# Fast and Slow Causation: An Interventionist Account of Speed of Change

Contributed Paper, *Philosophy of Science Association 28th Biennial Meeting,* Pittsburgh

Contact:
Brian J. Hanley
Merrimack College
brianjohnhanley@gmail.com

## Abstract

This paper elucidates an important feature of type-level causal relationships that is critical for understanding why disasters occur in sociotechnical systems. Using an interventionist theory, the paper explicates a concept, *causal delay*, to characterize differences between how rapidly or slowly interventions can make a difference to their effects. The paper then uses this explication to illuminate aspects of causal reasoning in everyday and scientific cases involving speed of change. In particular, the paper shows how *causal delay* clarifies why some systems are more prone to disasters than others. The paper closes by analyzing critical tradeoffs in choices between different interventions.

## Introduction

Traditionally, philosophers of causation focused on identifying criteria to distinguish causal from non-causal relationships. However, recently, philosophers have begun investigating other important distinctions in causal reasoning: distinctions among causal relationships. This paper contributes to this literature by adding a new causal concept to a growing list that includes *actual* and *potential-difference makers* (Waters 2007), *causal stability* (Woodward 2010), *causal specificity* (Waters 2007, Woodward 2010), and *reversible* and *irreversible causation* (Ross & Woodward forthcoming). The concept explicated here describes differences in the relative rapidity or lag between an intervention and a subsequent change to its effect. I call it *causal delay*.

Philosophers have primarily considered temporal aspects of causation within the traditional project of defining causation, per se. Here, I show this is neither necessary nor a preferable approach to this subject. First, I show how the strategy Woodward (2010) introduced to characterize causal features, such as *stability* and *specificity*, can also be applied to temporal proximity. Using this strategy, I distinguish two temporal dimensions of causal relationships: *causal delay* and *causal inertia*. To motivate why this strategy is fruitful, I apply my explication of *causal delay* to everyday and scientific cases and show how it illuminates causal reasoning. In particular, I show how *delay* clarifies why technological disasters occur and helps analyze crucial choices between interventions that may prevent future disasters. I conclude by combining *delay* with other causal concepts to reveal tradeoffs in these choices.

## 1. Woodward's Interventionist Strategy Towards Causal Features

At least two temporal features of causation can be distinguished within an interventionist theory. The first characterizes how quickly or slowly a change to cause can make a difference to its effect. I call this feature "causal delay," which is the focus of what follows. The second describes how quickly or slowly an intervention can change a causal variable. I call this feature "causal inertia" and discuss it elsewhere (BLIND 2021). To define these two features, I draw upon a strategy used by Waters (2007) and Woodward (2010) that uses a minimal criterion for causation as a basis to identify nuanced distinctions among causal relationships. I begin by describing this strategy and the interventionist theory in which it is framed.

Woodward's interventionist theory rests on what he calls a "minimalistic" and "undemanding" definition of a causal relationship (2010). For Woodward, $X$ is a cause of $Y$ if and only if there is some intervention that changes the value of $X$ that could lead to a subsequent change in $Y$ (or the probability distribution of $Y$). In this definition, the variable relata represent properties or events in a system. Different values of variables represent changes to the respective property or events. Roughly, an "intervention" is a manipulation of just the causal variable, whereby any subsequent change to the effect occurs only through the change to the cause. Though "minimal" and "undemanding," this definition provides a powerful basis for analyzing causation.

Consider a simple illustration. For Woodward, a match strike is an interventionist cause of a flame since intervening on whether a strike occurs or not would make a subsequent difference to whether a flame occurs or not. Notice, no actual strike must occur for the causal relationship to hold. For this reason, interventionism is a counterfactual "type-level" causal theory. To formalize the example, let a causal variable, *S*, represent a strike and let an effect variable, *F*, represent the occurrence of a flame. For simplicity, let both be binary variables, where *S* has two possible values (*no-strike*, *strike*) representing the absence or occurrence of a strike, and *F* has two possible values (*no-flame*, *flame*) representing an absence or ignition of a flame. Since interventions changing *S* (from *no-strike* to *strike*) make a subsequent difference to whether *F* will change (from *no-flame* to *flame*), then *S* is a cause of *F*.

According to this theory, there are many causes of a match flame. For instance, dampness is also a cause of a flame, since intervening on the presence or absence of moisture also would make a subsequent difference to whether a flame can occur. Let another variable, *D*, represent moisture surrounding the match, where different values represent different moisture levels. Since interventions changing this variable lead to a subsequent change in whether *F* could turn from *no-flame* to *flame*, then *D* is a cause of *F*.

Here, the minimalism of Woodward's criteria might lead some to suspect a problem. For example, some might find it awkward to say that *dampness caused a match to light*. Certainly, a match strike is importantly different from surrounding moisture, causally speaking. Such an observation might lead some to suggest emending Woodward's criteria to avoid such awkward results. After all, one might suppose causal theories should not conflict with everyday causal judgments. Woodward himself points out this worry is not illusory, but arises from his definition leaving out features presumed to be "characteristic of paradigmatic causal relationships" (2010, 290). Nevertheless, Woodward resists calls for emending his definition. Instead, he takes an alternative strategy that leverages the definition's minimalism to bring other "paradigmatic characteristics" of causation into causal analysis in deeply clarifying ways.

To illustrate Woodward's strategy, consider one important characteristic left out of his criteria: *invariability*. Woodward (2010) shows causal reasoning often considers the range of background circumstances within which a cause can continue to make a difference to its effect. In response, some philosophers try to include a particular "invariability" requirement in their definition of causation, per se. However, Woodward instead allows causes to come in degrees of invariability. To characterize this variable causal feature, Woodward explicates the concept "causal stability" built upon the minimal definition of causation. Interventionist causes have some degree of *stability*, determined by the relative range of circumstances where interventions on a cause can continue to make a difference to its effect. Causes that can make a difference across many circumstances are *strongly stable* causes. Those holding in fewer circumstances are *weakly stable* causes. I will draw upon *stability* later (**Sections 3, 6**).

Applying this distinction to the match example can clarify why a match strike is distinctive from dampness. Striking a match can make a flame across many circumstances, especially those where people typically manipulate matches. By contrast, changing surrounding moisture may only make a flame in fewer circumstances. For example, drying a match may lead to a subsequent flame only in circumstances where striking is already taking place. Both make a difference to this effect, but intervening on $S$ can turn $F$ from *no-flame* to *flame* in more circumstances than $D$. Hence, $S$ is a more stable cause of $F$ than $D$. This causal difference helps clarify why match-strikes are often more salient than dampness in causal reasoning about matches and why this salience can shift.

Because a match-strike is a *more stable* intervention, it offers a more reliable and useful means to achieve typical purposes for using a match, across more of the most common circumstances. In this way, match-strikes should be distinctively salient in causal reasoning about matches. Their increased stability makes them more reliable means of intervening to make a desired difference. Yet, as campers at wet campsites know, this salience can easily shift. This too makes sense, as dampness should be more salient in the narrower, but real, circumstances where it makes a critical difference to whether a flame can occur. In this way, concepts like *stability* illuminate why match-strikes are generally more salient and how causal reasoning adapts to different situations.

## 2. Causation and Temporal Proximity

In addition to invariability, Woodward (2010) discusses two other features philosophers often view as characteristic of paradigmatic causal relations: *specificity* and *proportionality*. This paper analyzes a fourth: *temporal proximity* between a cause and its effect. Woodward's definition includes no temporal requirements, except causes precede their effects. Yet, many philosophers have highlighted the significant role temporal proximity plays in causal reasoning. Like with invariability, philosophers traditionally tend to incorporate temporal requirements into their definitions of causation, per se.

In their influential theory of causation in the law, Hart & Honoré (1959) advocate for "natural limits" on the time that may pass between a cause and its effect. They observe that most outcomes depend on many necessary historical factors. For example, a train derailment may depend on a bent rail, but it also depends on the train's design, construction, and operations. However, Hart & Honoré claim "the cause" of the derailment must be a factor that is not "too remote" or "too close" in time to the accident. Consequently, they define causes as "abnormal" events: a cause is a deviation from "normal" conditions that precipitates an effect. Factors relating to design, construction, and operation may be necessary, but they are not causes, according to Hart & Honoré, because they lie outside their "natural" temporal criteria for causation.

Philosophers have attempted to define causation using a variety of temporal requirements. For example, compared to Hart & Honoré's flexible "natural limits" criteria, causal process theories adopt more stringent temporal requirements by defining causation in terms of *temporal continuity* (e.g., Salmon 1984, Dowe 2000). On the other hand, formal and informal practices of "root cause analysis" invert this requirement by associating greater temporal remoteness with more real causal influence. Given these disagreements, *temporal proximity* appears similar to *invariability*. It is an important causal feature, that comes in degrees (from continuity to "roots"), and influences causal reasoning in various ways. If so, then applying Woodward's strategy may be fruitful.

## 3. Two Types of Interventionist Speed

Woodward's minimal interventionist criteria leave room for at least two distinct temporal dimensions of causation. Each marks a different feature of causal relationships associated with speed of change. In other words, each concept describes a different aspect of how quickly or slowly a system can respond to interventions. The feature I call “causal delay” relates most closely to the subject of temporal proximity discussed in the previous section. Though the focus here is *delay*, it is helpful to contrast it with a distinct temporal concept I call “causal inertia” (Hanley 2021).

*Causal delay* characterizes differences in the time between a change to a causal variable and the subsequent change of an effect. Causes with *shorter delay* are marked by a more rapid change in an effect after an intervention on a cause occurs. Causes with *longer delay* are marked by longer lag between an intervention on a cause and the subsequent change in effect. In contrast, *inertia* characterizes differences in the amount of time required to change the value of a causal variable. *Causal inertia* characterizes differences in the time it takes to change a causal variable itself. Causes with *higher inertia* are those where changes to a causal variable require longer times before a change can occur. Causes with *lower inertia* are those where causal variables can be changed relatively quickly.

Causal relationships may have various combinations of *delay* and *inertia*. For example, causal relationships could have *short delay* and *high inertia* if it would take ample time to change a causal variable, but such a change could bring rapid subsequent differences in its effect. Alternatively, causal relationships could have *long delay* and *low inertia* if quick causal changes are possible but make subsequent differences after a substantial lag. In **Sections 5-6**, I demonstrate how distinguishing these two concepts and combining them with other causal concepts reveals tradeoffs in causal reasoning. For now, I mention these combinations to help isolate the precise temporal causal property *causal delay* identifies and to underline a critical aspect of how the concept works.

*Causal delay* should be analyzed as a feature of "type-level" causal relationships. For example, *delay* should not be conceived only as the time that does pass between a particular match strike and a particular flame. *Delay* should be understood as the time that would pass between possible match strikes that could make a difference to possible flames. In other words, the *delay* of an intervention does not merely describe how much time happens to pass between an actual intervention and its effect; it describes how much time tends to elapse between a potential intervention and the potential difference it would make. *Causal delay* is less like the reading on a stopwatch and more like a description of a "fuse" that begins to burn when a causal change occurs.

Construed as a type-level feature, *causal delay* can identify differences among interventionist causes much as *stability* does. To illustrate, consider another potential cause in the match example. Let a causal variable, *G*, represent a magnifying glass focusing sunlight on a match. Like a match strike, represented by *S* with two values (*no-strike*, *strike*), let *G* take on two possible values (*no-focused-sunlight*, *focused-sunlight*), representing whether sunlight is or is not focused through the glass. Since changes to *G* also make a difference to whether a flame occurs or not, represented as *F* (*no-flame*, *flame*), *G* is an interventionist cause of *F*. While *S* and *G* are both causes of *F*, a match strike and magnifying glass differ in terms of their *delay*.

For most matches in many circumstances, changes to *S* more rapidly lead to a subsequent change in *F* than changes to G. Striking a match (changing *S* from *no-strike* to *strike*) tends to be quickly followed by a flame (a change in *F* from *no-flame* to *flame*). On the other hand, focusing sunlight through a glass (a change in *G* from *no-focused-sunlight* to *focused-sunlight*) tends to make a similar difference but only after a longer lag. Hence, *S* causes *F* with *shorter delay* relative to the *longer delay* that *G* causes *F*. This distinction has a significant role in causal reasoning.

Learning to use a match requires knowing which interventions can make a flame. However, learning how to use a match safely and effectively often requires also knowing how quickly or slowly interventions tend to make their difference. For example, using a match to create a campfire requires knowing that there will be insufficient time to do any other activities campfires require (collecting and organizing kindling) once a match is struck. However, if the camper's only means of intervention was a magnifying glass, then they may need to accommodate the "longer fuse" by beginning their intervention at an earlier point, since the sun going down will introduce complications. Ignoring these differences would render control much more difficult, even in such a simple causal system used for a simple purpose.

Considering *delay* also illuminates choices between possible interventions, especially when combined with other features, such as *stability*. In most scenarios, striking a match is preferable because it often has a *shorter delay* and *stronger stability* than alternatives. For example, not only does a strike make a flame much more quickly than a magnifying glass, it also does so in many low-light circumstances a glass cannot. Together, *greater stability* and *short delay* can improve an intervention's reliability and readiness, both assets for a purpose like creating a campfire on a campsite at an uncertain time of day. This combination explains why campers justifiably tend to forgo magnifying glasses in favor of strike-strips. However, for a camper with a mischievous child, they may justifiably choose to forgo strike-strips in favor of a magnifying glass. Though a longer fuse has drawbacks for the camper, the additional time it takes to make flame can allow the camper to disrupt unsafe uses of a match. Hence, for safety purposes, interventions with *longer delays* can have particular advantages over those with *shorter delays* because it is more feasible to alter or reverse unsafe interventions before their harmful effects may occur.

Considering differences in *delay* in causal reasoning helps guide our expectations about how systems behave, coordinate activities with these expectations, and make informed choices between possible interventions by revealing essential tradeoffs. This added clarity and guidance are critical in contexts where these choices have higher stakes than a campfire.

## 4. Why Fast Systems are Prone to Disasters

To highlight *causal delay*'s importance for refining expectations and informing choices about interventions, consider the deep significance speed of change has in technological disasters demonstrated by the sociologist Charles Perrow's influential "normal accidents" model of safety. According to Perrow, how quickly a system responds to perturbations is a crucial determinant of why many disasters have occurred and why some systems are more prone to accidents than others. Consider an example to illustrate this insight.

The tragic *Challenger* shuttle disaster resulted from a rapid series of changes during its brief launch. Shortly after igniting the solid rocket boosters, a complex technical failure began inside one rocket's walls. This failure quickly compromised the rocket walls, leading to more structural damage to the system. In a little over a minute, the initial failure precipitated a catastrophic breakdown. Not unlike a match system, the shuttle system is characterized by changes making rapid differences to the system's subsequent behavior. As the tragic shuttle loss shows, such rapid change carries certain risks. Because of how quickly the technical failure spread through the system, little could have been done to prevent the disaster once this change occurred. In Perrow's model, the space shuttle is a paradigmatic "tightly-coupled" system: one marked by rapid change. The *Challenger* disaster is a striking example of the distinct risk these systems carry.

Perrow's model reveals that disasters like *Challenger* arise from a combination of complexity and rapid speed of change. Increased complexity alone does not carry the same risk. In fact, according to Perrow, complex "loosely-coupled" systems marked by slower change, are often resilient to disasters. Hence, how quickly a system responds to change often determines whether a disaster is likely to occur. To appreciate this insight, compare the shuttle program with another complex system: a school district.

In Perrow's model, both the shuttle program and public school district are highly "complex" systems. They both involve interactions between many people, serving in many different roles, pursuing multifaceted organizational objectives. Moreover, many of these interactions may lead to complex and unexpected behaviors. While both are complex, Perrow's model shows why a shuttle program is more prone to accidents than a school district. This disparity is due to, what Perrow calls, differences in the systems' respective "coupling."

The shuttle program is a more "tightly-coupled" system than a school district. In contrast with the sudden cascade of failures that led to *Challenger*'s loss, even significant changes to a school district can rarely induce such rapid change. For example, Perrow considers how a new literacy program affects a school system. Implementing the program may have some quick effects (perhaps on hiring, funding, reading lists), but most effects will occur after a significant lag. In fact, effects most closely related to the program's desired outcome, improving literacy, may have the most pronounced lag. While this delay may be undesirable, the slower rate of change lowers risk of harmful outcomes, since slower change allows altering or reversing interventions before any harmful effects occur. If changes in the shuttle program moved at a similarly slow speed, then catastrophic outcomes like *Challenger* could have been more readily averted. For Perrow, differences in coupling explain why systems like shuttle programs are riskier than those like public schools.

*Causal delay* clarifies Perrow's insight and allows philosophers to use an interventionist theory to elucidate the scientific reasoning surrounding them. Differences between "tightly" and "loosely coupled" systems can be understood as differences in the relative number of causal relationships with *short delay* or *long delay*. If a causal system contains more causes with *shorter delays*—a system with many "short fuses"—then their prevalence carries additional risk of accidents. If a causal system contains more causes with *longer delays*—a system with "longer fuses"—then this system is less prone to disasters, since changes that could lead to disaster are more readily altered or reversed before they do. Without considering *causal delay*, or assuming causation is defined by just one temporal requirement, the causal significance of these temporal differences is difficult to see. However, considering *delay* reveals these difference along with the challenges that speed of change raise for controlling complex systems and making them safer.

## 5. Tradeoffs Between Functionality and Risk

Though causal systems that undergo faster changes raise difficulties for control and safety, rapid speed of change also helps achieve important purposes. *Challenger* is a vivid illustration. While *short delays* in the solid rocket booster system allowed a technical failure to precipitate a disaster, this same speed also allowed the system to achieve flight. In fact, this technical failure occurred when changes unraveled too slowly during liftoff.

The shuttle's solid rocket booster design depends on a causal sequence to unfold at a particularly rapid speed. Upon ignition, changes to the rocket's interior pressure must quickly create a pressure gradient within the rocket walls capable of causing large, rubbery o-rings to expand and quickly seal joints. If this sequence occurs fast enough, then hot gas produced by the rocket is directed down, creating lift. However, if this process occurs too slowly, the system can quickly break down as it did during *Challenger*'s launch. Upon ignition, a pressure gradient formed too slowly in one joint, causing a sluggish O-ring expansion. As a result, hot gas entered the joint, damaging the O-ring and interior walls. This damage hampered the O-ring's ability to expand, further lengthening the time before a seal could be achieved. For *Challenger*, this extended lag between ignition and the joints sealing was so long that the joint could not seal. Damage from the joints led to a rapid cascade of failures throughout the system, leading to the shuttle's loss. Though fast change carries the risks Perrow highlights, *Challenger* demonstrates that improving safety is not simply a matter of slowing changes. Rather, it involves tradeoffs.

Causal systems often depend on changes occurring at what Ross (2018) calls a “particular speed” to achieve certain functions. Interestingly, she observes a similar dependence on speed as described in *Challenger*, but in biological contexts. Ross (2018) shows that some biological functions depend on molecular processes occurring at particular speeds. Much like *Challenger*, if these processes occur too slowly, the biological system can break down. Here, enzymes play a distinctively important causal role by influencing the speed of chemical reactions and, thus, helping to coordinate these processes’ speed and maintain the system’s functionality. If systems require the creation or coordination of changes causing their effects at certain speeds, such that introducing slower or faster changes leads to failures, then there are important tradeoffs between functionality and safety.

One tradeoff is between the higher risk associated with rapid change and the functionality of faster speed. As *Challenger* demonstrates, achieving space travel involves a tradeoff between a need to create and coordinate systems characterized by *extremely short delay* with the undesirably higher risk such a system poses. *NASA* knew of this tradeoff from the program’s outset (Vaughan 2006). Initially, engineers viewed the risk as acceptable, given the benefits of the shuttle’s innovative reusable design. However, after the loss of *Challenger* and *Columbia*, the program was ultimately scrapped because the tradeoff between safety and functionality became indefensible. Many vital decisions, such as whether abundant nuclear energy or agricultural chemicals are worth the risks posed by the complex, tightly-coupled systems they require, involve similar tradeoffs between risks and functionality associated with rapid speed of change.

Another tradeoff exists between lower risk and decreased functionality associated with slower speed of change. Consider Perrow's school district example. Interventions make their differences more slowly in such systems. Consequently, they are more immune to rapid losses of control that lead to disaster. However, if rapid change is desirable, achieving it is infeasible. For example, if interventions that improve literacy do so after a lag, this *delay* limits how quickly important outcomes can be achieved. Hence, slow-moving systems may be safer, but the *longer delay* involved in changing them creates constraints.

Overlooking these tradeoffs between safety and functionality can lead to flawed expectations about how systems behave and can be changed. For example, expecting interventions on literacy should work faster than a system allows may lead to undue pessimism about the intervention's effectiveness; even if it will make a desired difference, but only after an undesirably long lag. This false expectation could lead to abandoning one intervention for another. However, in a loosely-coupled system, this alternative intervention likely also has *long delay*, thus raising the same frustration and prolonging desired outcomes. Moreover, frustration with slow change may lead to a desire to replace loosely-coupled systems with more tightly-coupled ones. However, this amounts to a desire for a riskier system. For cases like school districts, this could lead to a dangerous choice.

## 6. Analyzing Intervention Choices

The desire for immediate effects can also compel choices of less effective interventions. To see why requires appreciating that tradeoffs between the functionality and dangers of *shorter delay* and safety and constraints of *longer delay* are often layered alongside other tradeoffs. I conclude by elaborating on the previous analyses by combining *delay* with other concepts.

Recall the camper's choice from **Section 3** between a match-strike, *S*, and a magnifying glass, *G*, as interventions to ignite a match. *S* had a combination of two features: *stronger stability* and *shorter delay*. A third can be added: *low inertia*, since match-strike interventions can occur quickly. These qualities each render *S* a functional intervention for creating a flame. Together, they render *S* exceptionally functional since it can rapidly make a flame across many circumstances. By contrast, *G* had two different features: *weaker stability* and *longer delay*. Again, a third can be added: *higher inertia*, since focusing sunlight through glass is cumbersome to achieve. These qualities render *G* less functional than *S* for creating a flame, but much safer. In the context of a mischievous kid, an intervention that takes more time to make, which only slowly creates a flame across fragile circumstances, will be easier to avert before a flame occurs. This everyday case provides a template for combining these concepts to illuminate choices of interventions more broadly.

Many traditional safety measures work by changing systems quickly. Emergency shutdowns, flare towers, and familiar examples like fire extinguishers rely on fast activation and rapid effects. In other words, they rely on combinations of *low inertia* and *short delay*. Conceived as such, their value and limitations as safety interventions can be further analyzed. For example, it may appear ill-advised for engineers to add further fast-acting relations to a system since this increases the risk of rapid losses of control. However, if a system is already tightly-coupled, one sees why adding fast-acting measures is necessary. After all, if one change can rapidly ripple through a system to create disaster, only readily available faster-acting interventions can avert disaster once this is underway. *Low inertia* and *short delay* explain safety measures' value. Adding *stability* reveals their limitations and why overreliance on them is deleterious to safety.

Safety scientists warn against a myopic focus on "proximate" interventions, e.g., safety measures and technical fixes (BLIND 2021). Instead, they argue, "systemic" interventions offer more robust and resilient safety improvements. Combining *delay*, *inertia*, and *stability* can explain their reasoning. Consider one example of a systemic intervention.

Safety researchers argue that creating more equitable decision-making in a system can reduce accidents across many circumstances (Hanley 2021). Roughly, the idea is systemic changes can lower the probability that accident-inducing events will occur and also whether they can precipitate disaster if they do. By contrast, intervening on a component's reliability lowers risk only in narrow circumstances where its failure precipitates accidents. Hence, systemic interventions have *stronger stability* than proximate ones. Moreover, while systemic changes take more time and effort to achieve, researchers argue they are more resilient improvements. Roughly, the idea is that once changed, safer decision-making structures are harder to undo. Hence, systemic interventions have *higher inertia* as well. Combined, the *stronger stability* and *higher inertia* of systemic causes can lower the risk of disasters across more circumstances for longer periods of time.

Because of their *higher inertia*, as well as often having *longer delay*, exploiting the advantages of systemic interventions require ample time. If safety concerns are dire, then the long delay and high inertia render them impotent and potentially distracting. If quick effects are needed, then proximate interventions will be preferable due to their *shorter delay and lower inertia*, but at the cost of being less robust and less lasting improvements. By employing the strategy and concepts offered here, such tradeoffs become clear, allowing choices of interventions best suited for the shifting situations causal reasoning often needs to navigate to achieve important goals like preventing disaster.

## Bibliography

Hanley, Brian J. 2021. *A Pragmatic Epistemology of Causal Selection in Safety Science*. Doctoral Dissertation, University of Calgary.

Dowe, Phil. 2000. *Physical Causation*. New York: Cambridge University Press.

Hart, Herbert L.A., and Anthony M. Honoré. 1959. *Causation in the Law*. London: Oxford University Press.

Perrow, Charles. 1984/1999. *Normal Accidents: Living with High-Risk Technologies*. Princeton: Princeton University Press.

Ross, Lauren. 2018. "Causal Selection and the Pathway Concept." *Philosophy of Science* 85 (4): 551-572.

Ross, Lauren, and James Woodward. Forthcoming. "Irreversible (One-hit) and Reversible (Sustaining) Causation." *Philosophy of Science*.

Salmon, Wesley. 1984. *Scientific Explanation and the Causal Structure of the World*. Princeton: Princeton University Press.

Vaughan, Diane. 2016. *The Challenger Launch Decision, Enlarged Edition*. Chicago: Chicago Press.

Waters, C. Kenneth. 2007. "Causes That Make A Difference." *The Journal of Philosophy* 104 (11): 551-579.

Woodward, James. 2010. "Causation in Biology: Stability, Specificity, and the Choice of Levels of Explanation." *Biology and Philosophy* 25 (3): 237-318.