

# What Caused the Bhopal Gas Tragedy? The Philosophical Importance of Causal and Pragmatic Details

Brian J. Hanley\*†

---

In cases in which many causes together bring about an effect, it is common to select some as particularly important. Philosophers since Mill have been pessimistic about analyzing this reasoning because of its variability and the multifarious causal and pragmatic details of how it works. I argue Mill was right to think these details matter but wrong that they preclude philosophical analysis of causal selection. I show that analyzing the pragmatic details of scientific debates about the important causes of the Bhopal Gas Tragedy can illuminate causal reasoning about disasters and shed new light on causality and causal selection.

---

**1. Introduction.** Why are some causes judged to be more important than others? Do some causes have features that make them more important? Or are these judgments unfounded? In cases in which many causes together bring about an effect, it is common to select some as particularly important

Received October 2019; revised June 2020.

\*To contact the author, please write to: Department of Philosophy, University of Calgary, 2500 University Drive NW, Calgary, AB T2N 1N4, Canada; e-mail: bhanley@ucalgary.ca.

†I received invaluable feedback and support from many individuals and communities in the course of writing this article. Thank you to Ken Waters for working with me through so many drafts, discussions, and difficulties; to my family, and to Mariesa, for everything; to my University of Calgary community, especially Ananya Chatteraj, Oliver Lean, and Timothy “TJ” Perkins, for help with several drafts; and to Niall Roe for countless formative discussions, as well as Beige Lussier, Jeremy Chapdelaine, Alison McConwell, Justin Caouette, and Denise Retzlaff. Finally, thank you to the anonymous referees, for comments that significantly strengthened the article, and to Susan Sterrett and Marc Johansen, for their thoughtful commentary at the American Philosophical Association Central Division meeting in Denver. Research for this article was supported in part by grant 50191 from the John Templeton Foundation: From Biological Practice to Scientific Metaphysics.

Philosophy of Science, 88 (October 2021) pp. 616–637. 0031-8248/2021/8804-0003\$10.00  
Copyright 2021 by the Philosophy of Science Association. This work is licensed under a Creative Commons Attribution-NonCommercial 4.0 International License (CC BY-NC 4.0), which permits non-commercial reuse of the work with attribution. For commercial use, contact journalpermissions@press.uchicago.edu.

and background the others. *Causal selection* is used regularly in scientific and everyday reasoning. Yet, philosophers tend to respond to it with either pessimism or retreat. Many doubt its merit as a philosophically interesting form of reasoning or simply ignore it in favor of more general notions of causality. The variability in how causes are selected and the diverse *pragmatic details* about different rules, reasons, and purposes involved in selecting important causes across different contexts lead many to these dismissive conclusions.

I argue that embracing the pragmatics of causal selection, something most philosophers avoid, is key to understanding how this reasoning works. Ignoring these details leaves important cases of causal selection intractable. While the reasoning behind causal selection does indeed vary, there are principled and philosophically interesting ways to analyze the pragmatics of causal selection. I show this by analyzing safety scientists' reasoning about the important causes of the Bhopal Gas Tragedy, the deadliest industrial disaster in history.

Before proceeding, it is important to distinguish two senses of "causal selection." Philosophers of science typically conceive causal selection in terms of distinctions among many causal factors. They analyze how scientists reason about *important causes* in cases with many causal factors. For example, Waters (2007) examines why genetic causes are more important than nongenetic causes for explanatory reasoning in biology. Philosophers of causation working in an analytic tradition are interested in a different sense of selecting causes. They analyze distinctions made between genuine causes and mere background conditions (Schaffer 2005, 2016). This sense of causal selection is set in terms of distinguishing causes from noncauses. For example, Schaffer (2005) asks why a spark is selected as "the cause" of a fire, while oxygen is a mere (noncausal) background condition. In this article causal selection is conceived in the former sense.

Despite significant discrepancies, a common thread binds together much of the literature on causal selection. The thread is John Stuart Mill. Waters (2007) uses Mill as a foil. Schaffer (2016) takes Mill to offer the "main argument" for the "standard view." Mill's (1843/1981) discussion of causal selection has a pervasive influence on how philosophers think about causal selection. His pessimism and dismissive conclusions about it continue to frame the philosophical problem.

**2. A Brief History of Causal Selection.** Since John Stuart Mill, causal selection has been rightly associated with variability. Mill demonstrates that even when selections follow a fixed rule, there is significant variability in how selections are made. A single rule can pick out various types of causes across different cases. Sometimes a rule makes sensible selections, and other times it selects causes "no one would regard as . . . principal" (Mill 1843/1981, 329). Mill observes that this depends on "the purpose we have in view" (329).

A rule may select causes useful for some purposes but not others. While his discussion is framed around one rule, Mill says there is no reason to suppose “this or any other rule is always adhered to” (329). Many types of causes are selected for many different purposes using many different rules. For Mill, this variability prompts pessimism.

The diverse causal and pragmatic details that vary across cases of causal selection lead Mill to conclude causal selection is “capricious,” unscientific, and outside the purview of philosophical analysis (1843/1981, 329). The severity of his conclusions is striking in the context of *System of Logic*, where Mill uses a method of refining our often-capricious everyday reasoning into more principled “ratiocination.” His method is pragmatic, involving analyzing how reasoning and logical “contrivances” are used toward their epistemic purposes (6). Yet, causal selection is uniquely problematic. Mill thinks selections are influenced by too many, and too diverse, causal and *pragmatic details*—details about different purposes, rules, and reasons for selecting causes—to be given adequate philosophical analysis.<sup>1</sup>

Mill resolves his concerns by retreating from causal selection entirely. He proposes restricting all causal reasoning to a general notion of universal and invariable causes that preclude any need for selection. This pattern of reasoning is not unique to Mill and enjoys a lasting legacy. Lewis follows a similar pattern, noting the “invidious principles” of causal selection before retreating to the “prior question” of a “broad” concept of causation (1973, 559). With this endorsement, Mill’s pessimism has “won the field” in many circles (Schaffer 2016). Consequently, causal selection and the reasoning behind it are widely ignored. However, Lewis himself recognizes that pessimism and retreat are not necessary. In a footnote, Lewis (1973, 559 n. 6) says he would be amenable to an account of causal selection like the one offered by White (1965). White’s account is similar to the more influential one formulated by Hart and Honoré (1959). Not all philosophers have agreed with Mill’s pessimism.

**3. Exceptions to Mill’s Pessimism.** Hart and Honoré’s (1959) analysis of causal selection is perhaps the most widely known challenge to Mill. Their approach represents a typical way philosophers break from the Millian pattern of pessimism and retreat. Noting that causal selection is an “inseparable feature” of causal reasoning in law, their area of interest, Hart and Honoré are compelled to challenge Mill’s view (11). Like Mill, they acknowledge

1. In *System of Logic*, Mill describes other types of everyday reasoning as “capricious.” Within his philosophical methodology this is not necessarily problematic. Capriciousness implies that philosophy is needed to refine and retool that reasoning to better serve certain epistemic ends. It does not imply pessimism or wholesale dismissal. However, Mill thinks the capriciousness of causal selection is uniquely problematic. Mill’s claims, and the problem he identifies, are more nuanced and challenging than most philosophers have appreciated (Hanley 2021).

the variability of selections. However, they contend that Mill misidentifies the source of it. They argue that only variable causal details are relevant for analyzing how selections are made, not diverse pragmatic details (17). By pushing to the background those variable purposes, rules, and reasons for selecting causes, they think analysis is possible.

Hart and Honoré develop a general rule for selecting *abnormal causes*, causes deviating from normal circumstances. Mill thinks such a rule is inadequate because many different rules are used to select causes for diverse purposes. To avoid this problem, Hart and Honoré argue that variability across selections is not due to these pragmatic details. Instead it is entirely due to how the same rule is applied to structurally different cases.

For example, in a factory where the presence of oxygen is normal, an abnormal spark is the important cause of a fire. In a factory where sparks are normal, the abnormal presence of oxygen is the important cause of a fire (Hart and Honoré 1959, 10). Important causes vary, but that variance is entirely due to “a subordinate aspect of a more general principle” as it is applied to cases with varying causal structures (17). They think causal details matter, but pragmatic details do not.

This approach offers a way to dissolve the source of Mill’s pessimism. However, it relies on a crucial assumption. Hart and Honoré claim causal selection is unique to the law. They claim that other areas, including all sciences, have only a “derivative interest” in it (Hart and Honoré 1959, 9). This circumscription is what justifies subordinating all the variability of selections to different applications of their single rule, which warrants disregarding the diverse pragmatic details Mill identifies. Only by narrowing causal selection to the law can they construe their rule as a general analysis and a solution to Mill’s problem. If causal selection is not unique to law, then their approach is a tenuous response to Mill founded primarily on neglecting the diverse pragmatics that are at work in causal selection.

Contra Hart and Honoré, causal selection is not unique to one discipline. It runs through important reasoning in science, engineering, and other areas of thought. Several philosophers have noted its significant role in explanation and investigation in biology and related sciences (Gannett 1999; Waters 2007; Woodward 2010, Franklin-Hall 2014, forthcoming; Stegmann 2014; Weber 2017; Baxter 2019; Lean 2020; Ross, forthcoming). While this undermines a key premise in Hart and Honoré’s approach, many of these philosophers of biology follow similar ones. Many have strong preferences for “ontological” analyses (Waters 2007; Stegmann 2014; Weber 2017), analyses in terms of a single general rule (Waters 2007, Franklin-Hall, forthcoming), or analyses that minimize the role of pragmatics (Franklin-Hall 2014).

Apparent exceptions to Mill’s pessimism confirm a key aspect of his concern. The diverse pragmatic details of causal selection pose problems for philosophical analysis. Mill almost certainly would deny that Hart and Honoré’s

approach addresses his problem. In fact, he rejects a very similar notion to their abnormal causes (Mill 1843/1981, 328). Mill thinks diverse pragmatics are inseparable from how causal selection works. Formulating a general rule that minimizes or ignores them is not a solution. It leads to deficient accounts that leave out much of the reasoning behind selection. If Mill is correct, then the problem is how to meaningfully analyze the pragmatic details of selections. But is he? There are good reasons to think so.

**4. Analyzing Pragmatics.** Several philosophers have acknowledged that pragmatics matter for analyzing causal selection (Gannett 1999; Waters 2006; Kronfelder 2014; Woodward 2014). While Waters (2007) offers a general rule for selecting important causes, Waters (2006) embraces pragmatic details (e.g., the purposes, strategies, and activities of biologists) to argue genetic causes are important because they are useful for manipulating biological processes. However, the most compelling evidence that pragmatics matter is the striking diversity among accounts of causal selection.

Hart and Honoré's concern with assigning liability, Sober's (1988) focus on apportioning causal contributions, Franklin-Hall's (2014) interest in how much causal information is needed to explain, and Collingwood's (1957) focus on manipulation each lead to remarkably different analyses of why some causes are selected as important. Each emphasizes what Mill would consider distinct purposes for selecting causes. If different purposes require different types of selections, and consequently require different analyses, then the variability among philosophical accounts implies pragmatics do matter.

I argue that Mill was right that pragmatic details matter for how causal selection works. However, I show that analyzing them is possible. My approach is similar to van Fraassen's (2008) way of analyzing the role of pragmatics in representation. Van Fraassen acknowledges the "variable polyadicity" of representation and embraces diverse pragmatic details about use, purposes, practices, and context in his analysis (29). Similarly, analyzing causal selection in terms of the different purposes it can serve, and the various ways selections are actually used to achieve these purposes, across varying contexts, can elucidate causal reasoning about important causes. My approach can be contrasted with other pragmatic approaches, such as Collingwood's account of causal selection in terms of what a particular human has an ability to prevent (1957, 302–4). This is a narrow sense of "pragmatic" compared to the one I demonstrate in what follows. My approach is more akin to the research program articulated by James Woodward in his PSA presidential address. It shows how to analyze causal reasoning in terms of the details about "various goals and purposes" and how causal concepts and causal knowledge "conduce" to their achievement (Woodward 2014, 693).

To argue my case, I analyze a dispute among safety scientists concerning the causes of the Bhopal Gas Tragedy. Stripped of pragmatics, the dispute

appears to involve an intractable disagreement over preferences for different methods, the resolution of which looks more like a question for sociologists than philosophers. However, when pragmatic details such as the purposes guiding the selection of important causes of disasters and the actual uses of causal knowledge are considered, the disagreement is clearly principled and philosophically interesting.

**5. Causal Selection in the Bhopal Gas Tragedy.** In 1984, a deadly disaster occurred in Bhopal, India. This disaster resulted from the release of a toxic chemical from a processing plant. The chemical spread through surrounding populated areas, killing thousands and permanently injuring hundreds of thousands more. The Bhopal Gas Tragedy continues to affect local populations today.

Many causal factors brought about this disaster, but investigators focused on a human error. During routine cleaning a worker failed to insert a device designed to prevent water from entering chemical tanks in case of valve failures. This error caused water to leak through the valve during cleaning, pouring into a chemical tank containing a large volume of methyl isocyanate. The mixing of water and methyl isocyanate created a chemical reaction with enough energy to vent tons of the toxic chemical into the air.

This account fits the traditional method for modeling disasters used by safety scientists. Typically, as in the findings of the investigators, disasters are modeled as chains of causal events. *Causal chain models* like these were developed at the dawn of scientific studies of accidents and safety (Heinrich 1959).<sup>2</sup> Experts rely on these models to understand accident causality across a range of sociotechnical systems including the plant at Bhopal. The method models disasters using chains of *proximate causes*, discrete causal events arising from relations among human agents and physical technologies within a system that occur spatially and temporally close to an accident. Examples include human errors, component failures, and energy-related events. The chain of proximate causes involved in the Bhopal disaster is the failure to insert the safety device, the resulting valve leak, and the chemical reaction. Had these causes not occurred, or occurred differently, then the disaster would have been less likely to occur if at all.

Recently, some safety scientists have argued against this methodological orthodoxy. Citing changes in complexities of systems since the beginning of safety science in the early twentieth century, they argue that causal chain models have outlived their usefulness: these models are no longer adequate

2. Safety experts sometimes refer to this class of accident models as “linear causal models.” However, as an anonymous reviewer notes, this term has a narrower technical meaning in causal modeling literature. For clarity, I refer to the models used by safety experts as causal chain models.

for understanding and learning from disasters like the one at the chemical plant in Bhopal. At the heart of their dissent is what they see as an incorrect emphasis on proximate causes, which they contend mislead scientists and engineers. Causal models that emphasize more important causes should be developed and replace the widespread use of chain models. Proximate causes are genuine causes, but they are not the most important causes of disasters.

Nancy Leveson is a leading voice against the effectiveness of causal chain models. She argues that the traditional focus on these models should be replaced by an emphasis on *systemic causes*.<sup>3</sup> Systemic causes of disasters are distinct from their proximate causes. Systemic causes are properties of a system that causally influence its behavior. They are associated with the overall design and organization of a system. In the case of Bhopal, there were many systemic causes of the disaster (Leveson 2012). One systemic cause was the operating conditions at the plant. Before the disaster, many safety devices were disabled to save money, early warning alarms and refrigerated chemical tanks among them. Poor operating conditions like these allowed the chemical reaction to occur at the strength it did while leaving the reaction undetected. Had operating conditions been different—had they been better—the disaster would have been much less likely to occur.

Leveson identifies a number of other systemic causes involved in the Bhopal disaster. A systems approach reveals design deficiencies also caused the disaster. Devices to minimize chemical releases were designed for smaller, less powerful events than the one at Bhopal. Vent scrubbers and flare towers that neutralize vented chemicals were designed for much smaller amounts of chemicals than were released during the disaster. Water curtains designed to minimize released chemicals only reached heights well below where venting actually occurred. The only devices able to reach that height were inefficient individually operated water jets. Safety culture at the plant was also very poor. Safety audits before the disaster were ignored. Alarms sounded erroneously during normal operations, making genuine alarms impossible to discern. Employees had sparse safety training and safety equipment, and there were few qualified engineers at the plant. Had these design deficiencies and the poor safety culture been remedied, the disaster would have been much less likely to occur.

Disasters in complex sociotechnical systems like in Bhopal are caused by systemic and proximate causes together. Poor safety culture, operating conditions, and design deficiencies as well as the maintenance error and leaky valve were all causes of the disaster. Since systemic and proximate causes

3. Conceptual issues involving systemic causes and complex sociotechnical systems merit further philosophical attention (Noy et al. 2015). The concept of a systemic causes can be given fuller analysis, but for this article they can be conceived in the rough way safety scientists discuss them: as causal properties of a system distinct from proximate causes.

are distinct causes of disasters, disagreements about whether proximate causal chain models or systems-based models identify more important causes are disagreements about causal selection, much as Mill defined it. One sense of the term “invidious” that Lewis (1973) uses to describe causal selection correctly applies here. Selection of the most important causes of the Bhopal disaster arouses discontent, resentment, or animosity.

To this day, corporate owners of the plant have largely avoided or denied responsibility. Contentious disputes among activists, the Indian government, and corporate entities endure over reparations and deficient cleanup. Legal battles continue steadily since 1984, with court hearings as recent as 2012. Serious questions of ethics are tightly interwoven with questions of what was the most important cause of the disaster. So too are methodological questions about how best to prevent such terrible tragedies from ever happening again.

The Bhopal disaster plays a motivating role in safety science and engineering. The tragedy demonstrates how critical it is for engineers to practice the best methods and reasoning practices available. Given the stakes, decisions over which causes are the most important for understanding and preventing disasters have aroused debate. However, while Lewis implies that all this invidiousness is reason for philosophers to retreat to generalities, the ethical and methodological significance of causal selection in cases such as the Bhopal disaster imply the opposite.

The material weight of how debates about causal selection are adjudicated gives philosophers undeniable reason to take them seriously. It also gives reason to worry that dismissing or retreating from these debates may be deleterious. Bad actors interested in avoiding responsibility surely would agree with Mill’s notion that “we have, philosophy speaking, no right” to make selections of important causes (1843/1981, 328). Cases like the Bhopal Gas Tragedy show that philosophers have, philosophically speaking, a duty to engage with causal selection deeply, not dismissively. But how can philosophers analyze disagreements about important causes, such as the methodological one in safety science?

On one side of the disagreement is a traditional method that emphasizes proximate causes. It is hard to deny that human errors and component failures are important causes of accidents in some sense. On the other side, dissenting safety experts argue against this traditional way of modeling accidents. They contend that the most important causes of disasters in complex systems are systemic ones, and the proximate causal chain models that exclude them are unsatisfactory. Systems-based approaches should be pursued instead. Again, it is hard to deny that issues with safety culture and design deficiencies were important causes of the Bhopal disaster.

Is one position more justified than the other? This question relates to another sense of the term “invidious” as unfairly discriminating, or unjust. This is the sense most assume Lewis intended. At this point it is not obvious what



a philosopher can say about which side is justified. Philosophers could analyze the causes themselves and their relations to accidents and to other causal factors. But this alone would not adjudicate the disagreement about which are more important or elucidate any reasons why. One could take the Millian approach and conclude that all a philosopher should say is, strictly speaking, all the factors taken together were the real cause. On this view, beyond that, the disagreement appears to be simply an expression of competing preferences. Engineers in supervisory or maintenance roles prefer thinking about human errors, while engineers who manage or design systems prefer systemic factors. What can philosophers say about preferences that would bear on causality?

The situation is reminiscent of Carnap's (1995) discussion of how different professionals would select different causes to explain a car crash.<sup>4</sup> Carnap attributes those differences to divergent preferences determined by professional interests. He promptly ends his analysis there and appeals to a general notion of "cause" nearly identical to Mill's. This abrupt end implies Carnap thinks there is nothing philosophical left to say about disagreements about important causes. Road engineers prefer causes related to their interest in roads, police prefer causes related to an interest in policing, but preferences like these are not philosophically interesting. How can a philosopher of causation give a deeper analysis to preferences?

**6. Pragmatics of Leveson's Reasoning.** There is more to the disagreement about the Bhopal disaster than competing preferences. This becomes clear when Leveson's (2012) reasoning is analyzed in terms of pragmatics. Leveson's reasons for selecting systemic causes over proximate ones relate to the purposes safety scientists have when evaluating the causes of disasters and to how they actually use causal knowledge to achieve these purposes. Analyzing Leveson's causal reasoning in terms of these pragmatics reveals that the disagreement over the relative importance of systemic and proximate causes is principled. Principles that are about which types of causes are more important than others for certain purposes.

Leveson argues that proximate causes are "misleading at best" (2012, 28). To narrowly focus on them leads to "ignoring some of the most important factors in terms of preventing future accidents" (33). In large, complex systems of humans and technology, proximate causes are less important

4. An anonymous reviewer correctly points out that this example is first found in Collingwood (1957). Unlike Carnap, Collingwood thinks there is a principle guiding selections of different causes of a crash. He explains different selection in terms of what the selector is "able to produce or prevent" (302–4). Like Carnap, Collingwood's analysis of selection does not go deeper. As the next sections shows, there is more to the pragmatics of selecting important causes for preventing accidents than Collingwood's notion of selecting what a person can "put right."

causes for preventing certain types of behavior. From a pragmatic perspective, there is a principled rationale behind claims about the relative importance of systemic causes over proximate causes.

As Leveson explains, proximate causes of disasters are unlikely events. The probability of a human error occurring in the way that it did at the plant in Bhopal is quite low. Given the complexity of the system it occurred within, anticipating when it would occur would be near impossible. Furthermore, directly preventing proximate causes such as human error is notoriously difficult (Reason 1990). Finally, whether a maintenance error leads to disaster is contingent on an unlikely confluence of many other factors. Proximate causes are unlikely, difficult to predict, problematic points of interventions that are only weakly related to disasters. Because of these features, focusing on them implies disasters like the Bhopal Gas Tragedy are unpredictable and unavoidable events.

Leveson and other safety experts argue disasters, including the one in Bhopal, can often be predicted and prevented. Given the poor state of safety culture, operating conditions, and so on, at the chemical processing plant, imminent disaster was likely and knowable. Had engineers intervened on these systemic factors, then disaster would have been less likely. Yet, this perspective is obscured by focus on proximate causes.

The possibility to predict and prevent disasters in complex sociotechnical systems requires causes that conduce to these goals. Safety experts who argue for the importance of systemic causes are not basing this on their preferred methods or what they happen to be able to control. Their argument is that systemic causes are more important because knowledge of these causes offers distinct epistemic and technical advantages over other causes for the purpose of preventing disasters such as the Bhopal Gas Tragedy.

Stopping where Carnap concluded his analysis of the car accident would render the disagreement among safety experts philosophically intractable. Digging into pragmatic details here has revealed important lines of reasoning behind claims like the ones made by Leveson. In turn, it raises further questions amenable to philosophical analysis. Philosophers can ask, what properties do systemic causes have that make them useful for prediction and prevention? What do proximate causes lack that makes them less important for these purposes?

**7. Philosophical Analysis of Pragmatic Details.** To begin analyzing the importance of systemic causes in the Bhopal Gas Tragedy, the proximate and systemic causes of the disaster can be framed in the same causal theory. To do this, I use Woodward's (2003) interventionist theory of causation. Roughly speaking, according to Woodward's account, causal relations hold between two variables when a change in the value of one variable would bring about a subsequent change in the value of the other (or an increase

or decrease in the probability of a change in value of the second variable). Woodward conceives this relation in terms of hypothetical changes and counterfactual dependence relations. There are many technical aspects to his account, but it is not necessary to draw on them here.

Recall, proximate causes are discrete causal events arising from relations among human agents and physical technologies within a system. At the plant in Bhopal, the human error causing the disaster was an employee failing to insert a physical safety device. This caused the subsequent component failure of water leaking through a valve. These proximate factors caused the release of a chemical from the plant. Safety experts represent this kind of causal process as a chain of proximate causal variables, as discussed in section 5. Each node of a causal chain model can be given a straightforward interpretation as interventionist causal variables, as represented in figure 1.

Human errors and component failures are understood as events, but they can be formulated as binary variables in the interventionist framework. For example, the variable HE in figure 1, representing a particular human error, can take on two values: a human error occurring or not. The variable CF, representing a particular component failure, also takes on two different values: a component failure occurs or not. The arrows connecting each variable represent a causal relationship, defined by the following dependence relations. Changes in the value of HE have counterfactual or actual control over the probability of a change in the value of CF. Changes in the value of CF in turn have counterfactual or actual control over the probability of a change in the value of the target effect variable E. There are a number of different ways to conceive of the target effect variable when representing the Bhopal disaster or any complex causal process. For the purpose of this analysis, it is easiest to formulate the variable as safety experts do in causal chain models, as a binary accident variable. In this case, the variable E can take on two different values: chemical is released or not.

Systemic factors, such as safety culture and operating conditions, also stand in an interventionist causal relation with accidents and disasters. Actual or hypothetical interventions on systemic properties do or could bring about subsequent changes to the behavior of the system and the likelihood of certain types of effects. The possibility of systems engineering and systems-based approaches to safety, such as Leveson's, depend on systemic factors having this



Figure 1. Release of chemicals, represented by E, depended on a component failure, CF, and a human error, HE. The values these variables take on are associated with the occurrence of the respective associated behavior in the system.




Figure 2. Systemic variables have causal control over the release of chemicals. Changes in the value of S are associated with the probability distribution of changes to the value of the same effect variable, E, as in figure 1.

kind of causal control over system behavior. Figure 2 offers one way to model a systemic cause of the Bhopal disaster.

A systemic causal variable S also stands in a counterfactual dependence relationship with the same target effect variable E from figure 1. Changes to the values of a systemic causal variable can be associated with changes made to a systemic property. To make this more concrete, in the case of the Bhopal disaster, changes to the value of S can be associated with changes to the safety culture at the plant. By changing safety culture, safety experts can change the probability of the effect associated with E occurring. Changes in the value of S lead to subsequent changes in the probability of a change in the value of E, that is, the probability of a chemical being released or not.

Interventions on systemic properties are one means of controlling behaviors of sociotechnical systems and the likelihood of certain types of events occurring within them. While systems approaches are founded on this possibility, experts can also intervene on systems in other ways. Before the development of systems-based approaches, most safety practices focused on proximate factors as the primary points of intervention. Safety was understood primarily in terms of controlling human errors and component failures. Approaches like Leveson's are an alternative to these traditional approaches based on emphasizing different causal factors to use controlling systems. Hence, figures 1 and 2 represent different causal relationships emphasized by alternative approaches.

So far it is still not clear why one of these models is more predictive or offers better interventions for prevention. The two models show that both proximate and systemic causes offer causal control over disasters. They do not indicate much else and do not elucidate why systemic interventionist causes may be more important than proximate ones for preventing disasters. To understand the causal selection surrounding cases such as the Bhopal disaster, analysis of these causal relationships needs to go deeper. More nuanced causal concepts like Woodward's notion of causal stability, and a new concept I begin to develop here, fill out a richer philosophical analysis. When considered together with the pragmatics of selection they give a principled account of why systemic causes are more important than proximate causes in complex sociotechnical systems.

In complex systems like the one in the Bhopal disaster, no proximate or systemic causes invariably cause accidents. For example, human errors or

component failures in sociotechnical systems, such as the ones that occurred in the plant at Bhopal, do not always lead to accidents. A wide range of other circumstances must have obtained in order for them to have actually brought about their effects. Consequently, proximate causes like those in figure 1 must be understood as causes strictly in relation to a background of other causal factors within which they cause their effects. The same goes for most causes involved in disasters, including systemic causes, and most causes of anything represented in Woodward's interventionist framework. Interventionist causes almost always require some range in the values of other causal variables to obtain in order to bring about their effect.

For Woodward, unlike Mill, this is not problematic. Causes need not be invariable but can come in degrees of invariance. For example, consider the causal relationship between the temperature of water and the effect of changes to phase state. The relation holds within certain ranges of temperature changes. The relation also only holds within particular ranges of values in other variables (e.g., ambient pressure). The variables of temperature and pressure, among others, create a multidimensional space. Within that space, there are regions within which changes to temperature bring about changes to the effect. This region in the multidimensional space under which the relation holds is the *invariance space* of the cause. Woodward (2010) sometimes refers to this concept as the 'background conditions' of a cause. I choose not to use this term because the term *conditions* can suggest that the background is not causal. While some parts of a background are not causally relevant, much of it is in Woodward's account. The causally relevant background is what I refer to when using the term *invariance space*. All the causal relations represented in figures 1 and 2 should be conceived as holding only within some invariance space of other causal factors.

The size of invariance spaces varies for different causal relations. In other words, some interventionist causes hold under larger invariance spaces than others. In Woodward's framework, the relative sizes of invariance spaces for different causal relations can be analyzed and compared using the concept he calls *causal stability* (Woodward 2010). Minimally stable causes hold under small regions. Maximally stable causes hold across wider ranges of circumstances. Causes are more or less stable, depending on the relative sizes of their respective invariance spaces. This kind of conceptual tool that distinguishes causes in terms of the features they possess, and to what degree, is essential for analyzing causal selection.

While Woodward develops the concept of causal stability for analyzing explanations of biological phenomena, the concept can also help elucidate causal selection that is aimed at preventing rather than explaining. The causes of the Bhopal disaster with the most stability are systemic causes. The proximate causes hold only under highly specific circumstances. As I explain below, systemic causes hold across a wide range of different circumstances.

Hence, the concept of stability offers a clear basis for analyzing why some causes are more useful for the purposes of predicting and preventing accidents than other causes.

Recall, a key consideration in Leveson's claims about the importance of systemic causes is that proximate causes make the Bhopal disaster appear contingent and unpredictable. Yet, she argues systemic causes make it clear that a disaster was bound to happen. For these reasons, Leveson thinks systemic causes are more important for the purpose of predicting disasters. Another key consideration in her reasoning was that systemic causes offer more important means of intervening on a system to prevent disasters, while proximate causes such as human error are less effective means. These pragmatic considerations can be analyzed in terms of differences in causal stability and used to give a principled analysis to the selection of systemic causes as more important for achieving these epistemic and technical aims.

Had the causal circumstances surrounding the maintenance error and leaky valve been even slightly different at the plant in Bhopal, then they likely would have had much different effects. For example, had the maintenance error occurred elsewhere in the plant, or at different time during its operation, then it would not have had the kind of disastrous effects safety experts are interested in preventing. Had the plant been organized and operating more safely, then these variables would no longer bear a causal relationship to disaster. The variables represented in figure 1 stand in a causal relationship under a relatively small invariance space compared to systemic variables, as I show next. Proximate causes are weakly stable causes.

Systemic causes increase the probability of a disaster in a system like the one in Bhopal across many different circumstances. For example, the poor safety culture could have caused a disaster at the plant under many different possible proximate causal events. Had the maintenance error occurred differently or not at all, the probability of disaster would have remained high because of the state of the systemic factors. Changes to these systemic factors would have lowered the probability of a disaster under the actual circumstances and many others. Excellent safety culture lowers the probability of disaster across many different proximate causal events. In sum, changes to value of systemic causal variables in figure 2 control the probability distribution of an effect variable under a large invariance space. Systemic causes of disasters have strong causal stability.

Relative differences in stability give systemic causes and proximate causes different epistemic and technical functions. Causes with less stability are useful for predicting effects in narrow sets of changes to a system. In contrast, more stable causes predict their effect across a larger set of changes a system does or may exhibit, even for changes that are difficult to predict or observe. Highly stable systemic causes offer an epistemic advantage for safety experts to achieve their goal of assessing the risk of disastrous effects

in large, complex systems that exhibit many unpredictable changes over time.

Causes with increased stability offer interventions on a system more conducive to the task of preventing disasters. Systemic causes can suppress their effects under many perturbations to a system. This kind of control is advantageous when the effects are disasters. Toward the technical ends of safety experts, systemic factors have an additional causal feature offering complementary advantages. The feature can be further analyzed by looking closer at the interrelations of proximate and systemic causes.

There are at least two ways proximate and systemic causes might relate in systems like the Bhopal disaster. Figure 3 represents one way they are related. Systemic factors can have control over the values of many different possible proximate causes of an accident effect. For example, improved safety culture at the plant in Bhopal would lower the probability of the maintenance error. It also would lower the probability of a number of other possible human errors. This models one aspect of Leveson's reasoning in which she points out that even if the exact maintenance error had not occurred in Bhopal, given the systemic causes involved, some other human error was likely to cause a disaster anyway. By controlling a systemic cause, engineers can decrease the probability of many different human errors occurring and consequently decrease the likelihood of an accident (see fig. 3). This feature explains why systemic causes are important for the purpose of preventing disasters.

In addition to having control over the value of proximate causes, there is reason to think systemic causes bear a second kind of causal relationship with the probability of accidents and disasters. Some safety scientists think human errors and component failures are inevitable in complex systems (Reason 1990, 409). If proximate causes are inevitable, then systemic causes do not have complete control over their occurrence or the values they take




Figure 3. One way systemic and proximate causes relate is through a systemic variable having control over the probability of many different proximate causal chains leading to the effect.

on. Nevertheless, these safety experts think systemic causes offer causal control needed to prevent accidents and disasters. Consequently, systemic causes do not just relate to proximate causes by controlling changes to their values (as represented in fig. 3). Furthermore, the causal relation between systemic variables and their effects (fig. 2) is not necessarily a mere abstraction of the intermediate control of proximate causes (fig. 3). The interrelationships of systemic and proximate causes, and their effects, are more complicated. To this point, I propose that systemic causes exhibit an interesting causal feature that has not been analyzed in terms of Woodward's interventionist framework.

Some changes to systemic factors, such as when safety culture is improved, weaken the stability of proximate causes. Yet, when the systemic factor is changed in another way, such as when safety culture is worsened, then the stability of proximate causes is strengthened. Systemic causes have control over the strength of other causal relations. For example, human errors cause disasters under a wider range of circumstances when systemic factors are at a certain value (e.g., when safety culture is poor). When those systemic factors are changed (e.g., when safety culture is improved), then the same proximate factors cause disasters under a smaller range of circumstances. Systemic factors bear a *damping* and *amplifying* relation, a relation not to their effects but to the strength of other causal relationships around them, that is, other causal relations in their invariance space. I represent this relationship in figure 4 with the wavy line and in figure 5 in the standard framework for directed acyclic graph.

This relation is not a relation of cause and effect, but it is nevertheless causal in nature. It can be understood similar to a *moderating variable* as discussed in the context of some social sciences.<sup>5</sup> Suppose occurrences of human errors can be reduced but not eliminated. If so, then changes in systemic factors do not produce changes in the values of these ineliminable human errors (as in fig. 3). Put in more interventionist language, if particular proximate causes are inevitable, then changes to systemic variables do not have control over the values that these proximate causes take on. Systemic factors are not interventionist causes of ineliminable proximate causes. Nevertheless, changes in systemic factors do influence whether these proximate causes can actually bring about accidents and disasters. One way they do this is through altering the background circumstances a proximate cause

5. I thank an anonymous reviewer for pointing out this similarity and for suggesting the form of representation in fig. 5. Amplifying/damping causes may be a type of moderating cause in the sense discussed by Kincaid (2012, 59–60), but here they have a narrower sense set in terms of Woodward's concepts of stability and invariance space. I agree with Kincaid that there are limitations representing moderating causes with acyclic graphs; hence, I offer the nonstandard representation in fig. 4.





Figure 4. Wavy line represents another way systemic and proximate causes interrelate. This amplifying/damping relation holds when changes in a variable influence the strength of other causal relations, rather than the values the variables take on. The wavy line is not part of standard method of constructing directed acyclic graphs.

continues to bring about its effects. Changes to some systemic factors can increase or decrease the range of circumstances in which the occurrence of a human error stands in a causal relation with disasters. Hence, changes to this type of systemic factor appear to change the size of the invariance space of some proximate causes. In other words, some systemic causes can influence the stability of other causes. I call this relationship *amplifying/damping*.

This causal feature further fills out the analysis of why systemic causes are important for preventing disasters by showing again how they are more useful means of preventing certain types of effects from coming about in complex systems. This causal concept also offers a way to enrich Woodward's framework further, raising an underexplored aspect of causality. More detailed exploration and analysis can be pursued beyond this preliminary presentation. However, this should point to the importance of digging deeper into details, rather than retreating to generalities or ending analysis when pragmatics arise.

The disagreement surrounding Bhopal is over what the most important causes are for the purpose of predicting and preventing disasters. All agree what the causes are but disagree which are most important. Stripped of pragmatic considerations, it was unclear how to analyze the disagreement. Once these details were considered, analysis could continue to show why more stable systemic causes that control many proximate causes, and amplify/dampen others, are better predictors and means of intervening for preventing disasters



Figure 5. Amplifying/damping relation can be represented as collider structures in a standard framework for directed acyclic graphs. In this simplified form, this structure does not evoke the distinctive nature of the amplifying/damping relation systemic causes bear on proximate causes.

in complex sociotechnical systems. As Mill thought, the pragmatic details of Leveson's reasoning matter for understanding how selections are made. However, he was wrong to think that this precludes philosophical analysis. As this section shows, analysis can go deeper. When embraced, the kind of multifarious pragmatic details of causal selection that worried Mill actually provide a basis for deeper analysis.

**8. The Philosophical Importance of Pragmatic Details.** The history of philosophical interest in causal selection is largely defined by skepticism about the pragmatics this reasoning involves. The literature tends to follow the general tenet that *causal details* (types of causes and their features, properties, structures, interrelationships, etc.) are important for analysis, but *pragmatic details* (different purposes, rules, reasons, activities involved in how people reason about and use causes) can and should be avoided. Many consider these details to be “mere pragmatics,” and their inclusion is assumed to give less rigorous analyses, or anything goes. Philosophers who acknowledge some role for pragmatics mostly do so narrowly, trying to minimize their significance. However, the Bhopal case study implies pragmatic details are crucial for analyzing causal reasoning about important causes.

Some readers might nevertheless still question the significance of pragmatics in my analysis. They might claim that the philosophically interesting aspects of this case of causal selection are exhausted by purely causal concepts. They might argue that the world contains more or less important causal features, stability among the important ones. They could accept the idea that considering pragmatics helps uncover these important features but argue that whether a cause has those features does not depend on whether humans are interested in using them for their purposes. They could conclude that pragmatics offer a window into interesting aspects of causality, but specific details about our purposes can ultimately be disregarded in favor of general concepts about objectively important causal details. I close by showing the adverse philosophical and social consequences of this view and how they are resolved by embracing pragmatics in their rich detail.

Recall, Hart and Honoré ignore the diverse pragmatic details that worried Mill by circumscribing causal selection to a problem unique to history and law. For Mill, the consequences of this approach are part of why causal selection is philosophically problematic. Construing a single principle as a general solution, Mill says, results in a lot of faulty causal reasoning. The Bhopal case study confirms the adverse epistemic consequence of disregarding pragmatic details in favor of generality. The case also shows potentially pernicious social consequences of ignoring Mill's warnings.

Hart and Honoré's abnormalism principle selects the maintenance error as the important cause of the Bhopal disaster, since it deviates from normal circumstances and is associated with a human act. Their principle deems

systemic causes quite unimportant, since they are remote, mostly fixed aspects of the causal process. Section 7 shows this is an ineffective selection for preventing disasters. However, even for questions about liability that Hart and Honoré develop their account to answer, their selection is dubious.

Assigning legal or ethical responsibility to a mistake made by a poorly trained and weakly managed worker is contentious. Even more so is minimizing responsibility for the systemic deficiencies of the plant. Maintaining that this analysis is endorsed by the general solution to this kind of question is more problematic still, easily bolstering denials of culpability by corporate actors involved in the tragedy and working against activists seeking justice and reparations for the suffering it created. These controversial results needlessly arise from ignoring pragmatics.

The Bhopal Gas Tragedy involves complex causal details and pragmatic details. Considering them in their rich complexities allows philosophers to ask and answer more precise questions about the case. For assigning liability, philosophers can ask whether responsibility should be afforded to those with better control over preventing disasters or those with less control. Causal and pragmatic details taken together provide more nuanced answers. While section 7 has an epistemological focus, the analysis could be expanded to give principled answer to legal and ethical questions. If liability ought to be afforded to the agents most capable of preventing tragic disasters, and systemic causes are most important for preventing them, then those closely associated with systemic causes should be held most responsible. This gives a principled basis for assigning greater responsibility to the corporate actors who controlled the systemic causes than to workers associated with proximate causes less important for prevention. Embracing causal and pragmatic aspects of causal selection in their fuller detail clarifies reasoning about these cases and helps philosophers and society avoid many needless, and potentially harmful, confusions.

While Hart and Honoré's account is not useful for analyzing which causes are important for preventing the Bhopal Gas Tragedy, or determining liability in the case, their analysis is based on cases in which it makes effective selections. Readers might correctly note the analysis in section 7 has similar limitations. The analysis of why systemic causes are important in the Bhopal case study does not necessarily provide a general analysis that applies across all kinds of systems or purposes.

In simpler systems than the one in Bhopal, systemic causes may be less important than proximate causes for prediction and prevention. Proximate causes also figure prominently in methods of "root-cause analysis" used to reconstruct how accidents occur. For engineers, such contexts in which proximate causes are important can undermine the need for the methodological changes Leveson proposes. For philosophers, cases in which proximate causes are important may appear to generate a contradiction or reintroduce the specter of

caprice. However, in a robustly pragmatic perspective apparent contradictions and caprice are illusory and ripe sources of insight.

Safety scientists use causal knowledge to assess the risk of certain behaviors and identify effective interventions to prevent them. For complex sociotechnical systems, they do so with limited means of making interventions and with limited knowledge of the types of behaviors a complex sociotechnical system will exhibit over time. Given their purposes, and the epistemic and technical methods and constraints they have for achieving them, the stability and amplifying/damping of systemic causes make them important tools. However, for different purposes and practices, it makes sense why proximate causes may be more useful tools.

In early stages of investigation, root-cause analyses are the principal method safety experts use to piece together what happened. They consist of ways to trace sequences of proximate causes backward from an accident. Because accidents are complex, difficult to observe, and often destroy physical evidence of their causes, chains of proximate causes occurring close to an accident are important causes in such investigatory practice. Knowledge of proximate causes that have occurred or are most likely to occur within a system are also important for inferring time-sensitive decisions about how to prioritize interventions. The importance of proximate causes for these practices does not contradict the analysis of their unimportance for prevention in systems like in the Bhopal Gas Tragedy. Rather, it is further evidence of the value of analyzing causal selection in terms of the diverse pragmatics guiding it.

Causes are important for some purposes and not for others. This depends on how the causal details of a case (the types of causes present, what features they have, their interrelations, etc.) relate to details about the types of activities (epistemic, technical, etc.) used to pursue certain purposes. Keeping these pragmatic details in the forefront of analysis transforms apparent contradictions and caprice into revealing contrasts for philosophers to examine. Contrasting how and why causes with certain features are useful for different purposes, and how this changes across different types of systems, offers insight for developing a richer picture of how causal selection works. Pragmatics illuminate how causal reasoning adapts to different circumstances. Ignoring pragmatics obscures this.

Finally, readers might challenge the analysis in section 7 with the following kind of counterexample. They might argue that, if increased causal stability is why causes are important for preventing disasters, then there is a more important cause than systemic causes. The cause with maximally stable control over the Bhopal Gas Tragedy is the existence of the plant in Bhopal. Intervening on its existence prevents disasters across the widest range of possible circumstance. Within some approaches to philosophy of causation, philosophers might claim this is a trivial consideration and is as a counterexample along the lines of a “problem of profligate causes” (Menziez 2004).

They might claim endorsing the existence of the plant as an important cause is an absurdity similar to endorsing the queen's failure to water my plants as causally relevant to their deaths and contend this forms a *reductio ad absurdum* to the analysis in section 7. However, accepting this requires unjustifiable oversimplification.

Safety experts like Leveson are tasked with developing causal knowledge and methods to improve the safety of systems without eliminating them. Outside of these pragmatic constraints, the existence of the plant is far from a trivial causal consideration. The Bhopal Gas Tragedy raises a real question of whether the best intervention to prevent horrific disasters is to eliminate chemical processing plants. This consideration is not an instance of a tried-and-true philosophical counterexample. It is a live option, depending on the details of what we choose to achieve and how best to pursue it. Ignoring pragmatic details obscures this perspective. It may also inadvertently ally philosophers with those who dismiss as meaningless questions about whether some system's existence is a cause of suffering. Analyses of causal selection that are sensitive to diverse pragmatic details make this clear and enable philosophers to play a more active role in socially important debates.

## REFERENCES

- Baxter, Janella. 2019. "How Biological Technology Should Inform the Causal Selection Debate." *Philosophy, Theory, and Practice in Biology* 11 (2): 1–17.
- Carnap, Rudolph. 1995. *An Introduction to the Philosophy of Science*. 2nd ed. New York: Dover.
- Collingwood, Robin G. 1957. *An Essay on Metaphysics*. Oxford: Clarendon.
- Franklin-Hall, Laura. 2014. "Explaining Causal Selection with Explanatory Causal Economy." In *Explanation in Biology*, ed. Pierre-Alain Braillard and Christophe Malaterre, 416–36. New York: Springer.
- . Forthcoming. "Causal Economy Account of Scientific Explanation." In *Minnesota Studies in the Philosophy of Science*, vol. 21, ed. Brian J. Hanley, C. Kenneth Waters, and James Woodward. Minneapolis: University of Minnesota Press.
- Gannett, Lisa. 1999. "What's in a Cause? The Pragmatic Dimensions of Genetic Explanations." *Biology and Philosophy* 14:349–74.
- Hanley, Brian J. 2021. "What Capricious Reasoning Means for Philosophy: Mill's Problem of Causal Selection." Unpublished manuscript, University of Calgary.
- Hart, Herbert L. A., and Anthony M. Honoré. 1959. *Causation in the Law*. Oxford: Oxford University Press.
- Heinrich, H. W. 1959. *Industrial Accident Prevention*. New York: McGraw-Hill.
- Kincaid, Harold. 2012. "Mechanisms, Causal Modeling, and the Limitations of Traditional Multiple Regression." In *The Oxford Handbook of Philosophy of Social Science*, ed. Harold Kincaid. New York: Oxford University Press.
- Kronfelder, Maria. 2014. "How Norms Make Causes." *International Journal of Epidemiology* 43 (6): 1707–13.
- Lean, Oliver. 2020. "Binding Specificity and Causal Selection in Drug Design." *Philosophy of Science* 87 (1): 70–90.
- Leveson, Nancy G. 2012. *Engineering a Safer World*. Cambridge, MA: MIT Press.
- Lewis, David. 1973. "Causation." *Journal of Philosophy* 70 (17): 556–67.
- Menzies, Peter. 2004. "Difference-Making in Context." In *Causation and Counterfactuals*, ed. John Collins, Ned Hall, and L. A. Paul. Cambridge, MA: MIT Press.

- Mill, John S. 1843/1981. *A System of Logic*. Toronto: University of Toronto Press.
- Noy, Y. I., L. J. Hettinger, M. J. Dainoff, P. Carayon, N. G. Leveson, M. M. Robertson, and T. K. Courtney, eds. 2015. "Editorial: Emerging Issues in Sociotechnical Systems Thinking and Workplace Safety." *Ergonomics* 58 (4): 543–47.
- Reason, James. 1990. *Human Error*. Cambridge: Cambridge University Press.
- Ross, Lauren. Forthcoming. "Causal Control: A Rationale for Causal Selection." In *Minnesota Studies in the Philosophy of Science*, vol. 21, ed. Brian J. Hanley, C. Kenneth Waters, and James Woodward. Minneapolis: University of Minnesota Press.
- Schaffer, Jonathan. 2005. "Contrastive Causation." *Philosophical Review* 114 (3): 297–328.
- . 2016. "The Metaphysics of Causation." In *Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta. Stanford, CA: Stanford University. <https://plato.stanford.edu/entries/causation-metaphysics/>.
- Sober, Eliot. 1988. "Apportioning Causal Responsibility." *Journal of Philosophy* 85 (6): 303–18.
- Stegmann, Ulrich. 2014. "Causal Control and Genetic Causation." *Noûs* 48 (3): 450–65.
- van Fraassen, Bas C. 2008. *Scientific Representation*. Oxford: Oxford University Press.
- Waters, C. Kenneth. 2006. "A Pluralist Interpretation of Gene-Centered Biology." In *Minnesota Studies in the Philosophy of Science*, vol. 19, ed. Stephen Kellert, Helen Longino, and C. Kenneth Waters, 190–214. Minneapolis: University of Minnesota Press.
- . 2007. "Causes That Make a Difference." *Journal of Philosophy* 104 (11): 551–79.
- Weber, Marcel. 2017. "Which Kind of Causal Specificity Matters Biologically." *Philosophy of Science* 84:574–85.
- White, Morton G. 1965. *Foundations of Historical Knowledge*. New York: Harper & Row.
- Woodward, James. 2003. *Making Things Happen*. Oxford: Oxford University Press.
- . 2010. "Causation in Biology: Stability, Specificity, and the Choice of Levels of Explanation." *Biology and Philosophy* 25 (3): 237–318.
- . 2014. "A Functional Account of Causation." *Philosophy of Science* 81 (5): 691–713.