

Folk intuitions of actual causation: a two-pronged debunking explanation

David Rose¹

Published online: 18 August 2016
© Springer Science+Business Media Dordrecht 2016

Abstract How do we determine whether some candidate causal factor is an actual cause of some particular outcome? Many philosophers have wanted a view of actual causation which fits with folk intuitions of actual causation and those who wish to depart from folk intuitions of actual causation are often charged with the task of providing a plausible account of just how and where the folk have gone wrong. In this paper, I provide a range of empirical evidence aimed at showing just how and where the folk go wrong in determining whether an actual causal relation obtains. The evidence suggests that folk intuitions of actual causation are generated by two epistemically defective processes. I situate the empirical evidence within a background discussion of debunking, arguing for a two-pronged debunking explanation of folk intuitions of actual causation. I conclude that those who wish to depart from folk intuitions of actual causation should not be compelled to square their account of actual causation with the verdicts of the folk. In the dispute over actual causation, folk intuitions deserve to be rejected.

Keywords Actual causation · Folk intuitions · Debunking

*“[I]f a stone falls suddenly in the brush near an adult, he will usually mutter ‘a spirit’”
Margaret Mead’s (1932) notes on the Manus people of Papua New Guinea*

✉ David Rose
drose@philosophy.rutgers.edu

¹ Department of Philosophy, Rutgers University, New Brunswick, NJ, USA

How do we determine whether some candidate causal factor is an actual cause of some particular outcome?¹ Many philosophers have wanted a view of actual causation which fits with folk intuitions of actual causation (e.g., Halpern and Hitchcock 2015; Hitchcock 2007; Hitchcock and Knobe 2009; Lewis 1986; McDermott 1995; McGrath 2005; Mellor 1995; Menzies 1996, 2009; Schaffer 2000, 2004, 2005). And those who wish to depart from folk intuitions of actual causation are often charged with the task of providing a plausible account of just how and where the folk have gone wrong.

Thus, one important task is for the revisionist is to meet the *challenge from folk belief* (Korman 2009): the revisionist should explain why the folk believe as they do when the resultant theory apparently conflicts with relevant folk beliefs. In doing so, the revisionist should offer up an explanation that is not (1) globally self-defeating (i.e., challenging the very ability to form true beliefs) or (2) locally self-defeating (e.g., if one locates the alleged source of error in mistaken intuitions in ways that suggests a general skepticism about intuitive judgments, then one cannot also rely on intuitions to support the premises in an argument for some revisionary view). As Korman puts it, “virtually everyone agrees that, even after having presented the arguments for their positions, proponents of revisionary philosophical theories—that is, those that deviate from the pretheoretical conception—are required to provide some sort of account of the conflict between their theories and the pretheoretical beliefs of non-philosophers (“the folk”)” (2009, p. 242).

Though many are pessimistic that the challenge from folk belief can be met—thinking for instance that “revisionists standardly delude themselves into thinking that they can plausibly explain why people make the mistakes they allege” (Hirsch 2002, p. 117; see also e.g., Korman 2009, p. 242; Paul 2012, p. 22)—I’m optimistic. My view is that, aided and guided by work from cognitive science, one can meet the challenge from folk belief by providing a targeted debunking explanation for the relevant folk intuitions in the target domain; that is, one that does not run afoul of the challenge from folk belief by succumbing to either global or local self-defeat.

My plan is to speak on behalf of the revisionist—in an empirically informed way—by providing a targeted debunking explanation for folk intuitions of actual causation. Specifically, I’ll provide a two-pronged debunking explanation for folk intuitions of actual causation. Both prongs target epistemically defective processes involved in generating folk intuitions of actual causation. One process is rooted in a motivation to blame; the other is rooted in primitive teleological considerations. Taken together, these two epistemically defective processes provide resources for helping meet the challenge from folk belief in the specific domain of actual causation and as such provide the revisionist with the resources for holding that measuring a

¹ There is a standard distinction made between actual (or token or singular) causation, on the one hand, and generic (or type) causation, on the other. Roughly, generic causation is typically thought to be a relation between types of events. Actual causation is typically taken to be a relation between individual events (Lewis 1986). For example, we might say that “running causes weight loss”. Or, we might say that “John’s running fifteen miles caused him to lose weight”. The former would be a case of generic causation while the latter would be an example of actual causation. My focus throughout is on actual causation.

theory of actual causation by its fit with folk intuitions of actual causation is not a wise policy. Folk intuitions of actual causation deserve to be rejected.

To clarify, I won't be casting a general pox on folk intuitions of actual causation by arguing that all aspects of causal cognition are infused with the motivation to blame or primitive teleological considerations. My debunking explanation for folk intuitions of actual causation is targeted—and thus tempered—and should be understood as follows: insofar as folk intuitions of actual causation are generated by a motivation to blame or primitive teleological considerations, then those intuitions deserve to be rejected. This is a targeted debunking explanation for folk intuitions of actual causation. It is thereby fit to meet the challenge from folk belief.

The Plan: I'll begin by briefly documenting the role of folk intuitions of actual causation in evaluating theories of actual causation. Then, in Sect. 2, I'll briefly discuss some background empirical work on folk intuitions of actual causation to set the stage for the two-pronged debunking explanation. In Sect. 3, I present empirical evidence supporting the two-pronged debunking explanation. Section 4 discusses debunking and situates the empirical evidence within a background discussion of debunking and the challenge from folk belief. Section 5 considers some objections.

1 Fitting folk judgments of actual causation

Perhaps one of the clearest statements that a theory of actual causation needs to respect folk intuitions comes from Lewis (1986):

When common sense delivers a firm and uncontroversial answer about a not-too-far-fetched case, theory had better agree. If an analysis of causation does not deliver the common-sense answer, that is bad trouble (p. 94).

Others have followed suit in thinking that folk intuitions about actual causation need to be respected (e.g., Halpern and Hitchcock 2015; Hitchcock 2007; Hitchcock and Knobe 2009; Mellor 1995; McDermott 1995; McGrath 2005; Menzies 1996, 2009; Schaffer 2000, 2004, 2005). Yet some have departed from this, thinking that folk intuitions deserve to be rejected.

For instance, Beebe (2004), in arguing that preventers and omissions are not causes, claims that the folk confuse causal explanation and causation and so argues for a dismissive take on alleged folk intuitions that preventers and omissions are causes.² And Dowe (2000, 2004), who is explicit about not placing “a premium on respecting folk intuitions”, realizes that “others do” and so takes up the task of explaining why we mistakenly treat preventers and omissions as causes: it's because we confuse causation and quasi-causation.³

So, fit with folk intuitions of actual causation is taken to serve as an important desideratum in evaluating theories of actual causation. This is true not only for those

² Though see Livengood and Machery (2007) for evidence that the folk do distinguish causation and causal explanation.

³ I've only offered a brief sampling of disputes over the common sense view of actual causation. For more, see the excellent discussion in Paul and Hall (2013).

who are engaged in conceptual analysis. It's also a constraint on theories which aim to produce causal concept(s) useful to scientists or metaphysicians (see e.g., Paul and Hall 2013). The constraint may be put as follows:

[I]f an analysis of causation does not deliver the common-sense answer, that is certainly *prima facie* trouble, since it is evidence that something of importance has been overlooked. So it may make sense—but only up to a point!—to proceed as if your analysis has been refuted, when it runs afoul of common sense. (Paul and Hall 2013, p. 3)

Even on this approach, “causal intuitions...[are] defeasible guides to potentially interesting and important features of our causal concept or a causal relation” and the philosopher needs to be “prepared to jettison those intuitions in the event that they are discovered to lead nowhere...” (p. 2).

Whether a theory of actual causation should be applauded for fitting folk intuitions depends on whether those intuitions should be respected. And empirical evidence on why the folk intuit as they do can help in deciding whether the relevant folk intuitions deserve respecting or rejecting. That said, I'll briefly discuss some empirical work on folk intuitions of actual causation to set the stage both for my own studies and for the two-pronged debunking explanation.

2 Empirical work on folk judgments of actual causation

There are two threads of empirical evidence I want to consider. The first—which is more developed—concerns the role of moral considerations in folk intuitions of actual causation. The second—understudied but taken up in the empirical studies below—concerns the role of primitive teleological considerations in folk intuitions of actual causation.

2.1 Moral considerations

The role of moral considerations in generating judgments of actual causation is well documented (e.g., Alicke 1992, 2000; Alicke and Rose 2010; Alicke et al. 2011; Sytsma et al. 2012). For instance, Alicke (1992) presents evidence that, with all other factors held fixed across cases, people are much more willing to assign blame and causation to an individual involved in a car accident when that individual was speeding home to hide cocaine as opposed to an anniversary present. And Alicke et al. (2011) model the causal relationship between judgments of blame and causation, finding that blame plays a direct causal role in generating judgments of actual causation.

According to Alicke and colleagues, in the realm of harmful and offensive actions, folk intuitions of actual causation are skewed by a desire to blame those who we evaluate negatively.⁴ We exaggerate an actor's causal role in bringing about

⁴ For ease, I'm only discussing the role of blame in causal judgment. But, as Alicke et al. (2011) argue, causal assessments can also be influenced by a desire to praise. Also, for evidence on the flip side of this—excuse validation—see Turri and Blouw (2015).

an outcome since doing so allows us to support our desire to blame the agent. Thus, our desire to blame an individual actually leads us to adjust our assessment of the agents causal role in the production of the outcome since doing so supports our desire to blame. On this view, the effect of moral considerations on folk intuitions of actual causation is an *error*,⁵ rooted in a motivational bias to blame those who engage in harmful or offensive actions (see e.g., Alicke 1992, 2000; Alicke and Rose 2010; Alicke et al. 2011).

This work suggests that there is a psychological process—which I'll call the *evaluative process*—that plays a direct role in generating folk intuitions of actual causation. This forms the basis for the first prong of my targeted debunking argument. Though the evidence in favor of this process is well developed, I move beyond extant work in two key ways. First, though much of the extant work on the role of moral considerations in causal judgments has focused on human action, I will provide evidence that the evaluative process extends beyond the realm of human action and do so by empirically integrating these results with results concerning promiscuous teleology (Sect. 3). And second, while some have claimed that the role of blame in causal judgment is a bias (e.g., Alicke and Rose 2010; Alicke et al. 2011) there has been no explicit discussion of the philosophical upshot of these empirical results. I will, however, philosophically integrate these results within a debunking framework (Sect. 4) and thus provide the resources to clearly depart from those who hold that a philosophical theory of causation should respect morally laden causal intuitions (e.g., Halpern and Hitchcock 2015; Hitchcock and Knobe 2009; McGrath 2005). In doing so, I'm taking a stand on how best to interpret the role of moral considerations in causal judgment.

Controversy remains as to how to best interpret the role of moral considerations in causal judgments.⁶ Where I see a glaring epistemically defective process, others see an epistemically appropriate process. For instance, on one leading view, Hitchcock and Knobe (2009) argue that norm violations directly impact judgments of actual causation. People assign heightened causation when a causal candidate deviates from its normal state: judgments of actual causation are guided by norm violations broadly construed. Blame plays no role at all in this. Given that the role of norm violations in causal judgment isn't driven by blame, perhaps the role of norm violations in causal judgment is entirely appropriate. However, a range of empirical evidence does not cohere well with the norm violation view. For instance, Sytsma et al. (2012) provide evidence that typical behaviors, as opposed to atypical ones, lead to heightened causal assignment, which is the exact opposite of what Hitchcock and Knobe predict. Moreover, Alicke et al. (2011)—utilizing the cases presented by Hitchcock and Knobe as well as some of their own—present a range of empirical evidence supporting the view that the desire to blame plays a direct role in generating judgments of actual causation. In light of this, I take the evidence to provide support for my interpretive stance: the role of moral considerations in causal judgment—captured in what I've called the evaluative process—is epistemically defective.

⁵ See Sect. 5.1 for further discussion.

⁶ Thanks to an anonymous reviewer for raising this.

That said, I would flag that the interpretive issue is far from settled. If the best interpretation of the role of moral considerations turns out to be that they are entirely appropriate, then this will undercut one prong of the debunking explanation I'll be offering. Even so, the second prong—to be discussed below—would still stand. I won't be trying to settle this interpretive issue here. Instead, I'm making a "judgment call" (Stich 2013, p. 156). But I take it that the evidence presented in Sect. 3 will further support the interpretive stance I'm taking. I'll also have a bit more to say in favor of my interpretation in Sect. 5.1.

2.2 Agentive considerations

The second process involved in generating folk intuitions of actual causation—which I will call the *agentive process*—has been almost entirely neglected. Some work suggests that agentive considerations—in particular, whether an individual's behavior is construed as intentional or accidental—impact folk intuitions of actual causation. For instance, work by Lagnado and Channon (2008), Channon et al. (2010) and Lombrozo (2010) has found that when negative outcomes are brought about, individuals who bring about the outcome intentionally are assigned a greater causal role in producing the outcome than individuals who bring about the outcome by accident. These studies have solely focused on the role intentionally and accidentally construed behaviors play in causal judgments about agents. But the specific process I'm targeting isn't restricted to the domain of human action. Rather, it extends to nature as a whole. Evidence in support of this can be provided by considering work on promiscuous teleology.

There is a wide range of evidence suggesting that people are promiscuous teleologists in that teleological considerations play a role not only in our conception of human actions, but also of artifacts, biological organisms and non-living natural things like rocks. A range of evidence supports the view that children are promiscuous teleologists (e.g., Kelemen and DiYanni 2005; Kelemen 1999a, b, 2004) in that they naturally accept such statements as "rocks are pointy to prevent animals from sitting on them". Other work suggests that even adults never fully outgrow their childhood tendencies toward promiscuous teleology (e.g., Kelemen and Rosset 2009; Kelemen et al. 2013; Lombrozo et al. 2007). For instance, Kelemen and Rosset (2009) found that college aged students indulge in accepting unwarranted teleological explanations—endorsing such statements as "The sun radiates heat because warmth nurtures life", "Fungi grows in forests to help with decomposition" and even "Lightening occurs to release electricity"—even in scientific contexts. Similarly, Kelemen et al. (2013) also found that even trained physical scientists show a similar pattern of accepting unwarranted teleological explanations when their cognitive resources were limited (when in a "speeded task").

Though there is a range of evidence that we're inclined toward promiscuous teleological thinking, the question arises as to *why* we're so inclined. As Bloom (2007, p. 150) proposes:

We have a bias to attribute an agent when we see nonrandom structure. When we see complex structure, we see it as the product of beliefs and goals and

desires. We chew over the natural world with our social mode of understanding, and it is difficult to make sense of it in any other way.

Support for the view that nature as a whole is viewed in agentive terms comes from work by Kelemen et al. (2013) who found that people's endorsement of background Gaia beliefs predicted their tendency toward accepting teleological explanations. Other work—specifically from the science education literature—suggests that people's tendency toward viewing nature as a whole in agentive terms is one of the primary obstacles in students' path to acquiring an adequate understanding of natural selection (see Galli and Meinardi 2011; Kelemen 2012 for an overview). For instance, students tend to think that a “personified “Mother Nature” responded to animal's functional needs by generating or conferring the functional part with a view to preserving the animal's survival” (Kelemen 2012, p. 4; see also Kampourakis and Zogza 2008; Moore et al. 2002; and Gregory 2009), such as by stretching a giraffe's neck so it could reach leaves on trees (e.g., Clough and Wood-Robinson 1985; Demastes et al. 1995; Evans et al. 2010; Kampourakis and Zogza 2008). Summing all this up, Kelemen (2012, p. 7) writes:

Findings suggesting that underlying beliefs about natural agency exert non-obvious influence on students' biological reasoning are potentially less surprising when considered in a broader context of research which suggests that such immanent agentive ideas influence adults' scientifically incorrect ideas about living and non-living nature more generally. For example, in contrast to their ratings of belief in God, students' ratings of the Gaia notion that “Nature is driven to preserve living things” has been found to strongly predict undergraduates promiscuous (but often covert) tendencies to teleologically explain not only living but also non-living natural phenomena in terms of a purpose: That is, an agentive construal of nature provides a significant reason why American undergraduates find scientifically inaccurate teleological statements such as “the sun makes light so that plants can photosynthesize” highly believable even after extensive high school and college level instruction in both the physical and life sciences (Kelemen et al. 2013; also Kelemen and Rosset 2009)

Taken together, the evidence suggests that we're inclined toward endorsing teleological explanations because we view nature as a whole in agentive terms. More specifically, this work suggest that folk teleology is best understood as *promiscuous teleomentalism*, which is rooted in a primitive, superstitious view of nature as a whole in agentive terms.⁷ Promiscuous teleomentalism represents a strong, robust tendency to *error*.⁸

⁷ Here I'm connecting *promiscuous* teleology (i.e., that teleological explanations extend beyond the artifact and biological domain and play a role in explaining non-living natural phenomena) with *teleomentalism* (the view that the teleology of psychological intentions, goals, and purposes is the primary model for understanding teleology outside the domain of human action e.g., in considering non-living natural phenomena). See Allen and Bekoff (1994), p. 13 for a discussion of teleomentalism in biology.

⁸ Teleomentalism is typically regarded as an error and so eliminable. See Allen (2009) and Allen and Bekoff (1994) for a discussion of the eliminability of teleomentalism in biology.

Though no work has looked at whether this agentic construal of nature—promiscuous teleomentalism—impacts folk intuitions of actual causation, there's reason to expect that it will. People's tendency toward teleological explanation is predicted by background agentic considerations. Insofar as causation backs explanation, we should expect agentic considerations to impact folk intuitions of actual causation. Thus, promiscuous teleomentalism—embodied in what I have called the agentic process—should have a direct impact on folk intuitions of actual causation. This forms the basis for the second prong of my targeted debunking argument. Now, on to some direct evidence.

3 Evidence for two processes

Two candidate processes—the evaluative process and the agentic process—are hypothesized to play a role in generating folk intuitions of actual causation. The first process involved in generating intuitions of actual causation—the evaluative process—enjoys more empirical support than the second process, the agentic process. For this reason, most of the focus will be on providing support for the role of the agentic process in generating folk intuitions of actual causation. But, I will provide some extensions to the extant work on the role of the evaluative process. In particular, most of the work supporting this has been confined to causal judgments involving human actions. The primary extension will be to investigate whether the evaluative process extends beyond the realm of human action.

3.1 Study 1

Study 1 was aimed at investigating whether manipulating an intentional construal of events affects causal judgments about human action and causal judgments for non-living natural objects like rocks. Recall from Sect. 2.2, that some work has shown that manipulating whether an individual's behavior is intentional or accidental affects causal judgments. I'll move beyond this work by doing two things (1) investigate whether an intentional or accidental construal of events affects causal judgments beyond the realm of human action and (2) model the causal relationships between the candidate variables under consideration.

The cases for Study 1 had the following structure: A bird, Cantup, receives an essential nutrient, Keterine, from eating Weeble worms. Weeble worms receive Keterine by feeding from a rock, Zenite, which produces the Keterine. Zenite produces Keterine by absorbing heat. As heat is absorbed, this initiates a chemical reaction in Zenite which produces Keterine. The chemicals involved in producing Keterine are densely packed in the upper surface of Zenite. Zenite is prevented from absorbing heat. As a result, Keterine is not produced and the Cantups slowly start to die.

In the version involving a person, John moves Zenite so that it can absorb heat; in the version involving just the rock, Zenite redirects the chemicals toward the heat.

This is what I'll call the Individual (John, Zenite) manipulation. I also varied whether Zenite began producing Keterine and so the Cantups survived or whether Zenite failed to produce Keterine and so the Cantups died. The purpose of this manipulation—which I'll refer to as the Condition manipulation—was to vary an intentional or accidental construal of the events under consideration. The guiding idea was that in trying to remedy a bad situation, sometimes we're successful and other times we're not. When the goal is to fix a bad situation and we're successful, the outcome should be more likely to be viewed as intentionally brought about in comparison to a case where we're unsuccessful. Together, the study was a 2 (Individual: John, Zenite) \times 2 (Condition: Intentional, Accidental) design (for full cases see "[Appendix](#)").

3.1.1 Participants and measures

A total of 154 participants were recruited through Amazon's Mechanical Turk and randomly assigned to one of the four above conditions. After reading the story, participants were given the following probes (in random order):

Causation Probe: [John/Zenite] caused the Cantups to [survive/die].⁹

Intentionality Probe: [John/Zenite] intentionally [helped/harmed] the Cantups.¹⁰

Evaluation Probe: How would you evaluate [John's/Zenite's] behavior?¹¹

3.1.2 Results

An ANOVA revealed a statistically significant large-sized effect of Condition (Intentional, Accidental) on causal judgments $F(1, 150) = 62.797$, $p < .001$, $\eta^2 = .295$, no statistically significant effect of Individual (John, Rock) on causal judgments, $F(1, 150) = .027$, $p = .869$, and a statistically significant small-sized interaction between Condition and Individual $F(1, 150) = 4.518$, $p = .035$, $\eta^2 = .029$ (Fig. 1). T-tests for both the John and rock cases can be found in Table 1.

Importantly, the results indicate that regardless of whether the individual was an agent or a rock, the outcome had a dramatic effect on causal judgments. The crucial question now is: *why* is this pattern in people's causal judgments arising? Given that there was an interaction between Condition and Individual, I'll analyze responses in the rock and John cases separately to determine why this pattern is arising.

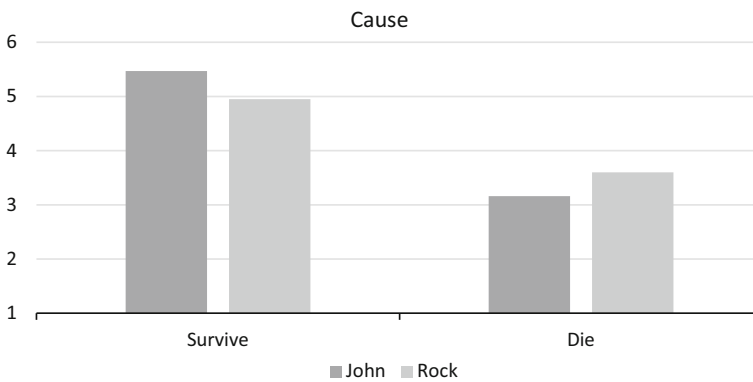
⁹ Responses were made on a 6-pt scale anchored with 1 = strongly disagree, 6 = strongly agree.

¹⁰ Responses were made on a 6-pt scale anchored with 1 = strongly disagree, 6 = strongly agree.

¹¹ Responses were made on a 6-pt scale anchored with 1 = extremely blameworthy, 6 = extremely praiseworthy.

Table 1 Study 1 t-tests


	Good outcome (intentional)	Bad outcome (accidental)	t value	p value
<i>John cases</i>				
Causation	5.47 (.646)	3.15 (1.85)	7.27	.000
Evaluation	5.31 (.701)	3.92 (1.04)	6.81	.000
Intentionality	5.71 (.459)	1.65 (1.45)	16.35	.000
<i>Rock cases</i>				
Causation	4.95 (1.31)	3.60 (1.59)	4.02	.000
Evaluation	4.03 (1.06)	3.24 (.942)	3.44	.001
Intentionality	2.54 (1.61)	1.48 (1.05)	3.45	.001

**Fig. 1** Causal judgments for John and the Rock

3.1.2.1 John cases I'll begin by looking at responses in the John cases. To determine the causal relationships among the candidate variables, I ran a causal search on the data, using Greedy Equivalence Search (GES).¹² GES returned the following model:¹³


¹² Roughly, GES operates by considering the possible models available given the different variables. GES begins by assigning an information score to the null model (i.e., a disconnected graph). GES then considers various possible arrows ("edges") between the different variables. It begins by adding the edge that yields the greatest improvement in the information score (if there is such an edge) and repeats the process until additional edges would not further improve the information score. GES then considers deletions which would yield the greatest improvement in the information score (if there is such an edge), repeating this procedure until no further deletions will improve the score. In all cases, the orientation of the edges is given by edge-orientation rules in Meek (1997). It has been shown by Chickering (2002) that, given enough data, GES will return the true causal model of the data. GES is often interpreted as returning the best fitting causal model, given the data (for further details and some applications, see Chickering 2002; Rose et al. 2011; Rose and Nichols 2013).

¹³ This model fits the data well, $df = 3$, $X^2 = 4.9871$, $p = .1727$, $BIC = -8.0051$.



Two things are worth noting about this model. The first is that evaluations of John's behavior make no contribution to causal judgments. Second, and more importantly, intentionality judgments screen off the effect of condition on causal judgments: whether John is viewed as a cause of the outcome depends on the extent to which he is viewed as intentionally bringing about the outcome. When the outcome is negative, participants are much less likely to view John as intentionally bringing about the outcome in comparison to the case where the outcome is positive.

3.1.2.2 Rock cases Next, to understand the causal relationships between the candidate variables in the rock cases, I ran a causal search on the data. GES returned the following model:¹⁴



In comparison to the cases involving John, in the rock cases, Intentionality had no effect on causal judgments. Moreover, whereas Intentionality screens off the effect of Condition on causal judgments for the cases involving John, here we see that people's evaluations screen off the effect of Condition on causal judgments.

Perhaps it's the case that the agentive process plays a role in generating causal judgments for human action but not non-living natural objects like rocks. For non-living natural objects like rocks, it may be that the evaluative process as opposed to

¹⁴ This model fits the data well, $df = 3$, $X^2 = 6.0381$, $p = .1098$, $BIC = -6.9933$.

the agentic process plays a role. If so, this would still provide support for the two-pronged debunking argument on offer. But it may be that people's explicit judgments of intentionality are not a good guide here. Indeed, it does seem quite unnatural to explicitly say that a rock was intentionally harming or helping. So, perhaps explicitly asking about intentionality for cases involving non-living natural objects is not the best way to probe for whether the agentic process is playing a role in people's causal judgments for these cases. A different measure would be preferable. I'll take this up in Study 2.

3.2 Study 2

3.2.1 Participants and measures

Endorsement of quasi-religious Gaia beliefs has been shown to significantly predict people's tendency to endorse teleological explanations (Kelemen et al. 2013). To determine whether the agentic process is having an effect on people's causal judgments in cases involving non-living natural objects, I ran the same cases involving a rock that were used above, used the same probes as above, but—borrowing from Kelemen et al. (2013)—added in measures to probe for Gaia Beliefs. They were:¹⁵

1. I believe Nature is driven to preserve living things
2. I believe the Earth is alive
3. I believe that Nature is a powerful being
4. I believe the Earth is driven to provide optimal conditions for Life

211 participants were recruited from Amazon Mechanical Turk and randomly assigned to either the good outcome or bad outcome cases.

3.2.2 Results

First, the results from the Rock cases above were replicated (see Table 2). Next, I examined the inter-correlations between the various Gaia Belief Probes, finding that they exhibited a high degree of internal consistency.¹⁶ So, I combined them together to form Gaia Composite.

I then ran a causal search on the data. GES returned the model below.¹⁷

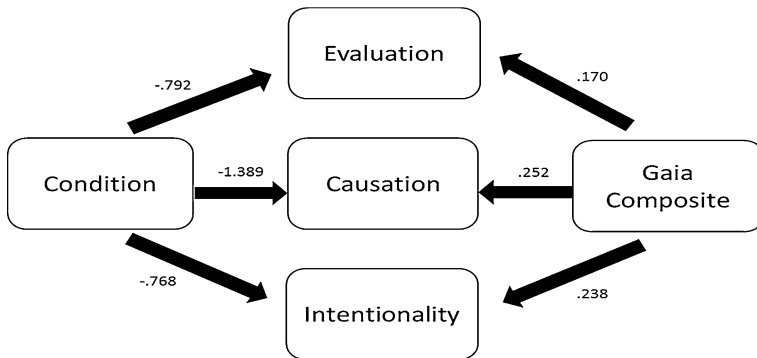
¹⁵ Ratings for each of these probes was made on a 6-pt scale anchored with 1 = strongly disagree, 6 = strongly agree.

¹⁶ Cronbach's Alpha = .847.

¹⁷ This model fits the data well, $df = 4$, $X^2 = 5.3028$, $p = .2576$, $BIC = -16.1047$.

Table 2 Study 2 t-tests

	Good outcome (intentional)	Bad outcome (accidental)	t value	p value
Causation	4.67 (1.29)	3.29 (1.71)	7.02	.000
Evaluation	4.14 (1.07)	3.29 (1.01)	5.90	.000
Intentionality	2.41 (1.52)	1.56 (1.12)	4.56	.000
Gaia	4.39 (1.28)	4.08 (1.34)	1.75	.081



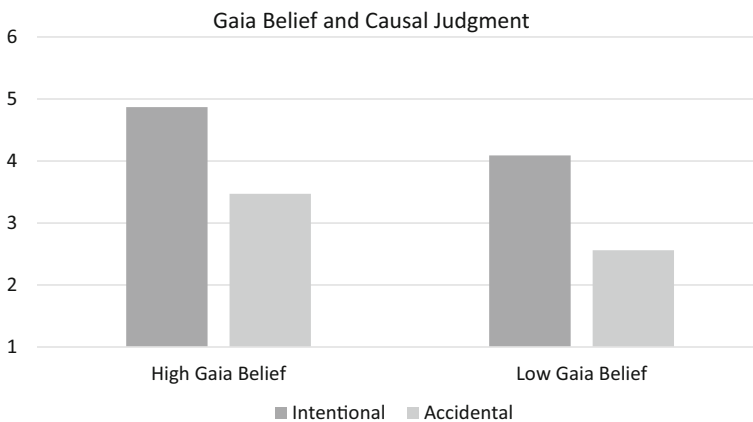
Just as with John, where it was found that judgments of intentionality directly caused causal judgments, so too with Zenite the rock, Gaia Beliefs directly caused causal judgments. This suggests that just as the agentive process directly affects causal judgments for human actions, non-obvious agentive considerations—as measured by Gaia beliefs—directly affect causal judgments for non-living natural things.

One of my main claims is that the agentive process plays a causal role in generating judgements of actual causation. Utilizing Gaia beliefs as a measure for whether the agentive process plays a causal role in generating judgments of actual causation beyond the realm of human action, the results from the causal modeling provide key, strong support for this main claim: background Gaia beliefs do indeed play a direct role in generating causal judgments. Indeed, the results from the causal modeling show that as Gaia belief increases, so too does causal judgment. On its own the results from the causal modeling are sufficient to provide support for the causal hypothesis that the agentive process generates causal judgments and moreover is sufficient to support one prong of the debunking explanation I’ll set out below in Sect. 4. That said, there is a secondary question regarding the way in which the intentional/accidental construal of events works among those with different background Gaia beliefs. Here are two main things we might expect on the present proposal: (1) the impact of the intentional/accidental construal of events will have a greater impact on causal judgments for those displaying High Gaia Belief and (2) when the event is construed intentionally causal judgments should be greater for those displaying High Gaia Belief.¹⁸ Taking those who had an overall score

¹⁸ One might think that for (2) we should instead expect that the intentional/accidental construal of events will never have an impact on causal judgments among those displaying Low Gaia Belief. But as

Table 3 High and Low Gaia Belief and causal judgments

	Good outcome (intentional)	Bad outcome (accidental)	t value	p value	Cohen's d
High Gaia Belief	4.87 (1.05)	3.47 (1.77)	5.70	.000	.962
Low Gaia Belief	4.09 (1.84)	2.56 (1.59)	2.95	.005	.884
t value	2.46	2.17			
p value	.016	.032			
Cohen's d	.521	.541			

**Fig. 2** High and Low Gaia Belief causal judgments in good outcome (intentional) and bad outcome (accidental) cases

between 1 and 3 on Gaia Composite as displaying Low Gaia Belief and those who had an overall score between 4 and 6 on Gaia Composite as displaying High Gaia Belief this is exactly what we find (Table 3).

First, among those who displayed High Gaia Belief and among those that displayed Low Gaia Belief, the intentional and accidental construal of events produced differences in causal judgments (left to right in the above table), though the magnitude of the difference was larger for those displaying High Gaia Belief. So (1) above is supported. Second, within the intentional condition, there were differences in causal judgments among those who displayed Low and High Gaia Belief (top to bottom in the above table), with those displaying High Gaia Belief being significantly more inclined to assign causation. So (2) above is supported. This can be visualized in Fig. 2.

Footnote 18 continued

the results of Kelemen et al. (2013) show, even those who display Low Gaia Belief sometimes display promiscuous teleological tendencies, although to a lesser extent than those with High Gaia Belief (see Rose 2015 for further discussion).

Putting all of this together, the main, key issue at hand—whether the agentive process plays a causal role in generating causal judgments—gains support from the causal modeling results. On its own, this finding provides strong support for the hypothesis that the agentive process plays a role in generating causal judgments. The secondary issue—how the intentional/accidental construal of events affects those with Low and High Gaia Belief—also gains support in that (1) the impact of the intentional/accidental construal of events was greater for those displaying High Gaia Belief and (2) when the event was construed intentionally, those with High Gaia Belief were significantly more inclined to assign causation.

One key limitation of the present study though is that an intentional/accidental construal of events is not varied within cases where the outcome is good or bad. Moreover, it would be useful to know if the agentive process plays a role in generating causal judgments in other kinds of cases. The next study takes up both of these issues.

3.3 Study 3

To extend the pattern of findings that the agentive process generates causal judgments, I decided to run a new set of cases, this time involving a plant. Moreover, since the extant research on the effect of agential considerations on causal judgment has only been conducted with cases where the outcome is negative (e.g., Lagnado and Channon 2008; Channon et al. 2010; Lombrozo 2010), I wanted to look at cases where an intentional/accidental construal is varied within cases with a positive and negative outcome.

Here is an overview of the cases: Suzy discovers a rare plant, called *Cerbolis*. Some of the plants live longer than others. The reason is that *Cerbolis* coats its leaves with a toxin to prevent insects from eating it. Some *Cerbolis* plants produce excess amounts of the toxin and poison themselves. She assigns her intern, Andy, to study the plants and find out why some produce excess amounts of toxin.

Across the cases, I varied whether a person—Andy—or a biological structure—KKM—was the target candidate causal factor. In the cases involving the person, Andy is responsible for administering the toxin to the plant, while in the cases involving the biological structure, KKM is responsible for administering the toxin. I also varied whether the outcome was good (an appropriate amount of the toxin was released and the plant survived) or bad (excess amounts of the toxin were released and the plant was killed) and whether administering too much or the right amount of the toxin was described as intentional or accidental. This resulted in a 2 (Case: KKM, Andy) × 2 (Behavior: Accidental, Intentional) × 2 (Outcome: Good, Bad) design (for full cases see “[Appendix](#)”).

3.3.1 Participants and measures

372 participants were recruited through Amazon’s Mechanical Turk and randomly assigned to one of eight conditions. After reading the cases, participants were given the following probes (in random order):¹⁹

¹⁹ The scales for the probes were the same as those used in Study 1.

Causation Probe: [KKM/Andy] caused the Cerbolis plant to [die/survive].

Evaluation Probe: How would you evaluate [KKM's/Andy's] behavior?

Intentionality Probe: [KKM/Andy] intentionally [killed/saved] Cerbolis.

Finally, in only the cases involving KKM participants were given the same Gaia Belief probes as used in Study 2.

3.3.2 Results

An ANOVA revealed a main effect of Behavior, $F(1, 364) = 17.836$, $p < .001$ $\eta^2 = .047$ but no main effect of Outcome, $F(1, 364) = 2.440$, $p = .119$ or of Case $F(1, 364) = .428$, $p = .514$. However, the main effect of Behavior was qualified by a two-way interaction between Behavior and Outcome, $F(1, 364) = 9.578$, $p = .002$, $\eta^2 = .026$. There were no other significant two way interactions nor was there a significant three way interaction. Below are t-tests for the KKM (Table 4) and Andy (Table 5) cases.

Table 4 T-tests for KKM cases with behavior as IV

	Intentional	Accidental	t value	p value
<i>Good outcome</i>				
Causation	4.61 (1.16)	4.51 (1.33)	.428	.669
Evaluation	4.09 (1.18)	4.04 (.988)	.256	.799
Intentionality	3.77 (1.62)	2.60 (1.49)	3.69	.000
Gaia	4.80 (1.06)	4.29 (1.31)	2.15	.034
<i>Bad outcome</i>				
Causation	4.71 (1.15)	3.80 (1.64)	3.37	.001
Evaluation	2.52 (1.13)	3.48 (1.01)	-4.76	.000
Intentionality	2.66 (1.30)	1.70 (1.15)	4.12	.000
Gaia	4.33 (1.22)	4.40 (1.30)	-.320	.749

Table 5 T-tests for Andy cases with behavior as IV

	Intentional	Accidental	t value	p value
<i>Good outcome</i>				
Causation	4.69 (1.12)	4.47 (1.46)	.777	.440
Evaluation	4.32 (.918)	3.84 (1.01)	2.26	.026
Intentionality	4.62 (1.51)	2.65 (1.68)	5.55	.000
<i>Bad outcome</i>				
Causation	5.05 (1.27)	3.78 (1.94)	3.37	.001
Evaluation	2.64 (1.11)	3.74 (.977)	-4.58	.000
Intentionality	4.15 (1.42)	1.31 (.702)	11.05	.000


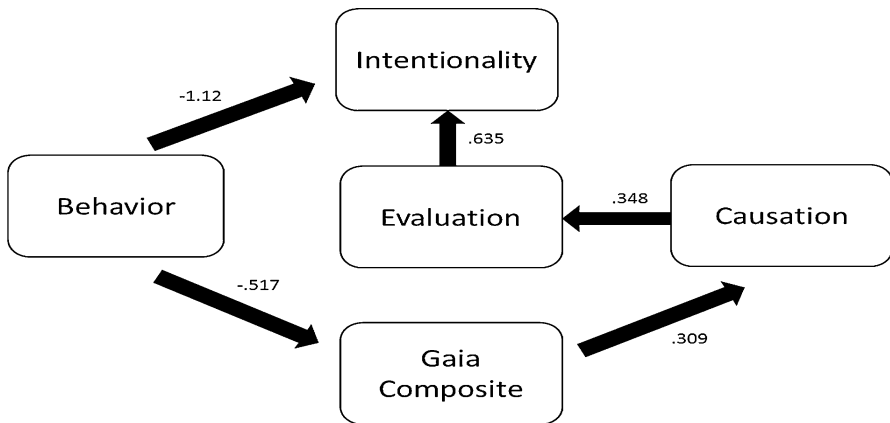


Fig. 3 Causal judgments in non-agent (KKM) cases

Analyzing just the data involving KKM, an ANOVA revealed a main effect of Behavior $F(1, 214) = 7.362, \eta^2 = .034$ and no main effect of Outcome $F(1, 214) = 2.557$. The main effect of Behavior was qualified by a two way interaction with the Outcome $F(1, 214) = 4.524, p = .035, \eta^2 = .021$. This can be seen in Fig. 3.

Given the interaction between Behavior and Outcome for the cases involving KKM, I'll run two separate causal searches, one for the cases where the outcome is good and another one for the cases where the outcome is bad.²⁰ Here is the model for the good outcome cases:²¹

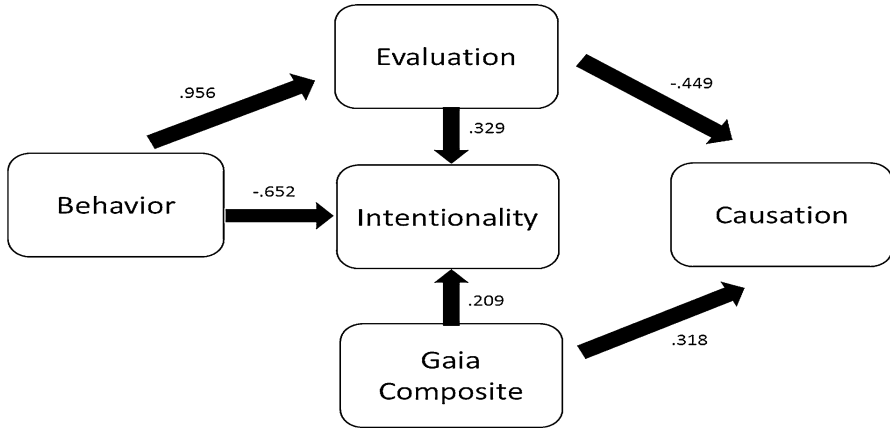


And, here is the model for the bad outcome cases:²²

²⁰ The inter-correlations between the various Gaia Belief probes was high, Cronbach Alpha = .863 for good outcome cases, Cronbach Alpha = .887 and so the items were combined into a single measure to yield Gaia Belief Composite.

²¹ This model fits the data well, $X^2 = 3.7328, df = 5, p = .5885, BIC = -19.2428$.

²² This model fits the data well $X^2 = 7.5765, df = 4, p = .1084, BIC = -11.4032$.



Again, one of my main claims—that the agentic process plays a causal role in generating judgements of actual causation—gains strong support from the causal modeling results. These results show that background Gaia beliefs do indeed play a direct role in generating causal judgments and that as Gaia belief increases, so too does causal judgment. To reiterate, and further emphasize, the results from the causal modeling are, on their own, sufficient to provide support for the causal hypothesis that the agentic process generates causal judgments. Moreover this result is sufficient to support one prong of the debunking explanation (Sect. 4). But, as with Study 2, there is a secondary question about the way in which the intentional/accidental construal of events works among those with different background Gaia beliefs. Again, as with Study 2, two main things are expected: (1) the impact of the intentional/accidental construal of events will have a greater impact on causal judgments for those displaying High Gaia Belief and (2) when the event is construed intentionally causal judgments should be greater for those displaying High Gaia Belief. To examine both (1) and (2) I followed the same procedure in Study 2 for grouping responses into High and Low Gaia Belief. The pattern of findings for those who display High and Low Gaia Belief in the Bad

Table 6 Causal judgments for High and Low Gaia Belief in bad outcome cases

	Intentional	Accidental	t value	p value	Cohen’s d
High Gaia Belief	4.95 (.986)	4.02 (1.67)	3.06	.003	.678
Low Gaia Belief	3.62 (1.76)	3.13 (1.80)	.559	.585	.275
t value	3.03	1.36			
p value	.004	.179			
Cohen’s d	.932	.512			

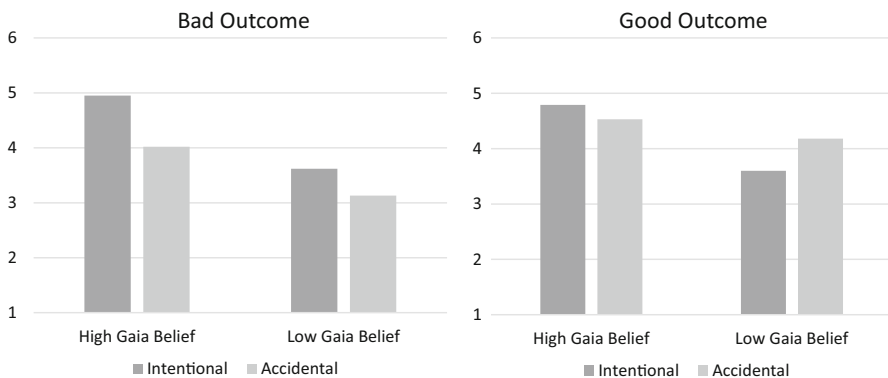
Table 7 Causal judgments for High and Low Gaia Belief in good outcome cases

	Intentional	Accidental	t value	p value	Cohen's d
High Gaia Belief	4.79 (1.09)	4.53 (1.29)	.915	.363	.217
Low Gaia Belief	3.60 (1.51)	4.18 (1.60)	-.684	.505	.372
t value	2.20	.727			
p value	.033	.472			
Cohen's d	.904	.240			

Outcome cases are presented in Table 6, while the pattern of findings for the Good Outcome cases are presented in Table 7.

In the Bad Outcome cases, among those who displayed High Gaia Belief and among those that displayed Low Gaia Belief, the intentional and accidental construal of events produced differences in causal judgments (left to right in Table 6), with the intentional/accidental construal of events having a much greater impact on causal judgments among those with High Gaia Belief. So, as with Study 2, (1) is supported. In addition, those with High Gaia Belief were significantly more inclined to assign causation when the outcome was viewed as being brought about intentionally (top to bottom in Table 6). Again, and as with Study 2, (2) is supported.

In the Good Outcome cases (Table 7), among those who displayed High Gaia Belief and among those that displayed Low Gaia Belief, the intentional and accidental construal of events did not produce differences in causal judgments (left to right in Table 7). So (1) doesn't gain support for the Good Outcome cases (more on this in Sect. 3.4). But there was a significant difference uncovered between those with High and Low Gaia Belief within the intentional condition, with those displaying High Gaia Belief being significantly more likely to assign causation (top to bottom in Table 7). So (2) is supported for these case. These results can also be seen in Fig. 4.

**Fig. 4** Gaia Belief and causal judgment for bad outcome cases (left) and good outcome cases (right)

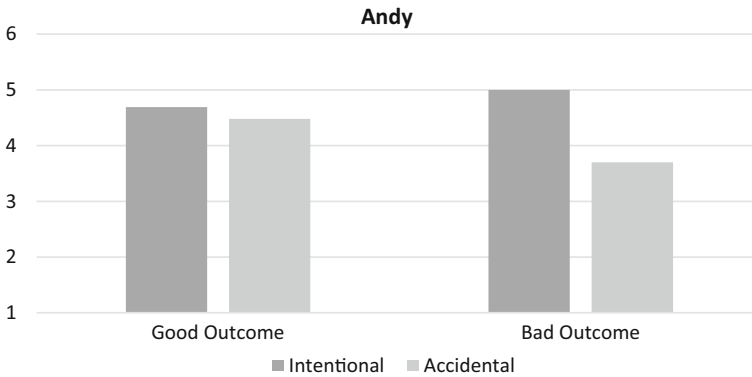



Fig. 5 Causal judgments in agent (Andy) cases

Analyzing the data involving Andy, an ANOVA revealed a main effect of Behavior $F(1, 154) = 10.003, p = .002, \eta^2 = .061$ and no main effect of Outcome $F(1, 154) = .495, p = .483$. The main effect of Behavior was qualified by a two way interaction with the Outcome $F(1, 154) = 4.880, p = .029, \eta^2 = .031$. This can be seen in Fig. 5.


As with the cases involving KKM, since there was an interaction between Behavior and Outcome for the cases involving Andy, I conducted two separate causal searches, one for the cases where the outcome is good and the other for the cases where the outcome is bad. Here is the causal model for the cases where the outcome is good:²³



And here is the causal model for the cases where the outcome is bad:²⁴

²³ This model fits the data well, $X^2 = 1.2709, df = 3, p = .7361, BIC = -11.9125$.

²⁴ This model fits the data well, $X^2 = .3155, df = 2, p = .8541, BIC = -8.3722$.



3.4 Summary of results

One of the crucial, key claims—that the agentic process causes causal judgments—gains strong support from the causal modeling results in Study 3. As with Study 2, explicit judgments of intentionality had no effect on causal judgments for the cases involving the biological mechanism, though the manipulation of whether the behavior was viewed as intentional or accidental, as revealed by differences on this measure, was successful (see Table 4). But interestingly, regardless of whether the outcome was good or bad, Gaia beliefs directly caused causal judgments. In line with Study 2, this suggests that non-obvious agentic considerations directly impact causal judgments beyond the realm of human action. Moreover, the model produced for the negative outcome cases involving the agent was comparable to the model produced for the biological mechanism in the negative outcome cases. This model showed that both intentionality and evaluative judgments have a joint effect on causal judgment, which was similar to the model involving the biological mechanism, where it was found that both Gaia Beliefs and evaluations produced a joint effect on causal judgment. For the positive outcome cases, the models for the agent and the biological mechanism were somewhat different. For the biological mechanism cases, Gaia Beliefs directly caused causal judgment while for the agent cases, intentionality judgments indirectly caused causal judgments via evaluative judgments. Taken together, the results from the causal modeling provide strong support for two of the main hypotheses: namely, that the agentic and evaluative processes play a causal role in generating causal judgments.

The secondary issue—how the intentional/accidental construal of events works among those with different background Gaia Beliefs—also gained support. Just as in Study 2, it was found that, for both the Good and Bad Outcome cases, when the event was construed intentionally (i.e., in the Intentional Conditions) causal judgments were greater for those displaying High Gaia Belief. It was also found that, just as in Study 2, the impact of the intentional/accidental construal of events had a greater impact on causal judgments for those displaying High Gaia Belief, though this result only obtained in the Bad Outcome cases. The main question now is why the intentional/

accidental construal of events did not have an impact on causal judgments among those with High and Low Gaia Belief in the Good Outcome cases.

One reason that this asymmetric effect may be arising is because we have a deep seated, implicit “intentionality bias” where the default is to view behavior as intentional and only by effortfully overriding this bias do we come to view behavior as accidental (Rosset 2008). In negative outcome cases, the intentionality bias may be overridden when doing so would excuse the person or object and “let them off the hook”. In positive outcome cases, since there’s no need to excuse, the implicit intentionality bias is not overridden. If this is right then though we shouldn’t always expect to find differences in causal judgments based on whether the events are construed intentionally or accidentally in positive outcome cases, we might nonetheless expect that those with High Gaia Belief will be more susceptible to the intentionality bias when considering good outcome cases. This is just what we find. Those with High Gaia Belief ($M = 4.68$, $SD 1.18$) were more inclined to assign causation overall than those with Low Gaia Belief ($M = 4.0$, $SD 1.54$), $t(88) = 1.96$, $p = .05$. This suggests that those with High Gaia Belief are indeed more susceptible to the intentionality bias when considering good outcome cases.

In sum: the pattern of results suggest that the influence of agentic considerations on causal attributions is direct, robust and similar regardless of whether causal judgments are made with respect to human actions or non-agents. And in line with previous research, the evaluative process also plays a role in generating folk judgments of actual causation. Importantly, this work has been extended. While one may have thought that the evaluative process is only operative when making causal judgments about human action, the results suggest the evaluative process extends beyond the realm of human action. The reason it extends beyond the realm of human action is because the folk take a perspective on reality whereby it is infused with agency. That is, the evaluative process extends beyond the realm of human action because it is connected to promiscuous teleoementalism. Thus I claim empirical support for the claim that two processes—the evaluative and agentic processes—play a role in generating folk intuitions of actual causation. On to debunking.

4 Debunking folk intuitions of actual causation

Given empirical support for two processes—the evaluative and agentic processes—playing a role in generating folk intuitions of actual causation, I now want to situate the findings within a background discussion of debunking. Having done that, I’ll then discuss the challenge from folk belief, showing how the targeted debunking explanation on offer is fit to meet the challenge.

The specific version of debunking that I will be offering invokes the following two level structure:²⁵

1. S’s belief that P is based on an epistemically defective process.
2. Insofar as S* relies on S’s belief that P as reason to accept P, accepting (1) serves as an undermining defeater for S*’s belief that P.

²⁵ This version of debunking grew out of separate work with Shaun Nichols.

Level one invokes a claim about the causal origins of a belief, where the causal origin of the candidate belief issues from an epistemically defective process. Level one only shows that S's belief is unjustified. It does not yet show that S*'s belief, at Level 2, is unjustified. Prior to learning about S's belief issuing from an epistemically defective process, S*'s belief is *prima facie* justified. After learning that S's belief issues from an epistemically defective process, this acts as an undermining defeater for S*'s belief. This is the second level of debunking.

An illustration. Suppose John suffers from a throbbing headache. He visits a doctor who displays various credentials in her office which attest to her medical expertise. The expert doctor tells John that his throbbing headache calls for special treatment. John must apply lipstick to his forehead to alleviate the headache. Since forming beliefs on the basis of expert testimony is typically a good way of forming beliefs, John's belief is *prima facie* justified. If this were all there were to the story, and thus there were no undefeated defeaters, John's belief might enjoy the status of *ultima facie* justification. But there's more. John learns that the alleged medical expert is actually an expert in astrological medicine. The doctor's belief that applying lipstick to the forehead alleviates headaches issues from an epistemically defective process. Upon learning this, John's belief suffers from an *undermining defeater* (Pollock 1987). This case invokes the above two level structure: (1) the testimony of S does not provide good evidence for S*'s belief that P (since S's belief issues from an epistemically defective process—astrological reasoning) and (2) after S* becomes aware of (1), S*'s belief that P on the basis of (1) suffers from an undermining defeater.

I won't attempt to provide an account of *what* makes a process epistemically defective.²⁶ For my purposes, all that is required is agreement on which processes are epistemically defective. Indeed, even among those who deeply disagree about what makes a process epistemically defective, there is agreement on *which* processes are epistemically defective. For instance, Goldman (1979)—who offers an externalist account of what makes a process epistemically defective—includes the following processes on his list of epistemically defective processes: “confused reasoning, wishful thinking, reliance on emotional attachment, mere hunch or guesswork, and hasty generalization”. Cohen (1984, pp. 282–283) agrees on which processes are epistemically defective, but offers an internalist account of what makes these processes epistemically defective. That said, all that is required is agreement that the evaluative process and agentive process are epistemically defective processes when they generate judgments of actual causation. I take it that these two processes are clear cases of epistemically defective processes. Now to the specific debunking argument for actual causation.

Both the evaluative and agentive processes seem, at least *prima facie*, to be epistemically defective bases for making judgments of actual causation.²⁷ Insofar as folk judgments of actual causation issue from either the evaluative or agentive processes, their judgments of actual causation are *prima facie* unjustified.

²⁶ I follow Kahane (2011) in associating undermining defeaters and debunking explanations (p. 106) which have a two-level structure. But I depart from Kahane in that while Kahane is focused on evolutionary debunking arguments, my focus is only on psychological debunking arguments (Nichols 2014).

²⁷ For more, see Sect. 5.1.

Philosophers who rely on folk intuitions of actual causation as support for a theory of actual causation are faced with an undermining defeater. Putting this together:

1. Folk intuitions of actual causation are based on epistemically defective processes (the evaluative and agentic processes).
2. Insofar as philosophers rely on folk intuitions of actual causation as a reason to accept a view of actual causation, accepting (1) would be an undermining defeater for the philosopher's belief.

To clarify, I don't take this debunking argument to undermine the usefulness of philosophers' intuitions in disputes about actual causation. And I'm not claiming that the philosopher who endorses a view of actual causation on the basis of her own intuitions or on considerations independent of folk intuitions of actual causation suffers from an undermining defeater. Rather, I take the psychological findings to provide the basis for an argument that debunks philosophical views which are based, at least in part, on fitting folk intuitions about actual causation.

I would also clarify that the debunking argument on offer is a two-pronged debunking argument. It's two-pronged since two epistemically defective processes have been put forward: the evaluative and agentic processes. The claim about the epistemically defective processes should thus be read disjunctively. Indeed, it is inclusive since as the empirical evidence suggests, in certain contexts, one or both processes may be generating folk intuitions of actual causation. Thus, the claim is not that *both* processes *always* play a role in generating folk intuitions of actual causation. Sometimes it is one; sometimes it is both.

I would emphasize that I take the two-pronged debunking argument to mark an advantage for the would-be debunker who might have wanted to go it alone on the basis of the evaluative process playing a role in generating folk intuitions of actual causation. This debunker—with only the evaluative process in hand—might be charged with offering a weak debunking argument. Those who would insist on following the folk might simply suggest that those who would debunk folk intuitions of actual causation on basis of the evaluative process alone have only shown that folk intuitions of actual causation are undercut in the realm of human action, where the evaluative process is most likely to be restricted. And so those who would hold that a philosophical theory of actual causation should be beholden to folk intuitions of actual causation might recommend that only those folk intuitions generated in response to events outside the realm of human action are worthy of respecting since—so the suggestion might go—outside the realm of human action, blame does not influence folk intuitions of actual causation.²⁸ But as the above results suggest, this is not the case. The evaluative process extends beyond the realm of human action because it is connected with promiscuous teleoemalism. Those who may have been attracted to debunking folk intuitions of

²⁸ Indeed something along these lines might explain why those who have investigated the role of evaluative considerations in generating intuitions of actual causation have not explicitly argued for a debunking explanation of folk intuitions of actual causation. Likewise, something along these lines might also explain why some have upheld respecting folk intuitions despite the fact that they are generated by evaluative considerations within the realm of human action.

actual causation on basis of the evaluative process alone should thus find the two-pronged debunking argument on offer appealing. It provides a stronger debunking argument: the two-pronged debunking argument extends the evaluative process beyond the realm of human action by connecting it with promiscuous telementalism, which is embodied in the agentive process.

Though the specific two-pronged debunking argument is stronger than a debunking argument based on the evaluative process alone, it is not indiscriminating and thus unfit for meeting the challenge from folk belief. Rather the two-pronged debunking argument on offer is *targeted* and thus fit for meeting the challenge from folk belief. Recall that the challenge was for the revisionist to explain why the folk believe as they do when the resultant theory apparently conflicts with relevant folk beliefs and do so in a way that is not (1) globally self-defeating or (2) locally self-defeating. As for global self-defeat, the two pronged debunking argument locates the mistake in two epistemically defective processes—the evaluate and agentive processes—and thus the two-pronged debunking argument is not premised on any claim about some general inability of the folk to form true beliefs about the world. Neither does the two-pronged debunking argument succumb to local self-defeat. My specific results on folk intuitions of actual causation do not call into question the general usefulness of *philosophers'* intuitions in the target domain of actual causation.²⁹ Indeed, I do not ever appeal to naive teleological reasoning or blame myself in debunking. In this way I hope to have illustrated a stable and targeted strategy for debunking folk intuitions.

Finally, a few clarifications. First, I would note that the two-pronged debunking argument is not aimed at casting a general pox on folk intuitions of actual causation. The claim is not that causal cognition, as a whole, is infected by either of the two epistemically defective processes. Indeed, concerning the evaluative process, Danks et al. (2013) provide evidence that whether moral considerations impact causal judgments depends on whether one learns about causal relations on the basis of experience or on the basis of description, as in the case of thought experiments. Specifically, they present evidence that moral considerations play a significant role in impacting causal intuitions when one learns about causal relations via a description—as in a typical thought experiment—but that moral considerations do not affect causal intuitions when one learns about causal relations via experience, as in observing candidate causes and effects covarying. Thus the evidence suggests that moral considerations significantly impact folk intuitions in the context of learning via description, as in the case of considering thought experiments. Though I'm not aware of any specific empirical evidence on whether the agentive process is operative in causal learning via experience, I suspect that just as with the evaluative process, the agentive process is not operative in all aspects of causal learning. What the evidence does suggest is that the evaluative and agentive processes are operative in some

²⁹ Though my own results do not themselves cast doubt on the usefulness of philosopher's intuitions on these matters, this isn't to say that philosopher's intuitions are useful on these matters. Indeed, when one looks at the range of conflicting claims about what is "intuitive" in the literature one finds a wide range of disagreement among philosopher's, which may cast doubt on whether philosopher's intuitions may be a helpful guide in these matters.

aspects of causal cognition; namely, in those aspects of causal cognition which are invoked when considering thought experiments. Insofar as these processes give rise to folk intuitions of actual causation, they're subject to debunking.

Second, focusing on the role of teleological considerations in causal judgment, I would note that I am not adopting the following extreme view: for all cases, people either do or do not view a causal process in teleological terms. When they do view the process in teleological terms, they view it as being more causal; when they do not view the process in teleological terms, they view it as less causal. This is a mistake. So all folk judgments of actual causation are mistaken.³⁰ I would emphasize that, as previously mentioned, I doubt that teleological considerations play a role in all aspects of causal cognition. As the work of Danks et al. (2013) suggests, there is good reason to suspect that teleology doesn't play a role in causal judgment when learning via experience. I also doubt that teleological considerations play a role in causal perception. Again, what the evidence does suggest is that teleological considerations play a role in some aspects of causal cognition; in particular, teleological considerations play a role in those aspects of causal cognition that are operative in instances of learning via description. That said, even in cases of learning via description—as in typical thought experiments—we don't currently have enough evidence to suggest that teleological considerations always play a role in this aspect of causal cognition. Instead what the extant empirical evidence indicates is that teleological considerations sometimes play a role in generating intuitions about actual causation. So a great deal more empirical work would need to be done before this extreme view gained strong empirical support. Until then, let me reiterate the more moderate view I am endorsing: insofar as folk intuitions of actual causation are generated by the agentive process, they are subject to debunking.

To further clarify, I'm not claiming that the fact that teleological considerations play any role at all in folk intuitions of actual causation is a mistake.³¹ Instead, my view is that the role of teleological considerations in folk intuitions of actual causation is illegitimate in cases where the folk are mistaken to view the events in question in teleological terms. Regarding rocks and clouds, I take it that it is *prima facie* clear that teleology is irrelevant. In other domains, such as the biological domain, whether teleology is legitimate is a contested issue (see e.g., Allen and Bekoff 1994). My purpose here isn't to settle the issue of whether teleology is legitimate in domains where it is contested. Instead, I'm only taking it that teleology is a mistake when people were wrong to attribute it. And to this, the evidence does indicate that the folk take teleological considerations to be relevant in assessing actual causation in connection with events involving rocks. Putting this altogether: I take it that teleological considerations don't always play a role in judgments of actual causation. Nor, for that matter, do teleological considerations play a role in all aspects of causal cognition. But when they do, and when the folk are mistaken to view the events under consideration teleologically, I take it that these intuitions are subject to debunking. The two-pronged debunking argument on offer is targeted—

³⁰ Thanks to an anonymous reviewer for raising this.

³¹ Thanks again to an anonymous reviewer for asking for clarification on this point.

and thus tempered—and so succumbs to neither global nor local self-defeat. It is thus fit for meeting the challenge from folk belief.

5 Objections

I presented evidence—in Sect. 3—that two epistemically defective processes—the evaluative and agentic processes—play a role in generating folk judgments of actual causation. And, in Sect. 4, I argued that these two processes give rise to a two-pronged debunking argument. I now want to briefly consider some natural objections.

5.1 The two processes are epistemically appropriate

The first objection I want to consider is that the two processes which I claim are clearly epistemically defective are not clearly epistemically defective. What the objector wants are some reasons for thinking that these two processes are epistemically defective.

First, I take it that an account of actual causation ought to cohere well with a background scientific picture of the nature of reality.³² According to current scientific methodology, spirited beings—for instance, agentic forces such as Gaia—play no role in understanding the actual world and causal processes within our world. To invoke spirited beings that make things happen in the world for a purpose is to adopt a muddled, pre-scientific, outmoded perspective on the natural world. To insist on following the folk is to buy into a Stone Age metaphysical perspective on the natural world, with all the crudity and superstition that comes along with it.

Second, and concerning the evaluative dimension of folk intuitions of actual causation, one might be attracted to something like the following plausible epistemic principle:

Recommendation: If E is an epistemically appropriate basis for reasoning about P, then one should be willing to recommend that others use E in reasoning about P.

Insofar as one is willing to accept Recommendation, I would only ask whether it would be a wise policy to recommend that one consult their desire to blame in reasoning about actual causation. Imagine, for instance, that judges instructed jury members to consult their desire to blame in reasoning about whether some defendant caused some outcome. Adopting such a policy would likely lead to disastrous consequences, especially for defendants being judged by jury members using such standards. Why? Because consulting ones desire to blame is an epistemically defective basis for reasoning about actual causation.

Third—and again concerning the evaluative process—one might follow Driver (2008) in thinking that it's natural to hold that “someone is morally responsible for

³² Here I join Paul (2012): after drawing on experience to develop a theory, in evaluating it we need to look back at the natural science just in case our ordinary experience of the world conflicts with what our best natural science says about the world. If it does conflict, then often the assumptions based on ordinary experience should be rejected (p. 17).

an event only when that person has caused the event” (p. 423). That is, a natural view is that moral responsibility entails causal responsibility. On this view, causation is determined independent of considerations of moral responsibility. As the evidence suggests, folk judgments run afoul of this reasonable principle.

5.2 Other processes are operative

The next objection is that other, epistemically appropriate processes are operative. There are two versions of this objection. The first is that an epistemically appropriate process actually underwrites the two processes identified here. The idea here is that the two epistemically defective processes which give rise to undermining defeaters are themselves defeated. The objector is thus proposing a *reinstater*, seeking to gain *ultima facie* justification in following the folk. The second is that though these two epistemically defective processes are operative, other “core” causal judgments—which are epistemically appropriate—are operative too and not infected by evaluative or agentive considerations.

A natural proposal in connection with the first version of this objection is that causal and counterfactual selection are driven by a “counterfactual-influences-cause” process (Mandel 2003a). Evaluative or agentive considerations operate by guiding one toward the selection of counterfactuals which in turn guide causal selection.³³ But a range of evidence suggests that counterfactual selection doesn’t necessarily guide causal selection (see e.g., Mandel 2003a, b; Mandel and Lehman 1996, 1998). For instance, the “counterfactual-influences-cause” process makes two predictions: (1) counterfactual judgments should facilitate causal judgments more than vice versa and (2) the correlation between importance ratings for counterfactual and causal selection should be greater when a counterfactual task precedes a causal task than vice versa (Mandel 2003a, p. 421). Against (1), Mandel et al. (2001) and Mandel (2003a) did *not* find evidence that counterfactual selection facilitated causal selection, while against (2) Mandel (2003a) did *not* find that the correlation between importance ratings for counterfactual and causal selection was greater when a counterfactual task preceded a causal task. Taken together, these results suggest, at best, a tenuous connection between counterfactual and causal selection and thus it is doubtful that evaluative or agentive considerations operate by guiding one toward the selection of counterfactuals which in turn guide causal selection. Perhaps the evaluative and agentive processes are correlated with some other reliable indicator of actual causation. But for this objection to work, we’d need some specific proposal to evaluate it and empirical evidence that the alleged epistemically appropriate process is underwriting the operation of the two epistemically defective processes.

For the second version of the objection—that other “core” causal judgments which are epistemically appropriate are operative too—a natural proposal—operating with the image of billiard balls colliding—is that billiard ball kinds of causal judgments are instances of core causal judgments not corrupted by either the

³³ See e.g., Hitchcock and Knobe (2009) for an argument that norm violations guide counterfactual selection which then guide causal selection.

evaluative or agentive processes. Following Michotte (1963) one might think, for instance, that causal judgments in response to motion events involving contact—such as launching or entraining—represent core causal judgments in that they are developmentally and conceptually prior to causal judgments based on agentive considerations. But as Saxe and Carey (2006) write, “the available data are...inconsistent with...[this] claim” (p. 145).

At the earliest ages at which infants show sensitivity to launching or entraining events, infants also show a keen sensitivity to agentive considerations. With entraining, Leslie (1984) presents evidence that infants attend to contact relations between a hand and inanimate object and view a hand and inanimate object moving together as causally interacting but do not view two inanimate objects moving together as causally interacting. For launching, Saxe et al. (2005) and Saxe et al. (2007) present evidence that infants infer a hidden agent as the source of and primary cause of an inanimate object being set in motion. Muentener and Carey (2010) showed that when a train approaches a box and the box collapses, infants do not differentiate between cases where the train and box are or are not in contact but yet they do differentiate between contact and non-contact cases when the train is replaced by an agent (a puppet). And they are surprised when the agent contacts the box and the box doesn't collapse. Taken together the candidate core causal judgments—embodied in Michottian billiard ball causation—seem to be influenced by agentive considerations at the earliest ages at which infants show sensitivity to launching or entraining events. Concerning the billiard ball model, and as Margaret Mead (1932) observed among the Manus people: “if a stone falls suddenly in the brush near an adult, he will usually mutter ‘a spirit’” (p. 118).

6 Conclusion

Philosophers often invoke the mantle of commonsense when evaluating theories of actual causation. But, if a philosophical theory is to be measured by its fit with commonsense, then it seems that empirically discerning why the folk believe as they do will help in deciding whether measuring a theory by its fit with commonsense is a wise policy.

I presented a range of evidence which suggests that two processes—the evaluative and agentive processes—are involved in generating folk judgments of actual causation. In light of the empirical evidence, I argued for a two-pronged debunking explanation, which operates at two levels. At level one, just as the medical astrologer's belief that applying lipstick to the forehead alleviates headaches issues from an epistemically defective process so too folk intuitions of actual causation—insofar as they are based on either the desire to blame or primitive teleological considerations—issue from an epistemically defective process. Level two invoked an undermining defeater. Just as John's belief, after learning that it was based on the testimony of a medical astrologer, suffers from an undermining defeater, so too the philosopher who relies on folk intuitions of actual causation as a reason to accept a theory of actual causation would suffer from an undermining defeater.

Taken together, I hold that discussion over actual causation should be liberated from any demanded conformity with folk intuitions: the revisionist should not be compelled to square her account with the verdicts of the folk. In the dispute over actual causation, folk intuitions deserve to be rejected. Thus, absent further empirical evidence, it seems that measuring a theory of actual causation by its fit with folk intuitions is not a wise policy.

Acknowledgments I would like to thank Wesley Buckwalter, Elizabeth Camp, Georgi Gardiner, Alvin Goldman, Josh Knobe, Shaun Nichols, Jonathan Schaffer, Stephen Stich, John Turri and an anonymous reviewer for helpful comments on previous versions of this paper.

Appendix

Study 1

John, bad outcome

John is an ecologist, studying a rare bird called Cantup. Cantups need a rare, essential nutrient, Keterine, which they receive only from eating Weeble worms. Indeed, if the Cantups do not receive Keterine, they will die. Weeble worms, however, do not naturally produce Keterine. Rather, Weeble worms receive Keterine by feeding from a rock, Zenite, which produces the rare, essential nutrient.

Zenite produces Keterine through a chain of chemical reactions which is initiated and facilitated only by adsorbing heat. For this reason, John keeps Zenite under a special heat lamp. The chemicals involved in producing Keterine are densely concentrated in the upper surface of Zenite, since this is the area that is directly heated by the lamp.

One day the power goes out in the whole town and so the heat lamp shuts off. This prevents the upper surface of Zenite from being heated. John knows that he must act quickly or else the Cantups will begin dying. He notices that the bottom of the glass case containing Zenite is still very warm. The only option he has is to turn Zenite over so that the densely concentrated chemicals can absorb heat from the bottom of the glass case. So, he turns Zenite over.

Unfortunately, there is a volatile reaction between the chemicals and heat from the bottom of the glass case. But John is completely unaware that this volatile reaction has occurred. The Weeble worms continue feeding from Zenite. The Cantups continue to eat the Weeble worms and John begins to realize that something went wrong as the Cantups slowly start to die.

John, good outcome

John is an ecologist, studying a rare bird called Cantup. Cantups need a rare, essential nutrient, Keterine, which they receive only from eating Weeble worms. Indeed, if the Cantups do not receive Keterine, they will die. Weeble worms,

however, do not naturally produce Keterine. Rather, Weeble worms receive Keterine by feeding from a rock, Zenite, which produces the rare, essential nutrient.

Zenite produces Keterine through a chain of chemical reactions which is initiated and facilitated only by adsorbing heat. For this reason, John keeps Zenite under a special heat lamp. The chemicals involved in producing Keterine are densely concentrated in the upper surface of Zenite, since this is the area that is directly heated by the lamp.

One day the power goes out in the whole town and so the heat lamp shuts off. This prevents the upper surface of Zenite from being heated. John knows that he must act quickly or else the Cantups will begin dying. He notices that the bottom of the glass case containing Zenite is still very warm. The only option he has is to turn Zenite over so that the densely concentrated chemicals can absorb heat from the bottom of the glass case. So, he turns Zenite over.

The Weeble worms continue feeding from Zenite. The Cantups continue to eat the Weeble worms and do not die.

Rock, bad outcome

In South Africa, there is a rare bird called Cantup. Cantups need a rare, essential nutrient, Keterine, which they receive only from eating Weeble worms. Indeed, if the Cantups do not receive Keterine, they will die. Weeble worms, however, do not naturally produce Keterine. Rather, Weeble worms receive Keterine by feeding from a rock, Zenite, which produces the rare, essential nutrient.

Zenite produces Keterine through a chain of chemical reactions which is initiated and facilitated only by adsorbing heat. For this reason, the chemicals involved in producing Keterine are densely concentrated in the upper surface of Zenite, since this is the area that is directly heated by sunlight.

One day a tree branch falls on top of Zenite. The branch prevents the sunlight from heating the upper surface of Zenite. However, the densely concentrated chemicals located in the upper surface of Zenite migrate toward the lower surface of Zenite and absorb heat from the ground. Unfortunately, there is a volatile reaction between the chemicals and heat from the ground. The Weeble worms continue feeding from Zenite. The Cantups continue to eat the Weeble worms and slowly start to die.

Rock, good outcome

In South Africa, there is a rare bird called Cantup. Cantups need a rare, essential nutrient, Keterine, which they receive only from eating Weeble worms. Indeed, if the Cantups do not receive Keterine, they will die. Weeble worms, however, do not naturally produce Keterine. Rather, Weeble worms receive Keterine by feeding from a rock, Zenite, which produces the rare, essential nutrient.

Zenite produces Keterine through a chain of chemical reactions which is initiated and facilitated only by adsorbing heat. For this reason, the chemicals involved in producing Keterine are densely concentrated in the upper surface of Zenite, since this is the area that is directly heated by sunlight.

One day a tree branch falls on top of Zenite. The branch prevents the sunlight from heating the upper surface of Zenite. However, the densely concentrated chemicals located in the upper surface of Zenite, migrate toward the lower surface of Zenite and absorb heat from the ground. The Weeble worms continue feeding from Zenite. The Cantups continue to eat the Weeble worms and do not die.

Study 2

Cases Rock, bad outcome and Rock, good outcome above

Study 3

Andy, intentional, good outcome

Suzy is a botanist. She recently discovered a rare plant, Cerbolis. She noticed that some Cerbolis plants lived much longer than others, even though the plants were subject to the same environmental conditions. So she studied the plants to try and determine why some were living longer than others.

Suzy discovered that Cerbolis emits a toxin which coats its leaves. The toxin prevents insects from eating the plant. She noticed that the longer living plants were producing much lower doses of the toxin than the plants that were suffering premature death. The plants that were producing excess amounts of the toxin seemed to be poisoning themselves. She noticed that one of the Cerbolis plants was not producing any of the toxin. So, she gave it to one of her interns, Andy, to experiment with in order to determine what level of the toxin is appropriate to ensure the plants survival.

Andy puts the Cerbolis plant in an aquarium and runs a hose into it. The hose administers the toxin and is connected to a button. When Andy wants to administer the toxin he simply presses the button. The amount of toxin released is determined by how long Andy holds down the button.

Andy places some insects in the aquarium and plans to administer some of the toxin when they begin swarming the plant. When the insects begin swarming the plant, Andy pushes the button which administers some of the toxin. An appropriate amount of the toxin is released and the plant is saved from being infested by insects.

Andy, accidental, good outcome

Suzy is a botanist. She recently discovered a rare plant, Cerbolis. She noticed that some Cerbolis plants lived much longer than others, even though the plants were subject to the same environmental conditions. So she studied the plants to try and determine why some were living longer than others.

Suzy discovered that Cerbolis emits a toxin which coats its leaves. The toxin prevents insects from eating the plant. She noticed that the longer living plants were producing much lower doses of the toxin than the plants that were suffering premature death. The plants that were producing excess amounts of the toxin seemed to be poisoning themselves. She noticed that one of the Cerbolis plants was

not producing any of the toxin. So, she gave it to one of her interns, Andy, to experiment with in order to determine what level of the toxin is appropriate to ensure the plants survival.

Andy puts the Cerbolis plant in an aquarium and runs a hose into it. The hose administers the toxin and is connected to a button. When Andy wants to administer the toxin he simply presses the button. The amount of toxin released is determined by how long Andy holds down the button.

Andy places some insects in the aquarium and plans to administer some of the toxin when they begin swarming the plant. When the insects begin swarming the plant, Andy walks over to the aquarium so that he can push the button to administer some of the toxin. As he is walking over he suffers a mild stroke and becomes confused and scrambled. Surprisingly, he pushes the button which administers some of the toxin. An appropriate amount of the toxin is released and the plant is saved from being infested by insects. Given that Andy was completely confused and scrambled from the stroke, it was a complete accident that he allowed appropriate amounts of the toxin to be released.

Andy, intentional, bad outcome

Suzu is a botanist. She recently discovered a rare plant, Cerbolis. She noticed that some Cerbolis plants lived much longer than others, even though the plants were subject to the same environmental conditions. So she studied the plants to try and determine why some were suffering a premature death.

Suzu discovered that Cerbolis emits a toxin which coats its leaves. The toxin prevents insects from eating the plant. She noticed that the longer living plants were producing much lower doses of the toxin than the plants that were suffering premature death. The plants that were producing excess amounts of the toxin seemed to be poisoning themselves. She noticed that one of the Cerbolis plants was not producing any of the toxin. So, she gave it to one of her interns, Andy, to experiment with in order to determine what level of the toxin is appropriate to ensure the plants survival.

Andy puts the Cerbolis plant in an aquarium and runs a hose into it. The hose administers the toxin and is connected to a button. When Andy wants to administer the toxin he simply presses the button. The amount of toxin released is determined by how long Andy holds down the button.

Andy places some insects in the aquarium and plans to administer some of the toxin when they begin swarming the plant. When the insects begin swarming the plant, Andy pushes the button which administers some of the toxin. He keeps the button pressed, administering a steady flow of the toxin. Large amounts of the toxin are released and Andy continues to keep the button pressed as the plant slowly starts to die.

Andy, accidental, bad outcome

Suzu is a botanist. She recently discovered a rare plant, Cerbolis. She noticed that some Cerbolis plants lived much longer than others, even though the plants were

subject to the same environmental conditions. So she studied the plants to try and determine why some were suffering a premature death.

Suzy discovered that Cerbolis emits a toxin which coats its leaves. The toxin prevents insects from eating the plant. She noticed that the longer living plants were producing much lower doses of the toxin than the plants that were suffering premature death. The plants that were producing excess amounts of the toxin seemed to be poisoning themselves. She noticed that one of the Cerbolis plants was not producing any of the toxin. So, she gave it to one of her interns, Andy, to experiment with in order to determine what level of the toxin is appropriate to ensure the plants survival.

Andy puts the Cerbolis plant in an aquarium and runs a hose into it. The hose administers the toxin and is connected to a button. When Andy wants to administer the toxin he simply presses the button. The amount of toxin released is determined by how long Andy holds down the button.

Andy places some insects in the aquarium and plans to administer some of the toxin when they begin swarming the plant. When the insects begin swarming the plant, Andy walks over to the aquarium so that he can push the button which administers the toxin. As he is walking over he suffers a mild stroke and becomes confused and scrambled. He pushes the button which administers some of the toxin. Large amounts of the toxin are released and the plant slowly starts to die. Given that Andy was completely confused and scrambled from the stroke, it was a complete accident that large amounts of the toxin were released.

KKM, intentional, good outcome

Suzy is a botanist. She recently discovered a rare plant, Cerbolis. She noticed that some Cerbolis plants lived much longer than others, even though the plants were subject to the same environmental conditions. So she studied the plants to try and determine why some were living longer than others.

Suzy discovered that Cerbolis emits a toxin which coats its leaves. The toxin prevents insects from eating the plant. She noticed that the longer living plants were producing much lower doses of the toxin than the plants that were suffering premature death. The plants that were producing excess amounts of the toxin seemed to be poisoning themselves. To determine why some of the plants are producing appropriate amounts of the toxin, she gave some of them to her friend Andy, who is a molecular biologist.

While studying Cerbolis, Andy noticed that the plants house a large store of the toxin. The amount of toxin released is regulated by KKM. KKM works as a “gate” which regulates the amount of toxin released. So Andy thinks that KKM may have something to do with some of the Cerbolis emitting appropriate doses of the toxin.

Looking more closely at one of the Cerbolis plants, Andy notices that KKM is carefully regulating the release of the toxin. As insects start swarming the plant, KKM releases some of the toxin. As a result, the insects leave the plant and it does not die.

Andy tries to figure out why KKM is allowing appropriate amounts of the toxin to be released. But, he can find no factor which contributed to KKM’s allowing

appropriate amounts of toxin to be released. Andy is baffled and tells Suzy that it looks to him like KKM just suddenly decided to release appropriate amounts of the toxin, as if it was trying to save the plant from being infested by insects.

KKM, accidental, good outcome

Suzy is a botanist. She recently discovered a rare plant, Cerbolis. She noticed that some Cerbolis plants lived much longer than others, even though the plants were subject to the same environmental conditions. So she studied the plants to try and determine why some were living longer than others.

Suzy discovered that Cerbolis emits a toxin which coats its leaves. The toxin prevents insects from eating the plant. She noticed that the longer living plants were producing much lower doses of the toxin than the plants that were suffering premature death. The plants that were producing excess amounts of the toxin seemed to be poisoning themselves. To determine why some of the plants are producing appropriate amounts of the toxin, she gave some of them to her friend Andy, who is a molecular biologist.

While studying Cerbolis, Andy noticed that the plants house a large store of the toxin. The amount of toxin released is regulated by KKM. KKM works as a “gate” which regulates the amount of toxin released. So Andy thinks that KKM may have something to do with some of the Cerbolis emitting appropriate doses of the toxin.

Looking more closely at one of the Cerbolis plants, Andy notices that KKM is carefully regulating the release of the toxin. As insects start swarming the plant, KKM releases some of the toxin. As a result, the insects leave the plant and it does not die.

Andy tries to figure out why KKM is allowing appropriate amounts of the toxin to be released. He finds that a chemical—DD12—collided with KKM. After the collision, KKM appeared to be confused and scrambled, unable to regulate the release of the toxin. Andy tells Suzy that it looks to him like KKM wasn't trying to save the plant from being infested by insects: the collision made KKM confused and scrambled. It was a complete accident that KKM allowed appropriate amounts of the toxin to be released.

KKM, intentional, bad outcome

Suzy is a botanist. She recently discovered a rare plant, Cerbolis. She noticed that some Cerbolis plants lived much longer than others, even though the plants were subject to the same environmental conditions. So she studied the plants to try and determine why some were suffering a premature death.

Suzy discovered that Cerbolis emits a toxin which coats its leaves. The toxin prevents insects from eating the plant. She noticed that the longer living plants were producing much lower doses of the toxin than the plants that were suffering premature death. The plants that were producing excess amounts of the toxin seemed to be poisoning themselves. To determine why some of the plants are producing excess amounts of the toxin, she gave some of them to her friend Andy, who is a molecular biologist.

While studying Cerbolis, Andy noticed that the plants house a large store of the toxin. The amount of toxin released is regulated by KKM. KKM works as a “gate” which regulates the amount of toxin released. So Andy thinks that KKM may have something to do with some of the Cerbolis emitting large doses of the toxin.

Looking more closely at one of the Cerbolis plants, Andy notices that KKM is carefully regulating the release of the toxin. But suddenly KKM just stops regulating the release of the toxin. As a result, the plant slowly starts to die.

Andy tries to figure out why KKM is allowing large amounts of the toxin to be released. But, he can find no factor which contributed to KKM’s allowing the large amounts of toxin to be released. Andy is baffled and tells Suzy that it looks to him like KKM just suddenly decided to release large amounts of the toxin, as if it was trying to kill the plant.

KKM, accidental, bad outcome

Suzy is a botanist. She recently discovered a rare plant, Cerbolis. She noticed that some Cerbolis plants lived much longer than others, even though the plants were subject to the same environmental conditions. So she studied the plants to try and determine why some were suffering a premature death.

Suzy discovered that Cerbolis emits a toxin which coats its leaves. The toxin prevents insects from eating the plant. She noticed that the longer living plants were producing much lower doses of the toxin than the plants that were suffering premature death. The plants that were producing excess amounts of the toxin seemed to be poisoning themselves. To determine why some of the plants are producing excess amounts of the toxin, she gave some of them to her friend Andy, who is a molecular biologist.

While studying Cerbolis, Andy noticed that the plants house a large store of the toxin. The amount of toxin released is regulated by KKM. KKM works as a “gate” which regulates the amount of toxin released. So Andy thinks that KKM may have something to do with some of the Cerbolis emitting large doses of the toxin.

Looking more closely at one of the Cerbolis plants, Andy notices that KKM is carefully regulating the release of the toxin. But suddenly KKM just stops regulating the release of the toxin. As a result, the plant slowly starts to die.

Andy tries to figure out why KKM is allowing large amounts of the toxin to be released. He finds that a chemical—DD12—collided with KKM. After the collision, KKM appeared to be confused and scrambled, unable to regulate the release of the toxin. Andy tells Suzy that it looks to him like KKM wasn’t trying to kill the plant: the collision made KKM confused and scrambled. It was a complete accident that KKM allowed large amounts of the toxin to be released.

References

- Alicke, M. (1992). Culpable causation. *Journal of Personality and Social Psychology*, 63, 368–378.
Alicke, M. (2000). Culpable control and the psychology of blame. *Psychological Bulletin*, 126, 556–574.

- Alicke, M., & Rose, D. (2010). Culpable control or moral concepts? *Behavioral and Brain Sciences*, *33*, 330–331.
- Alicke, M., Rose, D., & Bloom, D. (2011). Causation, norm violations and culpable control. *Journal of Philosophy*, *108*, 670–696.
- Allen, C. (2009). *Teleological notions in biology*. The Stanford Encyclopedia of Philosophy. <http://plato.stanford.edu/archives/win2009/entries/teleology-biology/>.
- Allen, C. & Bekoff, M. (1994). Function, natural design and animal behavior: Philosophical and ethnological considerations. In N.S. Thompson (Ed.), *Perspectives in Ethology*, Vol. 11: Behavioral Design (pp. 1–47). New York: Plenum Press.
- Beebe, H. (2004). Causing and nothingness. In L. A. Paul, E. J. Hall & J. Collins (Eds.), *Causation and counterfactuals* (pp. 291–308). Cambridge: MIT Press.
- Bloom, P. (2007). Religion is natural. *Developmental Science*, *10*, 147–151.
- Channon, S., Lagnado, D. A., Drury, H., Matheson, E., Fitzpatrick, S., Shieff, C., et al. (2010). Causal reasoning and intentionality judgments after frontal brain lesions. *Social Cognition*, *28*, 509–522.
- Chickering, D. (2002). Optimal structure identification with greedy search. *Journal of Machine Learning Research*, *3*, 507–554.
- Clough, E. E., & Wood-Robinson, C. (1985). How secondary students interpret instances of biological adaptation. *Journal of Biological Education*, *19*, 125–130.
- Cohen, S. (1984). Justification and truth. *Philosophical Studies*, *46*, 279–295.
- Danks, D., Rose, D., & Machery, E. (2013). Demoralizing Causation. *Philosophical Studies*, *171*, 1–27.
- Demastes, S., Settlage, J., & Good, R. (1995). Students' conceptions of natural selection and its role in evolution: Cases of replication and comparison. *Journal of Research in Science Teaching*, *32*, 535–550.
- Dowe, P. (2000). *Physical causation*. Cambridge: Cambridge University Press.
- Dowe, P. (2004). Causes are physically connected to their effects: Why preventers and omissions are not causes. In C. Hitchcock, (Ed.) *Contemporary debates in philosophy of science* (pp. 189–196). Malden: Blackwell.
- Driver, J. (2008). Attributions of causation and moral responsibility. In W. Sinnott-Armstrong (Ed.), *Moral psychology* (Vol. 2, pp. 423–439). MIT Press.
- Evans, E. M., Spiegel, A. N., Gram, W., Frazier, B. N., Tare, M., Thompson, S., et al. (2010). A conceptual guide to natural history museum visitors understanding of evolution. *Journal of Research in Science Teaching*, *47*, 326–353.
- Galli, L. G., & Meinardi, E. (2011). The role of teleological thinking in learning the Darwinian model of evolution. *Evolution Education Outreach*, *4*, 145–152.
- Goldman, A. (1979). What is justified belief? In G. Pappas (Ed.), *Justification and knowledge*. Dordrecht: Reidel.
- Gregory, T. R. (2009). Understanding natural selection: Essential concepts and common misconceptions. *Evolution Education and Outreach*, *2*, 156–175.
- Halpern, J. Y., & Hitchcock, C. (2015). Graded causation and defaults. *The British Journal for the Philosophy of Science*, *66*(2), 413–457.
- Hirsch, E. (2002). Against revisionary ontology. *Philosophical Topics*, *30*, 103–127.
- Hitchcock, C. (2007). Prevention, preemption, and the principle of sufficient reason. *Philosophical Review*, *116*, 495–532.
- Hitchcock, C., & Knobe, J. (2009). Cause and norm. *Journal of Philosophy*, *106*, 587–612.
- Kahane, G. (2011). Evolutionary debunking arguments. *Nous*, *45*, 103–125.
- Kampourakis, K., & Zogza, V. (2008). Students' intuitive explanations of the causes of homologies and adaptations. *Science & Education*, *17*, 27–47.
- Kelemen, D. (1999a). The scope of teleological thinking in preschool children. *Cognition*, *70*, 241–272.
- Kelemen, D. (1999b). Why are rocks pointy? Children's preference for teleological explanations of the natural world. *Developmental Psychology*, *35*, 1440–1452.
- Kelemen, D. (2004). Are children "Intuitive Theists"? Reasoning about purpose and design in nature. *Psychological Science*, *15*, 295–301.
- Kelemen, D. (2012). Teleological minds: How natural intuitions about agency and purpose influence learning about evolution. In K. S. Rosengren, S. K. Brem, E. M. Evans, & G. M. Sinatra (Eds.), *Evolution challenges: Integrating research and practice in teaching and learning about evolution*. Oxford: Oxford University Press.
- Kelemen, D., & DiYanni, C. (2005). Intuitions about origins: Purpose and intelligent design in children's reasoning about nature. *Journal of Cognition and Development*, *6*, 3–31.

- Kelemen, D., & Rosset, E. (2009). The human function compunction: Teleological explanation in adults. *Cognition, 111*, 138–143.
- Kelemen, D., Rottman, J., & Seston, R. (2013). Professional physical scientists display tenacious teleological tendencies: Purpose-based reasoning as a cognitive default. *Journal of Experimental Psychology: General, 142*, 1074–1083.
- Korman, D. (2009). Eliminativism and the challenge from folk belief. *Noûs, 43*, 242–264.
- Lagnado, D. A., & Channon, S. (2008). Judgments of cause and blame: The influence of intentionality and foreseeability. *Cognition, 108*, 754–770.
- Leslie, A. (1984). Infant perception of a manual pick-up event. *British Journal of Developmental Psychology, 2*, 19–32.
- Lewis, D. (1986). *Philosophical papers* (Vol. II). Oxford: Oxford University Press.
- Livengood, J., & Machery, E. (2007). The folk probably don't think what you think they think: Experiments on causation by absence. *Midwest Studies in Philosophy, 31*, 107–127.
- Lombrozo, T. (2010). Causal-explanatory pluralism: How intentions, functions, and mechanisms influence causal ascriptions. *Cognitive Psychology, 61*, 303–332.
- Lombrozo, T., Kelemen, D., & Zaitchik, D. (2007). Inferring design: Evidence for a preference for teleological explanation in patients with Alzheimer's disease. *Psychological Science, 18*, 999–1006.
- Mandel, D. (2003a). Judgment dissociation theory: An analysis of differences in causal, counterfactual and covariational reasoning. *Journal of Experimental Psychology: General, 132*, 419–434.
- Mandel, D. (2003b). Effect of counterfactual Thinking on causal judgments. *Thinking & Reasoning, 9*, 245–265.
- Mandel, D. R., & Lehman, D. R. (1996). Counterfactual thinking and ascriptions of cause and preventability. *Journal of Personality and Social Psychology, 71*, 450–463.
- Mandel, D. R., & Lehman, D. R. (1998). Integration of contingency information in judgments of cause, covariation, and probability. *Journal of Experimental Psychology: General, 127*, 269–285.
- Mandel, D. R., N'gbala, A., & Bonnefon, J. F. (2001). Tests of facilitation in causal and counterfactual judgment. Unpublished raw data.
- McDermott, M. (1995). Redundant causation. *British Journal for the Philosophy of Science, 46*, 523–544.
- McGrath, S. (2005). Causation by omission: A dilemma. *Philosophical Studies, 123*(1), 125–148.
- Mead, M. (1932). An investigation of the thought of primitive children, with a special reference to animism. *Journal of the Royal Anthropological Institute of Great Britain and Ireland, 62*, 173–190.
- Meek, C. (1997). Graphical models: Selecting causal and statistical models. Ph.D. Thesis, Carnegie Mellon University.
- Mellor, D. H. (1995). *The facts of causation*. London and New York: Routledge.
- Menzies, P. (1996). Probabilistic causation and the pre-emption problem. *Mind, 105*, 85–117.
- Menzies, P. (2009). Platitudes and counterexamples. In H. Beebe, P. Menzies, & C. Hitchcock (Eds.), *The Oxford handbook of causation* (pp. 341–367). Oxford: Oxford University Press.
- Michotte, A. (1963). *The Perception of Causality* (T. R. Miles, & E. Miles, Trans). London: Methuen. (Originally published in 1946).
- Moore, R., Mitchell, G., Bally, R., Inglis, M., Day, J., & Jacobs, D. (2002). Undergraduates understanding of evolution: Ascription of agency as a problem for student learning. *Journal of Biological Education, 36*, 65–71.
- Muentener, P., & Carey, S. (2010). Infants' causal representations of state change events. *Cognitive Psychology, 61*, 63–86.
- Nichols, S. (2014). Process debunking and ethics. *Ethics, 124*(4), 727–749.
- Paul, L. A. (2012). Metaphysics as modeling: The Handmaiden's tale. *Philosophical Studies, 160*, 1–29.
- Paul, L. A., & Hall, N. (2013). *Causation: A users guide*. Oxford: Oxford University Press.
- Pollock, J. (1987). Defeasible reasoning. *Cognitive Science, 11*, 481–518.
- Rose, D. (2015). Persistence through function preservation. *Synthese, 192*(1), 97–146.
- Rose, D., Livengood, J., Sytsma, J., & Machery, E. (2011). Deep trouble for the deep self. *Philosophical Psychology, 25*, 629–646.
- Rose, D., & Nichols, S. (2013). The lesson of bypassing. *Review of Philosophy and Psychology, 4*, 599–619.
- Rosset, E. (2008). It's no accident: Our bias for intentional explanations. *Cognition, 108*, 771–780.
- Saxe, R., & Carey, S. (2006). The perception of causality in infancy. *Acta Psychologica, 123*, 144–165.
- Saxe, R., Tenenbaum, J., & Carey, S. (2005). Secret agents: Inferences about hidden causes by 10- and 12-month-old infants. *Psychological Science, 16*, 995–1001.

- Saxe, R., Tzelnic, T., & Carey, S. (2007). Knowing who dunnit: Infants identify the causal agent in an unseen causal interaction. *Developmental Psychology*, *43*, 149–158.
- Schaffer, J. (2000). Causation by disconnection. *Philosophy of Science*, *67*, 285–300.
- Schaffer, J. (2004). Causes need not be physically connected to their effects: The case for negative causation. In Hitchcock (Ed.), *Contemporary debates in philosophy of science* (pp. 197–216). Malden: Blackwell.
- Schaffer, J. (2005). Contrastive causation. *Philosophical Review*, *114*, 327–358.
- Stich, S. (2013). Do different groups have different epistemic intuitions? A Reply to Nagel. *Philosophy and Phenomenological Research*, *87*, 151–178.
- Sytsma, J., Livengood, J., & Rose, D. (2012). Two types of typicality: Rethinking the role of statistical typicality in ordinary causal attributions. *Studies in History and Philosophy of Science Part C*, *43*, 814–820.
- Turri, J., & Blouw, P. (2015). Excuse validation: A study in rule-breaking. *Philosophical studies*, *172*(3), 615–634.