

Actual Causation and Simple Voting Scenarios¹

Jonathan Livengood

**History and Philosophy of Science
University of Pittsburgh**

Actual Causation and Simple Voting Scenarios

Abstract. Several prominent, contemporary theories of actual causation maintain that in order for something to count as an actual cause (in the circumstances) of some known effect, the potential cause must be a difference-maker with respect to the effect in some restricted range of circumstances. Although the theories disagree about how to restrict the range of circumstances that must be considered in deciding whether something counts as an actual cause of a known effect, the theories agree that at least some counterfactual circumstances must be considered. I argue that the theories are still too permissive in the range of counterfactual circumstances they admit for consideration, and I present simple counter-examples that make use of this over-permissiveness.

Harold enjoyed playing basketball. Knowing that he had a heart condition, Harold asked himself what effect playing basketball would (likely) have on his heart. The question Harold asked himself is about the singular effect of a given singular cause. If Harold had been a statistician, he might have estimated the expected effect his playing basketball would have on the condition of his heart by testing causal models on data drawn from a large population of diverse people.² Since he was not a statistician, Harold thought about his situation as best he could without data or models.

After deliberating, Harold decided to go ahead and play basketball with some friends. While playing, he collapsed from a heart attack. Harold was rushed to the emergency room, where the doctors saved his life. While he was recovering, Harold's wife, Helen, asked the doctors, "Was Harold's heart attack caused by his playing basketball?" Helen's question is not like Harold's question about the singular effect of a given singular cause; rather, Helen's question is about the singular cause of a given singular effect. In Harold's case, the (potential) cause is known, but the effect of applying it is uncertain. In Helen's case, the effect is known, but the (actual) cause of that effect is uncertain.

Structural equation models were developed (primarily by twentieth-century statisticians and econometricians) in order to answer questions like the one Harold asked himself about the effects of causes.³ By contrast, philosophers, lawyers, and historians have typically been interested in questions like the one Helen asked the doctors about the causes of effects.⁴ Recently, structural equation models have been adapted by Pearl (2000), Hitchcock (2001), Woodward (2003), Halpern and Pearl (2005), Glymour and Wimberly (2007), and Hall (2007) in order to answer questions about causes of effects.⁵ Such accounts are called theories of *actual causation*.

In order to answer questions about the causes of effects, an adequate theory of actual causation must solve two problems: (1) a metaphysical problem and (2) a logical, epistemological, or inferential problem. The metaphysical problem is to specify in some detail what it is for one thing to be an actual cause of another. The metaphysician aims to identify what we should call the actual cause(s) of some thing. In carrying out her task, the metaphysician assumes that she has perfect information about the world. The inferential problem is to specify in some detail how we can come to know that one thing actually causes another. The two problems are clearly related. In the limit of total information, the inferential problem collapses into the metaphysical problem. Progress on the inferential problem requires having some account of the metaphysical problem, though progress may be made on the inferential side without having a complete account on the metaphysical side.⁶

The project in the present paper is metaphysical, though there is a tight connection to the inferential problem. The aim of philosophical theories of actual causation is to reduce inferences about actual causation to inferences about causal structure plus the application of a (metaphysical) definition. I argue that several promising, contemporary theories of actual causation are defective. In order to show their defects, I apply these theories to some quite ordinary voting scenarios and note that the theories say rather strange things about them. Two *prima facie* examples of defects in the theories are these: (1) in all simple-majority elections that allow abstentions, the theories of actual causation under consideration count every abstention as an actual cause of the winning candidate's victory, regardless of whether the election is closely contested; and (2) in all simple-plurality elections involving three or more candidates, every theory of actual causation under consideration counts every vote as an actual cause of the winning candidate's victory, regardless of how the votes are actually distributed among the candidates.

Why do the theories say these things? The theories of actual causation under consideration accommodate the *individualist intuition* that when some outcome is over-determined by two or more occurrences, each occurrence is a cause of the outcome. In order to accommodate the individualist intuition, the theories have to take account of facts about difference-making in the actual circumstances *and* in counterfactual scenarios. However, the theories are too permissive about the range

of counterfactual scenarios they consider. The theories do not have the resources to block enough counterfactual scenarios from consideration (or at least they do not have the resources to block the right ones).

Every theory of actual causation under consideration begins with a structural equation model, which represents a collection of structural causation relations, and then adds something in order to represent the actual causation relations. Hitchcock, Woodward, and Halpern and Pearl add the values that the variables in the model take on in the actual circumstances. Hall adds both the values that the variables in the model take on in the actual circumstances and also a designation of some values of the variables as defaults for the model. I argue that these added constraints are not enough: we need to know the *conditional* default values of the variables in the model and possibly much else besides.

Here is how I will proceed. In Section 1, I distinguish between structural causation and actual causation. I briefly review some necessary technical machinery and set out two closely related examples. In Section 2, I describe theories of actual causation due to Hitchcock, Woodward, Halpern and Pearl, and Hall. In Section 3, I discuss the application of those theories to three simple voting scenarios: two-candidate, simple-majority elections *without* abstentions, two-candidate, simple-majority elections *with* abstentions, and three-candidate, simple-plurality elections without abstentions. (The results easily generalize to all simple-plurality elections with and without abstentions.) I argue from examples that the theories cannot be correct as they stand. Finally, in Section 4, I speculate about why the theories fail and how they might be repaired.

1. *Structural Causation and Actual Causation*

In this section, I will begin with a description of a simple case of early pre-emption and close with a case of over-determination. I will use the first case to introduce some necessary technical details and to distinguish between structural causation and actual causation. Let U , called *the universe (of discourse) or population*, denote an arbitrary set of *units*, u_i . Units might be people, states, universities, actions, events, processes, or anything else one might be interested in. A *random variable* (or simply a *variable*) is a measurable function from U into the real numbers. Random variables typically represent properties of

units, and the value of a variable X for u_i , denoted $X(u_i) = x$, represents the result of a measurement of the property represented by X taken with respect to the unit u_i . For example, a random variable might represent height in meters or annual operating budget in dollars. A unit may be regarded as having (or being) a collection of measurable properties. Whenever a variable takes a unit as its argument, the variable indicates which property value the unit has.

Suppose U is a collection of assassinations carried out by a pair of marksmen, Ralph and Lauren, working in tandem. Each unit u is a single assassination. Sometimes Ralph takes the lead and Lauren acts as backup. Sometimes Lauren takes the lead and Ralph acts as backup. Suppose that for each unit in the population, whoever takes the lead is successful. Let the variable $R(\cdot)$ represent Ralph's action such that for all u , if Ralph shoots, then $R(u) = 1$ and if Ralph does not shoot, then $R(u) = 0$. Similarly, let the variable $L(\cdot)$ represent Lauren's action. Moreover, let the variable $V(\cdot)$ represent the state of the victim such that for all u , if the victim is alive, then $V(u) = 1$ and if the victim is dead, then $V(u) = 0$. For present purposes, suppose that both Ralph and Lauren are perfect marksmen, so that if either one shoots, the victim dies. Hence, for each u , one may write $V(u) = R(u) + L(u)$, where '+' is the Boolean OR.

A *structural equation model* (SEM) is a collection of equations in which (1) the independent variables in a given equation are interpreted as causes of the dependent variable in that equation and (2) the dependent variable in one (or more) of the equations may appear as an independent variable in one or more of the equations in the model.⁷ Let $\langle \mathbf{V}, \mathcal{F} \rangle$ denote an arbitrary SEM, where \mathbf{V} is an ordered set (or vector) of random variables and \mathcal{F} is a set of equations involving the variables in \mathbf{V} . A *structural causation relation* is a relation between random variables, which are just measurable functions. As the name suggests, a structural equation model represents a collection of structural causation relations. Another way of thinking about what an SEM represents is in terms of possible experiments.

An idealized experiment consists in manipulating some causal system. In an experiment, some properties of a unit are set to specific values and the results of that manipulation are observed. Let the manipulation of a variable X to the value x for the unit u be denoted $do(X(u) = x)$. The result of manipulating a variable in an SEM is determined by replacing the equation in which the variable appears as a dependent variable with a new equation that makes the variable equal to a constant and then

propagating that change through all the equations in which the variable appears as an independent variable. An SEM, then, represents the results of a collection of possible experiments. Following Holland (1986) and Pearl (2000), let $Y_{X=x}(u)$ denote the value Y would have, for unit u , were one to manipulate the variable X to the value x with respect to unit u , i.e. if one were to $do(X(u) = x)$.

Consider an ordered set V of variables, partitioned into the variables X and Y along with the ordered set Z of variables obtained by removing X and Y from V . Say that $X(\cdot)$ is a *direct structural cause* of $Y(\cdot)$ relative to the population U and the ordered set V of variables if for each u in U , there exist values x_1 and x_2 of X and values z of Z such that $x_1 \neq x_2$ and $Y_{X=x_1, Z=z}(u) \neq Y_{X=x_2, Z=z}(u)$. In other words, the variable $X(\cdot)$ is a direct structural cause of the variable $Y(\cdot)$ if there is a pair of $do(\cdot)$ operations such that the value of $Y(u)$ given $do(X(u) = x_1, Z(u) = z)$ differs from the value of $Y(u)$ given $do(X(u) = x_2, Z(u) = z)$. Returning to the example of Ralph and Lauren, $R(\cdot)$ is a direct structural cause of $V(\cdot)$, since (i) $V(u) = 1$ given $do(R(u) = 1, L(u) = 0)$ and (ii) $V(u) = 0$ given $do(R(u) = 0, L(u) = 0)$. In the same way, $L(\cdot)$ is a direct structural cause of $V(\cdot)$, since (i) $V(u) = 1$ given $do(R(u) = 0, L(u) = 1)$ and (ii) $V(u) = 0$ given $do(R(u) = 0, L(u) = 0)$. Despite the fact that only one of Ralph and Lauren fires in any actual case, both assassins are structural causes of the state of their victim in each assassination.

Structural equation models may be (partially) represented by directed graphs. The graph corresponding to the pre-emption case is given in Figure 1.⁸




Figure 1

For present purposes, let a *directed graph* be an ordered pair $G = \langle V, E \rangle$, where V is a finite set of *vertices* and $E \subseteq (V \times V)$ is a finite set of directed *edges*. An edge $\langle V_1, V_2 \rangle$ is *directed from* V_1 *into* V_2 . Denote the directed edge $\langle V_1, V_2 \rangle$ by $V_1 \rightarrow V_2$. A *path* of length $n > 0$ from V_i to V_j , denoted $V_i \mapsto V_j$, is a

sequence $V_{(1)}, V_{(2)}, \dots, V_{(n+1)}$ of vertices such that $V_i = V_{(1)}$, $V_j = V_{(n+1)}$, and $V_{(k)} \rightarrow V_{(k+1)}$, for $k = 1, \dots, n$.

The graph in Figure 1 is not very complicated. It has two paths of length one—one from Ralph to Victim and one from Lauren to Victim.

Now, consider a specific assassination. Suppose that Ralph fires in such a way as to kill Victim. Lauren merely watches as Victim dies; however, Lauren was prepared to fire in case Ralph missed or some other mischance took place. Many people have the intuition that in cases where only one of two assassins shoots a victim so that he dies, only the assassin that actually shot caused the victim to die.⁹ (Since they are working together, both Ralph and Lauren might very well be morally responsible for Victim's death.) Structural causation does not generally seem to capture our intuitions about the cause(s) of a given effect. Whereas structural equation models treat causation as a relation between random variables, theories of actual causation specify conditions under which a random variable taking on some value causes another random variable to take on some (other) value. Current theories of actual causation restrict attention to a single unit. They tell us whether for some unit u , $V_c(u) = v_c$ counts as an actual cause of $V_e(u) = v_e$.¹⁰

Let $\langle \mathbf{V}(u) = \mathbf{v}, \mathcal{F} \rangle$ denote an *actualized structural equation model*. Actualized structural equation models assign to each variable V_i in the vector \mathbf{V} the value v_i that the variable actually takes for some unit u . In our example, $\mathbf{V} = \langle R, L, V \rangle$, $\mathcal{F} = \{V = R + L\}$, and $\mathbf{V}(u) = \mathbf{v}$ denotes the vector $\langle R(u) = 1, L(u) = 0, V(u) = 1 \rangle$, which gives the actual values for the variables in the model. If $V_c(u) = v_c$ counts as an actual cause of $V_e(u) = v_e$ in some actualized structural equation model, then V_c will be a cause of V_e in the corresponding structural equation model. However, the fact that V_c is a cause of V_e in some structural equation model *does not* guarantee that $V_c(u) = v_c$ counts as an actual cause of $V_e(u) = v_e$ for any unit or for any actual values of V_c or V_e . In models of causal systems with redundant backups for some outcome, like the example of Ralph and Lauren, the value of the backup variable will often fail to be an actual cause of the value of the outcome variable.

Actual causation relations place a stronger constraint on how one quantifies over $do(\cdot)$ operations with respect to some collection of variables. Quantification over $do(\cdot)$ operations is not equivalent to quantification over units: the two are orthogonal. A collection of structural causation relations might hold

for a large population or for a small population (maybe even a single unit). Similarly, a collection of actual causation relations might hold for a large population or for a small population (maybe even a single unit). Different populations may support different generalizations about actual causation. For example, if $R(u) = 1$ and $L(u) = 0$ for most $u \in U$, then it will be true, relative to U , that $R(u)$ is probably an actual cause of $V(u)$ and $L(u)$ is probably not. By contrast, if $R(u) = 0$ and $L(u) = 1$ for most $u \in U$, then it will be true, relative to U , that $L(u)$ is probably an actual cause of $V(u)$ and $R(u)$ is probably not. Although structural causation relations are usually assumed to be homogeneous over a population, they need not be so. Hence, one might ask how likely $R(\cdot)$ is to be a structural cause of $L(\cdot)$ relative to a population U .

The case considered earlier in this section is a case of early pre-emption. Ralph fires. Lauren waits to see whether Ralph's shot hits before firing, and in the end, she does not fire. Such cases are not too difficult to handle. However, suppose that Ralph and Lauren do not always act like lead and backup, but sometimes, both assassins fire. And suppose that when they both fire, sometimes they fire at different times and sometimes they fire at the same time. Consider the case where Lauren and Ralph fire simultaneously. Both shots hit Victim in the head such that either one would have been sufficient to kill him. The theories of actual causation that I consider in the present paper all accommodate the individualist intuition that both Ralph and Lauren count as actual causes of Victim's death. However, the way they go about satisfying this demand opens the door for the strange results I identify for simple voting cases.

2. Theories of Actual Causation

In this section, I will describe three distinct attempts to provide necessary and sufficient conditions for determining whether a given variable taking on some specific value for a specific unit is an actual cause of another variable taking on some (other) specific value for that unit.

2.1 Hitchcock and Woodward

Hitchcock (2001) and Woodward (2003) have developed very similar theories of actual causation.

Though they are not equivalent in general, they are equivalent with respect to the examples I develop in

Section 3 below. Thus, I provide a single statement representative of their theories with respect to my examples. Begin with the simple theory that $X(u) = x$ is an actual cause of $Y(u) = y$ iff the following two conditions are satisfied:¹¹

- (HW1) The actual value of X is x , and the actual value of Y is y , for unit u .
- (HW2) There exists a path P from X to Y and there exists a manipulation $do(X = x^*)$ for some $x^* \neq x$ such that $Y_{X=x^*}(u) \neq y$ whenever all variables not on the path P are held fixed at their actual values.

Unfortunately, (HW2) is not satisfied if the value of Y (the putative effect) is over-determined by independent causal mechanisms, as in the second Ralph and Lauren example from Section 1. Thus, there are no actual causes of over-determined events according to the simple theory proposed by Hitchcock and Woodward. Although some philosophers, notably Lewis (1986, Appendix E), have been willing to accept this consequence, Lewis' intuition is not very widely shared among theoreticians today. How widespread it is among ordinary speakers of English is less clear.¹² Hitchcock and Woodward both find it unsatisfactory. In order to formulate a replacement for (HW2), we need a bit more notation and another definition (due to Hitchcock). Let \mathbf{w} denote an ordered n -tuple of values of the ordered n -tuple \mathbf{W} of variables, and let $do(\mathbf{W} = \mathbf{w})$ denote the ordered collection of manipulations $do(W_1 = w_1), \dots, do(W_n = w_n)$. Say that the ordered n -tuple \mathbf{w} of values of the ordered n -tuple \mathbf{W} of variables is in the *redundancy range* of the path P if carrying out the manipulations $do(\mathbf{W} = \mathbf{w})$ leaves all of the variables on P at their actual values. Now, in order to handle cases of over-determination, replace (HW2) with¹³

- (HW2*) There exists a path P from X to Y and there exist manipulations $do(X = x^*)$ for $x^* \neq x$ and $do(\mathbf{W} = \mathbf{w})$ for \mathbf{w} in the redundancy range of P such that $Y_{X=x^*}(u) \neq y$ whenever all the variables in \mathbf{W} are fixed by the manipulation $do(\mathbf{W} = \mathbf{w})$.

In other words, $X(u) = x$ is an actual cause of $Y(u) = y$ if one can find some path P from X to Y and some choice of (possibly non-actual) values for all of the variables not on path P such that the variables on P retain their actual values but also such that some change in the value of X would result in a change in the value of Y , if one were to set the variables not on path P to those values.

In order to check whether $X(u) = x$ is an actual cause of $Y(u) = y$, apply the following algorithm. First, pick some path P from X to Y . Second, pick some values for all the variables not on the path P such that the variables on the path P keep their actual values. (That is what it means for the off-path values to

be in the redundancy range of P.) Third, set X to each of its alternative values in turn, checking whether any such change requires a downstream change in the value of Y . If any such change in the value of X results in a change in the value of Y , then stop, $X(u) = x$ is an actual cause of $Y(u) = y$. If no change in the value of X results in a change in the value of Y , then repeat the second and third steps above. Do this until all possible values for the off-path variables in the redundancy range of P have been tried. If at any stage changing the value of X results in a change in the value of Y , stop: $X(u) = x$ is an actual cause of $Y(u) = y$. Otherwise, repeat the above steps with a new path from X to Y . If no untried paths from X to Y exist, then declare that $X(u) = x$ is not an actual cause of $Y(u) = y$.

2.2 Halpern and Pearl

Halpern and Pearl (2005) offer a more complicated theory of actual causation. Not all of the complication matters for my examples; consequently, my presentation of their theory removes some excess. Halpern and Pearl produce two different definitions of “actual cause”; however, as was the case with Woodward’s theory and Hitchcock’s theory, the two definitions are equivalent with respect to my examples.

According to Halpern and Pearl, $X(u) = x$ is an actual cause of $Y(u) = y$ iff the following three conditions are satisfied:

- (HP1) The actual value of X is x , and the actual value of Y is y , for unit u .
- (HP2) There exists a path P from X to Y and there exist manipulations $do(X = x^*)$ for $x^* \neq x$ and $do(\mathbf{W} = \mathbf{w})$ for the variables in \mathbf{W} (all those variables not on path P) such that $Y_{X=x^*}(u) \neq y$ whenever all the variables in \mathbf{W} are fixed by the manipulation $do(\mathbf{W} = \mathbf{w})$.¹⁴
- (HP3) Let \mathbf{W}^* be an m -tuple, $m \leq n$, formed from \mathbf{W} by selecting m components of the \mathbf{W} n -tuple, and let \mathbf{w}^* be the m -tuple of values formed by selecting similarly indexed components of \mathbf{w} . For all possible \mathbf{W}^* , $Y_{X=x}(u) = y$ whenever all the variables in \mathbf{W}^* are fixed by the manipulation $do(\mathbf{W}^* = \mathbf{w}^*)$.¹⁵

Condition (HP1) is the same as the first condition for Hitchcock and Woodward. Condition (HP2) allows for actual causation in cases of over-determination, and condition (HP3) guarantees that the assignment of values to off-path variables is not completely responsible for the change in the value of the putative effect Y —the change in the value of Y is due at least in part to the change in the value of X .

In order to check whether $X(u) = x$ is an actual cause of $Y(u) = y$ under Halpern and Pearl's theory of actual causation, apply the following algorithm. Pick a path from X to Y . Set the variables not on the path P to (potentially) new values such that Y retains its actual value y and so that there is some value $x^* \neq x$ such that if we set $X = x^*$, Y will take on a new value not equal to y . If such a set of values exists, then $X(u) = x$ is an actual cause of $Y(u) = y$. If there are no such values for the off-path variables, pick a new path from X to Y and try again. If all possible paths from X to Y have been tried, then $X(u) = x$ is not an actual cause of $Y(u) = y$.

2.3 Hall

Hall (2007) criticizes structural equation theories of actual causation on the grounds that they do not make use of any intrinsic properties of causes or effects in determining what is an actual cause of what. The intrinsic property Hall thinks we should care about is the property of being a *default* (as opposed to a *deviant*) state. For Hall, the distinction between a deviant state and a default state maps pretty closely onto the distinction between an occurrence (deviant state) and an absence or non-occurrence (default state). With the notion of default state in hand, Hall provides necessary and sufficient conditions for an event to count as an actual cause of another event. Translating into the notation of the present paper, Hall proposes that $X(u) = x$ is an actual cause of $Y(u) = y$ iff the following three conditions are satisfied:

- (H1) The actual value of X is x , and the actual value of Y is y , for unit u .
- (H2) There exists a manipulation $do(\mathbf{W} = \mathbf{w})$ for the variables in some subset \mathbf{W} of the variables in $V \setminus \{X, Y\}$ where the values in \mathbf{w} are all default values of the variables in \mathbf{W} and such that $Y = y$ after the manipulation.
- (H3) There exists a manipulation $do(X = x^*)$ for $x^* \neq x$ and such that $Y_{X=x^*}(u) \neq y$ whenever all the variables in \mathbf{W} are fixed by the manipulation $do(\mathbf{W} = \mathbf{w})$ in (H2).

The central idea behind condition (H2) is that the only permissible manipulations are those that set variables to their default values. As Hall writes:

In one situation, lots of events occur—*that is*, various bits of the world exhibit *deviations from their default states*. In another situation, strictly fewer events occur—*that is*, some of the bits of the world that are in deviant states in the first situation are in their *default states* instead; and every other bit is in the same state as it was. That is what it is for one situation to be, as I will call it, a *reduction* of another. Letting the “null” reduction of a situation just *be* that situation, we can now say the [sic] C causes E iff there is some reduction of the C - E situation in which E depends on C . (129)

In the notation of the present paper, an actualized structural equation model $\langle V(u) = v^*, \mathcal{F} \rangle$, which is produced from another actualized model $\langle V(u) = v, \mathcal{F} \rangle$ by setting an arbitrary subset of the variables in V to their default values, is called a *reduction* of the original actualized model. Conditions (H2) and (H3) require that the value of Y depends on the value of X in some reduction of the actualized SEM.

3. *Over-Determination and Election Results*

I begin this section by applying the various theories of actual causation to the second example from Section 1. I then consider three voting scenarios: (1) two-candidate, simple-majority elections *without* abstentions; (2) two-candidate, simple-majority elections *with* abstentions; and (3) three-candidate, simple-plurality elections with or without abstentions. I claim that the deliverances of the theories are *prima facie* wrong. Hence, voting scenarios provide counter-examples to the theories. I illustrate how the theories go wrong in voting cases with a representative proof.

3.1 *Over-Determination*

Recall the second story involving Ralph and Lauren, the perfect assassins. Ralph and Lauren fire simultaneously and both shots hit Victim in the head, killing him instantly. The model has only one equation—the same equation from the first example in Section 1: $V(u) = R(u) + L(u)$, where ‘+’ is the Boolean OR. The graph is the same as in Figure 1, and the actual values of the variables are $R(u) = 1$, $L(u) = 1$, and $V(u) = 1$. Either shot would have sufficed to kill Victim. Neither Ralph’s shot nor Lauren’s shot was a difference-maker in the actual circumstances. Manipulating Ralph’s shot makes no difference to Victim’s state. Neither does manipulating Lauren’s shot.

What should one say, then, about whether Ralph’s shot or Lauren’s shot actually caused Victim to die? Two reactions are possible here: individualist and collectivist.¹⁶ The collectivist denies that the individual over-determining occurrences are actual causes of the over-determined outcome but asserts that the mereological sum of the over-determining occurrences *is* an actual cause of the outcome. On the collectivist view, no special moves are required in order to deal with cases of over-determination, since an

over-determined outcome counterfactually depends on the mereological sum of the over-determining occurrences in the actual circumstances.

By contrast, the individualist has the intuition that every over-determining occurrence is individually an actual cause of the over-determined outcome. Schaffer (2003) provides four arguments in favor of the individualist view. The individualist view makes better sense of theoretically important aspects of causation, like prediction, explanation, and attribution of responsibility. The individualist view is naturally consistent with the fact that each over-determining occurrence is connected to the outcome by an independent, completed (causal) process. The individualist view is better able to explain the collective causal powers of the over-determining occurrences. And the individualist view is better able to account for the pragmatics of ordinary causal discourse.

However, the individualist pays a modal cost in order to accommodate the intuition that every over-determining occurrence is an actual cause of the over-determined outcome.¹⁷ Instead of attending to counterfactual dependence only in the actual circumstances, the individualist must attend to counterfactual dependence *in counterfactual circumstances* as well. Consider the assassins again. Victim's death does not counterfactually depend on Ralph's shot or on Lauren's shot in the actual circumstances. But the individualist wants to count both shots as actual causes of Victim's death. In order to see that they are both actual causes, the individualist recommends that we imagine the counterfactual scenario in which Ralph *did not* make an accurate shot. In the *counterfactual scenario* (but not in the actual scenario), Victim's death counterfactually depends on Lauren's shot. (Similar reasoning works for Ralph's shot in the counterfactual scenario in which Lauren's shot is bad.)

All of the theories of actual causation that I have been considering accommodate the individualist intuition by taking into account difference-making in counterfactual scenarios. Every theory of actual causation under consideration allows variables not on a path from the putative cause variable to the effect variable to be set to non-actual values, as long as the new setting does not change the value of the effect and, in Hall's case, as long as the change is to a default value. In this way, the theories of actual causation recover the individualist intuition. However, the constraints on how non-actual values may be assigned to variables—namely, that the new values may only be assigned to off-path variables and that each new

value must be the default value of its respective variable—are too weak or permissive. The constraints do not exclude enough counterfactual scenarios from consideration.

3.2 Election Results

Voting scenarios are idealizations of cases important to historians, ethicists, legal theorists, and diagnosticians. Sometimes the idealization is not very noticeable: for example, deciding whether Olympia Snowe’s October 2009 vote in the Senate Finance Committee was an actual cause of the health reform bill going to the Senate floor. Sometimes the idealization is extreme: for example, deciding whether the Missouri Compromise was an actual cause of the American Civil War. Voting scenarios—even simple ones—capture a number of interesting cases for ethical and legal theory as well. Are all of the members of successful lynch mobs actual causes of someone being hanged? What about people who actively resist the mob’s actions or bystanders who simply watch? Real cases will typically not be as clean and simple as the cases considered below; however, if a theory fails on the simplest of cases, then we should presume against the theory succeeding (in general) for more complicated cases, and when it *does* succeed, we should treat the success as accidental and uninformative.

All of the election scenarios I consider in the present paper share the structure pictured in Figure 2 below. In this figure, each vote (or voter) is labeled with a “ V_i ,” and the outcome is labeled “Elect.” This greatly simplifies our work, since for any vote, there is only one path to the outcome and that path contains a single directed edge.




Figure 2

The real work of my examples is done by the number of values the variables are allowed to range over and the non-linearity of the functional relationships in the SEMs representing voting scenarios, rather than

by complicated graphical features.¹⁸ In nearly all philosophical discussions of causation, especially where the relata of the causal relation are assumed to be events, SEMs that represent the cases being discussed will involve only indicator variables (binary variables that have values, “Yes, event *e* occurred,” or “No, event *e* did not occur”). However, indicator variables are often not well-suited to real work in causal representation and inference, where variables typically have multiple values. The voting cases show that apparently novel problems arise when one considers models having variables with more than two values.

As we have seen, the individualist pays a cost to accommodate the intuition that every over-determining occurrence is an actual cause of the over-determined outcome: the individualist has to attend to counterfactual dependence relations both in the actual circumstances and in counterfactual circumstances as well. The theories run into problems with voting scenarios because they are too liberal in the range of counterfactual circumstances they consider. The deliverances of the various theories with respect to two-candidate, simple-majority elections with abstentions, two-candidate, simple-majority elections without abstentions, and three-candidate, simple-plurality elections are summarized in Table 1:

	Two-Candidates, No Abstentions		Two-Candidates, with Abstentions			Three-Candidates		
	A	B	A	B	Null	A	B	C
Hitchcock & Woodward	Yes	No	Yes	No	Yes	Yes	Yes	Yes
Halpern & Pearl	Yes	No	Yes	No	Yes	Yes	Yes*	Yes
Hall	N/A	N/A	Yes	No	Yes	Yes	Yes†	Yes†

* Yes, except when the number of votes for candidate *A* is strictly less than twice the number of votes for candidate *B* plus one, the total number of votes is odd, and there are no votes for candidate *C*.

† Yes, as long as there is at least one vote for each of the candidates.

Table 1

The table describes the three types of election being considered. Under each type of election is a listing of the choices available in that election: vote for A, vote for B, vote for C, or abstain (Null). Without loss of generality, the candidates are assumed to be listed in descending order of the number of votes they received. Specifically, candidate A is assumed to have received the most votes, and candidate B is assumed to have received at least as many votes as candidate C. Each row of the table marks a specific theory of actual causation. If a row contains a “Yes” entry in the column for candidate X (or Null), then

the theory counts every vote for candidate X (or every abstention) as an actual cause of the result of the election, with the provisions indicated in the notes to the table.

Two quick counter-examples may be constructed from these results. I leave extended discussion and further examples to Section 4. *First Counter-Example.* Suppose Jack and Jill live in a Congressional district that is overwhelmingly Republican. Jack and Jill both know that the Republican candidate is going to win the election (and, in fact, the Republican does win). Jill prefers the Democratic candidate, but she finds herself so disgusted with the first-through-the-gate election system that she abstains in protest. Jack, on the other hand, prefers the Republican candidate. Since Jack believes that the Republican will win easily without his vote and since he has lots of important things to do, Jack decides to abstain. The theories treat Jack's abstention and Jill's abstention as exactly alike, but they are clearly different. *Second Counter-Example.* Elizabeth's teacher is letting her students vote on whether to read *Pride and Prejudice* or *Sense and Sensibility* for their section on Jane Austen. The class votes 18-2 in favor of *Pride and Prejudice*, with Elizabeth on the losing side. Elizabeth does not like *Pride and Prejudice*, and she likes losing even less. So, she asks her teacher if the class can vote again with *Emma* included in the options. The teacher agrees, and the class votes 18-1-1 in favor of *Pride and Prejudice*. The theories say that not only has Elizabeth lost in the second vote, she actually caused *Pride and Prejudice* to be selected as the book for the Jane Austen section of her class.

The results for three-candidate, simple-plurality elections are easily extended to elections having four or more candidates with or without abstentions. All the theories endorse the claim that every vote and every abstention (with the exceptions already noted) is an actual cause of the result of a simple-plurality election with three or more candidates. The proofs of the results summarized in Table 1 exploit over-permissiveness about how the votes may be redistributed in counterfactual situations. In order to save space and not distract the reader with irrelevant details, I will only exhibit one simple, illustrative proof in the main text. I have left the other proofs to the appendix or as exercises for the reader.

3.3 An Illustrative Proof

I prove that for Hitchcock and Woodward, every vote in a three-candidate, simple-plurality election is an actual cause of the result of the election, whatever that result might be. Consider an election in which there are $2k$ votes. (The result is identical for elections involving an odd number of votes. The proof for the odd-numbered condition is left as an exercise for the reader.) Suppose that i votes are for candidate A , j votes are for candidate B , and l votes are for candidate C . Further, suppose without loss of generality that $i > j \geq l$.

To see that every vote for candidate A is an actual cause of candidate A 's victory, choose a vote $V_A = A$. Distribute the votes such that there are $k + 1$ votes for candidate A (including V_A), $k - 1$ votes for candidate B , and no votes for candidate C . Changing the value of V_A from A to B results in a tie (k votes for A against k votes for B). Hence, $V_A = A$ is an actual cause of $\text{Elect} = A$.

To see that every vote for candidate B is an actual cause of candidate A 's victory, choose a vote $V_B = B$. Distribute the votes such that there are k votes for A , one vote for B , and $k - 1$ votes for C . Changing the value of V_B from B to C results in a tie (k votes for A against k votes for C), so $V_B = B$ is an actual cause of $\text{Elect} = A$. Similarly, every vote $V_C = C$ is an actual cause of $\text{Elect} = A$.

Thus, according to Hitchcock's theory and according to Woodward's theory, every vote cast in an election having three (or more) candidates is an actual cause of the result of the election! The proof works because Hitchcock and Woodward's theories always allow votes to be re-distributed (in the counterfactual condition) such that the election is decided by a single vote and such that the single vote is for the candidate we are thinking about. Often, the theory allows a full third or more of the votes to be assigned arbitrary new values in the counterfactual condition. The result in no way depends on the actual number of votes cast for each candidate. Even if no one votes for candidate C , every vote for candidate B is an actual cause of candidate A 's victory.

4. Intuitions

I now want to explore the deliverances of the theories of actual causation with respect to the simple voting scenarios I have been considering and suggest some generic revisions. I point out a number of situations

where the deliverances of theory do not accord well with my intuitions. I identify several intuitions and cast them as constraints on a theory of actual causation. I then argue for two things: (1) what counts as an actual cause depends on *conditional* defaults, and (2) what makes something an actual cause is not an intrinsic feature of that thing.¹⁹

4.1 *Symmetry and Causal Production*

In the case of two-candidate, simple-majority elections without abstentions, the theories say that all and only the votes for the winning candidate are actual causes of the winning candidate's victory. Insofar as one has individualist intuitions, intuition appears to agree with the theories in this case. One reason to be an individualist about elections is that each vote is a *producer*. If an election involved only one vote, then that vote would produce or determine the result. Votes for candidate *A* produce victory for candidate *A*, assuming that candidate *A* wins the election. Votes for candidate *B* produce victory for candidate *B*. Since the winning votes are indistinguishable (i.e. all of the votes have equal weight in determining the outcome), symmetry requires that all of the votes be considered equally efficacious.

4.2 *Abstentions, Contrasts, and Conditional Defaults*

In the case of two-candidate, simple-majority elections where voters are allowed to abstain, the theories say that all votes for the winning candidate and all abstentions are actual causes of the winning candidate's victory (and no other votes are actual causes of the winning candidate's victory). Counting the votes for the winning candidate as actual causes makes sense in the same way that it made sense for the case without abstentions. But how should we understand the causal role of abstentions?

One thing is clear about the causal role of abstentions in voting scenarios: if abstentions have any causal role at all, they have it in virtue of facts about causal dependence, not in virtue of facts about causal production.²⁰ Abstentions (and absences generally) do not *produce* anything. From the perspective of causal production, we should be able to ignore abstentions altogether: an election in which candidate *A* receives six votes, candidate *B* receives three votes, and two people abstain is equivalent to an election in which candidate *A* receives six votes, candidate *B* receives three votes, and seventeen people abstain.

However, those two elections strike me as being importantly different. In the second case, but not in the first case, the abstentions seem to matter. My initial impression is that in order for any abstention to matter, it must be the case that the abstentions matter collectively. In other words, if an abstention is to count as an actual cause of the result of an election between two candidates, the number of abstentions must be greater than or equal to the gap between the winner and the loser. This intuition turns Schaffer's discussion of the source of collective causal powers on its head.²¹ I ask, "How could the *individual* abstentions matter if the collection of abstentions does not?" When the actual votes are six for *A*, three for *B*, and two abstaining, I want to say that the abstentions, as a group, do not matter. Therefore, I want to say that the individual abstentions do not matter. However, when the actual votes are six for *A*, three for *B*, and seventeen abstaining, I want to say that the abstentions as a group do matter, or at least, they might matter. So following Schaffer, I want to say that the individual abstentions (might) matter.

The individual abstentions *might* matter, but my confidence about the judgment that they *do* matter depends on what the voters who abstained would have done—who they would have voted for—had they decided (or been forced) to vote. For example, if everyone who abstained would have voted for candidate *A* (the winning candidate), then I would say that the abstentions were not actual causes of the result of the election. Rather, candidate *A* won despite the fact that some *A*-supporters abstained. On the other hand, if everyone who abstained would have voted for candidate *B* (the losing candidate), then I would say that the abstentions were actual causes of the result of the election (assuming that the number of votes for *B* plus the number of abstentions is greater than or equal to the number of votes for *A*). Abstentions count as actual causes or fail to count as actual causes in the light of some relevant contrast. Two examples will make this clearer.

Imagine that in a certain department, a search committee recommends job candidates one at a time and then the faculty votes on whether to hire the job-seeker. Professors may vote "Yes," "No," or "Abstain." If the vote is tied, then the question is tabled for two days of debate and then brought to a new vote. Now, imagine that the committee has recommended a controversial candidate named Steve. Dr. Smith has an unfavorable impression of Steve, but at the last minute, he decides to abstain instead of voting "No." The vote comes in at 4-3 in favor with Dr. Smith abstaining, and Steve is offered a job. I

have the clear intuition that Dr. Smith abstaining *rather than voting* “No,” caused Steve to be offered the job. But now imagine that two other professors, Dr. Crane and Dr. King, who had favorable impressions of Steve also decided to abstain. With the final vote at 2-3 against and three abstaining, Steve is not offered a job. I have the clear intuition that by abstaining, Dr. Crane and Dr. King prevented Steve from being offered the job. But I do *not* want to say that Dr. Smith’s abstaining prevented Steve from being offered the job; rather, I want to say that Steve was offered the job *despite* Dr. Smith’s abstention.²²

In the 2008 U.S. Presidential election, many out-of-state university students in Virginia and Colorado were told that they could not vote in their school’s state.²³ This was false. Since students overwhelmingly support the Democratic Party, had John McCain won the election, it might fairly have been said that many students not voting was an actual cause of McCain’s victory. On the other hand, since Obama won the election, student abstentions were not actual causes of anything.²⁴ Judgments about actual causation are contrastive—something hidden by the way actual causation is treated by the theories under consideration.²⁵ When I say that Jill’s abstention was an actual cause of the Republican candidate winning the election, I am saying that Jill’s abstaining rather than voting for the Democratic candidate was an actual cause of the Republican candidate winning the election.²⁶ Where do the relevant contrasts come from? I claim that they come from facts about *conditional defaults*.

But first, we need to rethink the notion of default. Hall’s conception of defaults comes out of the dominant view that causation is a relation between events, where an event is binary—something that either happens or does not happen. From this perspective, it makes sense to say that by default, nothing happens. The default is for whatever event actually occurred to have not occurred. Hall thinks about this operation as removing an event and leaving a void in its place. By contrast, I claim that the default for any occurrence is exactly what actually happened. If Jill actually abstains in an election, then the default is for Jill to have abstained. If Jill actually votes for a Democratic candidate, then the default is for Jill to have voted for a Democratic candidate. Given the defaults, we can ask a further question: how would Jill have voted if she had been prevented from doing what she did by default? In other words, we can ask about Jill’s conditional defaults. Suppose Jill actually abstained in some election. The question, “How would Jill have voted had she not been allowed to abstain?” is a question about a conditional default for

Jill. Such questions may be iterated, provided there are enough options available. If an election involves six candidates and allows abstentions, then we might ask how Jill would have voted had she been prevented from abstaining and also prevented from voting for either candidate *A* or candidate *B*.

Define the zeroth-order default for a variable *V* as the value that *V* takes on in the actual circumstances, and define the $(n + 1)$ th-order default for a variable *V* as the value that *V* would take on were it prevented from taking on any lower-order default value. Conceptualizing defaults and actual causation in this way requires some rethinking of the *do*(·) calculus. Ordinarily, the *do*(·) operator forces its target to take on a specific value (or when probabilities are involved, a specific distribution). Thus, the *do*(·) operator may be understood as completely constraining its target variable. On the present view of defaults, a more generic partial-constraint operator is required. Instead of writing *do*($X = x$), write *do*($X \in \{x_i\}$) for admissible values x_i .

The relevant contrasts for a theory of actual causation come from considering what an agent (or other target of investigation) would do by default (by its nature), *conditional on* some constraint. When one asks about whether some occurrence was an actual cause of some outcome without specifying a conditioning constraint, then one is asking whether that occurrence rather than its first-order default was an actual cause of the outcome. Determining the conditional defaults will often require consultation of some special science for its resolution. Philosophers of science (and especially philosophers of physics) often say that metaphysics must take stock of physics. In voting cases, it appears that metaphysics must take stock of psychology, neuroscience, and much else as well.

4.3 *Asymmetry, Stability, and Irrelevant Details*

In the case of elections involving three or more candidates (with or without abstentions), every theory of actual causation under consideration counts *every* vote (and *every* abstention) as an actual cause of the winning candidate's victory, with the rare exceptions mentioned in footnotes to Table 1.

On its face, this result is absurd. How could a vote *against* a candidate possibly count as an actual cause of that candidate being elected? Two relatively recent U.S. Presidential elections show how votes for third-party candidates might count as actual causes of the outcome of three-candidate simple-

plurality elections. For instance, when Clinton defeated Bush in 1992, some pundits suggested that Perot had siphoned off enough votes from Bush to give the election to Clinton. Similarly, when Bush defeated Gore in 2000, some suggested that Nader cost Gore the victory. However, I would be very surprised to hear anyone say that *Gore cost Nader* the 2000 election or that *Bush cost Perot* the 1992 election. My attitudes toward the two losing candidates in three-candidate, simple-plurality elections are not symmetric. Only the last-place finisher seems fit to count as a cause of the victorious candidate winning the election. The intuition extends to the voters for the candidates. I have witnessed many conversations recently in which one person expresses an interest in voting for a third-party candidate, and his or her interlocutor replies, “If you vote for a third party and the Democrats lose, it will be your fault.” (Given the political preferences of the people in these conversations, the person voting for a third party would not be considered blameworthy if the Democratic candidate won.) Call this the *asymmetry intuition*.²⁷

None of the theories of actual causation considered in this paper respects the asymmetry intuition. Hitchcock’s theory and Woodward’s theory also fail to capture two further intuitions, which I call the *strong and weak stability intuitions*. According to the weak stability intuition, what counts as an actual cause of the outcome of an election in which some candidate receives no votes should be the same as what counts as an actual cause of the outcome of an otherwise identical election in which the candidate that received no votes does not appear. Here is an example. Suppose a corporate board consisting of 23 members takes a vote to decide whether to build their new facility in New York or Los Angeles. The vote is 16 for New York and 7 for Los Angeles. In this case, Hitchcock’s theory and Woodward’s theory both tell us that the 16 votes to build in New York are actual causes of the company building their new facility in New York, while the other 7 votes are not. Now, suppose that one of the board members thinks the board should consider building in Chicago, though she herself does not think Chicago is really the best place to build. She presents this view to the board, and consequently, the members vote on whether to build in New York, Los Angeles, or Chicago. The vote is 16 for New York, 7 for Los Angeles, and none for Chicago. According to the weak stability intuition, what counts as an actual cause of the company building in New York in this scenario should be the same as what counts as an actual cause of the company building in New York in the previous scenario. However, according to Hitchcock and

Woodward's theories, in the second scenario, the seven votes to build in Los Angeles are actual causes of the company building its new facility in New York!

According to the strong stability intuition, what counts as an actual cause of the result of an election should be insensitive to partitioning the votes for a losing candidate among the losing candidate and some number of new candidates (where by "new" I mean candidates not on the ballot in the original election). In the previous example of the corporate board, the strong stability intuition says that if the seven votes for Los Angeles were re-distributed as votes for Los Angeles and Chicago (but the 16 votes for New York were left alone), then the actual causes should be the same as in the original case. Similarly, if the board decided to add Chicago and Houston for consideration and the new vote came out as 16 for New York, 4 for Los Angeles, 2 for Chicago, and 1 for Houston, then according to the strong stability intuition, only the votes for New York should count as actual causes of the company building its new facility in New York.

Hall's theory fails to capture the strong stability intuition but not the weak stability intuition. Halpern and Pearl's theory fails to capture both the weak and strong stability intuitions, though not for all cases. One might think that this is good news for Halpern and Pearl, but it is not. Halpern and Pearl's theory will not count votes to build in Los Angeles as actual causes of the company building in New York, when the 23-member corporate board votes 16-7-0 in favor of building in New York over Los Angeles and Chicago. However, the reason is very peculiar. In this case, the total number of votes is odd, there are no votes for the third-place option, and the winning option has more than twice as many votes as the second-place option. Under those conditions, Halpern and Pearl's theory only counts the first-place votes as actual causes of the outcome of the election. However, if any of those conditions is different, then Halpern and Pearl's theory will count all of the votes as actual causes of the outcome of the election. For example, if the corporate board has 24 members that initially voted 17 to 7 for New York, then it endorses the same conclusion that Hitchcock and Woodward's theories endorse. However, the conclusion that votes for Los Angeles are not actual causes when Chicago is not a live option but are actual causes when Chicago is a live option—even if no one picks that option—is no more compelling when the board has an even number of members than it was when the board had an odd number of

members. In virtue of paying attention to whether the total number of votes is even or odd, Halpern and Pearl's theory fails to capture another important intuition, which I call the *irrelevant details intuition*.

Any process that keeps track of overall counts of occurrences will work out similarly to voting. For example, imagine that the area around an apple tree is divided into two patches (*A* and *B*), and suppose that we count how many apples fall from the tree and come to rest on each patch. Suppose that more apples land on patch *A* than on patch *B*. All of the theories under consideration will say that all and only the apples that landed on patch *A* were actual causes of patch *A* being covered with more apples. However, if we subdivide patch *B* into two new patches *C* and *D*, then (with the exceptions already noted), every theory will say that all of the apples were actual causes of patch *A* being covered with the most. The theories pay no attention to how likely each apple was to land on a given patch. Nor do they pay attention to how the apples were actually distributed over the available patches.

Again, the problem is that these theories of actual causation are too permissive in the range of counterfactual scenarios they consider. Instead of allowing arbitrary re-distributions of the values of variables in a structural model, theories of actual causation ought to consider only re-distributions in line with first-order defaults, unless a specific contrastive question is at stake. Hence, in order for a third-party vote to count as an actual cause of the winning candidate's victory, there should be a counterfactual scenario—constructed by re-distributing votes according to their first-order defaults—in which the third-party vote is a difference-maker.

5. Summary

Current theories of actual causation fail to capture several intuitions about simple voting scenarios because they over-correct in order to accommodate individualist intuitions about over-determination. Facts about the actual distribution of votes in an election as well as facts about the conditional default values of the votes in an election (which are presumably determined by the preferences of the voters) matter to whether a given vote counts as an actual cause of the result of the election. Causal structure plus the actual values of the variables in the model and even some intrinsic facts about those values are not enough to pick out the actual causes in general.

Appendix: Proofs

A.1 *Two-candidate, simple-majority elections*

The scenario envisioned here is very simple. Everyone must cast a vote. Each vote cast is for exactly one of two candidates, A and B . If both candidates receive the same number of votes, then the election results in a tie. Otherwise, whichever candidate receives the most votes wins the election. Thus, the election may end in a victory for one or the other of the two candidates, or it may end in a tie.

A.1.1 *Hitchcock and Woodward*

Consider an election in which there are $2k$ votes. (I leave it to the reader to show that elections with an odd number of votes produce identical results.) Suppose that i votes are cast for candidate A , and suppose without loss of generality that $2k - i < i$. Thus, candidate A is the actual winner of the election.

Is a vote for candidate A an actual cause of candidate A 's victory? Yes. Choose a vote $V_A = A$ for candidate A .²⁸ There is only one path from V_A to the result of the election, and no other vote is on this path. Hence, we are free to change any of the other votes, so long as candidate A wins the election after the changes. Distribute the votes such that there are $k + 1$ votes for A (including V_A) and $k - 1$ votes for B . Now, change the value of V_A from A to B . Since such a change results in a tie (k votes for A against k votes for B), $V_A = A$ is an actual cause of A 's victory. Because V_A was chosen arbitrarily, the same reasoning applies to every vote for candidate A . Hence, every vote for candidate A is an actual cause of candidate A 's victory.

What about a vote for candidate B ? No. No vote for candidate B is an actual cause of candidate A 's victory. Choose a vote $V_B = B$ for candidate B . Again, there is only one path from V_B to the result of the election, and no other vote is on this path. Hence, we are free to change any of the other votes, so long as candidate A wins the election after the changes. However, there is no redistribution of the votes such that A is the winner of the election but would not have been the winner had V_B not been a vote for candidate B . Let r be the redistributed votes for candidate A . Since candidate A must be the winner after any redistribution, $2k - r < r$. For v_B to be an actual cause of candidate A 's election, there must be $k, r \geq 0$

such that $2k - r - 1 \geq r + 1$. That is, a change in vote V_B must result in a change in the election, either to a tie or to a victory for B . But $2k - r < r \Rightarrow 2k - r < r + 2 \Rightarrow 2k - r - 1 < r + 1$. So, $V_B = B$ is not an actual cause of candidate A 's election. Because V_B was chosen arbitrarily, the same reasoning applies to every vote for candidate B . Hence, no vote for candidate B is an actual cause of candidate A 's victory.

A.1.2 Halpern and Pearl

Consider an election in which there are $2k + 1$ votes. (I leave it to the reader to show that elections with an even number of votes produce identical results.) Suppose that i votes are cast for candidate A , and suppose without loss of generality that $2k + 1 - i < i$. Thus, candidate A is the actual winner of the election.

Is a vote $V_A = A$ for candidate A an actual cause of candidate A 's victory? Yes. The only path from V_A to the outcome is the direct edge connecting them. Thus, we are free to set all the votes however we like so long as candidate A still wins the election if an arbitrary subset of the votes were assigned the new values while V_A retains its actual value. To satisfy (HP3), assign new values to the votes by leaving $k + 1$ votes (including V_A) at their original values A and setting k votes to B . This can always be done, since there were originally i votes for candidate A , and by supposition, $k < i$. The required redistribution of votes is shown graphically in Figure 4 below.




Figure 4

Since setting V_A to B results in a victory for candidate B , (HP2) is satisfied as well. Hence, $V_A = A$ is an actual cause of candidate A 's victory. Because V_A was chosen arbitrarily, the same reasoning applies to every vote for candidate A .

Is a vote $V_B = B$ for candidate B an actual cause of candidate A 's victory? No. No assignment of values to the votes that satisfies (HP2) can also satisfy (HP3). Let $do(V = v^*)$ be a proposed manipulation satisfying (HP2). Since candidate A won the election and the original value of V_B was B , the change in the outcome of the election needed in order to satisfy (HP2) must be due *solely* to changes in the values of votes for candidate A . Let W be an ordered tuple of the votes for A that were changed to votes for B in carrying out the manipulation to satisfy (HP2). Leaving V_B at its actual value and changing all the variables in W to their manipulated values, candidate B wins the election in violation of (HP3). Hence, no vote for candidate B is an actual cause of candidate A 's victory.

A.1.3 *Hall*

In order to apply Hall's theory, we need to identify default values for the variables in the model. Neither a vote for candidate *A* nor a vote for candidate *B* can be thought of as the default choice. Hence, on one reading, Hall's theory can only be applied to the null reduction in forced-choice models. For null reductions of two-candidate simple-majority elections, Hall's theory reduces to simple counterfactual dependence. If the outcome of the election counterfactually depends on the actual value of some vote, then that vote is an actual cause of the outcome; otherwise, not.

However, on another reading, Hall's theory *can* be applied to the case. The issue is about how to model the scenario. One might think that even though votes must be for one of the two candidates in the actual scenario, one should *model* the scenario with three-valued variables just like in the case in Section 3.2 below, which includes abstentions. The contrast Hall needs to draw is not between voting for candidate *A* as opposed to voting for candidate *B*; rather, the contrast he needs to draw is between voting for some candidate and not voting at all. As we will see below, if we model the scenario with three-valued variables including abstentions, then for Hall, all and only votes for candidate *A* count as actual causes of candidate *A*'s victory.

A.2 *Two-candidate, simple-majority elections with abstentions*

In this scenario, every vote cast is for exactly one of the two candidates (just like in the previous scenario). However, voters are no longer obligated to cast a vote. As before, if both candidates receive the same number of votes, then the election results in a tie. Otherwise, whichever candidate receives the most votes wins the election. So again, the election may end in a victory for one or the other of the two candidates, or it may end in a tie.

A.2.1 *Hitchcock and Woodward*

Consider an election in which there are $2k$ votes. Suppose that i votes are for *A*, j votes are for *B*, and l votes are abstentions. Suppose without loss of generality that $i > j$.

Is a vote for candidate A an actual cause of candidate A 's victory? Yes. Choose a vote $V_A = A$. Distribute the votes such that there is one vote for candidate A (V_A itself), zero votes for candidate B , and $2k - 1$ abstentions. Changing the value of V_A from A to B results in a victory for candidate B (zero votes for A against one vote for B), so $V_A = A$ is an actual cause of A 's victory.

Is a vote for candidate B an actual cause of candidate A 's victory? No. Choose a vote $V_B = B$. There is no redistribution of the votes such that A wins the election but would not after $do(V_B \neq B)$. Let r be the redistributed votes for candidate A and a be the redistributed abstentions. Since candidate A must be the winner after any redistribution, $2k - r - a < r$. For $V_B = B$ to be an actual cause of candidate A 's election, there must be $a, k, r \geq 0$ such that either $2k - r - a - 2 \geq r$ (in case the vote for B is changed to a vote for A) or $2k - r - a - 1 \geq r$ (in case the vote for B is changed to an abstention). But $2k - r - a < r \Rightarrow 2k - r - a - 1 < r \Rightarrow 2k - r - a - 2 < r$. So, $V_B = B$ is not an actual cause of candidate A 's election.

Are abstentions actual causes of candidate A 's victory? Yes, they are. Choose an abstention, $V_{\text{none}} = 0$. Distribute the votes such that there is one vote for candidate A , zero votes for candidate B , and $2k - 1$ abstentions. Change the value of V_{none} from an abstention to a vote for B . Since such a change results in a tie (one vote for A against one vote for B), $V_{\text{none}} = 0$ is an actual cause of A 's victory.

A.2.2 Halpern and Pearl

Consider an election in which there are $2k + 1$ votes. Suppose that i votes are cast for candidate A , j votes are cast for candidate B , and there are l abstentions. Suppose without loss of generality that $j < i$.

Is a vote $V_A = A$ for candidate A an actual cause of candidate A 's victory? Yes. To satisfy (HP3), assign new values to the votes by leaving $j + 1$ votes (including V_A) for candidate A at their original value A , leaving all j votes for candidate B at their original value B and setting all the other votes to abstentions. The required redistribution of votes is shown graphically in Figure 5 below.

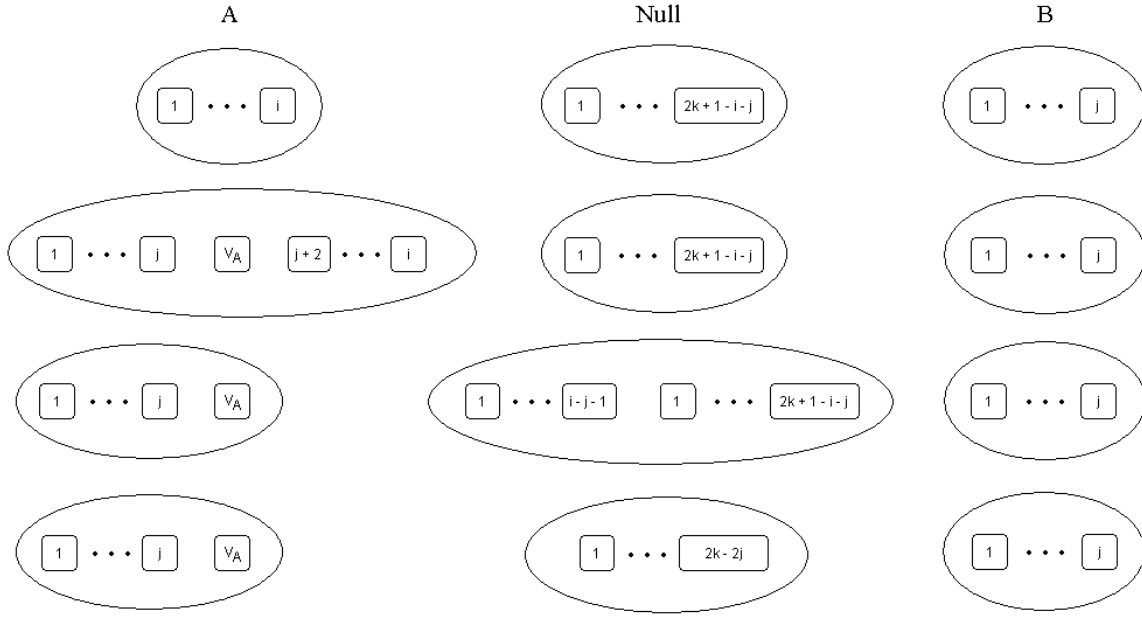


Figure 5

Since all of the votes being set to abstentions were actually either votes for A or abstentions, setting an arbitrary subset of the votes to their manipulated values leaves candidate A as the winner, so long as V_A has its actual value. Under the new assignment of values to the votes, if the value of V_A is changed, then the election results either in a tie or a victory for candidate B , which satisfies (HP2). Hence, $V_A = A$ is an actual cause of A 's victory.

Is a vote $V_B = B$ for candidate B an actual cause of candidate A 's victory? No. No assignment of values to the votes that satisfies (HP2) can also satisfy (HP3). Let $do(\mathbf{V} = \mathbf{v}^*)$ be a proposed manipulation satisfying (HP2). Since candidate A won the election and the original value of V_B was B , the change in the outcome of the election needed in order to satisfy (HP2) must be due *solely* to changes in the values of votes for candidate A . Let \mathbf{W} be an ordered tuple of the votes for A that were changed to votes for B in carrying out the manipulation to satisfy (HP2). Leaving V_B at its actual value and changing all the variables in \mathbf{W} to their manipulated values, either candidate B wins the election or the election results in a tie. Either way, (HP3) is violated. Hence, no vote for candidate B is an actual cause of candidate A 's victory.

Are abstentions actual causes of candidate A 's victory? Yes. Choose an abstention $V_{\text{none}} = 0$. To satisfy (HP3), leave all j votes for candidate B at their actual value, leave $j + 1$ votes for candidate A at their actual value, and set all other votes to abstentions. Since all the votes being set to abstentions were actually either votes for candidate A or abstentions, setting an arbitrary subset of those votes to their manipulated values leaves candidate A the winner, so long as V_{none} retains its actual value. Under the new assignment, if the value of V_{none} is changed to a vote for candidate B , then the election results in a tie rather than a victory for candidate A , which satisfies (HP2). Hence, $V_{\text{none}} = 0$ is an actual cause of A 's victory.

A.2.3 Hall

The obvious choice for a default value of a vote in a voting scenario is abstention—at least, when that value is available. (Consequently, the default value for Elect is a tie.) All reductions of elections involve setting some votes to abstentions. Consider an election in which there are $2k$ votes. Suppose that i votes are cast for candidate A , j votes are cast for candidate B , and there are l abstentions. Further, suppose without loss of generality that $i > j$.

Is a vote for candidate A an actual cause of candidate A 's victory? Yes. Choose a vote $V_A = A$ for candidate A . Set all the votes except V_A to abstentions. In this reduction, changing the value of V_A changes the result of the election. Hence, $V_A = A$ is an actual cause of $\text{Elect} = A$.

Is a vote for candidate B an actual cause of candidate A 's victory? No. To see this, choose a vote $V_B = B$ for candidate B . In every reduction that satisfies (H2) by leaving $\text{Elect} = A$, the number of votes for candidate A is strictly greater than the number of votes for candidate B . Elect does not depend on V_B in any of these reductions, since the result does not change for $do(V_B = A)$ or for $do(V_B = 0)$, which are the only possible manipulations of V_B . Hence, $V_B = B$ is not an actual cause of $\text{Elect} = A$.

Is an abstention an actual cause of candidate A 's victory? Yes. Choose a vote $V_{\text{none}} = 0$. Consider the reduction in which there is one vote for candidate A , zero votes for candidate B , and $2k - 1$ abstentions (of which V_{none} is one). In this reduction, changing V_{none} to B makes $\text{Elect} = 0$. Hence, $V_{\text{none}} = 0$ is an actual cause of $\text{Elect} = A$.

A.3 *Three-candidate, simple-plurality votes*

In this scenario, every vote cast is for exactly one of the three candidates. If all three candidates receive the same number of votes or if two candidates have the same number of votes as each other and more votes than the third candidate, then the election results in a tie. Otherwise, whichever candidate receives the most votes wins the election. (In other words, a candidate need not receive the majority of the votes, and there are no run-offs.) So again, the election may end in a victory for exactly one of the three candidates, or it may end in a tie. In order to save space, the proofs have been truncated in this section.

A.3.1 *Hitchcock and Woodward*

Consider an election in which there are $2k$ votes. (I leave it to the reader to show that elections with an odd number of votes produce identical results.) Suppose that i votes are for A , j votes are for B , and l votes are for C . Further, suppose without loss of generality that $i > j \geq l$.

To see that every vote for candidate A is an actual cause of candidate A 's victory, we proceed as before. Choose a vote $V_A = A$. Distribute the votes such that there are $k + 1$ votes for candidate A (including V_A), $k - 1$ votes for candidate B , and no votes for candidate C . Changing the value of V_A from A to B results in a tie (k votes for A against k votes for B). Hence, $V_A = A$ is an actual cause of $\text{Elect} = A$.

To see that every vote for candidate B is an actual cause of candidate A 's victory, choose a vote $V_B = B$. Distribute the votes such that there are k votes for A , one vote for B , and $k - 1$ votes for C . Changing the value of V_B from B to C results in a tie (k votes for A against k votes for C), so $V_B = B$ is an actual cause of $\text{Elect} = A$. Similarly, every vote $V_C = C$ is an actual cause of $\text{Elect} = A$.

Thus, according to Hitchcock's theory and according to Woodward's theory, every vote cast in an election having three (or more) candidates is an actual cause of the result of the election! Notice that this result does not in any way depend on the actual number of votes cast for each candidate. Even if no one votes for candidate C , every vote for candidate B is an actual cause of candidate A 's victory.

A.3.2 Halpern and Pearl

Of the theories considered in this paper, Halpern and Pearl's theory of actual causation is the most challenging to correctly apply. For the three-candidate case, I will provide general proofs for my claims about when specific votes are causes. Suppose that i votes are for A , j votes are for B , and l votes are for C . Suppose without loss of generality that $i > j \geq l$.

Is a vote $V_A = A$ an actual cause of candidate A 's victory? Yes. We need to pay attention to the difference in votes for candidate A and candidate B . We consider two cases: (1) $i - j = 2m$ for some $m \in \mathbf{N}$ and (2) $i - j = 2m + 1$ for some $m \in \mathbf{N}$. Case 1. Let $i - j = 2m$ for some $m \in \mathbf{N}$. Leave $j + m + 1$ votes for A (including V_A) at their original value, and leave all l votes for C at their original value. Set $j + m - 1$ votes to B . That is, move $m - 1$ votes from A to B . Given this distribution of votes, if we change V_A to B , the election results in a tie. Case 2. Let $i - j = 2m + 1$ for some $m \in \mathbf{N}$. Again, leave $j + m + 1$ votes for A (including V_A) at their original value, and leave all l votes for C at their original value. Set $j + m$ votes to B . That is, move m votes from A to B . Given this distribution of votes, if we change V_A to B , the election results in a win for candidate B .

Is a vote $V_B = B$ an actual cause of candidate A 's victory? Yes, with the following exception: A vote $V_B = B$ is an actual cause of candidate A 's victory unless candidate A wins by fewer than $j + 1$ votes, the total number of votes is odd, and there are no votes for candidate C .²⁹ Consider three cases: (1) $j + l > i$, (2) $j + l = i$, and (3) $j + l < i$.

Case 1. Suppose $j + l > i$. Let $m = i - l$. Leave all i votes for candidate A at their original value A and leave all l votes for candidate C at their original value C . Move $m - 1$ votes (but not including V_B) from B to C . Under the new distribution, there are i votes for A , there are $i - 1$ votes for C , and V_B is still a vote for B . If we change V_B to C , the election results in a tie.

Case 2. Suppose $j + l = i$. In this case, the total number of votes is $2i$. Leave all i votes for candidate A at their original value A and leave all l votes for candidate C at their original value C . Move all the votes for B except V_B from B to C . Under the new distribution, there are i votes for A , there are $i - 1$ votes for C , and there is one vote for B (V_B itself). If we change V_B to C , the election results in a tie.

Case 3. Suppose $j + l < i$. We need to consider two sub-cases: (a) the total number of votes is even, say $2k$, for some $k \in \mathbf{N}$ and (b) the total number of votes is odd, say $2k + 1$, for some $k \in \mathbf{N}$.

If the total number of votes is even, then $i - (j + l) = i - (2k - i) = 2i - 2k$. Let $m = 2i - 2k = 2p$ for some $p \in \mathbf{N}$. Leave all l votes for C at their original value. Move all the votes for B except V_B from B to C , and then move p votes from A to C . Under the new distribution, there are $j + l + p$ votes for A , there are $j + l + p - 1$ votes for C , and there is one vote for B (V_B itself). If we change V_B to C , the election results in a tie.

If the total number of votes is odd, then $i - (j + l) = i - (2k + 1 - i) = 2i - 2k - 1$. Let $m = 2i - 2k - 1 = 2p + 1$ for some $p \in \mathbf{N}$. Leave all l votes for C at their original value. Move $j - 2$ votes for B (not including V_B) from B to C , and then move $p + 1$ votes from A to C . As long as there was at least one vote originally cast for C , the number of votes for A is guaranteed to be at least $p + 2$ greater than the number of votes for B according to the original vote distribution, satisfying (HP3). Under the new distribution, there are $j + l + p$ votes for A , there are $j + l + p - 1$ votes for C , and there are two votes for B (including V_B). If we change V_B to C , the election results in a tie. This last case is depicted graphically in Figure 6.

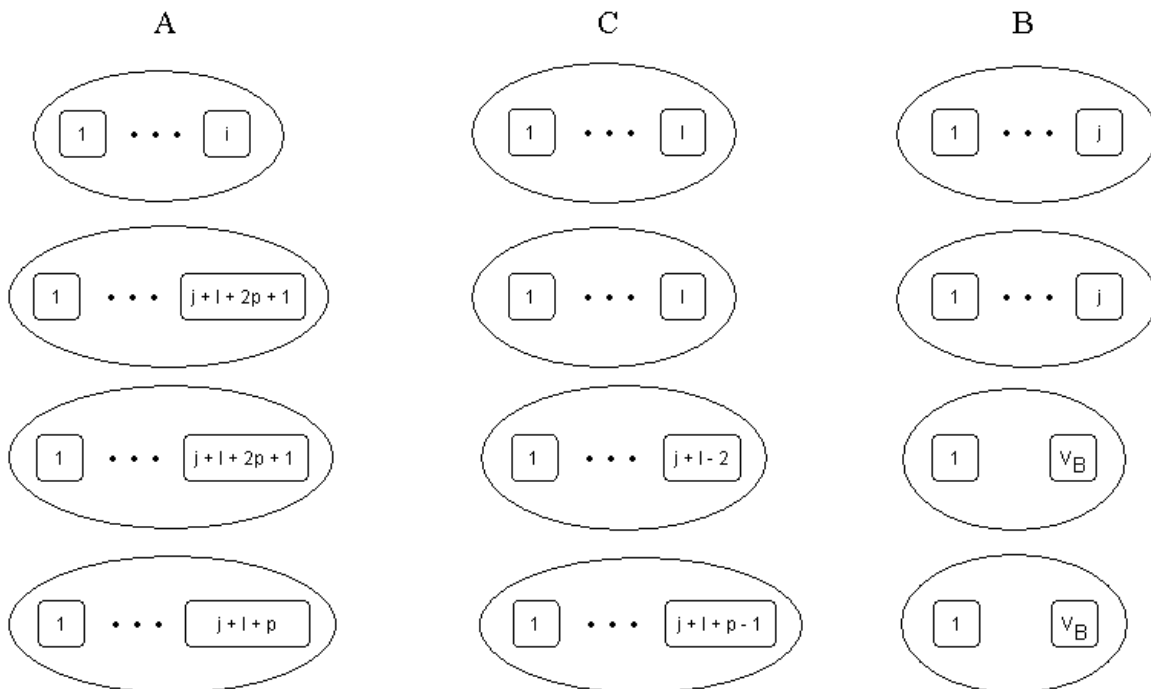


Figure 6

Is a vote $V_C = C$ an actual cause of candidate A 's victory? Yes. Again, consider three cases: (1) $j + l > i$, (2) $j + l = i$, and (3) $j + l < i$.

Case 1. Suppose $j + l > i$. Let $m = i - l$. Leave all i votes for candidate A at their original value A and leave all j votes for candidate B at their original value B . Move $m - 1$ votes (but not including V_C) from C to B . Under the new distribution, there are i votes for A , there are $i - 1$ votes for B , and V_C is still a vote for C . If we change V_C to B , the election results in a tie.

Case 2. Suppose $j + l = i$. In this case, the total number of votes is $2i$. Leave all i votes for candidate A at their original value A and leave all j votes for candidate B at their original value B . Move all the votes for C except V_C from C to B . Under the new distribution, there are i votes for A , there are $i - 1$ votes for B , and there is one vote for C (V_C itself). If we change V_C to B , the election results in a tie.

Case 3. Suppose $j + l < i$. We need to consider two sub-cases: (a) the total number of votes is even, say $2k$, for some $k \in \mathbf{N}$ and (b) the total number of votes is odd, say $2k + 1$, for some $k \in \mathbf{N}$.

If the total number of votes is even, then $i - (j + l) = i - (2k - i) = 2i - 2k$. Let $m = 2i - 2k = 2p$ for some $p \in \mathbf{N}$. Leave all j votes for B at their original value. Move all the votes for C except V_C from C to B , and then move p votes from A to B . Under the new distribution, there are $j + l + p$ votes for A , there are $j + l + p - 1$ votes for B , and there is one vote for C (V_C itself). If we change V_C to B , the election results in a tie.

If the total number of votes is odd, then $i - (j + l) = i - (2k + 1 - i) = 2i - 2k - 1$. Let $m = 2i - 2k - 1 = 2p + 1$ for some $p \in \mathbf{N}$. Leave all j votes for B at their original value. Move $l - 2$ votes for C (not including V_C) from C to B , and then move $p + 1$ votes from A to B . Since $j \geq l$ by assumption, the number of votes for A is guaranteed to be at least $p + 2$ greater than the number of votes for both B and C according to the original vote distribution, satisfying (HP3). Under the new distribution, there are $j + l + p$ votes for A , there are $j + l + p - 1$ votes for B , and there are two votes for C (including V_C). If we change V_C to B , the election results in a tie.

A.3.3 Hall

Consider an election in which there are $2k$ votes. Suppose that i votes are for A , j votes are for B , and l votes are cast for C . (In order to apply Hall's theory, assume that abstention is a possible value for any vote V .) Further assume without loss of generality that $i > j \geq l$.

Is a vote for candidate A an actual cause of candidate A 's victory? Yes. Choose a vote $V_A = A$. Set all the votes except V_A to abstentions. In this reduction, changing the value of V_A changes the result of the election. Hence, $V_A = A$ is an actual cause of $\text{Elect} = A$.

Is a vote for candidate B an actual cause of candidate A 's victory? As long as there is at least one actual vote for candidate C , the answer is "Yes." Choose a vote $V_B = B$. Consider the reduction in which there are two votes for A , one vote (namely, V_B itself) for B , and one vote for C . All other votes have been set to abstentions. In this reduction, changing V_B to a vote for C results in a tie. Hence, $V_B = B$ is an actual cause of $\text{Elect} = A$. The proof showing that a vote $V_C = C$ is an actual cause of $\text{Elect} = A$ as long as there is at least one actual vote for B is a mirror image of the proof for $V_B = B$.

References

- Baumgartner, M. (2008) "Regularity Theories Reassessed," *Philosophia* 36, 327-354.
- Bollen, K. (1989) *Structural Equations with Latent Variables*. John Wiley & Sons.
- Casella, G. and R. Berger (2002) *Statistical Inference*, Second Edition. Duxbury.
- Collins, J. et al. (2004) *Causation and Counterfactuals*. Cambridge: MIT Press.
- Dowe, P. (2000) *Physical Causation*. New York: Cambridge University Press, 2000.
- Freedman, D. (2005) *Statistical Models: Theory and Practice*. Cambridge: Cambridge University Press.
- Glymour, C. et al. (2010) "Actual Causation: A Stone Soup Essay," *Synthese* 175, 169-192.
- Glymour, C. and F. Wimberly (2007) "Actual Causes and Thought Experiments," in *Causation and Explanation*, 43-67. Edited by Campbell, O'Rourke, and Silverstein. Cambridge: MIT Press.
- Goldman, A. (1999) "Why Citizens Should Vote: A Causal Responsibility Approach," *Social Philosophy and Policy* 16(2), 201-217.
- Hall, N. (2007) "Structural Equations and Causation," *Philosophical Studies* 132, 109-136.
- Hall, N. (2004) "Two Concepts of Causation," in *Causation and Counterfactuals*, eds. Collins, Hall, and Paul, Cambridge: MIT Press.
- Halpern, J. and J. Pearl (2005) "Causes and Explanations: A Structural-Model Approach. Part I: Causes," *British Journal for the Philosophy of Science* 56, 843-887.
- Hart, H. and T. Honore (2002) *Causation in the Law*, 2nd Edition. Oxford: Clarendon Press.
- Hitchcock, C. and J. Knobe (forthcoming) "Cause and Norm," *The Journal of Philosophy*.
- Hitchcock, C. (2009) "Structural equations and causation: six counterexamples," *Philosophical Studies* 144, 391-401.
- Hitchcock, C. (2007) "Prevention, Preemption, and the Principle of Sufficient Reason," *The Philosophical Review* 116(4), 495-532.
- Hitchcock, C. (2001) "The Intransitivity of Causation Revealed in Equations and Graphs," *The Journal of Philosophy* 98(6), 273-299.
- Hitchcock, C. (1995) "The Mishap at Reichenbach Fall: Singular vs. General Causation," *Philosophical Studies* 78, 257-291.
- Holland, P. (1986) "Statistics and Causal Inference," *Journal of the American Statistical Association* 81, 945-960.
- Kallenberg, O. (2002) *Foundations of Modern Probability*, Second Edition. New York: Springer-Verlag.
- Kline, R. (1998) *Principles and Practice of Structural Equation Modeling*. New York: Guilford.

Livengood, J. (ms) "Structural Causation and Actual Causation"

Livengood, J. and E. Machery (2007) "The Folk Probably Don't Think What You Think They Think: Experiments on Causation by Absence," *Midwest Studies in Philosophy* 31, 107-127.

Mackie, J. (1965) "Causes and Conditions," *American Philosophical Quarterly* 2(4), 245-264.

Megill, A. (2007) *Historical Knowledge, Historical Error: A Contemporary Guide to Practice*. Chicago: University of Chicago Press.

Menzies, P. and H. Price (1993) "Causation as a Secondary Quality," *British Journal for the Philosophy of Science* 44, 187-203.

Norton, J. (2003) "Causation as Folk Science," *Philosophers' Imprint* 3(4)

Pearl, J. (2000) *Causality: Models, Reasoning, and Inference*. Cambridge: Cambridge University Press.

Salmon, W. (1998) *Causality and Explanation*. New York: Oxford University Press.

Sartorio, C. (2004) "How to be Responsible for Something without Causing It," *Philosophical Perspectives* 18, 315-336.

Schaffer, J. (2003) "Overdetermining Causes," *Philosophical Studies* 114, 23-45.

Tomer, A. (2003) "A short history of structural equation models," *Structural Equation Modeling*. B. Pugsek, editor. West Nyack: Cambridge University Press.

Wolff, P. (2007) "Representing Causation," *Journal of Experimental Psychology: General* 136(1), 82-111.

Woodward, J. (2003) *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press.

Notes

¹ I would like to thank Peter Spirtes, Christopher Hitchcock, Justin Sytsma, Peter Distelzweig, Karen Zwier, Jonah Schupbach, Edouard Machery, James Woodward, Josh Knobe, Benny Goldberg, and an anonymous referee for many useful comments and suggestions on earlier drafts. All errors are still owned by me.

² Ideally, the data would be collected in careful experiments. However, being a man of meager means, Harold would probably have to settle for data collected in observational studies rather than experiments.

³ See Tomer (2003) for an introduction to the fascinating history of structural equation models.

⁴ Hume and Kant were both concerned with identifying the cause(s) of a known effect, not with estimating the effect of a given cause. Ducasse (1926) and Anscombe (1971) were both concerned with observing the cause of an effect (or rather, observing that some event or action caused some event). Davidson (1967) and Mackie (1980) both aim to specify the conditions (including the laws of nature) under which one event counts as a cause of another. Suppes (1970) does the same thing using probabilities, and Lewis (1973) does the same thing using counterfactuals. See Collins et al. (2004) for a sample of more recent work on causation in philosophy. Hart and Honore (2002) present causation in the legal context, and Megill (2007), especially Chapter 4, discusses the role of causation for historians.

⁵ Strictly speaking, Hall (2007) does not adapt SEMs to his purposes, but his theory is compatible with SEMs.

⁶ Two other problems are often associated with the metaphysical problem. The first problem is the psychological problem of saying what the ordinary concept or concepts of actual causation are like insofar as such concepts exist. The second problem is the linguistic problem of specifying the semantics for ordinary discourse about actual causation. The psychological and linguistic problems are related insofar as one's concept(s) of actual causation together with pragmatic constraints give rise to ordinary discourse about actual causation. The psychological and semantic problems are related to the metaphysical problem insofar as one's concepts track the way actual causation really works, or to put it in a more pragmatic way, the psychological and semantic problems are related to the metaphysical problem insofar as they are fit for the work that we want from a concept of actual causation.

⁷ See Bollen (1989) and Kline (1998) for introductions to SEMs. Freedman (2000, 85-95) describes structural equation models in terms of *response schedules*. I have suppressed the statistical details in my treatment of SEMs, since the problems I am engaged with here do not involve probabilities.

⁸ An anonymous referee wondered in what sense the graph in Figure 1 is a graph of a pre-emption case. Rather, it seemed to the referee to correspond to a case of over-determination. In a sense, I agree. One cannot tell that the case is a case of pre-emption (as opposed to over-determination) on the basis of the graph alone. That is because pre-emption and over-determination are problems that belong specifically to *actual causation*, and causal graphs represent *structural causation*, not actual causation.

⁹ Such cases are sometimes called cases of *early* pre-emption to distinguish them from harder cases in which two or more causal processes, each of which would be sufficient to produce the actual effect, are initiated but only one runs to completion. The harder cases are called cases of *late* pre-emption.

¹⁰ In some accounts, for example Hitchcock (2007), $V_c(u) = v_c$ and $V_e(u) = v_e$ are understood as singular (or token) events. Although equations of random variables pick out events in the statistical sense, I claim that the semantics for random variables requires that $V_c(u) = v_c$ and $V_e(u) = v_e$ be understood as property-instances, not events. See Livengood (ms) for a discussion of this issue.

¹¹ See Woodward (2003, 74-77) and Hitchcock (2001, 286-287) for their descriptions of the simple theory.

¹² McDermott (1995) reports strong agreement with individualist intuitions among naïve undergraduates. Recent experiments that I have conducted with Justin Sytsma suggest that the picture is not so clear.

¹³ See Woodward (2003, 83-84) and Hitchcock (2001, 289-290) for their descriptions of the amended theory.

¹⁴ Some or all of the manipulated values of the variables in W may be the actual values of the variables in W .

¹⁵ Since Halpern and Pearl's definitions allow that an arbitrary vector X of variables may be an actual cause, they include a condition—minimality—that ensures that actual causes are as small as they can be. I have not included the minimality condition here because I will only be concerned with singleton sets below.

¹⁶ See Schaffer (2003) for an excellent review of the problem and an argument that individualism and collectivism exhaust the plausible reactions to over-determination cases.

¹⁷ An anonymous referee points out that some process, mark-transfer, or conserved-quantity accounts avoid paying any modal cost in cases of over-determination. I agree. Schaffer appears to have in mind counterfactual and regularity accounts, for which individualism/collectivism is a pressing issue, and I have followed his lead here.

¹⁸ Graphically more complicated voting scenarios can easily be imagined. One might include direct causal dependencies between pairs of votes. One might include common causes of votes (owing to the influence of a demagogue, for example) or common effects intermediate between the votes and the outcome of the election (in order to model voting machines or electors in the Electoral College, for example). Still, I think we ought to get clear about the simplest cases first, for if an account of actual causation cannot get the simplest cases right, it seems unreasonable to hold out hope that it will get more complicated variations right.

¹⁹ This is not the place for a lengthy discussion of the role of intuitions in philosophical research. However, some brief comments are in order. I rely on my intuitions as premisses in rather simple arguments. If the reader shares my intuitions, then the arguments go through. If not, then the intuitions become targets of further investigation. I do not intend to flag some epistemically privileged mental states with my use of “intuition.”

²⁰ See Hall (2004) on the differences between causal production and causal dependence.

²¹ See Schaffer (2003), Section 5.

²² Whether you think that the various professors are responsible for Steve getting or not getting the job is a different, though related question. See Livengood and Machery (2007) for a discussion of causation by absence and its relation to explanation. See Sartorio (2004) for a discussion of the relationship between causation and ascriptions of moral responsibility.

²³ See the New York Times article here <http://www.nytimes.com/2008/09/08/education/08students.html> and the Colorado Springs Gazette article here <http://www.gazette.com/articles/vote-40925-colorado-students.html>.

²⁴ Perhaps you think, like my friend and colleague Justin Sytsma, that if McCain had won the election under these circumstances, then *the voter fraud* perpetrated in Colorado and Virginia would have been the actual cause of his victory. I agree that it would have been an actual cause in such circumstances. Moreover, it would have been the most salient actual cause from the perspective of moral and legal responsibility attribution. Moreover, if a variable for fraud is included in the structural model, the theories will all say that the fraud was an actual cause of the outcome. However, the theories will continue to say the same things about the votes that they said before; adding a variable for voter fraud will not change the deliverances of the theories with respect to the causal role of the votes. Actual causation is widely agreed to be non-transitive. If $A(u) = a$ is an actual cause of $B(u) = b$ and $B(u) = b$ is an actual cause of $C(u) = c$, it might still be the case that $A(u) = a$ is not an actual cause of $C(u) = c$. On the other hand—and here is where the problem comes in—if the structural model looks like $A \rightarrow B \rightarrow C$, then every theory of actual causation supposes that if $A(u) = a$ is an actual cause of $C(u) = c$, then it must be the case that $A(u) = a$ is an actual cause of $B(u) = b$ and $B(u) = b$ is an actual cause of $C(u) = c$.

²⁵ Hitchcock (personal communication) reminded me of the importance of contrasts in causal judgments.

²⁶ I am making the simplifying assumption that there are only Republicans and Democrats in the election.

²⁷ Hitchcock (personal communication) replies to the asymmetry intuition as follows: “I agree that it is counterintuitive to say that a vote for Gore counts as a cause of Bush’s victory. But I’m not sure that is the same as saying that Gore cost Nader the election. We can imagine scenarios where that might sound true, even if the votes are the same. E.g., at first Nader and Bush are the only candidates, and slightly more voters favor Nader. Then Gore enters. Most Nader voters switch to Gore, but no Bush voters do. Bush narrowly wins. In this case, I don’t think it’s so unreasonable to say that Gore cost Nader the election. If Gore hadn’t run, Nader would have won.” I am sympathetic to this reply; however, I think it adds something to the structure of the election scenarios by bringing in time.

²⁸ Given my definition of random variable, we really should code the variables so that the values are real numbers. For the sake of clarity, I have simply used the letter of the candidate, instead.

²⁹ In the event that $i > 2j$, $l = 0$, and $i + j = 2k + 1$ for some natural number k , the difference between i and j is odd. In order for the outcome to depend on a vote for B , enough votes have to be moved from A to C such that if V_B had been a vote for C , then C would have won or tied the election. Otherwise, there will be a subset of vote re-distributions that will result in a change in the outcome without changing the value of V_B .