

UC Irvine

UC Irvine Previously Published Works

Title

Industrial Distraction

Permalink

<https://escholarship.org/uc/item/2zb028vw>

Authors

O'Connor, Cailin

Freeborn, David Peter Wallis

Publication Date

2025

DOI

10.1017/psa.2025.1

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Industrial Distraction

Cailin O'Connor and David Peter Wallis Freeborn

¹ *Department of Logic and Philosophy of Science, University of California,
Irvine*

² *Department of Philosophy, Northeastern University, London*

Abstract: There are myriad techniques industry actors use to shape the public understanding of science. While a naive view might assume these techniques typically involve fraud or outright deception, the truth is more nuanced. This paper analyzes industrial distraction, a common technique where industry actors fund and share research that is accurate, often high quality, but nonetheless misleading on important matters of fact. This involves reshaping causal understanding of phenomena with distracting information. Using case studies and causal models, we illustrate how this impacts belief and decision making even for rational learners, informing science policy and debates about misleading content.

Acknowledgements

Thanks to Ben Genta, Chris Torsell, Tori Cotton, Matthew Coates, Rebecca Korf, and Jim Weatherall for comments on this manuscript. Thanks to participants in the SKAT workshop at Columbia University for comments and feedback. Thanks to commentary from attendees at the PSA 2024 meeting in New Orleans, and to anonymous referees.

Funding Information and Declarations

The authors have nothing to declare. No funding sources were used in the preparation of this work.

Abstract

There are myriad techniques industry actors use to shape the public understanding of science. While a naive view might assume these techniques typically involve fraud or outright deception, the truth is more nuanced. This paper analyzes industrial distraction, a common technique where industry actors fund and share research that is accurate, often high quality, but nonetheless misleading on important matters of fact. This involves reshaping causal understanding of phenomena with distracting information. Using case studies and causal models, we illustrate how this impacts belief and decision making even for rational learners, informing science policy and debates about misleading content.

1 Introduction

Over the past few decades the Coca Cola company has engaged in an extensive campaign to fund and share research on the benefits of exercise to health, and especially its impacts on weight and diet-related diseases (Serodio et al., 2020; Wood et al., 2020; Nestle, 2015; O’Connor, 2015; Greenhalgh, 2024; Carpenter, 2025). In response, scientists have raised the alarm about the potential for negative health effects from this campaign. For example, in 2017 the Union of Concerned Scientists published a report documenting Coca Cola’s influence on the sciences of sugar, obesity, and exercise (UCS, 2017). Notably, though, these scientists made no accusations of fraud, questionable research practices, or lying. Neither did they suggest that the research funded by Coca Cola was itself bad or inaccurate. What, we might ask, is wrong with a company giving money to otherwise independent scientists to do research on a topic of interest to public health?

The worry is that even good science on exercise can shift blame for public health problems away from Coca Cola products, and towards sedentary lifestyles. This type of technique—funding and sharing accurate, often high quality, often independent, research with the goal of distraction—is one that has been used extensively in the history of industry influence on science (Proctor, 1995; Oreskes and Conway, 2011; Proctor, 2012). In this paper we analyze this sort of technique, which we call *industrial distraction*.¹ We use both case studies and causal models to show how and why industrial distraction works, and to identify a few variations of the technique.

At its heart, industrial distraction involves changing how targets understand some causal system in the world. Typically it shifts public understanding towards some distracting potential cause of a public harm, and away from a known industrial cause of the same harm. A second variation uses inaccurate information to introduce distracting mitigants of industrial harms. And a last variant shifts public beliefs about downstream effects of policies to focus on distracting harms they may cause.

¹Elsewhere Robert Proctor has referred to this technique as “distraction science”, but we wish to emphasize the role industry plays in it (Proctor, 1995, 2012; Kourany and Carrier, 2020).

One reason it is important to understand and analyze industrial distraction is that it does not fit with a naive understanding of how industry influences public opinion about science. A typical picture focuses on the production of fraudulent or influenced research, and/or the sharing of inaccurate, false, or deceptive scientific claims. While this does happen, it is far from the only method of industry influence (Lesser et al., 2007; Oreskes and Conway, 2011; Bes-Rastrollo et al., 2013; Proctor, 2012, 1995; O’Connor and Weatherall, 2019b). Industrial distraction does not work this way. Nonetheless, as our models will illustrate, it can shift public belief in harmful ways, and, as a result, shift policy decisions in harmful ways. As our models also show, this sort of harm need not depend on human fallibility—even fully rational learners and decision makers can err in the presence of industrial distraction.

Recent research has highlighted a suite of industry techniques that avoid moral and legal censure by technically “playing by the rules” (Oreskes and Conway, 2011; Holman, 2015; Holman and Bruner, 2017; Weatherall et al., 2020; Greenhalgh, 2024). In order to properly regulate industry influence, then, policy makers must be able to recognize how industrial actors can skirt current norms and regulations and nonetheless influence policy outcomes. Industrial distraction is one more technique in this vein. We argue that, given the presence of these techniques, policies are needed to more stringently separate industry from science, and to regulate how industry communicates with the public about science.

This paper will also be relevant to both philosophical and policy debates about how to understand misinformation, disinformation, and misleading content. While this kind of content is often defined as “false” or “inaccurate”, it is increasingly recognized that true and accurate content can mislead, industrial distraction arguably providing one example (Fallis, 2015; Wardle and Derakhshan, 2017). The ubiquity of accurate but misleading content online leads to thorny questions about how best to regulate both social and traditional media. Relatedly our analysis will be relevant to philosophical debates about how to characterize and identify illegitimate scientific dissent.

On one last note, there has been a great deal of excellent historical investigation into the details of industrial influence on public health.² Many of these investigations carefully outline various details of industrial strategy. What philosophers of science and social epistemologists have added to this research are systemic analyses of the epistemic impacts of industrial propaganda. These are formal and theoretical understandings of just how and why propaganda of various sorts can impact belief. This paper follows in this vein.

The paper will proceed as follows. Section 2 will introduce Bayesian causal models, giving the background information necessary to model various types of industrial distraction. Section 3 will discuss cases where industry shifts beliefs about causes of an industrial harm, and develop causal models that illustrate how this sort of industrial distraction works. The next section, 4, looks at

²See, for example, Proctor (1995); Brownell and Warner (2009); Oreskes and Conway (2011); Proctor (2012); Greenhalgh (2024); Carpenter (2025).

cases where industry introduces spurious mitigants of industry harms. And 5 analyzes cases where industry shifts understandings of the effects of policy. As will become clear, these three varieties of industrial distraction all work differently, though they all can be effective. In section 6 we discuss what this means for policy regulation of industry influence on science and public belief, and for thinking about misleading content more generally.

2 Causal Models

Causal models provide a useful framework for analyzing the various techniques of industrial distraction both because they illuminate the logic of these strategies, and because they make clear how even rational learners are misled by them. In fact, recent work in philosophy and the social sciences has demonstrated how this sort of model is useful to understanding a suite of phenomena related to false belief, propaganda, and polarization (Freeborn, 2023, 2024; Eliaz et al., 2022; Jern et al., 2014; Eliaz and Spiegler, 2024; Spiegler, 2020).

Causal models offer formal representations of systems with multiple stochastic variables and causal relationships between them. For example, when studying obesity in humans these variables could represent the events that some population 1) drinks sugary drinks, 2) has high rates of sedentary lifestyles, and 3) exhibits high levels of obesity. Causal models allow us to reason about cause-and-effect relationships between these variables, to predict how changes in one variable might influence others, and to estimate the effects of specific interventions. In addition, as we will see, they allow us to represent how an ideal learner might update their beliefs about such a causal system in light of new evidence.

2.1 Causal Bayesian Networks

Bayesian networks are one popular type of causal model, which allow for consistent probabilistic reasoning (Pearl, 2009; Spirtes et al., 2000). A Bayesian network represents a probabilistic system using a directed acyclic graph (DAG). These graphs consist of nodes and directed edges (arrows) between them. (They are “acyclic” because these arrows never form closed loops between the nodes, as will become clear shortly.) We can fully specify a Bayesian network by,

1. A set of n random variables $\mathbf{X} = \{X_1, \dots, X_n\}$. For example, these variables could be obesity, a sedentary lifestyle, and intake of sugar. Each variable is associated with a node on the graph.
2. A set of directed edges, $\mathbf{E}$, between nodes. Each edge represents a probabilistic relationship between the variables. For example, if sedentary lifestyles increase the probability of obesity, then there could be an edge pointing from sedentary lifestyles to obesity. If there is a directed edge from node X_i to node X_j , we call X_i a “parent” of X_j , and X_j a “child” of X_i .

3. Conditional probability distributions $P(X_i \mid \text{Pa}(X_i))$ for each random variable X_i , where $\text{Pa}(X_i)$ denotes the parents of X_i .³

These probability distributions determine how nodes are probabilistically related to each other. For example, they might specify a strong link between sugar intake and obesity, or else a weak one. Together, these conditional distributions must be probabilistically consistent with each other.⁴ Note that in the following we will label the two possible values for any binary variables, **true** or **false**. For instance, $P(X = \text{true} \mid Y = \text{true})$ will give the probability that variable X is true conditional on Y being true. Occasionally, it will be convenient to omit the values of variables, for instance when discussing independencies. For example, $P(X) = P(X \mid Y)$ means that variable X is independent of variable Y .

When we learn some new piece of information, E , the probabilities in the network can remain consistent by updating through Bayesian conditionalization, $P_{\text{new}}(X_i) = P(X_i \mid E = \text{true})$. As such, Bayesian networks can provide a model of rational learning. The nodes represent events that might hold, the edges their probabilistic relationships, and the constraints of the model specify how a rational agent should update their beliefs about all these events.

For example, suppose that high pollen count (P) and colds (C) are two independent causes of a bout of sneezing (S). Then, we can represent this situation with the Bayesian network in figure 1. Both variables increase the probability that one experiences a bout of sneezing, according to the conditional probabilities given in the corresponding table.⁵ Then, learning either that the pollen count is high or that I have caught a cold should increase my credence that I will have a bout of sneezing today. Alternatively, experiencing a bout of sneezing should increase my credences that the pollen count is high and that I have a cold.

To give an example, according to this Bayesian network, if I start with a prior belief of 0.5 that the pollen count is high, and a prior belief of 0.5 that I have a cold, then my prior degree of belief that I will experience a bout of sneezing should be 0.65. Suppose that I do start experience such a bout of sneezing. Then I can use this observation, plus Bayesian inference, to update my degree of belief that I have a cold, $P_{\text{new}}(C = \text{true}) \approx 0.65$.

We are often interested in knowing which variables are statistically dependent or independent of others. We say that variables X and Y are independent of each other, conditional on a set of variables, $\mathbf{Z}$, if $P(X \mid \mathbf{Z}) = P(X \mid Y, \mathbf{Z})$, or equivalently $P(Y \mid \mathbf{Z}) = P(Y \mid X, \mathbf{Z})$.⁶ For example, in the graph in figure 1 the

³We assume the *Causal Markov assumption*: each variable X_i is conditionally independent of its non-descendants given its parents, $\text{Pa}(X_i)$.

⁴They will form a factorized representation of a joint probability distribution, $P(\mathbf{X}) = \prod_{i=1}^n P(X_i \mid \text{Pa}(X_i))$.

⁵The table should be read as follows. If P is false and C is false, the probability of sneezing given both of these facts, $P(S = \text{true} \mid P, C)$, is .1, and so on.

⁶In a Bayesian network, we can generally identify a property of the graph structure, d-separation, to determine whether two variables must be statistically independent. If two variables are d-separated relative to a set of variables $\mathbf{Z}$ in a directed acyclic graph, then they




Figure 1: A causal graph and associated conditional probability table representing two possible causes, high pollen count (P) or a cold (C), of sneezing (S). We assume that these two causes are independent.

two possible causes, high pollen count P and a cold C , are independent of each other. Although they are connected by the path $P - S - C$, it is blocked by a “collider” at S . Roughly, we can understand this as saying that whilst both P and C might inform us about S , they do not inform us about each other. However, P and C are not independent conditional on S .⁷ If we assume that a sneezing bout is taking place, then each of the other variables can inform us about the other. For instance, if the pollen count is high, that might explain the sneezing, so it is less likely I have a cold. Or if I know I have a cold, this can already explain the sneezing, so it is less likely that the pollen count is high. This sort of conditional dependence will be relevant to cases we discuss below.

3 Distracting Causes

As noted industrial distraction involves attempts to reshape the way targets understand causal relations in the world, and thus avoid undesirable outcomes for industry. We divide these attempts into several sorts—those aimed at shifting beliefs about causes of some harmful phenomenon, those aimed at (falsely) shifting beliefs about factors mitigating harmful effects, and those aimed at shifting

must be statistically independent conditional on $\mathbf{Z}$ in all possible probability distributions that the graph can represent. The reverse does not hold. Two d-separated variables in a joint probability distribution might still be numerically independent given some other variables. See Pearl (2009) for further details.

We say that X and Y are d-separated by $\mathbf{Z}$ if there are no unblocked undirected paths through $\mathbf{G}$ that connect them. An undirected path between two nodes X_1 and X_n is a sequence of nodes $(X_1, X_2, \dots, X_n)$ such that for each pair of consecutive nodes X_i and X_{i+1} , there is an edge between them in either direction. An undirected path is blocked by a set of nodes, $\mathbf{Z}$, if the path contains a collider that is not in $\mathbf{Z}$ and has no descendants in $\mathbf{Z}$, or if the path contains a non-collider that is in $\mathbf{Z}$. A node X_i on an undirected path $(X = X_1, X_2, \dots, X_n = Y)$ is a collider if it has two incoming edges from its neighbors on the path, i.e., both $X_{i-1} \rightarrow X_i$ and $X_{i+1} \rightarrow X_i$, $X_{i-1} \rightarrow X_i \leftarrow X_{i+1}$. If variables X and Y are d-separated must be independent.

⁷They are d-connected given S , as the collider is now found in the set of dependent variables on which we are conditioning.

beliefs about effects of policy interventions.

The Coca Cola case described above is an excellent example of the first sort of industrial distraction. We have an undesirable phenomenon from the point of view of public health—obesity and obesity-related disease.⁸ We have clear scientific evidence connecting the consumption of sugar-sweetened beverages, such as sodas, to weight gain, diabetes, and heart disease (Ludwig et al., 2001; Malik et al., 2006; Schulze et al., 2004; Malik et al., 2010; Yang et al., 2014). We have increasing public attention to this connection, and increasing action by policy makers to regulate soda (Greenhalgh, 2024; Carpenter, 2025).

These events create pressure on industries producing soda to disrupt public belief about its health effects, and prevent policy regulation. However, in a case like this, enough scientific evidence has accumulated to make it difficult for Coca Cola to outright deny the causal connection between soda consumption and obesity. One way forward is to distract the public and policy makers from this connection by focusing on some other causal factor that contributes to obesity—in this case, sedentary lifestyle. By strengthening beliefs about the connection between a distraction (D) and an undesirable outcome (U), propagandists decrease beliefs that industry (I) is a relevant or important cause of U .

There are several ways that Coca Cola emphasized this distracting causal pathway. First, they funded research into exercise, for example through the Global Energy Balance Network—a Coca Cola funded research group promoting the idea that the best way to lose weight is through exercise. Second they widely shared research on exercise and obesity, whether or not they had funded that research. The variations in how they fund, and promote, this sort of research are many and complicated. They go beyond the scope of this paper, but interested readers can learn more in Greenhalgh (2024) or Carpenter (2025).

It is important to recognize that industrial distraction as used by Coca Cola is very far from an isolated case. Another notable case involved the tobacco industry, which spent enormous resources sowing doubt about the connection between tobacco and diseases like lung cancer and emphysema. (As Oreskes and Conway (2011) convincingly show, tobacco pioneered many industry techniques for influencing scientific belief, so this is, in fact, an early and important example of industrial distraction.) Notably, they promoted research about alternative causes of lung disease, including asbestos exposure, air pollution, coal smoke, and even early marriage (O'Connor and Weatherall, 2019b).⁹ Later, when fighting consensus on the dangers of second-hand smoke, tobacco publicized alternative causes for lung disease in spouses of smokers such as, “microorganisms, allergens, pesticides, herbicides, household chemicals, insect and rodent products, nitrogen and sulfur dioxides, ozone, formaldehyde, respirable

⁸We, the authors, are not making or supporting any claims about the desirability of fatness, but are describing here the way it has been understood by policy makers and the general public.

⁹For example, the Tobacco Industry Research Committee—a propaganda body funded by major US tobacco firms—publicized the work of Willhelm Hueper, a cancer researcher who appeared regularly as an expert witness arguing that lung illnesses of patients were caused by asbestos rather than smoking (Oreskes and Conway, 2011).

dusts, radon”.¹⁰

The sugar industry has been criticized, similarly, for funding research on the link between dietary fat and heart health in the mid-20th century (Kearns et al., 2016). Ironically, at the same time various industry groups connected to fatty foods, like the British Egg Marketing Board and the National Dairy Council—were funding research into the link between sugar and heart disease, and thus also attempting industrial distraction (Johns and Oppenheimer, 2018).

Industrial distraction sometimes involves poor science, but not necessarily so. For example, Johns and Oppenheimer (2018) argue that in the sugar case, the industry funded mainstream researchers doing high quality work. They argue there is little evidence that the nutrition research itself was directly impacted by industry funding. Notably, there is often no need, in industrial distraction, to promote low quality work. There are typically multiple, real causes of some undesirable outcome, and revealing these links constitutes important research. It is just when this research is funded and communicated cynically as a distraction strategy that it tends to harm public belief.

With these cases in hand, we now turn to causal models to illuminate how this sort of technique works generally, and to illustrate how learners updating on accurate and relevant data can be misled by it.

3.1 Distracting Causes Model

As noted, this version of industrial distraction involves promoting an alternative cause (D) to distract from the industry’s own causal role (I) in an undesirable outcome (U). Let us use the Coca Cola case to ground our analysis. If we regard the two possible causes (e.g. a sedentary lifestyle and intake of sugary sodas) as statistically independent, one way to represent this type of distraction is with a simple causal network like the one shown in figure 2 (note that this has the same structure as the sneezing example in figure 1).




Figure 2: A causal graph in which the effect U has two independent possible causes, an industrial product I and a distracting cause D .

Suppose that we encounter evidence that the distraction D is a cause of U .

¹⁰See the pamphlet “Environmental Tobacco Smoke and Health”, available at the UCSF’s Truth Tobacco Industry archive (Env, 1986).

How should that affect our beliefs about the industrial cause, I ? Well, although the variables I and D are marginally independent (i.e. $P(I) = P(I | D)$), they are not conditionally independent given U (i.e. $P(I | U) \neq P(I | D, U)$).¹¹ In many instances we might already know that the undesirable effect U is taking place. Or alternatively, we might acquire evidence about the causes that does not alter our beliefs about whether the effect is taking place. In either case, if D can account for some or all of the effect of U , then I does not need to account for as much. Thus we should often rationally lower our degree of belief in I being a cause of U .

There are at least two different ways we could model this effect using the Bayesian network structure. In the first approach, we use the conditional probabilities to represent changes in beliefs about the causal effect of one variable on another. In other words we change the strength of the “edges” between nodes, i.e., the entries in our conditional probability tables. In the second approach, we assume a change in our marginal probabilities (the “node” itself), whilst keeping the conditional probabilities fixed. Mathematically, we can achieve the same effect either way. However, each modeling choice will require slightly different interpretations of each of the variables. Different choices will be more natural in different cases. We explore both options in turn.

3.1.1 Updating only the Conditional Probabilities

Suppose we use the Bayesian network and conditional probabilities in figure 2. We use the variables to represent these events,

- I : The population has a high intake of sugary drinks
- D : The population has high rates of sedentary lifestyles
- U : There is an increase in obesity levels.

Suppose we begin with the following prior probabilities,

$$\begin{aligned} P(I = \text{true}) &= 0.8 \\ P(D = \text{true}) &= 0.8. \end{aligned}$$

Then from the conditional probability tables, it follows that $P(U = \text{true}) \approx 0.836$. Now suppose that we learn new information that increases our credence that sedentary lifestyles cause obesity,

$$\begin{aligned} P(U = \text{true} \mid D = \text{true}, I = \text{false}) &= 0.9 \\ P(U = \text{true} \mid D = \text{true}, I = \text{true}) &= 0.95, \end{aligned}$$

¹¹Although I and D are d-separated (i.e., they are independent), they are not d-separated given the outcome U . In other words, I and D become d-connected when conditioned on U .

but which does not alter our beliefs in the marginal probabilities ($P(I)$, $P(D)$ and $P(U)$) regarding whether obesity, rates of sugary drinks, and sedentary lifestyles are high. Furthermore, we assume that it does not alter the probability that obesity arises if neither the intake of sugary drinks nor rates of sedentary lifestyles are high, $P(U = \text{true} \mid I = \text{false}, D = \text{false})$.¹² Then, in order to keep the probabilities consistent, we are forced to revise our beliefs about whether sugary drinks cause obesity to arise (if sedentary lifestyles are not at high rates). Now $P(U = \text{true} \mid I = \text{true}, D = \text{false}) = 0.5$, which is substantially lower than our prior belief.

Note that this is a rational case of consistently updating beliefs in the light of evidence. Thus, if we become more persuaded that the distracting cause (D) can explain some or all of the undesirable outcome (U), we have less reason to ascribe some of that effect to the industrial product (I). The result is that we rationally decrease our degree of belief that the industrial product, I causes the undesirable effect, U . This is sometimes known as the *explaining away* effect in Bayesian epistemology (Kim and Pearl, 1983; Wellman and Henrion, 1993).

3.1.2 Updating only the Marginal Probabilities

In a causal modeling framework, it is often more mathematically natural to update the marginal probabilities, whilst leaving conditional probabilities fixed. This provides an alternative way to model the distracting causes scenario; however, it necessitates a different, less straightforward, interpretation of the variable—we include causal effects *within* the variables.

For example, we might use the variables to represent the following propositions,

- I : High sugary drink intake leads to obesity
- D : High rates of sedentary lifestyles lead to obesity
- U : There is an increase in obesity levels.

Let us suppose that at first, we treat the two causes as independent, and we believe that sugar-sweetened beverages are the most likely cause, whilst sedentary lifestyles are less likely, adopting these prior probabilities:

$$P(I = \text{true}) = 0.6$$

$$P(D = \text{true}) = 0.4.$$

If we are sure that there really is an increase in obesity, i.e. $P(U = \text{true}) = 1$, then by Bayesian conditionalization, we should increase our degree of belief in each of these two possible causes: $P_{\text{new}}(I = \text{true}) = P(I = \text{true} \mid U = \text{true}) \approx$

¹²Note that, without making this many assumptions about which beliefs the evidence does or does not affect, the problem would be unconstrained. It is also important to note again that this model assumes statistical independence between the industrial cause (I) and the distracting cause (D). In reality, these causes might be correlated, which would require a more complex model.

0.77 and $P_{\text{new}}(D = \text{true}) = P(D = \text{true} \mid U = \text{true}) \approx 0.52$. However, these conditional probabilities are not independent: if sedentary lifestyles can explain some of the known effect, U , then sugary drinks need to explain less. If we then learn that the distracting cause is true, i.e., that $P(D = \text{true}) = 1$, then we should decrease our degree of belief in I : $P_{\text{new}}(I = \text{true}) = P(I = \text{true} \mid U = \text{true}, D = \text{true}) \approx 0.63$.

Once again, we can think of this as a case of the explaining away effect. We can express this with the inequality $P(I = \text{true} \mid U = \text{true}, D = \text{true}) < P(I = \text{true} \mid U = \text{true})$. This effect will arise in the simple model as long as the two possible causes, I and D , are probabilistically independent, are the only two possible causes, and both always positively increase the probability of U being true.¹³

3.2 Accurate Sharing and Inaccurate Beliefs?

Before continuing to the next version of industrial distraction, we will take a moment to address a possible worry here. One might think that if industry is actually sharing accurate scientific data, recipients will develop accurate causal pictures of the world. In other words, although they might strengthen beliefs in a distracting cause, they will only do so in an accurate way, and thus are not harmed.

There are a few things to note here. First, as we will emphasize later, industry is often supporting and spreading real scientific information but in a cherry-picked way. Targets are receiving too much information about distracting causes, and not enough information about relevant industry causes. Even rational learners can develop inaccurate pictures of the world on the basis of good data that is cherry picked or curated (Mohseni et al., 2022).

Second, industry is often picking distracting causes to highlight that are not currently a public focus. In other words, they cynically select distracting causes where accurate information can decrease beliefs in the strength of industry causes. It is in this sort of context that the sharing of such distracting information functions as a type of misleading content (even if it improves beliefs about a distracting cause). It misleads by shaping beliefs in such a way as to purposefully prevent effective policy.¹⁴

Third, although we are emphasizing the role that accurate scientific information can play in industrial distraction, there is no reason that inaccurate, false, hyperbolic, or fraudulent information cannot play the same role. Furthermore, it is often the case that media coverage of science overstates the strength of results, meaning that the public may get an inaccurate picture of the strength of a distracting cause.

¹³I.e., if the condition $\frac{P(U=\text{true}|I=\text{true},D=\text{true})}{P(U=\text{true}|I=\text{true},D=\text{false})} < \frac{P(U=\text{true}|I=\text{false},D=\text{true})}{P(U=\text{true}|I=\text{false},D=\text{false})}$ holds (Wellman and Henrion, 1993).

¹⁴There are formal accounts in formal epistemology and philosophy of science of what accurate beliefs consist in, and what counts as deception. Here we do not ground claims about what is “misleading” using any such account. Instead we will argue that whatever notion of “misleading” we develop should be broad enough to include cases like this one.

4 Distracting Mitigations

The next sort of case occurs when industry promotes distracting mitigations to some industrial harm. To give some examples, the sugar industry promoted and publicized research into enzymes that would disrupt dental plaque, and into a tooth decay vaccine (Kearns et al., 2015). The plastic industry widely shared false claims about the effectiveness of plastic recycling (Singla, 2022; Allen et al., 2024). Tobacco invented “healthier cigarettes”, like those with filters (Cummings et al., 2007).

This kind of technique again reworks the public’s causal picture. Instead of thinking that industrial product (I) is necessarily connected to undesirable effect (U), the public now thinks there is some mitigating factor (M) that interrupts that causal connection. Unlike the last technique, though, this one typically must involve sharing spurious or false claims. If some mitigating factor actually could prevent industrial harms, then no industrial propaganda would be needed. Instead, because no such mitigating factors exist, industry must mislead observers as to their abilities to prevent harm. (Filters do not prevent harms from smoking, plastic recycling is mostly a myth, and there is no tooth decay vaccine.)

There are some similar cases where industry over-emphasizes the potential mitigating impacts of future technologies. In these cases, it may turn out that these technologies actually can disrupt the link between an industrial product and harms. For example, it is possible that carbon capture technologies might someday greatly mitigate the harms of fossil fuel use. But even in these cases industrial communication about these benefits should be understood as a harmful distraction technique. The benefits of these technologies are not yet clear, and they are being shared cynically to shape policy with little regard for public health.

4.1 Distracting Mitigations Model

To model distracting mitigation we can use a network with the same structure as in section 3.1. Here, the undesirable effect (U) may be causally influenced by two variables, one representing the presence of an industrial product (I), the other representing the presence of a mitigating factor (M). For example, we could interpret the variables as follows,

- I : High sugary drink intake leads to tooth decay
- M : There is an effective tooth decay vaccine
- U : There is an increase in tooth decay levels.

A Bayesian network representation and possible conditional probability table is shown in figure 3. The main difference here is in the conditional probabilities.

Without the mitigating factor in play, the presence of the industrial product (e.g. sugar) increases the probability that the undesirable effect (tooth decay)




Figure 3: A causal graph in which the effect U is influenced by two causal factors, the industrial product I and a mitigating factor M . The conditional probability table for U shows that M reduces the causal effect of I on U .

will arise. However, if the mitigating factor is in play, the effects of the industrial product on the undesirable effect are greatly reduced. For instance, suppose we hold the following prior probabilities,

$$P(I = \text{true}) = 0.6 \quad (1)$$

$$P(M = \text{true}) = 0.1 \quad (2)$$

If, say, we learn that the undesirable effect is taking place, then we should rationally update our credence in the industrial product being the cause, $P_{\text{new}}(I = \text{true}) \approx P(I = \text{true} \mid U = \text{true}) = 0.93$. After all, with this setup, the industrial product is our only likely (and therefore best) explanation of the undesirable effect. As such, the existence of the undesirable effect is itself good evidence that the industrial product is causing it.

However, suppose that we also come to believe that the mitigating variable is true (i.e. the mitigating factor is present). Now, the industrial product is a much weaker explanation. In this case, we should rationally alter our credences, $P_{\text{new}}(I = \text{true}) = P(I = \text{true} \mid U = \text{true}, M = \text{true}) \approx 0.75$. The industrial product may still be a cause, in spite of the mitigating factor, but it is a less convincing one. (Alternatively, in this case, we might be unsure about whether U will occur in the future as a result of I . If we learn that M is true we decrease our belief in U .)

This effect is highly analogous to the explaining away effect discussed in section 3.1. Once again, the mitigating factor and the industrial cause are no longer statistically independent once the undesirable effect is known. However, in this case, the mitigating factor serves to reduce some of the explanatory strength of the industrial product, rather than serving as a separate explanation in itself.

4.2 Distracting Causes and Mitigations

The effect of the mitigating factor was quite weak in this example, because we had no alternative good explanations of the undesirable effect. Notice, though, that in some of the cases above industry introduced both distracting causes and distracting mitigants. The Tobacco industry emphasized the harms of asbestos, and also the mitigating hope of filters, with respect to lung cancer, for example.

Assume that a distracting explanation D and a mitigating factor M are both in place. Now the undesirable effect U is influenced by three causal factors, the presence of the industrial product, I , the mitigating factor, M , and the distracting cause, D . Then the false mitigating factor might cause us to further rationally reduce our degree of belief that the industrial product I is responsible for the effect, analogous to the shifting causes model in 3.1. We can represent this in a hybrid model, shown in figure 4.




Figure 4: A causal graph in which the effect U is influenced by three causal factors: the industrial product I , a false mitigating factor M , and a distracting cause D . The conditional probability table for U shows that M reduces the causal effect of I on U .

For example, suppose that we adopt the following initial probabilities.

$$\begin{aligned}
P(I = \text{true}) &= 0.6 \\
P(M = \text{true}) &= 0.1 \\
P(D = \text{true}) &= 0.4
\end{aligned}$$

These lead to a prior expectation of the undesirable effect of $P(U = \text{true}) \approx 0.63$. Suppose we learn that the undesirable effect does take place and there is a public harm to worry about, i.e. $P(U = \text{true}) = 1$. Then by Bayesian conditionalization, we should update our degrees of belief as follows,

$$\begin{aligned}
P_{\text{new}}(I = \text{true}) &= P(I = \text{true} \mid U = \text{true}) &= 0.76, \\
P_{\text{new}}(M = \text{true}) &= P(M = \text{true} \mid U = \text{true}) &= 0.066, \\
P_{\text{new}}(D = \text{true}) &= P(D = \text{true} \mid U = \text{true}) &= 0.54.
\end{aligned}$$

Now we think that both causes are more likely to be acting to produce U . However, suppose we then come to believe the mitigating variable is true (i.e. the mitigating factor is present), $P(M = \text{true}) = 1$. Then the industrial cause is less able to explain the effect of U . Consequently, we should rationally increase our degree of belief in the alternative explanation, D , as a likely cause of the undesirable effect, $P_{\text{new}}(D = \text{true}) = P(D = \text{true} \mid U = \text{true}, M = \text{true}) \approx 0.77$. Likewise, we should rationally decrease our degree of belief in the industrial product, I as the cause, $P_{\text{new}}(I = \text{true}) = P(I = \text{true} \mid U = \text{true}, M = \text{true}) \approx 0.63$. In this case the false mitigating factor works to reduce our rational credence that the industrial product causes the undesirable effect. This is again similar to the explaining away effect.

5 Distracting Effects

The last variety of industrial distraction involves influencing beliefs about distracting effects of policy interventions. Compared to the first two variants, this one is more straightforward to understand. But it, too, involves industry using accurate data to shape a target’s causal understanding of the world, to their own benefit. And it has been an important technique employed in real cases of industrial distraction. For these reasons, we analyze it here.

There are typically multiple downstream effects of policy given the complexity of many social, natural, and economic systems. When industry propagandists wish to counter policy proposals, and when they cannot plausibly deny the relevance of such proposals to mitigating the harms of their products, one solution is to emphasize negative causal outcomes instead.

Consider the recent transition from fossil fuels to wind power, intended to prevent the harms of global warming. The oil and gas industry spent decades obfuscating the link between fossil fuels and global warming, but their ability to plausibly do so is waning (Oreskes and Conway, 2011). Instead a number of prominent Republican lawmakers in the United States—backed by powerful

oil and gas interests—have blamed off-shore wind turbines for the deaths of whales (Hu, 2023). Legitimate scientists are indeed worried about impacts of these installations on cetaceans, and have produced studies of these impacts (Quintana-Rizzo et al., 2021; Thompson et al., 2010). But their worries are being shared cynically to distract from the more important benefits of wind energy. Others connected to the Republican party, and funded by oil and gas, have emphasized the impacts of wind turbines on birds, despite evidence of fossil fuel’s much more serious impacts on bird life (Katovich, 2023; Sovacool, 2013; Bateman et al., 2020). Republicans have also focused on wind power as a cause of power outages and shortages, even in cases where it is a less important cause than outages in traditional energy sources (Benshoff, 2022).

In a similar case, a 2017 report by the US Chamber of Commerce—produced with money from companies like Exxon Mobile—seriously overstated the economic impacts to the US from complying with the Paris agreement (Bernstein et al., 2017; Negin, 2020). The report was debunked, but was used by politicians like then US President Donald Trump to justify inaction on climate change (Greenberg, 2017; Biesecker and Wiseman, 2017).

In all of these cases industry, and their political allies, introduce and/or emphasize distracting downstream effects of unwanted policy. In other words, they argue that policy (P), while causing a desirable outcome (O), also causes some other harmful outcome (H). Once again, this involves reworking the causal picture policy makers have of the world, using data that may be perfectly good. Now in assessing some policy proposal, their causal picture involves a harmful outcome, as well as a desirable one.

5.1 Distracting Effects Model

This type of case is easy to understand even without a model, so we keep this section brief. Unlike the previous case, it is natural to assume that the two outcomes, O and H , are not independent: they both have a common cause, the policy or product, P . However, we assume O and H are independent, conditional on P . If we already know for certain that a policy intervention is happening, learning about one effect does not give us further information about the other effect. In that case, we can represent the situation with the causal graph model in figure 5.

This technique works by changing our overall estimates of the likelihoods of positive and negative effects of the policy. The key is that changing our beliefs about one outcome does not directly affect our beliefs about the other outcome once we know that P is happening. Suppose that we learn that the negative outcome is more likely (perhaps $P(H = \text{true} \mid P = \text{true})$ increases to 0.9) while our beliefs about the positive outcome remain unchanged (i.e., $P(O = \text{true} \mid P = \text{true})$ stays fixed). Then this should decrease how positively we feel about the policy intervention overall.¹⁵ And, importantly, our shift in beliefs about effects can impact our subsequent decision making.

¹⁵This is a case of what Kim and Pearl (1983) term “inter-causes independence”.




Figure 5: A causal graph in which the common cause, policy P , has two possible effects, a desirable outcome O and a harmful outcome H , which are independent conditional on P .

For example, we might interpret the variables as the following events,

- P : There is increased use of wind power
- O : There are reduced effects of global warming
- H : There are harmful effects for birds.

We might initially be focused on the positive effects of preventing the harms of global warming (O). Upon learning evidence that harmful effects to birds, (H), can be caused by wind farms, (i.e. $P(H = \text{true} \mid P = \text{true})$ is high), we now think that H is a more likely outcome of P , although information does not alter how likely it is that O will occur. We might thus revise our overall judgement of whether we should expand wind power.¹⁶

6 Discussion

As we have seen there are a series of related techniques where industry can use distracting information to reshape causal beliefs to their benefit. One of these variants (distracting mitigations) relied on false or misleading information.

¹⁶To make this point more clearly, we could adopt a decision theoretic framework with explicit utilities or payoffs for these outcomes. Suppose that we assign a utility of 1 to $O = \text{true}$ and -1 to $H = \text{true}$. With $P(O = \text{true} \mid P = \text{true}) = 0.8$ and $P(H = \text{true} \mid P = \text{true}) = 0.5$, the expected utility of implementing P is 0.3. If new evidence increases $P(H = \text{true} \mid P = \text{true})$ to 0.9 (while $P(O = \text{true} \mid H = \text{true})$ remains at 0.8), then expected utility from H being true equals -0.1 , making the policy much less desirable. In other words, increasing the probability that a harm occurs from some policy can shift the expected utilities from implementing that policy.

Notably, though, the two others (causes and effects) could function perfectly well with accurate or true information (not that they always do). And all three techniques are perfectly capable of impacting decision making. If, for example, we do not think Coca Cola is an important cause of public health problems, we should not work to regulate it. If filters prevent tobacco deaths, we do not need to decrease smoking.

One thing to note is that while all our models represented rational learners, real world learners may sometimes be even more vulnerable to industrial distraction. For example, humans are known to strengthen their beliefs upon repeated exposure to a claim, even when it is not reasonable to do so, and even when that claim is known to be false (Hassan and Barber, 2021; Fazio et al., 2015; Udry and Barber, 2023). In cases where industry can flood media, advertisements, and social media with some claim—say that wind farms kill birds, or that filtered cigarettes are safe—repeated exposure to these claims may have a stronger impact than our Bayesian models would predict.

One upshot of our analysis is that policy aimed at protecting public belief should not be limited to industrial propaganda that promotes scientific fraud or shares false information. Such policy misses the harms of techniques like industrial distraction. In thinking about science policy, a nuanced understanding of the many and subtle ways industry influences belief and decision making is necessary to prevent harms from this influence. This is especially true because industrial distraction is far from the only subtle influence technique used by industry.

Holman and Bruner (2017) use a model to illustrate what they call *industrial selection* where industry promotes researchers who happen to already be producing favorable research. Doing so involves taking advantage of natural variation in the background beliefs, assumptions, focus, or methodology of different scientists, and then, through funding and other amplification methods, making some subset of work more productive or more salient. Notably, many instances of industrial distraction are also instances of industrial selection. In these cases, industry is selecting researchers to fund or promote based on the fact that they are working on a causal connection favorable to industry. For example, as Serodio et al. (2020) point out, Coca Cola promoted the careers of many academic researchers already friendly to their “energy balance” message.¹⁷ Whether industrial selection uses distraction or not, though, it is another technique where industry technically plays by the rules, but can nonetheless seriously impact the course of science.

Others have emphasized the role of cherry picking in industry misinformation. This involves selecting just some biased subset of independent research to share and promote. For example, the tobacco industry widely shared studies that happened to spuriously find no link between tobacco and disease (Oreskes and Conway, 2011). Both Weatherall et al. (2020) and Lewandowsky et al. (2019) use models to show how this sort of selection can influence rational learn-

¹⁷Earlier on sugar funded independent researchers already looking at the link between fat and heart disease, while fat funded researchers already looking at sugar as a cause of heart disease (Johns and Oppenheimer, 2018; O’Connor and Weatherall, 2019a).

ers to form false beliefs favorable to some propagandist.¹⁸ As noted, industrial distraction can involve a form of cherry picking when only research relevant to a limited part of a full causal picture is shared. When engaged in industrial distraction, propagandists cynically select just some areas of research to promote, and in doing so distort the importance of causes and effects, thus distorting the beliefs of their targets. But again, whether or not cherry picking involves distracting information or straightforwardly misleading information, this sort of industrial technique works within the rules of science and policy to impact decision making in ways that harm public health.

Given these influence techniques, what should the policy response be? We think it necessary to create a greater separation between industry and science funding, especially in cases where there is a potential conflict of interest between industry incentives and public health concerns.¹⁹ It is clear that as long as industry is incentivized to get around the rules, they will find ways to do so. Relatedly, Holman (2015) describes the arms race occurring between pharmaceutical companies and officials seeking to regulate their influence on science. In this history, policy aimed to protect public health was repeatedly, creatively dodged by industry. Industry is an important funder of new science, but it is clear that current policy to prevent harms from industry funding of science is inadequate given these creative techniques.

One solution could be centralized bodies, under public control, which funnel industry money for some research area to the scientists and labs deemed best given public interest. In such a case, industry cannot choose which labs to fund based on their methods, and cannot dramatically over-fund just some part of the causal picture. We are not the first to suggest something along these lines (O'Connor and Weatherall, 2019b; Pinto and Pinto, 2023). This is not necessarily an easy policy to implement given the complex involvement of industry in current research funding. Furthermore Holman and Bruner (2017) suggest that in the presence of industrial funding, centralized funding can sometimes exacerbate industry influence because it often rewards those who have already been rewarded. To work, such an agency would have to itself avoid significant influence from industry, which may not be easy given the (discussed) industry incentives to find creative ways to influence science. Pinto and Pinto (2023) suggest a greater reliance on lottery funding as a way to avoid industrial selection in such cases, which may be a useful tool.²⁰

Another relevant policy area concerns industry communication about science. In some cases industrial distraction functions mostly via communication rather than funding. Given free speech protections, it is tricky to regulate industry sharing of accurate scientific information. Relevant laws, though, could

¹⁸See also Eliaz and Spiegler (2024) and Mohseni et al. (2022) for models of how news media, by sharing just some accurate content, can likewise mislead.

¹⁹Both Resnik and Elliott (2013) and Elliott (2014) discuss the differences between cases where industry is incentivized to fund and share accurate science, versus the sort of cases we focus on here.

²⁰Others have argued in favor of lottery funding for different potential benefits (Avin, 2015; Gross and Bergstrom, 2019; Smaldino et al., 2019; Avin, 2019; Shaw, 2023; Wu and O'Connor, 2023).

require sharing appropriate context along with distracting information. Under this policy Coca Cola could share information about sedentary lifestyles only when also sharing information about the relationship between soda and diabetes. This proposal is related to journalistic “balance” norms—that reporters should share information with context and balance. The idea is to apply similar balance rules to industry publicized science.

There is a related debate in philosophy of science. The question is when and whether it is right to suppress inappropriate scientific dissent—dissent that seems to be grounded in industrial or political interests rather than scientific doubt. Some authors argue that it is too difficult to delineate appropriate from inappropriate dissent, and that to suppress dissent without a clear delineation is too risky (de Melo-Martín and Intemann, 2014; de Melo-Martín and Intemann, 2018; Coates, 2024). On the other side are those who think it appropriate to identify and suppress this sort of dissent (Nash, 2018; Oreskes, 2017; Cook, 2017; Biddle and Leuschner, 2015; Biddle et al., 2017; Leuschner, 2018). Analyses like ours, and those described above, looking into specific industry techniques do highlight difficulties for this sort of delineation. For example, as noted, Coca Cola often funds legitimate scientists who are doing important work on exercise. It can be hard to say whether such work is either propaganda or normal science—it straddles the fence. On the other hand, though, understanding these techniques gives us a deeper ability to identify and fight them. Given the clear harms of industrial manipulation, and a track record of researchers successfully identifying and analyzing this manipulation, there will be many cases where inappropriate dissent can be identified and managed.

Recently, a great deal of work in philosophy and the social sciences has sought to define or delineate various sorts of misleading content, including misinformation, disinformation, malinformation, and fake news (Fallis, 2016; Weatherall and O’Connor, 2024). A typical claim, especially earlier in this literature, was to define terms like misinformation and disinformation as involving false or inaccurate content (Floridi, 1996, 2011; Fetzer, 2004). But increasingly it is recognized that much content is true or accurate, but nonetheless misleading (Fallis, 2015; Wardle and Derakhshan, 2017). And, in addition, misinformation and disinformation take many, varied forms, and can have many different sorts of impacts on belief and decision making (Harris, 2023; Simion, 2023; Habgood-Coote, 2019). Analysis of industrial propaganda can helpfully inform this discussion (O’Connor and Weatherall, 2019b). Techniques used by industry, as noted, mislead in a variety of creative ways, not all of which involve falsehoods. Ultimately, it is unlikely that it will be possible to derive definitions capturing all the types of content we might like to label as misinformation, disinformation, or industrial propaganda. Instead, specific analyses, like the one here, can help us better understand the variety of misleading content out there. And a thorough understanding of this variety can guide and shape successful policy aimed at regulating misleading content.

Before finishing, one last note. We focus in this paper on purposeful attempts to reshape causal understandings of the world, with the goal of shaping public behavior and policy. But there are going to be many similar cases where other

sorts of factors bias 1) the list of causes and effects the public is aware of and 2) their understanding of the relative strengths of these causes and effects. For example, it is widely recognized that the values scientists hold end up shaping what they choose to study and thus, often, what results exist on which topics (Haraway, 1991; Longino, 1990). The values of science journalists, as well as incentives they face, shape what they communicate and when (Mohseni et al., 2022). Algorithms on social media, and the public values and cognitives tendencies that shape these algorithms, determine who sees what scientific results. All these factors determine what evidence members of the public and policy makers see, and thus what their causal picture of the world looks like. The sorts of effects we outline here can happen as an accidental result of endogenous social forces, rather than the purposeful results of propaganda. This means that in thinking about promoting good public belief, attention is needed not just to the quality of information shared, but to its distribution and frequency.

Altogether, we take it to be very important to provide clear analyses of industrial propaganda techniques like industrial distraction. Doing so makes clear how and when industry harms public belief, and how and when industry can sway policy in their favor. As is clear, this analysis illuminates the workings of industrial distraction, highlights its relevance to current discussions in philosophy and the social sciences, and suggests policy responses.

Acknowledgements

Thanks to Ben Genta, Chris Torsell, Tori Cotton, Matthew Coates, Rebecca Korf, and Jim Weatherall for comments on this manuscript. Thanks to participants in the SKAT workshop at Columbia University for comments and feedback. Thanks to commentary from attendees at the PSA 2024 meeting in New Orleans, and to anonymous referees.

References

- Environmental Tobacco Smoke and Health: the Consensus. Standard, The Tobacco Institute, Washington, DC, (1986).
- Allen, David, Chelsea Linsley, Naomi Spoelman, and Alyssa Johl. The Fraud of Plastic Recycling. Standard, Center for Climate Integrity, (2024).
- Avin, Shahar (2015). “Funding science by lottery.” *Recent developments in the philosophy of science: EPSA13 Helsinki*, 111–126.
- Avin, Shahar (2019). “Centralized funding and epistemic exploration.” *The British Journal for the Philosophy of Science*.
- Bateman, Brooke L, Chad Wilsey, Lotem Taylor, Joanna Wu, Geoffrey S LeBaron, and Gary Langham (2020). “North American birds require mitiga-

- tion and adaptation to reduce vulnerability to climate change.” *Conservation Science and Practice*, 2(8), e242.
- Benshoff, Laura (2022). “Renewable energy is maligned by misinformation. It’s a distraction, experts say.” *National Public Radio*.
- Bernstein, Paul, David Montgomery, Barat Ramkrishnan, and Sughandu Tuladar. Impacts of Greenhouse Gas Regulations On the Industrial Sector. Standard, NERA Consulting, (2017).
- Bes-Rastrollo, Maira, Matthias B Schulze, Miguel Ruiz-Canela, and Miguel A Martinez-Gonzalez (2013). “Financial conflicts of interest and reporting bias regarding the association between sugar-sweetened beverages and weight gain: a systematic review of systematic reviews.” *PLoS medicine*, 10(12), e1001578.
- Biddle, Justin B, Ian James Kidd, and Anna Leuschner (2017). “Epistemic corruption and manufactured doubt: The case of climate science.” *Public Affairs Quarterly*, 31(3), 165–187.
- Biddle, Justin B and Anna Leuschner (2015). “Climate skepticism and the manufacture of doubt: Can dissent in science be epistemically detrimental?.” *European Journal for Philosophy of Science*, 5, 261–278.
- Biesecker, Michael and Paul Wiseman (2017). “AP Fact Check: Trump’s shaky claims on climate accord.” *Associated Press*.
- Brownell, Kelly D and Kenneth E Warner (2009). “The perils of ignoring history: Big Tobacco played dirty and millions died. How similar is Big Food?.” *The Milbank Quarterly*, 87(1), 259–294.
- Carpenter, Murray (2025). *Sweet and Deadly: How Coca Cola Spread Disinformation and Makes Us Sick*. MIT Press.
- Coates, Matthew (2024). “Does it Harm Science to Suppress Dissenting Evidence?.” *Philosophy of Science*.
- Cook, John (2017). “Response by Cook to “beyond counting climate consensus”.” *Environmental Communication*, 11(6), 733–735.
- Cummings, K Michael, Anthony Brown, and Richard O’Connor (2007). “The cigarette controversy.” *Cancer Epidemiology Biomarkers & Prevention*, 16(6), 1070–1076.
- Melo-Martin, Inmaculadade and Kristen Intemann (2014). “Who’s afraid of dissent? Addressing concerns about undermining scientific consensus in public policy developments.” *Perspectives on Science*, 22(4), 593–615.
- Melo-Martín, Inmaculadade and Kristen Intemann (2018). *The fight against doubt: How to bridge the gap between scientists and the public*. Oxford University Press.

- Eliasz, Kfir, Simone Galperti, and Ran Spiegler (2022). “False narratives and political mobilization.” *arXiv preprint arXiv:2206.12621*.
- Eliasz, Kfir and Ran Spiegler (2024). “News Media as Suppliers of Narratives (and Information).” *arXiv preprint arXiv:2403.09155*.
- Elliott, Kevin C (2014). “Financial conflicts of interest and criteria for research credibility.” *Erkenntnis*, 79(Suppl 5), 917–937.
- Fallis, Don (2015). “What is Disinformation?.” *Library Trends*, 63(3), 401–426.
- Fallis, Don (2016). “Mis- and Dis- Information.” *The Routledge Handbook of Philosophy of Information*. Ed. Luciano Floridi. New York: Routledge, 332–346.
- Fazio, Lisa K, Nadia M Brashier, B Keith Payne, and Elizabeth J Marsh (2015). “Knowledge does not protect against illusory truth..” *Journal of experimental psychology: general*, 144(5), 993.
- Fetzer, James H. (2004). “Disinformation: The use of false information.” *Minds and Machines*, 2(14), 231–240.
- Floridi, Luciano (1996). “Brave.net.world: The internet as a disinformation superhighway?.” *Electronic Library*, 14, 509—514.
- Floridi, Luciano (2011). *The Philosophy of Information*. Oxford, UK: Oxford University Press.
- Freeborn, David Peter Wallis (2023). *Polarization and factionalization for agents with multiple, related beliefs*. University of California, Irvine.
- Freeborn, David Peter Wallis (2024). “Rational factionalization for agents with probabilistically related beliefs.” *Synthese*, 203(2), 46.
- Greenberg, Jon. Fact-checking Donald Trump’s statement withdrawing from the Paris climate agreement. Standard, PolitiFact, (2017).
- Greenhalgh, Susan (2024). *Soda Science: Making the World Safe for Coca-Cola*. University of Chicago Press.
- Gross, Kevin and Carl T Bergstrom (2019). “Contest models highlight inherent inefficiencies of scientific funding competitions.” *PLoS biology*, 17(1), e3000065.
- Habgood-Coote, Joshua (2019). “Stop talking about fake news!.” *Inquiry*, 62(9-10), 1033–1065.
- Haraway, Donna (1991). *Simians, cyborgs, and women: The reinvention of nature*. Routledge.
- Harris, Keith Raymond (2023). “Beyond Belief: On Disinformation and Manipulation.” Forthcoming. *Erkenntnis*.

- Hassan, Aumyo and Sarah J Barber (2021). “The effects of repetition frequency on the illusory truth effect.” *Cognitive research: principles and implications*, 6(1), 38.
- Holman, Bennett (2015). “The Fundamental Antagonism.” *PhD Dissertation*.
- Holman, Bennett and Justin Bruner (2017). “Experimentation by industrial selection.” *Philosophy of Science*, 84(5), 1008–1019.
- Hu, Akielly (2023). “Republic Donors are Funding Misinformation About Off-shore Wind.” *Canary Media*.
- Jern, Alan, Kai-Min K Chang, and Charles Kemp (2014). “Belief polarization is not always irrational..” *Psychological review*, 121(2), 206.
- Johns, David Merritt and Gerald M Oppenheimer (2018). “Was there ever really a “sugar conspiracy”?.” *Science*, 359(6377), 747–750.
- Katovich, Erik (2023). “Quantifying the Effects of Energy Infrastructure on Bird Populations and Biodiversity.” *Environmental Science & Technology*, 58(1), 323–332.
- Kearns, Cristin E, Stanton A Glantz, and Laura A Schmidt (2015). “Sugar industry influence on the scientific agenda of the National Institute of Dental Research’s 1971 National Caries Program: a historical analysis of internal documents.” *PLoS medicine*, 12(3), e1001798.
- Kearns, Cristin E, Laura A Schmidt, and Stanton A Glantz (2016). “Sugar industry and coronary heart disease research: a historical analysis of internal industry documents.” *JAMA internal medicine*, 176(11), 1680–1685.
- Kim, J.H. and J. Pearl (1983). “A computational model for combined causal and diagnostic reasoning in inference systems.” *Proceedings of the 8th International Joint Conference on Artificial Intelligence*. Karlsruhe, Germany: (IJCAI-83), 190–193.
- Kourany, Janet and Martin Carrier (2020). *Science and the production of ignorance: When the quest for knowledge is thwarted*. Mit Press.
- Lesser, Lenard I, Cara B Ebbeling, Merrill Goozner, David Wypij, and David S Ludwig (2007). “Relationship between funding source and conclusion among nutrition-related scientific articles.” *PLoS medicine*, 4(1), e5.
- Leuschner, Anna (2018). “Is it appropriate to ‘target’inappropriate dissent? On the normative consequences of climate skepticism.” *Synthese*, 195, 1255–1271.
- Lewandowsky, Stephan, Toby D Pilditch, Jens K Madsen, Naomi Oreskes, and James S Risbey (2019). “Influence and seepage: An evidence-resistant minority can affect public opinion and scientific belief formation.” *Cognition*, 188, 124–139.

- Longino, Helen E (1990). *Science as social knowledge: Values and objectivity in scientific inquiry*. Princeton university press.
- Ludwig, David S, Karen E Peterson, and Steven L Gortmaker (2001). “Relation between consumption of sugar-sweetened drinks and childhood obesity: a prospective, observational analysis.” *The lancet*, 357(9255), 505–508.
- Malik, Vasanti S, Barry M Popkin, George A Bray, Jean-Pierre Després, and Frank B Hu (2010). “Sugar-sweetened beverages, obesity, type 2 diabetes mellitus, and cardiovascular disease risk.” *Circulation*, 121(11), 1356–1364.
- Malik, Vasanti S, Matthias B Schulze, and Frank B Hu (2006). “Intake of sugar-sweetened beverages and weight gain: a systematic review.” *The American journal of clinical nutrition*, 84(2), 274–288.
- Mohseni, Aydin, Cailin O’Connor, and James Owen Weatherall (2022). “The Best Paper You’ll Read Today.” *philosophical topics*, 50(2), 127–153.
- Nash, Erin J (2018). “In Defense of “Targeting” Some Dissent about Science.” *Perspectives on Science*, 26(3), 325–359.
- Negin, Elliott. ExxonMobil Claims Shift on Climate But Continues to Fund Climate Science Deniers. Standard, Union of Concerned Scientists, (2020).
- Nestle, Marion (2015). *Soda politics: taking on big soda (and winning)*. Oxford University Press, USA.
- O’Connor, Anahad (2015). “Coca-Cola Funds Scientists Who Shift Blame for Obesity Away From Bad Diets.” *The New York Times*.
- O’Connor, Cailin and James Weatherall (2019a). “How Powerful Interests Use Science to Sway Public Opinion.” *Zocalo Public Square*.
- O’Connor, Cailin and James Owen Weatherall (2019b). *The misinformation age: How false beliefs spread*. Yale University Press.
- Oreskes, Naomi (2017). “Response by Oreskes to “Beyond counting climate consensus”.” *Environmental Communication*, 11(6), 731–732.
- Oreskes, Naomi and Erik M Conway (2011). *Merchants of doubt: How a handful of scientists obscured the truth on issues from tobacco smoke to global warming*. Bloomsbury Publishing USA.
- Pearl, Judea (2009). *Causality: Models, Reasoning and Inference*. 2nd edition. USA: Cambridge University Press.
- Pinto, Manuela Fernández and Daniel Fernández Pinto (2023). “Epistemic diversity and industrial selection bias.” *Synthese*, 201(5), 182.
- Proctor, Robert N. (1995). *Cancer Wars: How Politics Shapes What We Know and Don’t Know About Cancer*. New York: Basic Books.

- Proctor, Robert N (2012). *Golden holocaust: origins of the cigarette catastrophe and the case for abolition*. Univ of California Press.
- Quintana-Rizzo, E, S Leiter, TVN Cole, MN Hagbloom, AR Knowlton, P Nagelkirk, OO Brien, CB Khan, AG Henry, PA Duley, et al. (2021). “Residency, demographics, and movement patterns of North Atlantic right whales *Eubalaena glacialis* in an offshore wind energy development area in southern New England, USA.” *Endangered species research*, 45, 251–268.
- Resnik, David B and Kevin C Elliott (2013). “Taking financial relationships into account when assessing research.” *Accountability in research*, 20(3), 184–205.
- Schulze, Matthias B, JoAnn E Manson, David S Ludwig, Graham A Colditz, Meir J Stampfer, Walter C Willett, and Frank B Hu (2004). “Sugar-sweetened beverages, weight gain, and incidence of type 2 diabetes in young and middle-aged women.” *Jama*, 292(8), 927–934.
- Serodio, Paulo, Gary Ruskin, Martin McKee, and David Stuckler (2020). “Evaluating Coca-Cola’s attempts to influence public health ‘in their own words’: analysis of Coca-Cola emails with public health academics leading the Global Energy Balance Network.” *Public Health Nutrition*, 23(14), 2647–2653.
- Shaw, Jamie (2023). “Peer review in funding-by-lottery: A systematic overview and expansion.” *Research Evaluation*, 32(1), 86–100.
- Simion, Mona (2023). “Knowledge and Disinformation.” Forthcoming. *Episteme*.
- Singla, Veena (2022). “Recycling Lies.” *Chemical Recycling of Plastic Is Just Greenwashing Incineration*, NRDC.
- Smaldino, Paul E, Matthew A Turner, and Pablo A Contreras Kallens (2019). “Open science and modified funding lotteries can impede the natural selection of bad science.” *Royal Society open science*, 6(7), 190194.
- Sovacool, Benjamin K (2013). “The avian benefits of wind energy: A 2009 update.” *Renewable Energy*, 49, 19–24.
- Spiegler, Ran (2020). “Can agents with causal misperceptions be systematically fooled?.” *Journal of the European Economic Association*, 18(2), 583–617.
- Spirtes, Peter, Clark Glymour, Scheines N., and Richard (2000). *Causation, Prediction, and Search*. 2nd edition. Mit Press: Cambridge.
- Thompson, Paul M, David Lusseau, Tim Barton, Dave Simmons, Jan Rusin, and Helen Bailey (2010). “Assessing the responses of coastal cetaceans to the construction of offshore wind turbines.” *Marine pollution bulletin*, 60(8), 1200–1208.
- UCS. How Coca Cola Disguised its Influence on Science about Sugar and Health. Standard, Union of Concerned Scientists, (2017).

- Udry, Jessica and Sarah J Barber (2023). “The illusory truth effect: A review of how repetition increases belief in misinformation.” *Current Opinion in Psychology*, 101736.
- Wardle, Claire and Hossein Derakhshan. Information Disorder: Toward an interdisciplinary framework for research and policy making. Technical report, Council of Europe, (2017).
- Weatherall, James O and Cailin O’Connor (2024). “Fake News!.” *Philosopher’s Imprint*.
- Weatherall, James Owen, Cailin O’Connor, and Justin P Bruner (2020). “How to beat science and influence people: Policymakers and propaganda in epistemic networks.” *The British Journal for the Philosophy of Science*.
- Wellman, M.P. and M. Henrion (1993). “Explaining ‘explaining away’.” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(3), 287–292.
- Wood, Benjamin, Gary Ruskin, and Gary Sacks (2020). “How Coca-Cola shaped the international congress on physical activity and public health: an analysis of email exchanges between 2012 and 2014.” *International Journal of Environmental Research and Public Health*, 17(23), 8996.
- Wu, Jingyi and Cailin O’Connor (2023). “How should we promote transient diversity in science?.” *Synthese*, 201(2), 37.
- Yang, Quanhe, Zefeng Zhang, Edward W Gregg, W Dana Flanders, Robert Merritt, and Frank B Hu (2014). “Added sugar intake and cardiovascular diseases mortality among US adults.” *JAMA internal medicine*, 174(4), 516–524.