Variable Causation

1000: Introduction

In the middle of the twentieth century, Polio was a summer threat to every American child. Many died from the virus, and tens of thousands were left permanently paralyzed. Before 1950, there was no cure and there was no known method, short of totally isolating oneself from the rest of society, to avoid getting it. Working at the University of Pittsburgh around 1950, Jonas Salk developed a vaccine which provided immunity through the injection of "dead" polio virus. A short time later, a vaccine that contained disabled but live Polio virus that could be taken orally was developed, and was eventually given en masse to American children. Tragically, the live virus vaccine actually infected some children with the disease.

Imagine the debate that must have raged between those who favored using Salk's "dead" vaccine and those who favored using the "live" version.

- Mr. Smith: "My neighbor Mrs. Jones took her son Doug to the local clinic to be immunized with the experimental "live" Polio vaccine. He was fine when he went to that clinic, but was dead from Polio two weeks later. We can't be in the business of killing children! Use the dead vaccine."
- Mr. Wright: "We aren't in the business of killing children, Mr. Smith, we're in the business of saving them. We believe that the "live" vaccine is more effective, for two reasons. One, the "live" vaccine is better at promoting a truly protective immune response, partly because it is made of live (but disabled) virus. Two, more children will get the vaccine if it only requires swallowing a sugar cube, then if they have to get a shot (the dead vaccine). By using the dead vaccine, we will be killing the children who got Polio because the dead vaccine didn't fully immunize them, or those who got Polio because they were too scared of getting a shot and didn't think Polio could happen to them so they avoided going to the clinic at all."

This debate makes the difference between event and variable causation vivid. Mr. Smith and Mr. Wright have no dispute about what caused the event of Doug Jones' death. They both agree that Doug was infected with Polio from a live vaccine, and died from the infection. They both agree that there will be events of the following types if the "live" vaccine is deployed:

1. Child is immunized with "live" vaccine, exposed to Polio, resists it because of the vaccine.
2. Child is immunized with "live" vaccine, exposed to Polio, gets it anyway.
3. Child is immunized with "live" vaccine, not exposed to Polio, gets Polio from the vaccine.
4. Child is immunized with "live" vaccine, not exposed to Polio, doesn't get it.

and events of the following types if the "dead" vaccine is deployed:

1. Child is immunized with "dead" vaccine, exposed to Polio, resists it from vaccine.
2. Child is immunized with "dead" vaccine, exposed to Polio, gets it anyway.
3. Child is immunized with "dead" vaccine, not exposed to Polio, doesn't get it.

What they do not agree about is what constitutes the best **policy** regarding using the "live" or the "dead" vaccine. Their dispute is about which of two **interventions** made on an entire population of individuals is more desirable. Although ethical issues abound and cannot be ignored, part of their dispute should be settled with statistical evidence about which vaccine is more causally effective at combatting Polio. If it turned out that the "live" vaccine gave Polio to one or two children each year, but prevented approximately two thousand more cases of deadly Polio than did the "dead" virus, then the gain might arguably be worth it.

In this module we make the jump from causation among particular events to causation among variables (kinds of events) in populations. Causation among variables connects us to the sort of statistical and experimental data collected by social and biological scientists.

2000: Events to Kinds of Events

The sinking of the Titanic was a single event, and one of its causes was another single event: hitting an iceberg. (This is a bit of an idealization. sinking is not an instantaneous event, nor does it occur at a single point in space). Likewise, when we say: "The live Polio vaccine given to Doug Jones on May 1st, 1955 at noon caused a Polio infection which killed him," we mean that the particular vaccine administration event was responsible for that case of Polio. Each of these examples involves two particular events: one for the cause and one for the effect. We want more than particular claims concerning one cause event and one effect event, however. We want to generalize to how vaccine adminstration events cause Polio or its absence in childen overall.

We generalize from particular causal events to general causal claims in two steps. First we generalize from a particular cause and effect event involving a single individual to **events of a kind** involving a single individual. Then we move to events of a kind involving a **group of individuals**.

Consider a single cause event: Joel Smith rubs against poison ivy as he was walking through the woods on May 20th, 1999, and a single effect condition: Smith has an itchy rash that starts the next day and lasts a week.

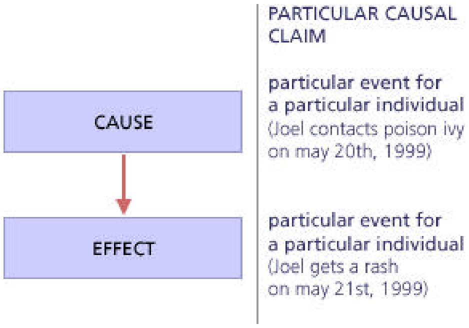

FIGURE 2000-1

The fundamental strategy for generalizing is to categorize separate events as being the same **kind of event**. For example, Smith might have rubbed against different poison ivy plants at different times. Although each event is separate, they are **events of a kind**: Smith rubbing against poison ivy. Similarly, the rash that appeared on Smith the evening of May 21st, 1999 was one event, but rash appearances like it on Smith are events of a kind for Smith. The individual remains the same, Smith, but similar events across different times can be lumped together as events of a kind. In so doing, we can make a general causal claim about an individual: "Whenever Joel Smith contacts poison ivy, it causes him to get a rash soon thereafter."

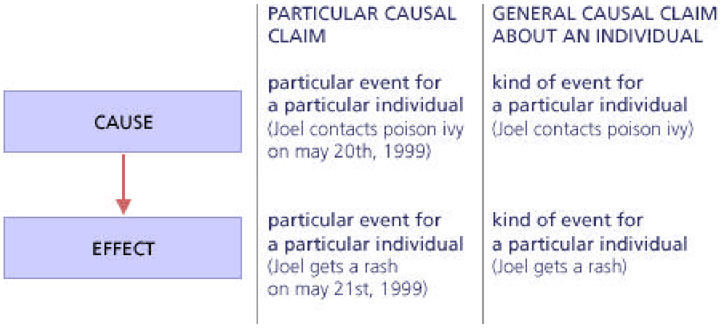

FIGURE 2000-2

A single event may be categorized as many different kinds of events. For example, the particular event: Joel Smith rubbed against poison ivy as he was walking through the woods on May 20th, can be categorized as: Smith rubbed against poison ivy, or Smith rubbed against a plant, or Smith walked through the woods.

The next move is to generalize from a single individual like Smith to a group of individuals, like adult males. Contacting poison ivy is a kind of event that any adult male can experience at any time, as is getting an itchy rash. What policy really requires, however, is a causal claim about a population (a group): "Contact with poison causes itchy rashes among adult males."

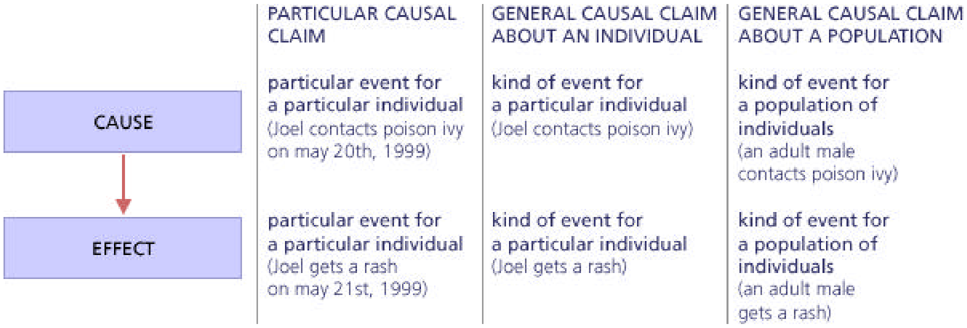

	PARTICULAR CAUSAL CLAIM	GENERAL CAUSAL CLAIM ABOUT AN INDIVIDUAL	GENERAL CAUSAL CLAIM ABOUT A POPULATION
CAUSE	particular event for a particular individual (Joel contacts poison ivy on may 20th, 1999)	kind of event for a particular individual (Joel contacts poison ivy)	kind of event for a population of individuals (an adult male contacts poison ivy)
EFFECT	particular event for a particular individual (Joel gets a rash on may 21st, 1999)	kind of event for a particular individual (Joel gets a rash)	kind of event for a population of individuals (an adult male gets a rash)

FIGURE 2000-3

So there are two steps in moving from event causation to general causal claims:

1. A **particular event** that occurs to a particular individual --> a **kind of event** that occurs to a particular individual at many times
2. A kind of event that occurs to a **particular individual** at many times --> a kind of event that occurs to many **different individuals** at many times.

In the module "Event Causation", we focused on particular causal claims about a single individual. Here we focus on general causal claims about a population of individuals.

In the table below, each row has a particular event, and then a kind of event for that individual of which the particular event is an instance, and finally the kind of event generally. We have left parts of two rows blank, and ask that you complete them in the exercises below.

TABLE 2000-1: PARTICULAR EVENTS AND KINDS OF EVENTS

Particular Event	Kind of Event for an Individual	Kind of Event Generally
Scheines climbs Giant Mountain in the Adirondacks on August 8th, 1999	Scheines climbing a mountain	Climbing a mountain
Jones watches the Steelers play the Cowboys on September 7th, 1997	Jones watches a football game	Watching a football game
The Exxon Valdez has an oil spill in Alaska in 1982	An oil spill for the Exxon Valdez	An oil spill
Bill Clinton's campaign for political office in 1992		
The U.S. Stock Market crash of October 17th, 1987		
Glymour purchasing a new car in August, 1999	Glymour purchases a new car	A used car purchase

< A link to exercises in the interactive version of this module. >

3000: Variables and Statistical Data

In the legal setting, when we are disputing the cause of a shooting accident, or the cause of a particular airplane crash, causes and effects are typically single events or conditions. In the policy setting, when we are debating whether to use the "live" Polio vaccine or the "dead" version, for example, causes and effects are typically **kinds of events**. Some part of a dispute involving policy about Polio vaccine ought to involve statistical data about the frequency with which the "live" vaccine accidentally causes infection. What does statistical data look like, and how does it relate to events, conditions, kinds of events, and a population of individuals?

When scientists collect data for statistical analysis, they don't write down events or conditions, they typically construct a data table in which each column is a **variable** and each row an individual. For example, in studying smoking and lung condition, they might collect data on four individuals and store the data in the following table:

TABLE 3000-1: DATA TABLE

Individual	# of Cigarettes Per Day	Lung Condition
Smith	5	Good
Jones	25	Fair
Gonzales	10	Good
Wright	30	Poor

Connecting this sort of statistical data table to general causal claims is, in some sense, the essence of "Causation and Statistics."

In statistical data tables, each row refers to a particular individual in some **population**. The population need not be people; it could be countries, cities, dogs, trees, or families, etc. Whether it is people or not, however, it should be a definite group. For example, all adult U.S. citizens is a population and all countries with any territory north of the equator is a population, but good looking people is not. The data tables actually analyzed rarely include every individual in population, however. The set of individuals the table does include is called a **sample** from the population. For example, if CNN declares that 48% of U.S. men favor tax cuts, they are describing a population that has over 100 million members. They usually go on to say that their numbers come from a random sample of about 1,000 men, which makes their overall results only accurate to within 3 or 4 %. Their data table had about 1,000 rows, and at least one column with a variable that indicates the individual's preferences about tax cuts.

The columns in the table are variables: that is, quantities that can take on different **values** for different individuals. The connection between events, conditions, kinds of events and variables is this: **kinds of events (conditions)** can be represented as **values of a variable**.

Variables have a range of possible values. For example, **Hair Color** is a variable with possible values of red, blonde, brunette, etc. Something as simple as a person's sex is a variable: it has two possible values: male and female. Some variables have only two possible values in no particular order, like **Sex**. Others have a small number of values that have a natural order, like **Drink Size**, with values = small, medium, or large, and others have an infinity of values, like **Height** in inches. Variables ought to have a range of values that correspond to a natural breakdown of the kinds of events or conditions that might obtain and that might make a difference.

For example, consider three kinds of events that relate to Polio vaccines:

1. Had "live" vaccine [live]
2. Had "dead" vaccine [dead]
3. Had no vaccine [none]

We might create a variable called **Vaccine Received** that can take on exactly one of these values, which we coded in the list above with the word in the square brackets, for example, [live]. An example of a simple data table involving five individuals in which we record the kind of Polio vaccine each received is as follows.

TABLE 3000-2: DATA TABLE

Individual	Variable: Vaccine Received
Lopez	live
Mitchell	live
Shaffer	none
Goodstein	dead
Sufez	dead

< A link to exercises in the interactive version of this module. >

What we really care about, of course, is connecting the kind of vaccine the child received with their health outcome, so a better table records two variables for each individual: **Vaccine Received** and **Health Outcome**. What values should **Health Outcome** range over? Perhaps [OK] and [Polio]. Then a table that recorded **Vaccine Received** and **Health Outcome** might look as follows.

TABLE 3000-3: DATA TABLE

Individual	Variable: Vaccine Received	Variable: Health Outcome
Lopez	live	OK
Mitchell	live	Polio
Shaffer	none	OK
Goodstein	dead	OK
Sufez	dead	OK

Lots of different quantities can be represented as variables. Lots of different kinds of events, or kinds of conditions, can be represented as a value of a variable. A person's age is a variable. The average amount of vitamin C a person consumes a day is a variable. The number of people who live in a city is a variable.

With data tables like the ones above, scientists aim to see if associations exist between events of one type (getting the live Polio vaccine), and events of another (not having Polio).

Enduring conditions can be treated in the same way. In the case of the Apollo 1 fire, for example, we might construct the following table to show the changes instituted between Apollo 1 and Apollo 2. The concentration of oxygen is a condition, not an event, yet it can easily be represented as the value of a variable.

TABLE 3000-4: DATA TABLE

Individual	Electrical Short	Oxygen Concentration	Fire
Apollo 1	True	100%	True
Apollo 2	False	18%	False

< A link to exercises in the interactive version of this module. >

4000: Causal Relativity

So far, we have gone from events to types of events and variables, and shown how data tables represent a sample of individuals and their values for some set of variables. We now move to general causal claims.

Individual causal claims, like "Hitting an iceberg caused the Titanic to sink," involve single events, and some sort of counterfactual relationship between them, for example: "If the Titanic hadn't hit an iceberg, it wouldn't have sunk." General causal claims, like "Smoking causes cancer," involve variables, like **Smoked** [heavily, moderately, no] and **Got Cancer** [yes, no]. What remains is to describe the sort of relationship that should hold between two variables in order for us to say that one variable is a cause of the other.

Claims about causation among variables are relative in at least two respects. First, causal claims are relative to a set of background conditions that are usually left unstated. Second, claims about the causal relations between two variables **X** and **Y** are relative to the set of variables **Z** that are explicitly under discussion (we put sets in bold face).

Causal claims are always to be understood as **relative to a particular set of background conditions**. For example, a claim about whether the level of violence in the mass media is a cause of the level of violence in our society will be viewed quite differently if the background is today's society as opposed to Egypt in 1800 BC. Today the average child spends something like 5 hours a day watching TV or connected to the internet, and in 1800 BC there was no mass media.

Similary, consider the claim: "Eating red meat several times a week reduces the length of one's life." In America today, where diets tend to contain lots of calories and lots of calories from fat, we might accept this causal claim as true. In Mongolia around 1500, however, eating red meat more than once a week would almost certainly increase life span, mostly because the normal diet was incredibly spare in calories and in fat. Whats the difference? The background conditions, which in this case include the "normal" diet for the time and place. Although the background conditions are typically not even mentioned, causal claims take on meaning only against some set of background conditions.

< A link to exercises in the interactive version of this module. >

Claims about the causal relationshiops between one variable **X** and another **Y** must also be judged relative to the set of variables **Z** that are explicitly under consideration. For example, consider four variables about matches: **Match Color** [blue, red], **Match Struck** [yes, no], **Match Tip Temperature** [above 300, below 300], and **Match Lit** [yes, no]. Asked whether striking a match is a direct cause of the match lighting, that is, whether the variable **Match Struck** is a direct cause of the variable **Match Lit**, the answer depends on the set of variables considered.

If the set is either:

- {**Match Color**, **Match Struck**, **Match Lit**}, or
- {**Match Struck**, **Match Lit**}

then the answer is yes. If the set includes **Match Tip Temperature**, however:

- {**Match Color**, **Match Struck**, **Match Tip Temperature**, **Match Lit**}, or
- {**Match Struck**, **Match Tip Temperature**, **Match Lit**}

then the answer is no. In these sets, **Match Struck** is only an indirect cause of **Match Lit**. **Match Struck** is a direct cause of **Match Tip Temperature**, which in turn is a direct cause of **Match Lit**.

We call the set of variables explicitly under consideration the **causal system**. The variables one includes in the causal system are in some sense arbitrary, but in order to be useful, a causal system should include as many potential causes as possible, and not include variables known to be irrelevant.

< A link to exercises in the interactive version of this module. >

So causal claims are relative to:

1. the background conditions, usually left tacit, and
2. the causal system, that is, the set of variables explicitly under consideration.

5000: General Causal Claims

5100: Causal Assignments

To build an account of causation among variables, we need to introduce two pieces of terminology: **causal assignments** and **response structures**.

Imagine that you went to your friend's house for dinner, and you are taking out the garbage, which leads you through her garage. When you enter the garage, it's pitch black, but next to the door you notice two light switches and a button. You need to get a light on and then open the electric garage door. You have two causal questions - what causes the garage light to go on, and what causes the garage door to go up. Lets consider both in turn, and build the ideas behind **causal assignments** and **response structures** as we do so.

Assume that the background conditions include working light bulbs in all her lights, and working electricity throughout her property. If you assume that the button is either for the garage opener or something else, but completely irrelevant to the light, then the first causal question can be represented with three variables: **Switch 1** [up, down], **Switch 2** [up, down], **Garage Light** [on, off].

< A simulation in the interactive version of this module. >

In this situation, and many like it, one of the variables is the "effect" we care about. Here it is the **Garage Light**, and the event we desire is: **Garage Light** = on. From common sense, we have narrowed the potential causal variables of this effect to **Switch 1** and **Switch 2**. We don't know, however, how the switches should be set to produce the effect. Put in the language of variables, we don't know which switch is a cause, and we don't know the **values of the causal variable(s)** that will produce the **value of the effect variable** that we want. We use the phrase **causal assignment** to refer to how we **set** the causes. More formally, if a causal system is just any set of variables, then we define causal assignment as follows.

Definition: Causal assignment

A **causal assignment** is one particular assignment of values to all the variables in a given causal system except one variable designated as the effect.

Since there are two switches, and each has two positions, there are four possible causal assignments in the causal system involving the garage light and the two light switches:

TABLE 5100-1: Causal Assignments for Switch 1 and Switch 2

Assignment	Switch 1	Switch 2
1	down	down
2	down	up
3	up	down
4	up	up

Although the system involves three variables, we designated the garage light as the effect. We call the other two **potential causal variables**.

The number of different causal assignments for a given effect grows with both the number of potential causal variables in the system and the number of values for each. In our example, there are only two potential causal variables, each with two values, so the number of possible causal assignments is four. If there were three switches, the number would be eight, with four switches it would be 16, and so on. If the two switches had three positions each, there would be nine different assignments. In general, if there are N variables, then the number of causal assignments is equal to:

(# of values for Var. 1) X (# of values for Var. 2) X ... X (# of values for Var. N)

A causal assignment is not merely the observation that particular variables have particular values. Rather, it crucially involves the idea of **forcing** the variables to have a particular set of values. We are **assigning values, not observing them**, and the difference is crucial to causal reasoning.

In the case of the light switches, assigning values to the switch positions is easy to imagine and to carry out. Consider a different example, however, in which we want to know the causes of malaria, a tropical disease that produces high fever and sometimes death. The effect variable is: **Got malaria** [yes, no], and we want to know what variables cause it. At a certain point in history, the British believed that quinine, which comes in tonic water, prevented malaria. Suppose that, according to our background knowledge, the following variables are candidates for causes of malaria:

- **Bitten by a mosquito** [yes, no]
- **Had a malaria inoculation** [yes, no]
- **Has the sickle cell gene** [yes, no]
- **Drank gin and tonics regularly** [yes, no]

What are the possible causal assignments for these variables?

TABLE 5100-2: Causal Assignments for Malaria

Assignment	Bitten by a mosquito	Had a malaria inoculation	Has the sickle cell gene	Drank gin and tonics regularly
1	yes	yes	yes	yes
2	yes	yes	yes	no
3	yes	yes	no	yes
4	yes	yes	no	no
5	yes	no	yes	yes
6	yes	no	yes	no
7	yes	no	no	yes
8	yes	no	no	no
9	no	yes	yes	yes
10	no	yes	yes	no
11	no	yes	no	yes
12	no	yes	no	no
13	no	no	yes	yes
14	no	no	yes	no
15	no	no	no	yes
16	no	no	no	no

To learn which causal assignment(s) produces malaria, we must be able to intervene and produce each assignment. Although easy to do with light switches, its not so clear how we would assign someone to have the sickle cell gene. Luckily, computers can simulate almost anything, so we programmed an imaginary individual named Sam that you may change as if you were all powerful. When you click on the button below, a window will display with the different settings Sam can have for each variable. Intervene to give Sam a particular value for a variable by clicking on the circle next to the value, and the bar at the bottom will tell you whether Sam gets malaria. Try to figure First, figure out which settings of these factors give Sam malaria. Then try to figure out which variables cause or prevent malaria.

< A link to a Java applet in the interactive version of this module. >

By enumerating all of the possible causal assignments for Sam, and testing each one, you could have discovered the entire response structure (the causal laws) we programmed into the software to simulate the cause-effect system for malaria. Here is what you should have found.

TABLE 5100-3: RESPONSE STRUCTURE FOR THE MALARIA CASE

Assignment	Variable 1: BITTEN	Variable 2: INNOCULATED	Variable 3: HAS GENE	Variable 4: DRINKER	Effect: MALARIA
1	True	True	True	True	False
2	True	True	True	False	False
3	True	True	False	True	False
4	True	True	False	False	False
5	True	False	True	True	False
6	True	False	True	False	False
7	True	False	False	True	True
8	True	False	False	False	True
9	False	True	True	True	False
10	False	True	True	False	False
11	False	True	False	True	False
12	False	True	False	False	False
13	False	False	True	True	False
14	False	False	True	False	False
15	False	False	False	True	False
16	False	False	False	False	False

< A link to exercises in the interactive version of this module. >

5200: Response Structures

In the simulation of Sam, you can produce each of the 16 possible causal assignments among the four potential causal variables, and in each case observe whether or not Sam got malaria. The simulation allows you to examine how Sam responds to each causal assignment, that is, the structure of his response. By building a table that lists his malarial response for each possible causal assignment, you can uncover the causes of malaria. In fact, the route to a general account of causation among variables goes through tables like this, tables that we call response structures.

The **response structure** for an effect describes the effect in every possible causal assignment of the other variables in the system. For example, consider the example of the light switches and the garage light again. We listed the possible causal assignments in a table with a column for each potential cause, as so.

TABLE 5200-1: Causal Assignments for Switch 1 and Switch 2

Assignment	Switch 1	Switch 2
1	down	down
2	down	up
3	up	down
4	up	up

To turn this table into a response structure for the system in which the two light switches are the potential causes and the garage light is the effect, simply add a column for the effect, as so.

TABLE 5200-2: Empty Response Structure for the Garage Light

Assignment	Switch 1	Switch 2	Effect: Garage Light
1	down	down	???
2	down	up	???
3	up	down	???
4	up	up	???

We put "???" into the table to indicate that we don't know the response yet, thus we call it an empty response structure. Suppose that after experimenting with the switches, we discovered that the response structure is as follows.

TABLE 5200-3: Response Structure for the Garage Light

Assignment	Switch 1	Switch 2	Effect: Garage Light
1	down	down	Off
2	down	up	On
3	up	down	Off
4	up	up	On

The response structure contains all the information needed to determine what variable causes the garage light.

Switch 2 controls the garage light. In every causal assignment (2 and 4) in which it is "up," the light is on, and in every assignment in which it is off (1 and 3) the light is off. For Switch 1, however, there are two causal assignments in which it is up (3 and 4), but in one the light is on and in another the light is off.

Now you get a chance to practice constructing response structures. Consider the response structure for a hypothetical person named Willie, with the effect "Willie gets Lyme disease," and these three potential causes:

- **Bitten by a tick** [yes, no]
- **Has a low red blood cell count** [yes, no]
- **Innoculated for Lyme disease** [yes, no]

Clicking on the top button in the applet below will bring up a window with buttons to let you assign particular values to thevariables, at which point you can then find out whether Willie has Lyme disease. Give Willie each of the eight possible causal assignments, and then see whether he has Lyme disease in each assignment. After you have done the eight "experiments" on Willie, click on the bottom button to bring up a window with a matrix for the response structure for whether Willie gets Lyme disease. To fill it in, go through each possible causal assignment, and click on the box in the Effect column for that assignment to appropriately fill in the response structure.

< A link to a Java applet in the interactive version of this module. >

< A link to a Java applet in the interactive version of this module. >

Response structures contain no information about the mechanisms by which causes influence their effects. The response structure for Willie, for example, leaves out information about the mechanisms by which Willie contracts Lyme disease. The response structure for the garage light also leaves out any description of the mechanism. Nevertheless, if you want knowledge of how to control the light, or how to prevent malaria or Lyme disease, you don't necessarily need to know the mechanisms. You need to know what variables cause the effect, and what values to assign those variables so that you can produce the value of the effect you want.

< A link to exercises in the interactive version of this module. >

Consider the same ideas for a new case. Suppose I tell you that Tim is a friend of mine who currently has no stomach problems whatsoever. But I want to know the circumstances in which Tim might get an ulcer, and those circumstances where he won't get an ulcer. Given Tim's body and the nature of the world he lives in, there is a set of causal mechanisms that determines whether Tim will get ulcers (or not) depending on the causal assignment that happens to him. That is, there is some set of physical and chemical and biological states of affairs that determines whether or not Tim gets an ulcer. However, we almost certainly cannot figure out all of the details of those mechanisms.

Instead, to gain an understanding of when Tim will get an ulcer, we use our representation of the response structure: an enumeration of whether the effect occurs in **all possible causal assignments**. First, we want to pick out a set of variables that contains all of plausible potential causal factors. Once we determine this set, then we can examine all of the possible causal assignments for a given effect, and catalogue the value of the effect in each. For the purposes of this exercise, we will assume that our set contains four causal variables and one effect. Of course, there are many other potentially causal factors that have been left out. However, to simplify the discussion, we will assume the variables are:

TABLE 5200-4: FACTORS INVOLVED IN GETTING AN ULCER

Variable	Values
Tim takes aspirin each day	[True/False]
Tim takes Tagamet each day	[True/False]
Tim drinks scotch each day	[True/False]
Tim drinks coffee each day	[True/False]
Tim gets an ulcer	[True/False]

Use the top applet to try to figure out how Tim responds in each causal assignment (by clicking on the button, and then using the familiar buttons). Then, use the bottom applet to fill in the entire response structure for Tim relative to this set of variables (by clicking on the button, and then filling in the matrix).

< A link to a Java applet in the interactive version of this module. >

< A link to a Java applet in the interactive version of this module. >

5300: Causation Among Variables

Causal assignments and response structures provide the tools we need to give an account of variable causation. In this section we finish the job.

General causal claims have the following form:

+ Variable **X** is a cause of variable **Y** relative to the set of variables **Z** under background conditions **B**.

Leaving the background conditions aside, when are claims like this true and when are they false? The answer, informally, is that if there are any two causal assignments that are identical **except** for what value of **X** is assigned, and there is a difference in the effect **Y**, then **X** is a cause of **Y**.

Consider again the simple response structure for the garage light, where we said that the variable **Switch 2** is a cause of the variable **Garage Light**. In this example, the variable **Switch 2** takes the role of the variable **X** in the informal answer above. The variable **Garage Light** takes the role of the variable **Y**, and the set of variables {**Switch 1**, **Switch 2**, **Garage Light**} takes the role of the set **Z**. So the question to answer is:

Are there any two causal assignments that are the same except for what value of **Switch 2** is assigned such that there is some difference in the **Garage Light**?

Let's first unpack what we mean by two causal assignments that are the same except for what value of **Switch 2** is assigned. Consider again the response structure for this system.

TABLE 5300-1: Response Structure for the Garage Light

Assignment	Switch 1	Switch 2	Effect: Garage Light
1	down	down	Off
2	down	up	On
3	up	down	Off
4	up	up	On

< A link to exercises in the interactive version of this module. >

Causal assignments 1 and 2 are the same **except** for the value of **Switch 2** assigned. Everything else in the system is assigned the same value (**Switch 1** = down), but **Switch 2** is down in assignment 1 and up in assignment 2. The same goes for causal assignments 3 and 4, except that in this pair **Switch 1** = up. **Switch 2** is down in assignment 3 and up in assignment 4. The idea, first made clear by the 18th Century philosopher John Stuart Mill, is to hold everything else in the system constant, and **only** vary the cause. If there is a difference in the effect, then we have causality.

The idea of two causal assignments that match except for one variable is important enough to name.

Definition: Test Pair of Causal Assignments

If two causal assignments C1 and C2 are identical except for the values assigned to variable **X**, then C1 and C2 are a **test pair of causal assignments** for **X**.

Now we can be a little more precise about variable causation.

Definition: Direct Cause

If, in a system of variables **S**, there are any test pair of causal assignments for **X** in which there is a difference in the effect **Y**, then **X** is a **direct cause** of **Y** relative to **S**.

In the garage light system, we agreed that **Switch 2** is a cause of the **Garage Light** but **Switch 1** is not. Applying our definition of variable causation ought to yield at least this result.

We already established that **Switch 2** is a cause of the **Garage Light**. Why? Because there is a test pair of causal assignments for **Switch 2**, (assignments 1 and 2) such that the effect (the light), is different in these assignments. Hopefully, applying the definition to **Switch 1** ought to tell us that we do **not** have causation.

To establish the lack of causation, it must be the case that for **every** test pair of causal assignments for **Switch 1**, there is **no** difference in the effect.

< A link to exercises in the interactive version of this module. >

Causal assignments 1 and 3 are the same **except** for the value of **Switch 1** assigned, as are assignents 2 and 4. So the effect had better be the same for pair 1 and 3, and for pair 2 and 4, and it is. For both 1 and 3, the light is off, and for both 2 and 4 the light is on. Thus, there are no test pairs for **Switch 1** such that there is any difference in the garage light, and thus the variable **Switch 1** is not a cause of the variable **Garage Light** in this causal system.

In the definition of variable causation, we say there is causation if there are **any** test pairs that make a difference to the effect. Why do use **any**? Why not demand that the effect be different for all test pairs? Consider a simple causal system involving a switch, a battery, and a light bulb.

Consider the light bulb below. Click on the red battery to toggle between charged (solid red battery) and uncharged (broken line through the battery), and click on the switch to open or complete the circuit.

< A simulation in the interactive version of this module. >

Considering the light bulb to be the effect, here is the response structure for this system.

TABLE 5300-2: Response Structure for the Garage Light

Assignment	Switch	Battery	Effect: Light Bulb
1	down	charged	Off
2	down	uncharged	Off
3	up	charged	On
4	up	uncharged	Off

< A link to exercises in the interactive version of this module. >

There are two test pairs for the **Switch**, 1 and 3, and 2 and 4. For the pair 1 and 3, the battery is charged, and for the pair 2 and 4, the battery is uncharged. The effect is different in pairs 1 and 3, because the switch makes the light bulb go on and off when the battery is charged. The effect is not different in pairs 2 and 4, because when the battery is uncharged the switch can't change the state of the light bulb. Nevertheless, we still want to say that the switch is a cause of the light bulb. To do so, we must allow that if there are **any** test pairs that make a difference to the effect, we have causation.

Lets restate the definitions for test pairs of causal assignments and for variable causation, and practice applying them to the malaria system.

Definition: Test Pair of Causal Assignments

If two causal assignments C1 and C2 are identical except for the values assigned to variable **X**, then C1 and C2 are a **test pair of causal assignments** for **X**.

Definition: Direct Cause

If, in a system of variables **S** there are any test pair of causal assignments for **X** in which there is a difference in the effect **Y**, then **X** is a **direct cause** of **Y** relative to **S**.

Here is the response structure for malaria.

TABLE 5300-3: RESPONSE STRUCTURE FOR THE MALARIA CASE

Assignment	Variable 1: BITTEN	Variable 2: INNOCULATED	Variable 3: HAS GENE	Variable 4: DRINKER	Effect: MALARIA
1	True	True	True	True	False
2	True	True	True	False	False
3	True	True	False	True	False
4	True	True	False	False	False
5	True	False	True	True	False
6	True	False	True	False	False
7	True	False	False	True	True
8	True	False	False	False	True
9	False	True	True	True	False
10	False	True	True	False	False
11	False	True	False	True	False
12	False	True	False	False	False
13	False	False	True	True	False
14	False	False	True	False	False
15	False	False	False	True	False
16	False	False	False	False	False

< A link to exercises in the interactive version of this module. >

6000: Populations

6100: Individuals and Populations

Often we want to answer questions about particular instances of causation, for example: What caused my child to have musical talent? Why did the Apollo 1 capsule burn explosively? Why did Sam get malaria? The answer to these questions depends on:

1. The **response structure** governing the individual.
2. The **particular causal assignment** that occurred.

In other situations, the questions we want to answer involve generalizations about causation in a population of individuals. Does lead in paint cause brain damage among children? Do school vouchers help children from poor households learn more from school? Does the death penalty deter criminals from committing crimes? Is a particular insecticide effective at controlling fruit flies?

The answers to these questions, which are questions about variable causation, depend on only the response structures for the individuals in the population of concern. For example, consider the question about insecticide and fruit flies. In California in the early 1990s, a fruit fly epidemic wiped out a substantial portion of the fruit crop and had to be controlled, lest California's economy suffer long term disaster. The fruit fly population in California consisted of millions of fruit flies. Each might have the same response structure, or they might vary. In either case, the state government of California needs to know something about the **response structure of the population**.

In the real world, at any given moment, one individual can have only one causal assignment. For that individual at that moment, the effect that occurs in the causal assignment depends on the response structure for that individual. Consider an example:

A fruit fly in a petri dish at a corporate lab in St. Louis (call him "Jeff") dies 22 hours after being sprayed with a new insecticide. For the chemical firm, two important pragmatic questions are:

- Was the spraying with the insecticide the cause of Jeff's death?
- If it was the cause of Jeff's death, will it also cause other fruit flies in California to die?

The first question involves Jeff's response structure. The second question is about whether the response structure for Jeff is representative of the response structure of fruit flies in general when put in this environment. We explore how we might answer both questions.

Question 1: Was the spraying with the insecticide the cause of Jeff's death?

Think back to the lessons learned in the module on Event Causation to establish what this question is really asking.

< A link to exercises in the interactive version of this module. >

Here's another way to describe what the company should want to know using the language we defined above:

They should want to know if there is a test pair of causal assignments for the spraying of insecticide such that when Jeff is not sprayed he lives but when he is sprayed he dies.

Representating Jeff's response structure requires enumerating all of the possible causal assignmenst for Jeff. Observing every possible causal assignment for one individual like Jeff is not possible, however. Having observed the death of Jeff in one assignment, we can't start over and see what happens in another causal assignment. In our simulations, we **can** start over, but in the real world we cannot.

Use the same artificial device of computer simulation we have been using all along to discover Jeff's response structure for dying with respect to these causal factors:

- **Sprayed with insecticide** [yes, no]
- **Handled with tweezers** [yes, no]
- **Have a parasite** [yes, no]
- **Exposed to low temperature** [yes, no]

< A link to a Java applet in the interactive version of this module. >

< A link to exercises in the interactive version of this module. >

6200: Response Structure Uniformity

Our simulations tell us about Jeff's response structure, but they don't tell us about fruit flies in general. What scientists do, however, is extrapolate from what happens to individual fruit flies, or from what happens to a small sample of fruit flies, to what happens in the whole population. In this case:

Question 2: If the insecticide was the cause of Jeff's death, will it also cause other fruit flies to die?

This question can be described as a question about the relationship between the response structure that governs Jeff in the corporation's St. Louis lab environment, and the response structure that governs each other fruit fly in the California environment. On the simplest level, the question is: Are all of the other fruit flies governed by the same response structure as Jeff? If they are, then they will die in the same causal assignments as Jeff. So, if they are all governed by the same response structure, then they will die if sprayed with the insecticide.

The concept we are after is called response structure uniformity.

Definition: Response Structure Uniformity

A population has **response structure uniformity** for a given effect if every individual in the population has the same response structure for that effect.

Let's investigate this concept using a very abstract world:

In this abstract world there are two squares that can each have one of two values of a variable color: one lighter than the other. You can change the colors from one state to the other by clicking on either of these two "causal factor" squares. They are causal factors for the color of a third square. The color that the third square has is determined by the values of the other two: the color of the third square is the "effect" in this abstract causal system. In this case, the causal factors are the red and green squares, and the effect variable is the blue square. Let's call this entire system of three squares a "triad" and call the particular triad in the question below "Triad #37".

Work with this simulation to determine the response structure of Triad #37.

< A simulation in the interactive version of this module. >

< A link to exercises in the interactive version of this module. >

Even though the triads look as if they are the same kinds of things, they clearly can have different underlying response structures.

The first population of triads (Triads #32, #10, #25, #12), for example, has **response structure uniformity**. Indeed, they have the same response structure as the first trial you worked with (Trial #37). The second population (Triads #18, #19, #114, #2) is not uniform with respect to the response structure in question. In the real world, we also have this same challenge in moving from what we learn about individuals (like Triad #37) to a population of such individuals. Hence our question:

If the insecticide was the cause of Jeff's death, will it also cause other fruit flies to die?

The answers are:

- Yes, among those fruit flies that are governed by the same response structure as Jeff.
- Among those fruit flies that are not governed by the same response structure as Jeff, we don't know.

Response structure uniformity is **not** the assumption that all individuals in a population are in the same causal assignment. It is the assumption that those individuals who do find themselves in the same causal assignment will have the same response.

We expect individuals in any population to be distributed among a variety of causal assignments. The response structure uniformity assumption says that they all have the same lawlike connection between causes and effects, regardless of their individual causal assignments. Therefore, if response structure uniformity holds, and the causes fully determine their effects, we can easily compute how many individuals in a population will have a particular effect, if we know how many individuals are in each causal assignment. We simply add together all of the individuals who will exhibit the effect (given their causal assignment and the response structure).

Scientifically, the response structure uniformity assumption is sometimes plausible, and sometimes not. Its plausibility depends on the population we're investigating, the background conditions, the effect we're interested in, and so on. We'll talk more later about the violation of this assumption. For now, use the questions below to help focus your thinking about the plausibility of this assumption for various populations in various circumstances.

< A link to exercises in the interactive version of this module. >

In some ways, what is at stake here is what we choose to include in a population of the "same" individuals. It makes sense to identify a population of Honda Civics (or fruit flies) as being the "same" kind of individuals, but they are of course the same in some ways and different in others. Civics are the same in basic structure, but different in the material used for their seats or the color of their paint.

One way to talk about this issue is to talk about different sub-populations (e.g. the sub-population of Civics with vinyl upholstery) in which response structure uniformity does hold even when it doesn't in a larger population of such individuals (like the entire population of Civics). The following section discusses this way of viewing the problem.

7000: Distributions of Response Structures

The response structure uniformity assumption is easily violated. For example, consider the following two-variable system: whether a person shaves their face during a four-week period, and whether a beard grows on their face. If we consider the entire U.S. population, the response structure uniformity assumption clearly fails. For men, not shaving for four weeks (usually) leads to growing a beard, whereas women won't grow a beard regardless of whether they shave. In other words, half of the population is governed by a response structure in which not shaving causes beard growth, and the other half is governed by a response structure in which neither variable causes the other.

It's important to understand the difference between differences among individuals with respect to response structures, and differences with respect to causal assignments. If two individuals are in different causal assignments, then there must be at least one variable whose value differs between them. But if two people are governed by different response structures, then there must be at least one causal assignment that leads to one response in one individual but a different response in the other. In other words, two individuals can differ either in their causal assignment, or in their response structure, or both, or neither. Knowing that two individuals are in the same (or different) causal assignment tells us nothing about whether they are governed by the same response structure. And likewise, knowing that they are governed by the same (or different) response structures tells us nothing about their particular causal assignments.

Causal assignments need not be distributed evenly throughout a population. Some causal assignments might occur frequently, and some rarely. For example, if the pesticide company decides, from experiments in their lab, to do a trial test of their pesticide in the field, they might spray one or two farms in California with the pesticide, but leave the great majority of farms alone. In that case, only the fruit flies on the test farms will be in a causal assignment that includes being sprayed with insecticide.

Likewise, response structures need not be distributed evenly throughout a population. Some response structures might occur frequently, and some rarely. For example, if the beard growing experiment was done only on males, but included 5% of Asian or Native American descent (who do not grow facial hair), then the population includes 95% of its individuals with one response structure, and 5% with another.

Although the response structure uniformity assumption might be violated in a given population, we can often pick out sub-populations in which the assumption holds. In the beard growing example, there are clearly (at least) two sub-populations (men and women), and the response structure uniformity assumption almost certainly holds in the sub-population of women. For this set of variables, all (or almost all) women have the same response structure. For another example, consider allergic reactions. Two people can be put into the same causal assignment (e.g., being around a long-haired cat), and one might suffer an allergic reaction while the other is unaffected. Assuming the response structure fixes what will happen, we can naturally conclude that there are different response structures governing the two people. And it is quite natural for us to group together all those people that do suffer an allergic reaction, and put into a separate group those people who do not have a reaction. In other words, we divide the population into two sub-populations, and the response structure uniformity assumption then holds in each sub-population.

8000: Summary

Particular causal claims involve particular events among particular individuals. For example, the Titanic sank because it hit an iceberg. The most general causal claims involve similar kinds of events or conditions among populations of individuals. For example, higher levels of education cause higher levels of income (in general) among Americans. Variables are quantitites which take on one of a possible range of values for a given individual. For example, **Height** is a variable which takes on a value in inches for every person in a given population. Kinds of events can be represented by values of a variable. Thus general causal claims in a population involve relationships between variables.

Causal reasoning takes place relative to background conditions as well as a specified set of variables (called the causal system). For a given effect variable, a **causal assignment** is an assignment of a particular value to each variable in the causal system besides the effect. A **response structure** for an effect is a description of what valuethe effect variable takes on in every possible causal assignment.

If two causal assignments C1 and C2 are identical except for the values they assign to a variable **X**, then we say C1 and C2 are a test pair of causal assignments for **X**. If there is **any** test pair of causal assignments for **X** that produce a different response in an effect variable **Y**, then **X** is a cause of **Y**.

A population consists of a set of individuals. For a given set of variables, each individual in the population can be given one of many posssible causal assignments. So individuals in a population might vary as to their causal assignment. Individuals may also differ, however, in their response structure. In the real world it is often the case that even when we can achieve experimental control sufficient to effectively make causal assignments, we cannot put the same individual in more than one causal assignment. Since, in order to observe enough about an individual to deduce their response structure we must be able to put them into every possible causal assignment, this threatens to severely limit our ability to infer anything about response structures. The way out is an assumption we call **response structure uniformity**: every individual in the population is governed by the same response structure. This assumption allows us to observe many individuals in many different causal assignments, and treat it the same as if we had observed a single individual in many different causal assignments.