthesize the subcategories of evidence into an overall evaluation on a five-category scale ranging from carcinogenic to probably not carcinogenic.⁴³ Figure 5 depicts the IARC scheme.

Each of the categories has natural counterparts to a Bayesian posterior, as explicitly described in the Institute of Medicine’s 2007 report, “Improving the Presumptive Disability

⁴¹ NATIONAL TOXICOLOGY PROGRAM (NTP), 11TH REPORT ON CARCINOGENS (U.S. Dep’t of Health and Human Servs. 2005).

⁴² They of course accompany this categorization with a long and inaccessible report.

⁴³ See IARC MONOGRAPHS, *supra* note 37, at 22-23.

Decision-Making Process.”⁴⁴ This sort of categorical output is accessible and comprehensible to a Judge and to a jury, and the steps to reach a consensus on the output are much better handled by relatively impartial scientists than by jurors trying to weigh evidence described to them by highly partial and well-paid experts.

The output of these panels regards the general causal claims that must be invoked in a legal case in which an individual causal claim is at stake, but the court must still decide the individual causal claim. In calculating the Probability of Causation (“PC”), or the Years of Life Lost (“YLL”), for example, the plaintiff must appeal to the general causal laws connecting the purported exposure and the injury it allegedly caused, but the court must still decide whether the estimate of PC of YLL is high enough, even allowing for uncertainty in this estimate, to warrant liability. This seems like more than enough complexity for juries to be asked to handle.

⁴⁴ DISABILITY DECISION-MAKING, *supra* note 17.