Degrees of Token Causality: Understanding and Reasoning

Phyllis Illari, David Lagnado, and Samantha Kleinberg

Abstract

Token causes in science and everyday reasoning are often described with degrees, allowing communication about how much each cause contributed to an outcome. Yet some philosophers argue that token causality does not admit of degrees. We respond to these arguments by reviewing a vast literature in psychology to argue that people think about causality in degrees, and also argue that degrees provide valuable information for science, aiding in both causal selection and prioritizing information.

Word count: 4479

Preprint of paper accepted for presentation at the 2026 PSA Biennial Meeting

1. **Introduction**

Token causal questions are pervasive across science and daily life: why did a non-smoker develop lung cancer, why did a plane crash, why is an individual's blood glucose high. These are questions about specific events at specific times and places, rather than what causes cancer, plane crashes, or glucose spikes in general (at the type level). Often many factors contribute to bringing about an outcome, so questions arise about the degree to which different factors contribute. That is, it may be correct to say a person's low blood glucose Tuesday at 10am was caused by their Type 1 diabetes, exercise the day before, insulin dose that morning, and breakfast, among others. However, this makes it difficult to prevent this outcome in the future and communicate what actually occurred (lest we are reduced to merely listing the dozens of factors that contributed). Even to find only whether or not something token-caused an effect, in practice many methods for finding causes quantify their contribution (Merck and Kleinberg, 2016) or assess their statistical significance, raising the question of what degree counts as a cause.

Despite the pervasiveness of degrees of token causality, some philosophical work has regarded token causality as a binary (cause or not-cause), arguing that token causality cannot admit of degrees. This leads to a puzzle, where degrees of token causality appear common in everyday reasoning and in science, yet papers argue that they are not necessary and are somehow undesirable. In this paper, we respond to these arguments and offer our own positive arguments that token causality can come in degrees. Our aim is not to take a metaphysical position about degrees of token causality, but rather to focus on cognitive and pragmatic arguments: we will argue that people think about causality in degrees, and degrees provide valuable information for science, aiding in both causal selection and prioritizing information.

In section 2 we examine philosophical objections to degrees of token causality, addressing the narrowness of the literature which focuses on simple examples, and devotes extensive argument to 'explaining-away' the appearance of degrees of token causality. Despite the philosophical objections, in section 3 we argue that extensive literature in psychology studying causal reasoning shows that people naturally attribute degrees of token causality. In section 4, we reject the explaining away strategy, and draw out the key disagreement we have with Sartorio (2020) and Baron et al. (2025): there does not need to be a unique way of establishing degrees for them to exist and be useful. In section 5, we argue that degrees of token causality are important to reasoning in science more generally, and can play the roles we have identified even with such conflict. Thus philosophical arguments against degrees of token causality would have to explain away both human reasoning and scientific practice, in an attempt to avoid a conflict that we can instead embrace. We conclude in agreement with Mumford (2013, p111) "I would hope that all future philosophical theories of causation take as a desideratum that the scalarity of causes must be accommodated." Specifically, any account of token causality should incorporate the important reasoning functions of degrees in both everyday and scientific reasoning.

2. **Objections to degrees of token causality**

Type-level causes are quantified using data from multiple instances, like collecting data on a number of people eating the same meal and measuring their glucose before and after to understand its effect. Degrees of type causality are less controversial than degrees of token causality, with many philosophers adopting different philosophical theories and developing ways of measuring degrees of causality within them (see for example Northcott, 2005; Braham and van Hees, 2009; Sprenger, 2018). However, key works accept degrees in type causality but reject them for token causality or consider it more of a problem (Baron et al. 2025, p2; Braham and Hees, 2009, p323-4). At the token-level, degrees of causality are about the specific instance at hand: how much did her lunch on Tuesday raise Jane's blood glucose? Given recent developments in personalized nutrition, we know that the same meal can have completely opposite effects for different people (Zeevi et al., 2015), so we cannot simply assign the population average for a meal. As a result, quantifying degrees of contribution for a token cause is a harder task in practice, but it is not impossible, as work in personalized nutrition has shown.

To address objections to degrees, we focus on two recent papers, as they illuminate a key reason why philosophers have rejected degrees of causality. First, Sartorio (2020) is primarily interested in moral responsibility and liability, particularly the case of just war, but her view is that degrees of causality cannot do this job, which she argues for using an extended puzzle. She compares one scenario where two assassins simultaneously fire two bullets at a victim, where both bullets were sufficient for death, with a contrasting scenario where each of two simultaneous bullets was necessary for death. Her objection to degrees of token causality, in a nutshell, is that if we understand degrees according to closeness to necessity, we will have a different judgment of causal contribution for each scenario than if we understand degrees according to closeness to sufficiency. There might even be further conflicting criteria for deciding when a cause makes a greater contribution than another, and Sartorio's key objection to degrees of token causality is that we cannot have a general resolution for these conflicting criteria. Second, Baron et al (2025) argue against degrees of token causality, while raising no concerns about type-level degrees. While their paper ranges over more cases than Sartorio's, their core argument is still that the "diversity of options" (Baron et al., 2025, p5) for measuring degrees undermines the case that degrees play a role in practice.

As Demirtas (2022) notes, what is meant by degrees is not extensively discussed in the literature we will criticise. Demirtas identifies three options: a relative sense of degrees, where we can compare one cause of an effect as more significant than another; a more absolute sense (following Northcott, 2005); and causal strength (p4). We might add 'causal contribution', as that is extensively used by both Sartorio and Baron et al. For our purposes, nothing turns on these distinctions, as we argue that degrees are *allowed* in some form or another. We reserve for future work examination of the distinction between causal contribution and causal strength.

We wish to draw attention to three features that Sartorio and Baron et al's arguments share. First, both papers agree that there *appear* to be degrees of causality, at least in some cases, and spend considerable time explaining away these appearances. A standard example is that of two companies polluting the same river, company X, which emits 80% of the total pollution, and company Y, which emits 20%. Sartorio and Baron et al. both agree that in such a case X appears to be a stronger cause of, or make a larger causal contribution to, the resulting environmental damage or disease. Second, both papers primarily argue using relatively simple cases such as the lake, bludgeoning someone to death, or two independent assassins shooting simultaneously. The extensive efforts spent 'explaining-away' the appearance of degrees of causality in both papers is for the most part targeted at such simple examples. Baron et al. acknowledge psychological literature that shows people attribute degrees of token causality, but quickly dismiss this as a mere artifact of the experimental process or indicator of degrees of another quantity like belief. Finally, both works assume that for apparent degrees to capture degrees of causality, there cannot be conflicting judgments about degrees of token causality, since degrees of causality must be unique. For example, Baron et al. (2025, p10ff) argue that their key objection to degrees of token causality is that there are a diversity of approaches to determining it, and we cannot all agree on the criteria for such a determination. They acknowledge that there could be positive reasons for accepting degrees of token causality. However, they claim that with multiple 'degree-theoretic' accounts of causality, and the fact that even within some theories there are still different ways to determine degrees, there is no reason to think that one is privileged (p10), and they maintain that, “No single measure of degree of causation will capture all of the work that the degrees of causation-adjacent features might be needed for” (p11). Baron et al. (2025, p12ff) claim that multiple dimensions of causality combined into one measure cannot yield a unique account, a “single aggregated comparative ranking” (p12). For them, this was the final possibility for saving degrees, it fails, and so they reject them.

In contrast to these views, we maintain that there does not need to be a unique account of degrees of causality, in order for degrees of causality to do the jobs we have identified for them. In this manner, we agree with Goh (2023), who defends the ‘incommensurable’ approach to degrees of causation that Sartorio’s criticisises. However, our analysis goes beyond Goh’s to address the psychological literature (section 3) and scientific practice (section 5).

3. **Degrees of Causality are a Feature of Human Reasoning**

To advance this debate, we begin with a substantial examination of what the literature in psychology has shown about degrees of causality. Baron et al. suggest that degrees may be a proxy for confidence, or may be an artifact of unnatural experimental procedures, specifically due to priming effects of the scales used. To the contrary, we argue that a large body of research supports the claim that degrees are an intuitive part of how lay people reason about token causality, and thus should be addressed by philosophical theories of token causality.

Notably, the cases examined in psychology are broader and richer than the simple intuitive examples that Sartorio and Baron et al. explain away as a mere appearance, and are much harder to explain away.

Degrees are not proxies for confidence: First, Baron et al. argue that participants may be reporting their belief, so that high degrees of causality merely indicate high confidence in whether something was a cause rather than the extent to which something was causal. However, the psychological literature has addressed this issue, showing that people understand the difference between questions of confidence and causal strength, and readily give judgments that track their beliefs in the latter. Although Baron et al cite a set of studies by O'Neill et al (2022) as evidence in favour of the confidence account, these studies actually show that people's graded judgments are a function of both confidence and causal strength: "Overall, we found that causal judgments were multimodal and that people make graded judgments both when they think a cause is weak and when they are uncertain about its causal role" (O'Neill et al., 2022). In Experiment 2 of that paper they systematically explored this with scenarios with varying number of causes (businesses dumping waste) and structure (disjunctive, where one business dumping waste leads to contamination; conjunctive where other numbers are required for contamination). With more causes of an outcome people considered each candidate less causal, yet this did not affect confidence, leading to cases where participants reported 100% confidence in their judgment that something was a weak cause. In other cases people were unsure but thought that if something was a cause, it was a major contributor.

Degrees are stable and replicable: Next, if reported degrees of causality were an artifact of experimental processes in psychology or due to priming (i.e., seeing continuous scales rather than binary choices making people more attuned to degrees) or question wording, we would expect to see significant variation between study results (due to sensitivity to study materials) and no consistent pattern in how degrees track other features of problems. Instead we find the contrary: people's causal attributions reliably track factors such as how much a cause contributed to the likelihood or intensity of an event. For example, Spellman (1997) found that when events had different contributions to the probability of an outcome, people assigned causal ratings that tracked those contributions and ratings strictly increased as probability did. Here the causal contribution was making the effect more likely to have occurred. In this case, participants had full information about the probabilities, what events occurred, and the outcome. Thus, degrees do not merely reflect uncertainty or lack of information.

Across a wide variety of experimental paradigms, and using a range of scales (e.g., 0-100 sliding scales, 1-7 Likert scales) and measures (e.g., 'to what extent did one event cause another', agreement with statements saying 'one event caused another'), lay people readily give graded judgments of causality (see Waldmann, 2017). Moreover, graded judgments vary systematically with key experimental manipulations. For example, when making judgments about physical events, people's causal ratings are sensitive to factors such as robustness or stability (how likely it is that the cause will still produce the effect even if background conditions are perturbed) and how much (quantitative) changes in the cause lead to (quantitative) changes in the effect (Grinfeld et al., 2020; Gerstenberg et al, 2021; Vasilyeva et al., 2018). In group contexts, people's causal ratings to individual members for a team outcome vary according to each member's contribution and role (Gerstenberg et al. 2023, Lagnado et al 2013; Zultan et al., 2012). For example, in team competition settings (e.g., a quiz, or cooking show) participants rate how much each team member causally contributes to the final team outcome using a 0-100 scale. Across a range of contexts, participants give ratings that vary systematically according to the task structure, such as whether all team members need to attain a high score, or whether only one high score is sufficient. Thus, in a cooking contest where the task structure is disjunctive (only one member needs a high score), but three members score highly, they are each assigned a lower causal rating than when only one player succeeds. Indeed participants' graded ratings are closely predicted by a formal counterfactual model that allows for degrees of causal responsibility in terms of closeness to a team member being pivotal for the team outcome (Chockler & Halpern, 2004). In the context of human agency, people's causal judgments vary according to the agent's level of intention and the degree to which the agent foresaw the outcome (Kirfel & Lagnado, 2021; Lagnado & Channon, 2008). People's causal judgments also vary systematically with the perceived normality of the causes (Icard et al., 2017; Kominsky et al., 2015; O'Neill et al., 2022).

On the other hand, if the results were due to priming by scales, we would instead expect the results to track differences in the scales used, which is not what we see. While various types of scales have been used (Likert agreement ratings, 0-100 rating scales) people's ratings remain consistent and show the same systematic variation with key experimental changes. For example, one major finding in the psychological literature is that when multiple causes combine to yield an outcome, people's causal judgments are influenced by the prior probabilities of the causes. For example, if winning a prize requires both rolling a six on a die and tossing heads on a coin, people attribute higher causal ratings to the die landing six than the coin landing heads (Kominsky et al., 2015). But the direction of the effect depends crucially on the nature of the causal interaction. In the conjunctive case, where both causes are necessary for the effect, people reliably judge the less probable cause to have higher causal strength, whereas when either cause is sufficient (and both occur), people judge the more probable cause to have higher strength. This is a well-established result across a range of studies,contexts, and measurement scales (Icard et al., 2017; O'Neill et al, 2022; Quillien & Lucas, 2024), and shows how causal ratings change systematically with prior probabilities. These patterns of ratings cannot be explained by priming, confidence, or deviant scale use.

Thus, the consensus view in psychology is that degrees of causality are a fundamental aspect of people's cognition, including in the token cases we have surveyed here. Given that people's behavior is consistent, systematic, and leads to effective interventions (as we discuss in section 5), this is not a cognitive quirk but rather a useful way to reason.

4. Why explain away degrees?

A core strategy of Sartorio and Baron et al. is to 'explain away' the appearance of degrees of causality as really degrees of something else such as blame, moral and legal responsibility, confidence in our causal ascriptions, necessity, sufficiency, or effort. Baron et al. (2025) and Sartorio accept degrees of other features, but question whether we can have degrees of causality, and spend time explaining away their appearance. Having reviewed the psychology literature, we see the true extent of what must be explained away. Degrees appear in simple intuitive cases, but also across a wide range of more complex and nuanced scenarios, as an extensive experimental literature shows. Baron et al. do not engage at all with the positive reasons why it might be worth retaining degrees. In our view, the crucial question is not *can* we explain away degrees, but *should* we? We have given important positive reasons to retain degrees drawing on the literature in psychology, and will give further scientific reasons in section 5.

In the face of those reasons, the puzzle, then, is why Sartorio, and Baron et al., are so keen to explain away degrees of token causality at all. Ultimately for both papers their key objection to degrees of causality seems to be that they cannot yield a unique account of degrees of causality to underlie an account of degrees of responsibility. Sartorio is explicit on this, as her whole paper is framed within the context of just war theory. In Baron et al. it is less clear, but at a pivotal point in the argument, we get, “Degrees of causation must generate a stable overall ranking among causes of *e* to govern the allocation of responsibility for *e*” (Baron et al., p11). In spite of this, though, the conclusion of their paper is framed as being much more general, that degrees of causality are not needed.

We agree with Sartorio and Baron et al. that degrees of causality cannot set degrees of moral responsibility. Bernstein (2017) surveys reasons to be sceptical that a metaphysics of causality can “account for the kind of causal differentiation relevant to moral responsibility”, (p181). This has also been addressed in the psychological literature, which has found that the relationship between people’s judgments about causality and other things like blame and responsibility is complex, although causal judgments often serve as a critical precursor to blame judgments (Lagnado & Channon, 2008; Malle et al., 2014, Shaver, 1985). However, as we will see next, degrees are highly useful for both understanding and intervention in science.

5. Degrees are a Feature of Scientific Reasoning

Degrees are intuitive to how people think about token causality, and are crucial to scientific practice. Thus in this section we expand our argument beyond people's ordinary causal judgments to the role of degrees of causality for reasoning in science. We think there are many places where degrees of causality are important to scientific practice, but focus here on two key ones: understanding and intervention. The core reason why degrees of causality are so important for understanding and intervention, though, is complexity.

Currently, the frontiers of many sciences are struggling with complexity, ranging across economics, medicine, epidemiology, climate science, ecology, and many more. Scientists are not searching for a needle in a haystack, hoping to locate a cause. Instead, the key feature of complexity we wish to draw out is this: causes are not scarce. The problem is instead the opposite, that scientists are swimming in possible causes, and we need to select from among many causes the most useful one(s). Take the Foresight Obesity Map, a systems map of the many factors influencing obesity (Foresight, 2007). It contains 108 nodes spanning 7 domains and about 300 edges (causal links) connecting them (McGlashan et al., 2018). Thus, one could rightly list over 100 factors contributing to obesity rates at a given time and place, but without further information, it is difficult to understand why obesity rates rose or why a policy to improve them was successful. Thus we are dealing with a kind of 'causal overload' analogous to information overload, making causal selection a critical problem. While most of the work has been at the type level, philosophers have examined selection for specificity, proximal variables, or variables that can be intervened on (Woodward, 2010; Ross, 2018; Quillien and Lucas, 2024; Frühstückl, 2026). However, we argue that degrees of causality are a very important one. Further, the need to select among a vast array of causes even in token cases shows a benefit of degrees that has been obscured by the overly simplified examples used frequently in the philosophical literature, as they typically only include a small number of causes.

The first reason scientists need to select among many token causes is to understand phenomena. Philosophers have recognized that not all factors that contributed to an outcome are equally appealing as causal explanations, and this is exacerbated in science by the complexity of many phenomena. In the Foresight obesity map, it would be hard to meaningfully understand rates of obesity in an area without information on the relative contribution of these factors, which are unavailable as degrees are not included. In medicine, many factors could have a nonzero contribution to a patient's symptoms, and so all count as token causes for their symptoms, but the goal is to find the strongest cause or one that contributes most to symptoms, not every relevant factor to better guide treatment. Importantly, we can achieve understanding and meaningful selection among token causes even with disagreement about how to measure degrees. One example can be found in nutrition research: dietary intake is complex so to aid communication and understanding there are different representations for diet such as ultra-processed foods (with disagreements on what counts and how to quantify), and scores such as the healthy eating index. None of these are objective and different groups have different scoring systems, but the simplification helps communicate what we think the cause is, and scientists still agree that a poor diet (regardless of exact definition) is a cause of numerous chronic diseases (Frongillo et al., 2023). Thus, even without unique degrees, knowing which factors contributed the most to an outcome helps communicate why it happened more clearly.

While understanding alone may be a goal of science, we also seek to use the knowledge gained to intervene, which is a second reason for selection. While pragmatic and other factors certainly affect our ability to intervene (Sloman & Hagmayer, 2006), we suggest causal contribution is also important. Recalling our earlier diabetes example, a person makes many decisions about what to eat, how much insulin to administer, and whether to exercise while aiming to ensure their glucose stays in a healthy range. To actually control glucose, though, the person needs to know the degree to which different factors brought about an outcome to ensure they are intervening on the right thing. Again, there would be an overwhelming number of causes involved if all were to be shared, and could lead to suboptimal interventions if people selected which to intervene on in the absence of degrees of contribution. Merck & Kleinberg (2016) developed methods to quantify impact of token causes on event probability, timing and intensity, allowing these to be determined for token causes. Other works such as by Kleinberg (2013) and Zheng & Kleinberg (2017) developed methods for using type-level causal knowledge to quantify token causes. Work on how people use causal information to make decisions is in line with the importance of degrees.

Kleinberg & Marsh (2023) showed that providing only the causal information needed for a decision (meaning, the cause most likely to produce a given outcome) led to better choices than a complex causal model in 16 different domains. Further, decision accuracy with the most complex models was equivalent to accuracy with no information at all. Kleinberg et al. (2022) demonstrated that when weak causes are omitted from a causal model, lay people trusted it more, while omitting stronger causes reduced trust in the model. Taken together this work suggests that degrees of causality can help prioritize which causes to share when explaining token cases, and deciding where to intervene. Merck & Kleinberg (2016), demonstrated how one can determine the effect of physical activity and meals on blood glucose, quantifying the effects on probability, timing (e.g., earlier/later), and intensity of hypoglycemia. It is up to the user which aspect to focus on but all provide intervention targets that can aid in preventing, delaying, or reducing the intensity of hypoglycemia. Not having a single unique measure does not impede understanding of token causes or ability to use them to intervene. Both scientists and lay people take in multiple dimensions of information, which they prioritize and integrate to act on.

In both science and daily life, there are many factors that contributed in some way to each token event. Philosophical theories of causality have long struggled with the idea of separating causes and background or enabling conditions (Hart & Honoré, 1985). We are not claiming that causal contribution is the only way of selecting, however people reliably and successfully prioritize some kinds of causes over others. If we were unable to distinguish degrees of causality for factors contributing to low blood glucose, we would have to ensure many factors are all present to prevent it. On the other hand, if we know that extended moderate exercise caused an instance of low blood glucose, we have now learned a useful way of preventing future instances.

Because most sciences are mired in complexity, where causes are not scarce and causal selection is crucial, degrees of token causality are an important means of selection for both understanding and intervening. Crucially, none of the uses we have laid out require consensus about degrees for specific cases. Just as people can make dietary decisions without a consensus on how to measure diet quality, degrees of token causality enable understanding and action without needing to be unique. We propose that degrees reflect an important tool in science and daily life, where causes are plentiful but cognitive resources are not.

References

Baron, Sam, Helen Beebee, David Braddon-Mitchell, Antony Eagle, and Kristie Miller. 2025. "How much did each of the authors of this paper causally contribute to its writing?" *Philosophers' Imprint*.

Bernstein, Sara. 2017. “Causal Proportions and Moral Responsibility” in David Shoemaker (ed.) *Oxford Studies in Agency and Responsibility*, Volume 4; 165-182.

Braham, Matthew, and Martin van Hees. 2009. "Degrees of causation." *Erkenntnis* 71, no. 3: 323-344.

Chockler, H., & Halpern, J. Y. 2004. “Responsibility and blame: A structural-model approach”. *Journal of Artificial Intelligence Research*, 22: 93-115.

Demirtas, Huzeyfe. 2022. “Causation comes in degrees”. *Synthese* 200:0, pp 1-17.

Foresight. 2007. Tackling obesities: future choices–Project report. Government Office for Science.

Frongillo, Edward, Jennifer Coates, Eric O. Verger, Mathilde Savy, Yves Martin-Prével, Lynnette Neufeld, Bridget Holmes et al. 2023. *Healthy diet metrics: a suitability assessment of indicators for global and national monitoring purposes*. World Health Organization.

Frühstückl, Robert. 2026. “Interactive Causes and Causal Selection.” *Philosophy of Science*, January 13, 2026, 1–32. https://doi.org/10.1017/psa.2025.10189.

Gerstenberg, Tobias, Noah D. Goodman, David A. Lagnado, and Joshua B. Tenenbaum. 2021. "A counterfactual simulation model of causal judgments for physical events." *Psychological Review* 128, no. 5: 936.

Gerstenberg, Tobias, David A. Lagnado, and Ro’I. Zultan. 2023. "Making a positive difference: Criticality in groups." *Cognition* 238: 105499.

Goh, Joshua. 2023. "It's Difficult to Explain Away the Appearance That Causation Comes in Degrees: A Reply to Sartorio." *Journal of Applied Philosophy* 40, no. 1: 109-122.

Grinfeld, Guy, David Lagnado, Tobias Gerstenberg, James F. Woodward, and Marius Usher. 2020. "Causal responsibility and robust causation." *Frontiers in Psychology* 11: 1069.

Hart, Herbert Lionel Adolphus, and Tony Honoré. 1985. *Causation in the Law*. Oxford University Press, Oxford.

Icard, Thomas F., Jonathan F. Kominsky, and Joshua Knobe. 2017. "Normality and actual causal strength." *Cognition* 161: 80-93.

Kirfel, Lara, and David Lagnado. 2021. "Causal judgments about atypical actions are influenced by agents' epistemic states." *Cognition* 212: 104721.

Kleinberg, Samantha. 2013. *Causality, probability, and time*. Cambridge University Press.

Kleinberg, Samantha, Eren Alay, and Jessecae K. Marsh. 2022. "Absence makes the trust in causal models grow stronger." In *Proceedings of the Annual Meeting of the Cognitive Science Society*, vol. 44, no. 44.

Kleinberg, Samantha, and Jessecae K. Marsh. 2023. "Less is more: information needs, information wants, and what makes causal models useful." *Cognitive Research: Principles and Implications* 8, no. 1: 57.

Kominsky, Jonathan F., Jonathan Phillips, Tobias Gerstenberg, David Lagnado, and Joshua Knobe. 2015. "Causal superseding." *Cognition* 137: 196-209.

Lagnado, David A., and Shelley Channon. 2008. "Judgments of cause and blame: The effects of intentionality and foreseeability." *Cognition* 108, no. 3: 754-770.

Lagnado, David A., Tobias Gerstenberg, and Ro'I. Zultan. 2013. "Causal responsibility and counterfactuals." *Cognitive Science* 37, no. 6: 1036-1073.

Malle, Bertram F., Steve Guglielmo, and Andrew E. Monroe. 2014. "A theory of blame." *Psychological inquiry* 25, no. 2: 147-186.

McGlashan, Jaimie, Josh Hayward, Andrew Brown, Brynle Owen, L. Millar, Michael Johnstone, Douglas Creighton, and Steven Allender. 2018. "Comparing complex perspectives on obesity drivers: action‐driven communities and evidence‐oriented experts." *Obesity Science & Practice* 4, no. 6: 575-581.

Merck, Christopher, and Samantha Kleinberg. 2016. "Causal explanation under indeterminism: A sampling approach." In *Proceedings of the AAAI Conference on Artificial Intelligence*.

Mumford, Stephen. 2013. “Causes for Laws”. Helen Beebee, Stephen Mumford, Dowe Phil & Michael S Moore. Legal Responsibility and Scalar Causation, *Jurisprudence*, 4:1, 102-137,

Northcott, Robert. 2005. “Pearson’s Wrong Turning: Against Statistical Measures of Causal Efficacy.” *Philosophy of Science*, Vol. 72, No. 5: 900-912.

O'Neill, Kevin, Paul Henne, Paul Bello, John Pearson, and Felipe De Brigard. 2022. "Confidence and gradation in causal judgment." *Cognition* 223: 105036.

Quillien, Tadeg, and Christopher G. Lucas. 2024. "Counterfactuals and the logic of causal selection." *Psychological Review* 131, no. 5: 1208.

Ross, Lauren N. 2018. “Causal Selection and the Pathway Concept.” *Philosophy of Science* (Chicago) 85, no. 4: 551–72. https://doi.org/10.1086/699022.

Sartorio, Carolina. 2020. "More of a Cause?" *Journal of Applied Philosophy* 37, no. 3: 346-363.

Spellman, Barbara A. 1997. "Crediting causality." *Journal of Experimental Psychology: General* 126, no. 4: 323-348.

Shaver, Kelly G. 1985. "Attributions of responsibility." In *The Attribution of Blame: Causality, Responsibility, and Blameworthiness*, pp. 87-113. New York, NY: Springer New York.

Sloman, Steven A., and York Hagmayer. 2006. "The causal psycho-logic of choice." *Trends in cognitive sciences* 10, no. 9: 407-412.

Sprenger, Jan. 2018. "Foundations of a probabilistic theory of causal strength." *Philosophical Review* 127, no. 3: 371-398.

Vasilyeva, N., Blanchard, T., & Lombrozo, T. 2018. "Stable causal relationships are better causal relationships". *Cognitive Science*, 42, no. 4: 1265-1296.

Waldmann, Michael R. 2017. "Causal reasoning: An introduction." The Oxford Handbook of Causal Reasoning: 1-9.

Woodward, James. 2010. "Causation in biology: stability, specificity, and the choice of levels of explanation." Biology & Philosophy 25, no. 3: 287-318.

Zeevi, David, Tal Korem, Niv Zmora, David Israeli, Daphna Rothschild, Adina Weinberger, Orly Ben-Yacov et al. 2015. "Personalized nutrition by prediction of glycemic responses." *Cell* 163, no. 5: 1079-1094.

Zheng, Min, and Samantha Kleinberg. 2017. "A Method for Automating Token Causal Explanation and Discovery." In *FLAIRS*, pp. 176-181.

Zultan, Ro'I., Tobias Gerstenberg, and David A. Lagnado. 2012. "Finding fault: Causality and counterfactuals in group attributions." *Cognition* 125, no. 3: 429-440.