PROMISCUOUS TELEOLOGY AND FOLK METAPHYSICS

by

DAVID ROSE

A Dissertation submitted to the

Graduate School-New Brunswick

Rutgers, The State University of New Jersey

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

Graduate Program in Philosophy

written under the direction of

Jonathan Schaffer

and approved by

_____

_____

_____

_____

New Brunswick, New Jersey

May, 2017

ABSTRACT OF THE DISSERTATION

Promiscuous Teleology and Folk Metaphysics

By DAVID ROSE

Dissertation Director:
Jonathan Schaffer

The primary aim of my dissertation is to set out one important role that cognitive science can play in debates in metaphysics. The focus of my dissertation is the appeal to intuitions in metaphysics. In certain metaphysical disputes, we often see common sense or folk intuitions being invoked in evaluating competing metaphysical theories. But when metaphysical theories are charged with violating common sense or folk intuitions, the charges tend to be met by claiming that the target theory is not, after all, at odds with what the folk think on the matter. As a result, we often end up with conflicting claims about what the folk think.

Here is one of the key ways that I see cognitive science contributing to metaphysics: Cognitive science can enter into the discussion by helping to identify the content of the relevant folk view under consideration. But the usefulness of cognitive science to metaphysics doesn't end there. Having identified the content of common sense on the target issue, a question remains as to how much weight, if any, should be placed on the folk view under consideration in evaluating target metaphysical theories. In other words: does the relevant folk view deserve to be respected or rejected in evaluating metaphysical theories? Empirical evidence from cognitive science can enter again here and play an

important role in evaluating whether the relevant aspect of the folk view is respectable or based on an unreliable or otherwise poor cognitive process. Taken together, one of the key ways that cognitive science can contribute to metaphysics is by identifying what the relevant folk view under consideration is and whether the relevant folk intuitions deserve to be taken seriously.

I focus on metaphysical disputes over composition, persistence and causation and the role of ordinary metaphysical beliefs in serving as a constraint on theories in these areas. Drawing on work in psychology, most notably the work of Deborah Kelemen, I find that the folk take a promiscuous teleological outlook on reality: all of nature—every rock and cloud—is viewed as being infused with agency and purpose. I bring the background psychological evidence into the discussion and also conduct various studies of my own, providing a range of evidence suggesting that this tendency toward promiscuous teleological thinking plays a significant role in folk intuitions about composition, persistence and causation. In light of this, I argue that folk intuitions about these matters deserve to be dismissed, that there is a debunking explanation for these folk intuitions. In short: metaphysical theories of composition, persistence and causation should not be beholden to folk intuitions since the folk view on these matters is tied into a primitive teleological view of nature.

I take myself to have illustrated a key way in which cognitive science can contribute to metaphysics. When folk intuitions are invoked as a constraint on metaphysical theories, cognitive science can enter in to help determine not only what the content of the relevant folk view is but also to help decide whether the relevant folk intuitions deserve to be taken seriously. In this way, cognitive science can help provide empirically informed

debunking explanations and thus liberate certain metaphysical theories from the demand that they conform to folk intuitions.

Acknowledgements

I would like to thank Mark Alicke, Wesley Buckwalter, David Danks, Alvin Goldman, Joshua Knobe, Edouard Machery, Shaun Nichols, Laurie Paul and John Turri for their invaluable discussion, feedback and support. I am especially grateful to Jonathan Schaffer and Stephen Stich for their mentorship, friendship, support and encouragement. Finally, I would like to thank my parents, Jennifer, Regan and my wife for their support.

Dedication

For Ashley

Table of Contents

List of Tables

List of Illustrations

Introduction

Contemporary experimental philosophy has crept into virtually every quarter of philosophical inquiry. Experimental philosophy has, for instance, found its way into epistemology (e.g., Weinberg et al, 2001), language (e.g., Machery et al, 2004), ethics (e.g., Sarkissian et al 2011), mind (e.g., Sytsma and Machery, 2010) and even logic (e.g., Ripley, 2016). Yet, one area that has received relatively little attention is metaphysics.

This isn't to say that experimental philosophers haven't done any work that is relevant to metaphysics. There has been work on causation (e.g., Alicke et al, 2011) and free will (e.g., Nichols and Knobe, 2007), two topics traditionally within the purview of metaphysics. But the implications of the empirical work are typically not explicitly connected to the relevant metaphysical issues.

It is surprising that experimental philosophers have paid relatively little attention to how empirical work might bear on issues in metaphysics. Several decades ago, Alvin Goldman (1987) suggested that cognitive science might be relevant to metaphysics. And over the years, Goldman has built on and refined those initial suggestions (see e.g., Goldman, 2007, 2015). Though Goldman has largely been focused on setting out a general rationale for how cognitive science can be evidentially relevant to metaphysics, others have explicitly connected work in cognitive science to traditional issues in metaphysics. For instance, Laurie Paul (2010) has connected work in cognitive science related to temporal experience to debates in the philosophy of time.

My aim is to take up the question of how experimental philosophy can contribute to metaphysics. Maybe this question strikes you as having an obvious answer: experimental philosophy *cannot* contribute to metaphysics. Perhaps you think that experimental philosophy has no philosophical significance whatsoever, that the whole practice has been a "big mistake" (e.g., Cappelen, 2014). This isn't the place to defend the philosophical significance of experimental philosophy (see e.g., Nado, 2016). But if "experimental philosophy" gives you conniptions, then perhaps it is worth pointing out that nothing much hinges on whether the main theme is centered on the question of how *experimental philosophy* can contribute to metaphysics. Experimental philosophers typically use the tools of cognitive science. So the question might instead be: how might cognitive science contribute to metaphysics? Work in experimental philosophy would then be one example of a more general kind of interaction between cognitive science and metaphysics.

For some, even on this way of posing the question—how might cognitive science contribute to metaphysics?—there is an obvious answer: cognitive science can't contribute to metaphysics. How, for instance, could what we *think* about causation inform us about what causation *really is*? Clearly it can't. So obviously cognitive science can't bear on issues in metaphysics. The essays that follow, however, suggest that it is not at all obvious that cognitive science can't contribute to metaphysics. Indeed, these essays suggest that experimental philosophy might play a significant role in metaphysics. So one shouldn't dismiss the idea that there is a significant connection between cognitive science or experimental philosophy and metaphysics because the question itself looks to have an obvious answer. The essays that follow may convince

you that cognitive science or experimental philosophy does indeed have a philosophically significant role to play in metaphysics.

Perhaps there are those who think it is obvious that experimental philosophy *can* contribute to metaphysics. A full blown philosophical naturalist or experimental philosopher, for instance, might think that empirical evidence has a role to play in virtually every area of philosophical inquiry. Cognitive science delivers empirical evidence and so do experimental philosophers given that they conduct psychological experiments. So experimental philosophy is obviously relevant to some issues in metaphysics. But even those who take the answer to the question of whether experimental philosophy can contribute to metaphysics to be obviously yes may want to know *how* it can contribute. What kind of specific proposals are there? What kind of concrete illustrations are there of experimental philosophy contributing to issues in metaphysics? The essays that follow provide some specific proposals and some concrete illustrations of how experimental philosophy can contribute to metaphysics.

So instead of taking it that the question—how can experimental philosophy (or cognitive science for the averse) contribute to metaphysics?—has an obvious answer, I would instead ask that you, as John Turri (forthcoming) has put it, "cleanse your mental palette". And with that, consider each essay on a case by case basis.

Here is an overview of the essays that follow: In Chapter 1, I argue that folk intuitions of actual causation are generated by two epistemically defective processes—what I call the evaluative and agentive processes. I provide empirical support for this hypothesis and, using causal modeling techniques, show that these two processes do indeed play a significant role in generating folk intuitions of actual causation. I then set out a two-

pronged debunking argument on the basis of this evidence. In Chapter 2, I consider folk intuitions of material object persistence. I provide evidence that the folk operate with a teleological view of persistence and go on to set out a debunking argument in light of the evidence. Chapter 3, co-authored with Jonathan Schaffer, takes up the flip side of persistence: composition. We find that, just as with persistence, the folk operate with a teleological view of composition. We trace out the implications of our findings, arguing that folk intuitions of composition should be dismissed: metaphysical theories of composition should not be beholden to folk intuitions. Finally, in Chapter 4, I set out how cognitive science can help the revisionary metaphysician discharge an explanatory burden that is typically lodged by the naïve realist. This burden is to provide a plausible explanation of how the folk could be so badly mistaken. I illustrate how this burden can be discharged by focusing on persistence and composition.

All of the material included is either published or forthcoming. Significant portions of the Introduction are forthcoming in a volume I am editing, which is entitled "Experimental Metaphysics". Chapter 1 is forthcoming in *Philosophical Studies* while Chapter 2 is published in *Synthese*. Chapter 3 is forthcoming in "Experimental Metaphysics" and an abridged version is forthcoming in *Noûs*. Finally, Chapter 4 is forthcoming in a volume edited by Alvin Goldman and Brian McLaughlin on cognitive science and metaphysics.

Chapter 1

Folk Intuitions of Actual Causation: A Two-Pronged Debunking Explanation

"[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

1.1. Introduction

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

[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.

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 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 Section 1.3, 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 Section 1.4, I present empirical evidence supporting the two-pronged debunking explanation. Section 1.5 discusses debunking and situates the empirical evidence within a background discussion of debunking and the challenge from folk belief. Section 1.6 considers some objections.

1.2 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, Beebee (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 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

[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).

> 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.

1.3. 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.

1.3.1 Moral Considerations

The role of moral considerations in generating judgments of actual causation is well documented (e.g., Alicke, 1992; Alicke, 2000; Alicke and Rose, 2010; Alicke, Rose and Bloom, 2011; Sytsma, Livengood and Rose, 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, Rose and Bloom (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 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; Alicke, 2000; Alicke and Rose, 2010; Alicke, Rose and Bloom, 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

[4] For ease, I'm only discussing the role of blame in causal judgment. But, as Alicke, Rose and Bloom (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 (forthcoming).
[5] See Section 1.6 for further discussion.

evaluative process extends beyond the realm of human action and do so by empirically integrating these results with results concerning promiscuous teleology (Section 1.4). And second, while some have claimed that the role of blame in causal judgment is a bias (e.g., Alicke and Rose, 2010; Alicke, Rose and Bloom, 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 (Section 1.5) 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, Livengood and Rose (2011) 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.

[6] Thanks to an anonymous reviewer for raising this.

Moreover, Alicke, Rose and Bloom (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.

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 Section III will further support the interpretive stance I'm taking. I'll also have a bit more to say in favor of my interpretation in Section 1.6.

1.3.2. 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, Lagnado, Drury, Matheson and Fitzpatrick (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, 1999b, 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, Rottman and Seston, 2013; Lombrozo, Kelemen and Zaitchick, 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, Rottman and Seston (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 and 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 & 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 & Wood-Robinson, 1985; Demastes, Settlage, & Good, 1995; Evans et al., 2010; Jensen & Finley,1995; 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 & 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]

Though no work has looked at whether this agentive 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 agentive considerations. Insofar as causation backs explanation, we should expect agentive considerations to impact folk intuitions of actual causation. Thus, promiscuous teleomentalism—embodied in what I have called the agentive 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.

1.4. Evidence for Two Processes

[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.

Two candidate processes—the evaluative process and the agentive 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 agentive process. For this reason, most of the focus will be on providing support for the role of the agentive 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.

1.4.1. 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 Section 1.3, 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) x 2(Condition: Intentional, Accidental) design (for full cases see Appendix A).

1.4.1.1. 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]

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

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

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

1.4.1.2. Results

Here are the t-tests for both John and the rock:

John Cases	Good Outcome (Intentional)	Bad Outcome (Accidental)	t-value	p-value
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
Rock Cases				
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

Table 1: Study 1 t-tests

An ANOVA revealed a statistically significant large-sized effect of Condition (Intentional, Accidental) on causal judgments $F(1, 150)=62.797$, $p<.001$, $\eta p^2=.295$, no statistically significant effect of Individual (John, Rock) on causal judgments, F(1,

[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

150)=.027, p=.869, and a statistically significant small-sized interaction between Condition and Individual F(1, 150)=4.518, p=.035, ηp^2=.029.

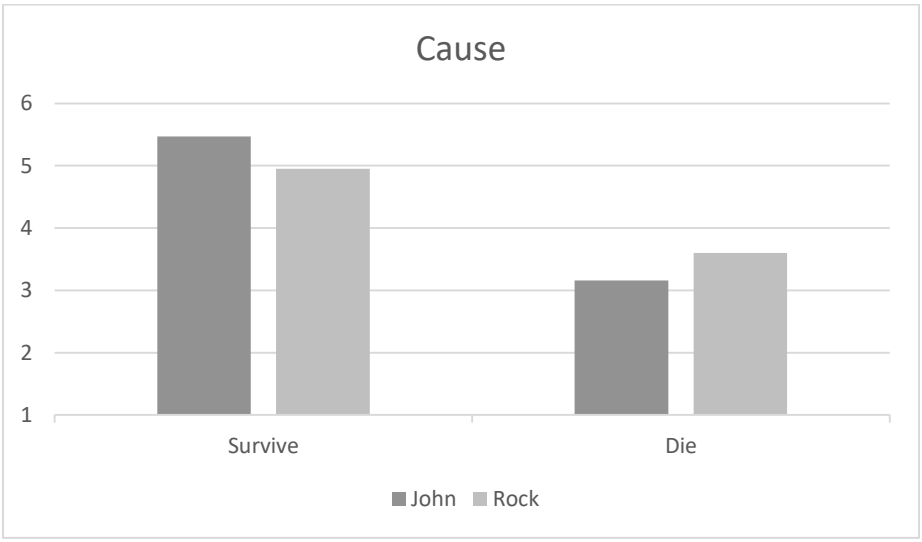

Figure 1: Causal Judgments for John and the Rock

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.

1.4.1.3. 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]

[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

GES returned the following model:[13]

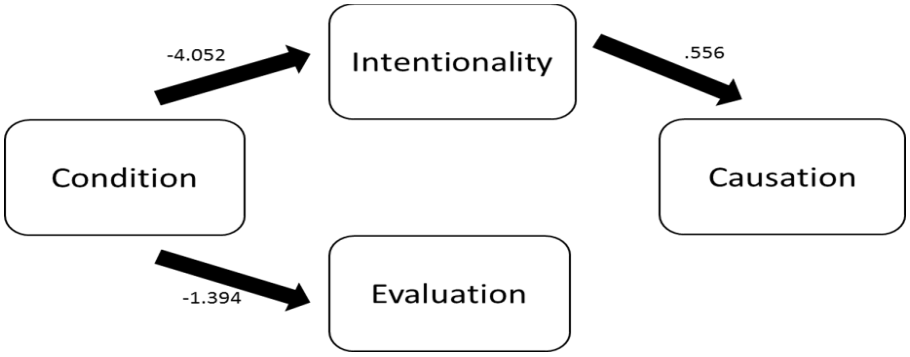

Figure 2: Causal Model for John Cases

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.

1.4.1.4. 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]

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

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

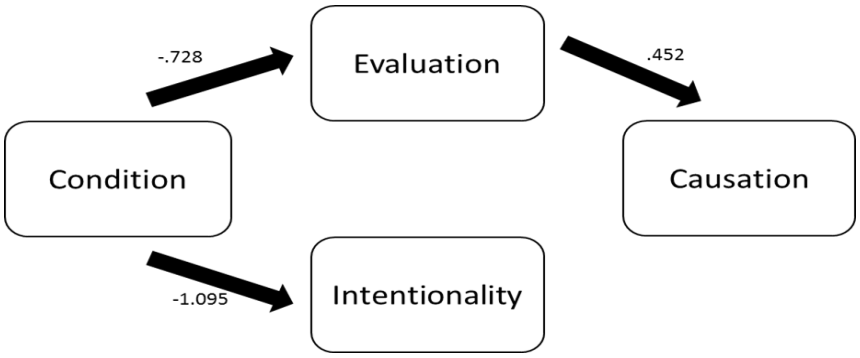

Figure 3: Causal Model for Rock Cases

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 the agentive 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 agentive 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.

1.4.2. Study 2

1.4.2.1. Participants and Measures

Endorsement of quasi-religious Gaia beliefs has been shown to significantly predict people's tendency to endorse teleological explanations (Kelemen, Rottman, and Seston 2013). To determine whether the agentive 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, Rottman, and Seston (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.

1.4.2.2. Results

First, here are the t-tests:

	Good Outcome (Intentional)	Bad Outcome (Accidental)	t-value	p-value
Causation	4.67(1.29)	3.29(1.71)	7.02	.000

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

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

Table 2: Study 2 t-tests

Second, the results from the Rock cases above were replicated. 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 following model:[17]

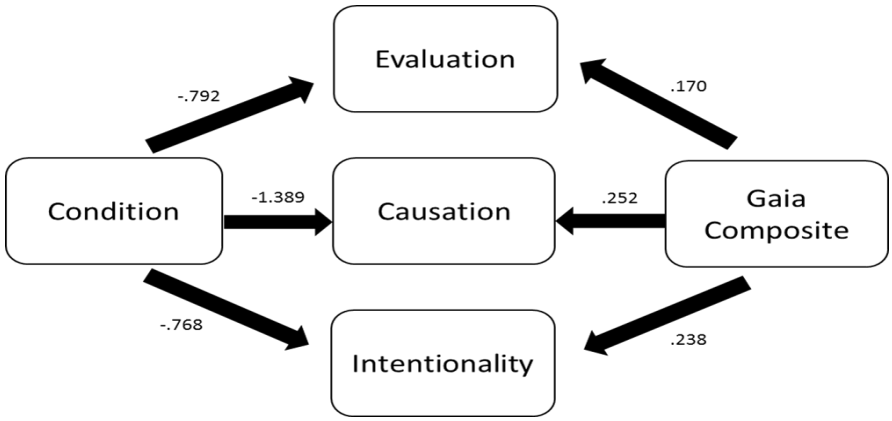

Figure 4: Causal Model for Gaia Beliefs in Rock Cases

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

[16] Cronbach's Alpha=.847
[17] This model fits the data well, df=4, X^2=5.3028, p=.2576, BIC=-16.1047

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 Section IV. 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 between 1 and 3 on Gaia Composite as displaying Low Gaia Belief and those who had an

[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 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).

overall score between 4 and 6 on Gaia Composite as displaying High Gaia Belief this is exactly what we find:

	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			

Table 3: High and Low Gaia Belief and Causal Judgments

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 the following graph:

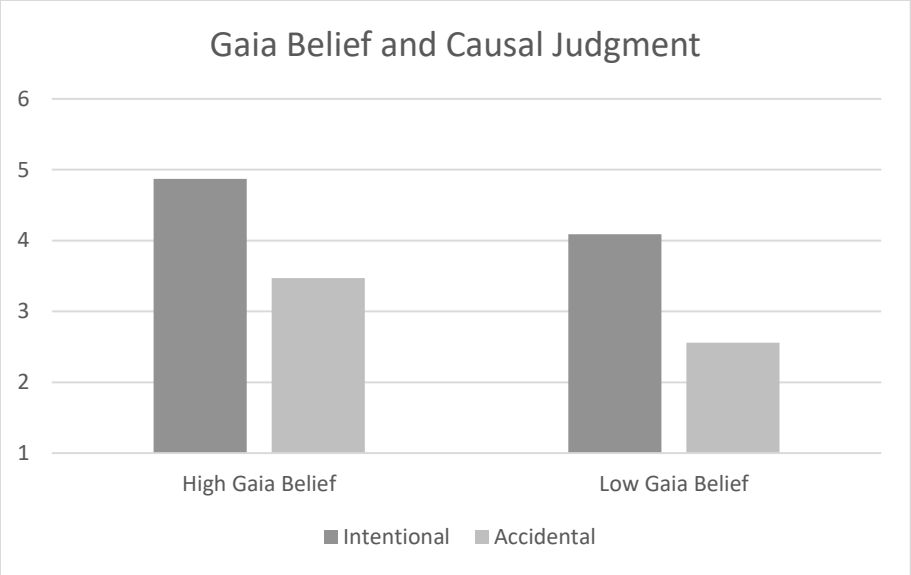

Figure 5: High and Low Gaia Belief Causal Judgments in Good Outcome (Intentional) and Bad Outcome (Accidental) Cases

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.

1.4.3. 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, Lagnado, Drury, Matheson and Fitzpatrick, 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) x 2(Behavior: Accidental, Intentional) x 2(Outcome: Good, Bad) design (for full cases see Appendix A).

1.4.3.1. 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]

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.

1.4.3.2 Results

Here are the t-tests for the KKM cases:

Good Outcome	Intentional	Accidental	t-value	p-value
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
Bad Outcome	Intentional	Accidental	t-value	p-value
Causation	4.71(1.15)	3.80(1.64)	3.37	.001
Evaluation	2.52(1.13)	3.48(1.01)	-4.76	.000

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

Intentionality	2.66(1.30)	1.70(1.15)	4.12	.000
Gaia	4.33(1.22)	4.40(1.30)	-.320	.749

Table 4: **T-tests for KKM cases with Behavior as IV**

And here are the t-tests for the Andy cases:

Good Outcome	Intentional	Accidental	t-value	p-value
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
Bad Outcome	Intentional	Accidental	t-value	p-value
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

Table 5: **T-tests for Andy cases with Behavior as IV**

An ANOVA revealed a main effect of Behavior, $F(1, 364)=17.836$, $p<.001$ $\eta p^2=.047$ 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, $F(1, 364)=9.578$, $p=.002$, $\eta p^2=.026$. There were no other significant two way interactions nor was there a significant three way interaction.

Analyzing just the data involving KKM, an ANOVA revealed a main effect of Behavior $F(1, 214)=7.362$, $\eta p^2=.034$ 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 $F(1, 214)=4.524$, $p=.035$, $\eta p^2=.021$. This can be seen in the following graph:

Figure 6: Causal Judgments in Non-Agent (KKM) Cases

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]

[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

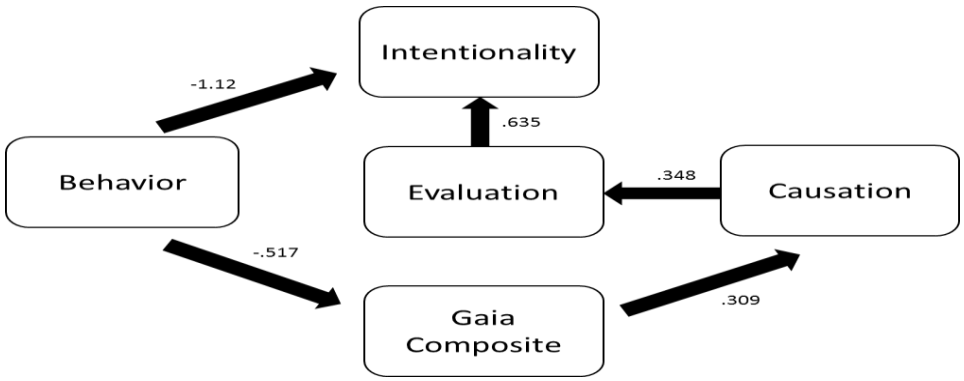

Figure 7: Causal Model for KKM Good Outcome

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

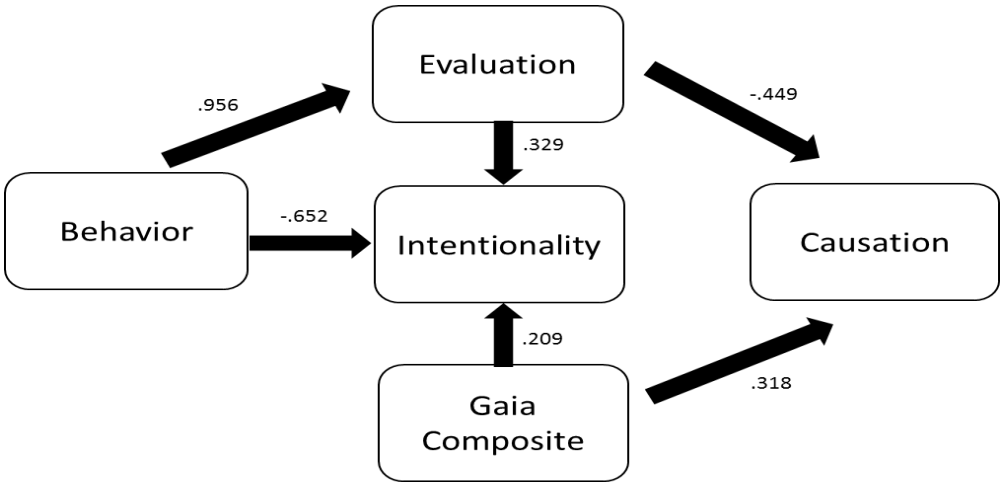

Figure 8: Causal Model for KKM Bad Outcome

Again, one of my main claims—that the agentive 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

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

modeling are, on their own, sufficient to provide support for the causal hypothesis that the agentive process generates causal judgments. Moreover this result is sufficient to support one prong of the debunking explanation. 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 Outcome cases are presented in Table 6, while the pattern of findings for the Good Outcome cases are presented in Table 7.

	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 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	4.79(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			

Table 7: Causal Judgments for High and Low Gaia Belief in Good Outcome Cases

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 Section 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 the following graph:

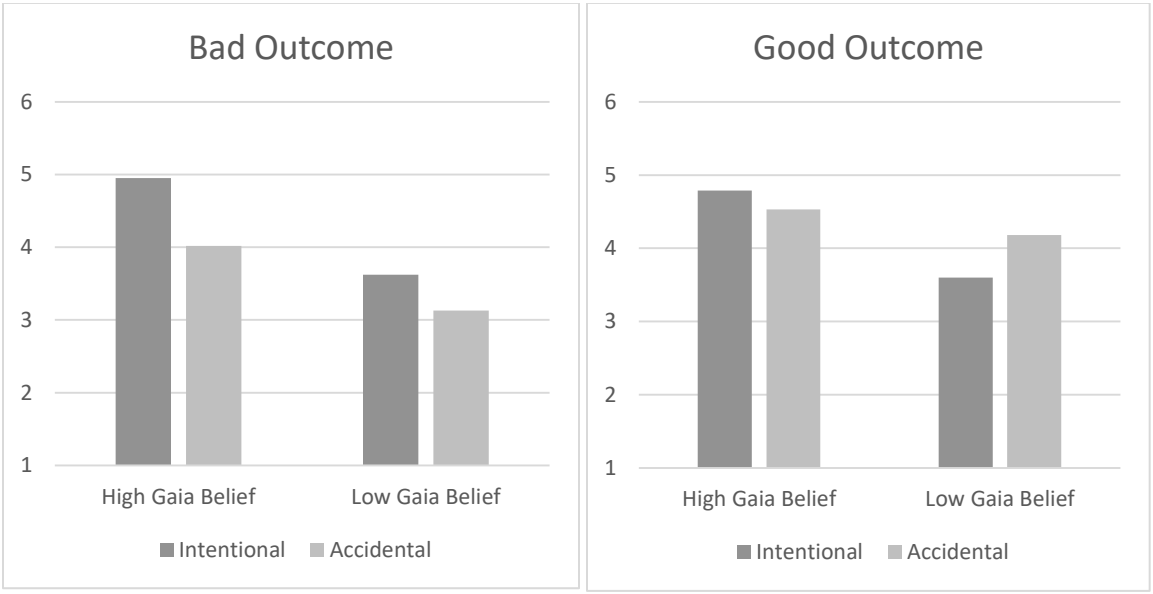

Figure 9: Gaia Belief and Causal Judgment for Bad Outcome Cases (Left) and Good Outcome Cases (Right)

Analyzing the data involving Andy, an ANOVA revealed a main effect of Behavior $F(1, 154)=10.003$, $p=.002$, $\eta p^2=.061$ 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 $F(1, 154)=4.880$, $p=.029$, $\eta p^2=.031$. This can be seen in the following graph:

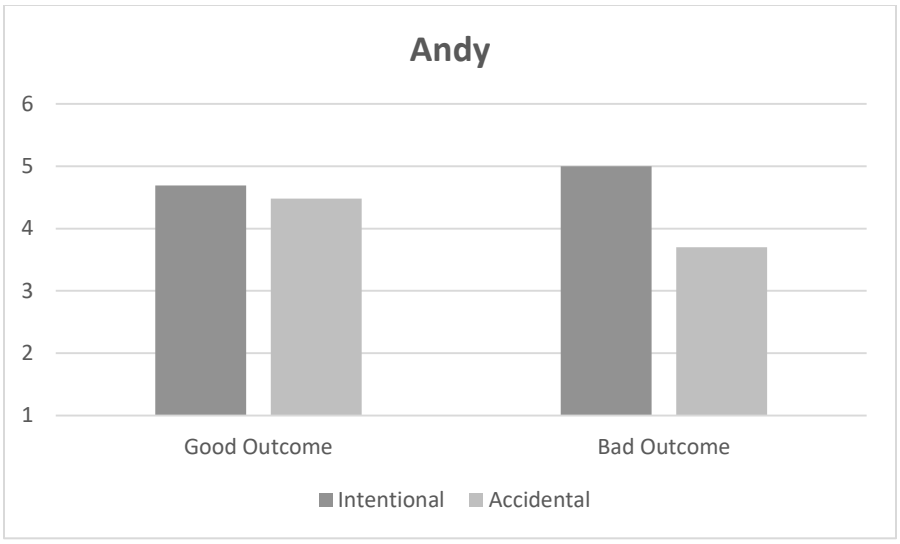

Figure 10: Causal Judgments in Agent (Andy) Cases

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]

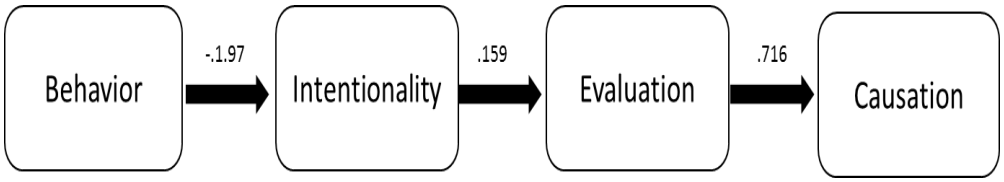

Figure 11: Causal Model for Andy Good Outcome

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

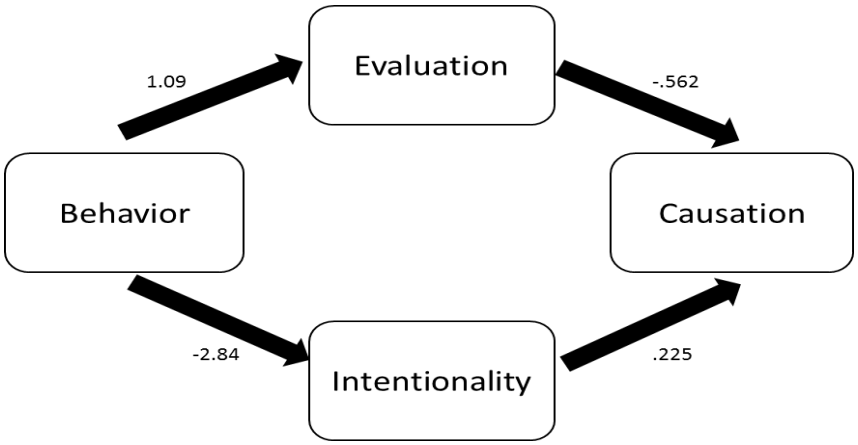

Figure 12: Causal Model for Andy Bad Outcome

1.4.4. Summary of Results

One of the crucial, key claims—that the agentive 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 agentive 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 agentive 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 agentive 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 teleomentalism. Thus I claim empirical support for the claim that two processes—the evaluative and agentive processes—play a role in generating folk intuitions of actual causation. On to debunking.

1.5. Debunking Folk Intuitions of Actual Causation

Given empirical support for two processes—the evaluative and agentive 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:

> (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.

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.[25] 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, p. 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

[25] 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).

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.[26] Insofar as folk judgments of actual causation issue from either the evaluative or agentive processes, their judgments of actual causation are prima facie unjustified. 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 agentive 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.

---

[26] For more, see Section 1.6.

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 agentive 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.[27] But as the above

[27] 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

results suggest, this is not the case. The evaluative process extends beyond the realm of human action because it is connected with promiscuous teleomentalism. Those who may have been attracted to debunking folk intuitions of 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 teleomentalism, 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 undiscriminating 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.[28] Indeed, I do not ever appeal to naive teleological reasoning or blame

why some have upheld respecting folk intuitions despite the fact that they are generated by evaluative considerations within the realm of human action.

[28] 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

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, Rose and Machery (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 aspects of causal cognition; namely, in those aspects of causal cognition which are invoked when

---

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.

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.[29] 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, Rose and Machery (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

[29] Thanks to an anonymous reviewer for raising this.

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

1.6. Objections

[30] Thanks again to an anonymous reviewer for asking for clarification on this point.

I presented evidence—in Section 1.4—that two epistemically defective processes—the evaluative and agentive processes—play a role in generating folk judgments of actual causation. And, in Section 1.5, I argued that these two processes give rise to a two-pronged debunking argument. I now want to briefly consider some natural objections.

1.6.1. 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.[31] According to current scientific methodology, spirited beings—for instance, agentive 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:

[31] 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).

> *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 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.

1.6.2. 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.[32] But a range of evidence suggests that counterfactual selection doesn't necessarily guide causal selection (see e.g., Mandel, 2003a, 2003b; Mandel and Lehman, 1996; Mandel and Lehman, 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, N'gbala and Bonnefon (2001) and Mandel (2003a) did *not* find evidence that counterfactual selection facilitated causal section, 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

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

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 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, Tenenbaum and Carey (2005) and Saxe, Tzelnic, and Carey (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).

1.7. 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.[33]

[33] 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.

Chapter 2

Persistence Through Function Preservation

2.1 Introduction

How do ordinary material objects persist? For instance, if a rock is hit with a hammer and chipped, does the rock survive? Or if a rock is smashed to pieces by a hammer, does the rock survive the smashing? Many metaphysicians have wanted a view of persistence that fits with folk intuitions and have charged leading views with failing to do so. Yet, there is disagreement about what the folk intuit and no empirical discipline to the discussion. For a debate so heavily centered on folk intuitions, it seems that some empirical evidence might help advance the discussion. Indeed, it seems that empirically discerning the folk view of persistence will help decide whether it deserves to be taken seriously.

So my question is: what is the folk view of persistence against which metaphysical accounts might be measured? My view is that the folk view of persistence is teleological in that the folk tend to intuit that a material object survives alterations when its function is preserved. As such, I hold that the folk view of persistence is tied into a benighted view of nature and thus deserves to be dismissed. Given an empirically informed understanding of the folk view of persistence, I hold that the discussion over how ordinary objects persist should be liberated from any demanded conformity with folk intuitions.

The Plan: I'll begin, in Section 2.2, by briefly documenting some conflicting claims about the folk view of persistence and charges of failing to fit common sense. In Section 2.3, I will present evidence from psychology which suggests that the folk are promiscuous teleologists. Given the background of promiscuous teleology, I will then go on, in Section 2.4, to present a range of evidence in support of the claim that the folk view of persistence is teleological. Having achieved sufficient empirical understanding of the folk view of persistence to judge its credentials, I will, in Section 2.5, argue for a dismissive take on folk intuitions about material object persistence. I will situate the discussion within the background of what Dan Korman (2009) calls the challenge from folk belief and discuss my results in the context of debunking arguments in order to show how the challenge can be met, concluding that in the specific case of persistence, the folk do not deserve to be taken seriously. They deserve to be ignored.

2.2. Persistence and Common Sense

How do ordinary material objects persist? There are two general answers. The first is the answer of the four dimensionalist. According to the four dimensionalist, ordinary material objects persist by having temporal parts, in addition to spatial parts, which are spread out across regions of spacetime. Ordinary objects persist by having distinct temporal parts at more than one time: they persist by *perduring*. The second answer, that of the three dimensionalist, denies that ordinary material objects persist by having temporal parts. Instead, the three dimensionalist holds that ordinary material objects persist by being "wholly present" at more than one time, sweeping across spacetime with

the whole of the object occupying different regions at different times.[1] Ordinary objects persist by *enduring*.

The four dimensionalists view of how ordinary material objects persist is typically charged with being "wildly counterintuitive" (Sider, 2001, p.3), "radically at odds with common sense" (Paul, 2002, p. 587), and "crazy metaphysics" (Thomson, 1983, p. 213). Three dimensionalists typically claim that their view better fits "common sense" (e.g., Merricks, 1994, p. 165; Paul, 2002, p. 586; Wiggins, 1980, p. 25). But four dimensionalists often dispute the charges. Some have suggested that "it is not clear that common sense offers any direct verdict" (Hawthorne, 2008, p. 3) against four dimensionalism, having never considered the question of whether ordinary objects have temporal parts (see also Benovsky, 2006). Indeed it has been suggested that "even though it is at first glance counterintuitive, [four dimensionalism] becomes more palatable even for common sense once one thinks more about it". And so having thought more about four dimensionalism, the folk will come to see that it's "the competing views that suffer from counterintuitive consequences" (Benovsky, 2006, p. 104).

So, we see a dispute between three and four dimensionalists about whether the three or four dimensionalist view of how ordinary objects persist fits with common sense. Does three dimensionalism better fit with common sense? Or rather would the folk, after having thought sufficiently about four dimensionalism, come to view it as more palatable,

[1] I will, along with others (e.g., Hawley, 2001, Lewis, 1986, Sider, 2001), take one of the main distinguishing features of three and four dimensionalism to be the acceptance of temporal parts (though see Parsons, 2000).

thinking that it's the competing three dimensionalist conception which suffers from being counterintuitive?

Aside from the dispute over whether the common sense view of how ordinary objects persist better fits with three or four dimensionalism, we also see a dispute over the common sense view of persistence in connection with one of the most standard accounts of persistence: the sortal based account (e.g., Wiggins, 1980; Hirsh, 1982; Lowe 1995).

Sortal based accounts hold that we trace the career of ordinary objects by tracing under a sortal. Sortals, on this view, answer the question of "what is it" and determine the persistence conditions for objects falling under the sortal. Many who hold a sortal based view allow that material objects can coincide. For instance, a statue and lump of copper can occupy exactly the same place at the same time and yet remain numerically distinct. The reason why, according to this account, is because the copper and statue fall under different sortals and accordingly have different persistence conditions associated with them. The statue, if pounded by a hammer, would not survive; yet the copper could survive being pounded by a hammer. So according to this account, coincidence can occur among objects that belong to different kinds (Wiggins, 1968).

Though the sortal based account has been recommended for fitting common sense, the account has also been charged with violating common sense. Burke (1997) recommends the standard sortal based account since it is "consisten[t] with the metaphysics implicit in ordinary ways of thinking" (p. 11). But, Burke continues, "There is one consequence of the standard account that many have found uncongenial, if not intolerable: that it is possible, indeed common, for one object to coincide with another" (p. 11). Burke holds that this is "at odds with commonsense" and thus avoids coincidence by arguing that a

piece of copper is *destroyed* upon being fashioned into a statue. In its place there comes to exist a new piece of cooper, which is identical to the resultant statue.

But Lowe (1995) claims, "it most certainly is…contrary to common sense to claim that a piece of copper must cease to exist merely upon assuming a certain shape—as any seven-year old child will confirm, if the question is put to him in terms of the fate of a piece of plasticene which is fashioned into the figure of a man…. I would most strenuously deny that his position is more in keeping with common sense than the standard account can claim to be" (1995, p. 176). Burke (1994), however, disputes the charges, claiming that once the distinction between the original lump and piece of copper is made clear, we see that it "is not the violation of common sense it initially seems" (p. 598)

Fit with common sense is an important desideratum in disputes over persistence.[2] Three dimensionalists tend to claim that their view fits better with common sense; four dimensionalists dispute the charges, claiming that their view fits common sense. And the standard sortal based account has been recommended for its (apparent) fit with common sense. But if fit with common sense is to serve as a desideratum in theory choice, then it seems that a precondition for deciding whether common sense is a good metric is determining what the folk view is and whether it deserves to be taken seriously. And empirical evidence on the folk view can help in deciding whether the measure of common sense, in specific domains, is a good measure. By empirically discerning what the folk view is in a specific domain, one is thereby in a position to decide whether

[2] I've only offered a brief sampling of disputes over the common sense view of persistence. For more, see the excellent discussion in Sidelle (2002).

measuring a metaphysical theory by its fit with the specific folk view under consideration is a wise policy.

My view is that the folk view of persistence is teleological and as such I see the view as being fit for dismissal in the debate over persistence. But before getting on to my own studies in support of the hypothesis that the folk view of persistence is teleological, I want to first briefly motivate the hypothesis by situating it within a background discussion of promiscuous teleology.

2.3. Psychological Background on Folk Teleology

2.3.1. Selective Teleology

It's widely agreed that we're at least *selective teleologists* in that our artifact and organism concepts are infused with considerations of purpose and function. Most work supporting this has come from studies investigating principles of object categorization (*what a thing is*) with the general result being that we tend to determine what a thing is by determining what it is for (*what is its purpose*).

With artifact concepts, German & Johnson (2002) note that it is well established that people take "the design stance, in which an entity's properties, behavior, and existence is explained in terms of its having been designed to serve a particular purpose" (pp. 279-280). Likewise Bloom (1996), reviewing earlier work by Rips and by Keil, notes: "This has suggested to many scholars that the psychological 'core' of artifact concepts is that their members share a common intended function" (p. 63).

As with artifact concepts, it is also agreed that our organism concepts are infused with teleological considerations. For example, Frank Keil (1995) writes:

> Historically there have been many arguments for a 'design' stance, which can include teleological interpretations and tool construction and use... Notions of functional architecture are among the most cognitively compelling ways of approaching the biological world and much of the artificial world as well.... (p. 245)

In a similar vein, Atran (1998) speaks of the folk idea of a "biological essence" as "an intrinsic...teleological agent, which physically... causes the biologically relevant parts and properties of a generic species to function and cohere 'for the sake of' the generic species itself" (pp. 550-551).

2.3.2 Promiscuous Teleology

Though it is widely agreed that we're at least selective teleologists in that our artifact and organism concepts are infused with teleological considerations, Kelemen (1999) suggests that the more psychologically plausible view is that of *promiscuous teleology*. On this view, teleological considerations not only influence our conception of artifacts and organism but also affect our general conception of an object and extend to even non-living natural things like rocks. She writes:

> [T]hroughout history, non-living natural objects have…been considered in… [teleological] terms... The earth, its climates, landforms, water sources, and elements, were seen as intentionally designed to create a habitat for, and meet the needs of, people. In other words, natural objects of all kinds—particularly those fulfilling a significant function in people's lives—were candidates for construal as quasi-artifacts. (p. 245)

Though there is a range of evidence supporting the view that children are promiscuous teleologists (e.g., Kelemen and Diyanni, 2005; Kelemen, 1999a, 1999b, 2004), other work suggests that even adults never fully outgrow their childhood tendencies toward promiscuous teleology (e.g., Kelemen and Rosset, 2009; Kelemen, Rottman and Seston, 2013; Lombrzo, Kelemen and Zaitchick, 2007). For instance, Kelemen and Rossett (2009) found that even college aged students revert to unwarranted teleological

explanations, even in scientific contexts, in conditions where their cognitive resources are limited. Similarly, Kelemen *et al* (2013) also found that even trained physical scientist show a similar pattern of reverting to unwarranted teleological explanations in similar conditions.

Indeed, recent work by Rose and Schaffer (forthcoming), investigating the role of teleological considerations in ordinary intuitions about mereological composition, provides evidence that the folk tend to intuit that a plurality has a fusion when the result has a purpose. In a wide range of cases, they found that the folk tended to judge that the plurality composed a further object when the result had a purpose and that the folk tended to judge that the plurality did not compose a further object when the result did not have a purpose. They found this basic pattern for cases involving artifacts, organisms and non-living natural objects like rocks. The background psychological literature on promiscuous teleology and principles of object categorization suggests that what something is (*sortal*) is given by what function it has; the results from Rose and Schaffer suggest that whether something is (whether this is a fusion) is determined by whether there is a function. Extending this pattern to the folk view of persistence: whether something persists is given by whether it continues to serve its function.

2.4. The Folk View of Persistence

I'll begin with gradual-replacement-of-parts cases, examining whether teleological considerations influence ordinary judgments about persistence for artifacts, organisms and rocks undergoing gradual part replacement. Having done this, I'll move away from gradual-replacement-of-parts cases and look at cases where a rock is hit with a hammer. The reason for considering diverse cases is twofold. First, by considering a range of

different cases, I hoped that the results might prove to be robust. As it turns out, this is the case. Across a range of cases, it turns out that a consistent pattern emerges: folk intuitions about material object persistence are largely driven by teleological considerations. Second, even though any single study may be questioned and open to diverse interpretations, by considering a range of different cases, I hoped that a consistent pattern would emerge. The overall pattern that appears to emerge is that the folk operate with a teleological view of persistence. That said, the results below, as with any empirical work, are defeasible in light of future inquiry. I take this to be the first, though not the last, word on the folk view of persistence and encourage further empirical work.

2.4.1 Gradual Replacement: Rowboat, Organism and Rock

I created three different cases—involving a rowboat, an organism, and a rock[3]—and for each case varied whether the replacement object or the object with the original parts preserved the function of the original object or whether no function was mentioned at all. For each case—rowboat, organism and rock—there were three conditions: No Function, Replacement Preserves Function and Original Parts Preserves Function.

The rowboat cases were modeled after the classic "Ship of Theseus" cases. In the No Function version, participants were told about John, a woodworker and sailor, who built his first rowboat "Drifter" thirty years ago. Over the years there was wear and tear on the boat and he eventually replaced all of the original planks with new planks. He kept all of the original planks and one day constructed a rowboat with them.

[3] For full cases, see Appendix B.

In both the Replacement Preserves Function version and the Original Parts Preserves Function version, participants were told that the original rowboat that John built was excellent and functioned perfectly as a rowboat. They were then told that either Replacement or Original Parts continued to function perfectly. After reading the stories, participants were told that two people—Andy and Suzy—wanted to borrow "Drifter" for an outing and disagreed about which boat is "Drifter". In all cases, Andy thinks that original parts is "Drifter" while Suzy thinks that replacement is "Drifter". Participants were asked to indicate the extent to which they agreed with Andy or Suzy, on a 7-pt scale anchored with 1=Suzy is right and 7-Andy is right.

The organism cases were similar to the rowboat cases. In the No Function case, participants were told that John discovered a new organism, which he named "Gollywag". Over the years, John experiments with it and each time he does, he removes a part from it and replaces it with a new part from the same type of organism until it has completely new parts. He keeps all of the original parts and one day assembles them.

In the function versions, participants were told that John suffered from eczema and that when handling the original organism he noticed that his eczema started to disappear. In one version, the Replacement preserved the function of relieving John's eczema while in the other, Original Parts preserved the function of relieving John's eczema. In all cases, participants were told that Andy and Suzy are interns, given the job of experimenting on "Gollywag". But Andy and Suzy disagree over which object is "Gollywag". Participants indicated the extent to which they agreed with Andy or Suzy on the same seven point scale used in the rowboat case.

The final case involved a rock. In the No Function case, John found a new rock which he named "Zenyte". Over the years, he experiments with the rock and each time he conducts an experiment he breaks off a piece of the rock and replaces it with a new part from the same type of rock until it has completely new parts. He keeps all of the original parts and one day assembles them.

As in the organism cases, in the function versions, participants were told that John suffered from eczema and that when handling the original rock he noticed that his eczema started to disappear. In one version, the Replacement preserved the function of relieving John's eczema while in the other, Original Parts preserved the function of relieving John's eczema. As in the organism cases, Andy and Suzy are given the job of experimenting on "Zenyte" but disagree about which object is "Zenyte". As in the other cases, participants indicated the extent to which they agreed with Andy or Suzy on the same seven point scale.

The results indicated that whether or not Replacement or Original Parts preserved the function of the original object had a strong effect on people's persistence judgments. And this effect was present, regardless of the type of object that underwent gradual part replacement, as can be seen in the images below:[4]

---

[4] A total of 330 participants were recruited from Amazon's Mechanical Turk. Participants were randomly assigned to one of nine conditions in a 3 (Function: None, Replacement Preserves Function, Original Parts Preserves Function) x 3 (Object Type: Rowboat, Organism, Rock) design. After reading the case, participants rated the extent to which they agreed with either Andy or Suzy. They were then taken to a separate screen where they were asked comprehension questions (see appendix for details). 10 people were excluded from the data analysis for missing one or more comprehension questions, leaving a total of 320 responses. The results indicated that there was a large-sized effect of Function, $F(2, 311)=33.568$, $p<.001$, $\eta p^2=.178$; no effect of Object Type, $F(2, 311)=2.832$, $p>.05$; and no interaction between Function and Type of Object, $F(4, 311)=258$, $p>.05$.

Rowboat: Planned pairwise comparisons in the Rowboat condition revealed a large-sized significant difference between Replacement Preserves Function ($M=2.79$, $SD=1.77$) and Original Parts Preserves Function ($M=5.03$, $SD=2.22$), $t(67)=4.60$, $p<.001$, $d=1.11$; a medium-sized significant difference

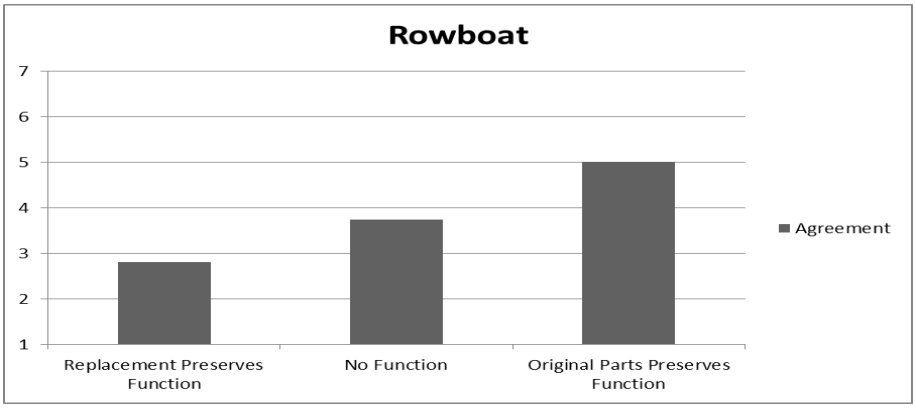

Figure 13: Rowboat

between No Function (M=3.75, SD=2.06) and Original Parts Preserves Function, t(68)=2.50, p<.01, d=.61; and a medium-sized significant difference between No Function and Replacement Preserves Function, t(67)=2.04, p<.05, d=.50.

Organism: Planned pairwise comparisons in the Organism condition revealed a large-sized significant difference between Replacement Preserves Function (M=2.94, SD=1.63) and Original Parts Preserves Function (M=4.81, SD=1.95), t(72)=4.45, p<.001, d=1.04; a medium-sized significant difference between No Function (M=3.72, SD=2.09) and Original Parts Preserves Function, t(67)=2.24, p<.05, d=.54; and a small-sized marginally significant difference between No Function and Replacement Preserves Function, t(67)=1.72, p=.09, d=.41.

Rock: Planned pairwise comparisons in the Rock condition revealed a large-sized significant difference between Replacement Preserves Function (M=3.27, SD=2.06) and Original Parts Preserves Function (M=5.75, SD=1.81), t(71)=5.45, p<.001, d=1.28; a large-sized significant difference between No Function (M=4.16, SD=2.12) and Original Parts Preserves Function, t(71)=3.43, p<.01, d=.81; and a small-sized marginally significant difference between No Function and Replacement Preserves Function, t(72)=1.83, p=.07, d=.43.

Throughout, I will be reporting effect sizes for significant effects. I'll be reporting partial Eta squared (ηp^2) and Cohen's d. ηp^2 indicates the amount of variance in the dependent variable explained by a given independent variable while Cohen's d indicates the magnitude of the mean difference between two groups. I'll follow Ellis (2010) for interpreting the magnitude of effect sizes. For ηp^2 I'll interpret values greater than or equal to .14 as large, greater than or equal to .06 but less than .14 as medium, and greater than or equal to .01 but less than .06 as small. And for Cohen's d I'll interpret values greater than or equal to .8 as large, greater than or equal to .5 but less than .8 as medium, and greater than or equal to .2 but less than .5 as small.

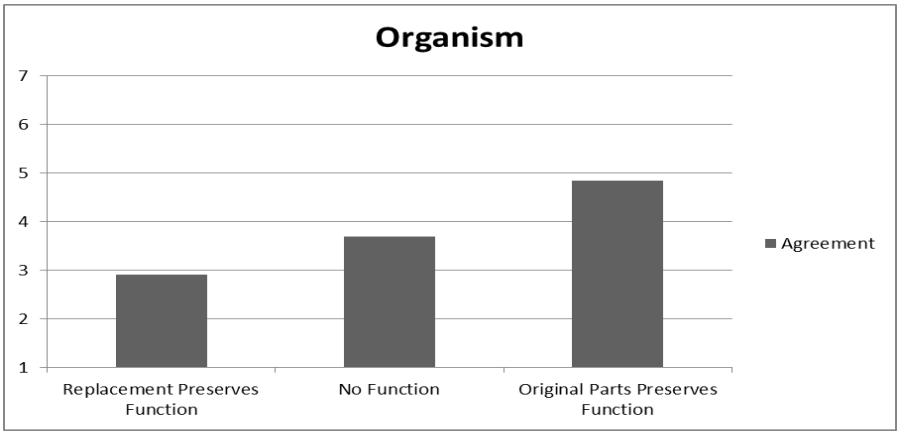

Figure 14: Organism

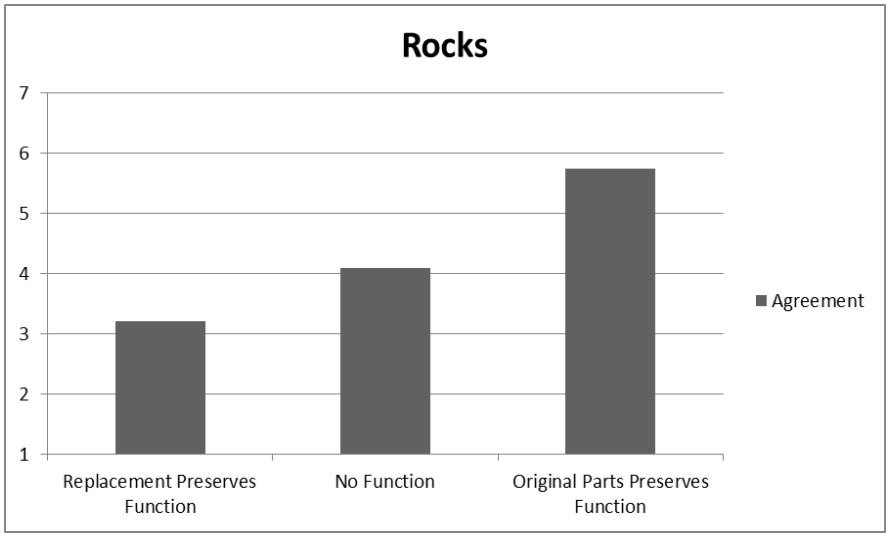

Figure 15: Rocks

Though it is perhaps natural to think that teleological considerations should play a role in determining whether artifacts or organisms persist (e.g., Wiggins,1980; Hirsh, 1982; Van Inwagen, 1990), the above results indicate that teleological considerations play a role in persistence judgments for non-living natural objects like rocks. As far as I am aware, no metaphysician has ever suggested that teleological considerations should or even expected that teleological considerations would play a role in folk persistence judgments

for non-living natural objects like rocks. But, against a background of promiscuous teleology, these results are perhaps unsurprising. These initial results fit with the view that we are promiscuous teleologists and suggest that teleological considerations even infuse our judgments about whether ordinary material objects persist.

Though these initial results begin to suggest that the folk operate with a teleological view of persistence, there are two natural concerns with the above studies. The first is that, instead of this pattern of judgments suggesting that the folk view is teleological, a better explanation of this pattern of findings is that the folk are merely operating on the assumption that John, in the rowboat case for example, is not a jerk and so would only lend a rowboat that worked. So, perhaps the pattern of judgments in the rowboat case has nothing to do with teleology but rather is better explained on the natural assumption that people's judgments track the functioning rowboat since they assume that John would not lend a broken rowboat to his friends. And similarly, for the organism and rock cases, perhaps it's natural to think that John cares about whether the organism or rock continues to function and so wouldn't assign Suzy and Andy to experiment with the non-functioning object. So perhaps these considerations better explain the pattern of judgments found above.

A second concern with the above studies pertains to the organism and rock cases. In both of those cases, the relevant function specifically benefitted John. Perhaps in a suitably modified story, where the relevant function didn't benefit people whatsoever, teleological considerations would cease to play a role in ordinary persistence judgments. Such a finding would be interesting in its own right since it would suggest that there is an

interesting restriction on the type of function that ends up playing a role in folk persistence judgments. The next set of studies takes up both of these concerns.

2.4.2. Gradual Replacement: Rowboat, Organism and Rock Modified

To address the concern that in the rowboat cases the pattern of judgments is better explained by participants assuming that John would not lend a broken rowboat, I modified the cases so that rather than wanting to borrow the rowboat for an outing, Andy and Suzy want to paint a portrait of the rowboat. Thus, the stories were similar to the ones above except participants were told that Andy and Suzy wanted to paint a portrait of the rowboat for an art class they were taking. John told them, over the phone, that they could paint "Drifter" and told them that he would be out of town and would have no phone or internet access. Participants were given a No Function version and a Replacement and Original Parts Preserves Function versions (for details, see Appendix B).

To address the concerns about the Organism and Rock cases, I did two things. First, to address the concern that the pattern of judgments found in Study 1 is better explained by participants assuming John would not assign his interns to experiment with a non-functioning object, I had John explicitly say that he did not care about the function of the object. Second, to address the concern that the pattern of judgments observed in Study 1 was only due to the object having a purpose for humans, I altered the function of the objects so that it did not benefit humans whatsoever.

So, in the Organism case, participants were told that John discovered a new microorganism, "Gollywag". In the No Function version, John disassembled it over the

years and replaced parts from the same type of microorganism. When he had accumulated enough of the original parts, he fashioned them into a microorganism which was exactly like the replacement.

In both the Replacement and Original Parts versions, participants were told that John gave the microorganism to his friend, Frank, who is a biochemist, to examine. Frank keeps the microorganism for several days and tells John that the microorganism is very delicate and must maintain a body temperature between 60 and 65 degrees. To maintain its temperature, it emits a unique sequence of chemicals that signal other microorganisms to group around it and heat it up. But John isn't interested in the chemicals emitted by the microorganism. Rather he is only interested in investigating and studying the various parts of the microorganism. So, John experiments with the microorganism, eventually replacing all of its parts and later assembles the original parts into a microorganism which is exactly like the replacement.

Function was varied by either having the Replacement preserve the function of the original by continuing to emit the chemicals or by having Original Parts preserve the function of the original by continuing to emit the chemicals. At the end of the stories participants were told that John gave two of his interns, Andy and Suzy, the job of experimenting on "Gollywag" (see Appendix B).

The Rock cases were similar to the Organism cases, except that the rock had a different function: it created a perfectly hospital environment for some worms. In addition to a No Function version, participants were given a version where the Replacement preserved the function of the original or Original Parts preserved the function of the original. As in the Organism case, John assigns two interns, Andy and Suzy, the job of experimenting on

"Zenyte". In this as well as the other cases, Andy and Suzy disagreed over which object was the original. Participants, in all versions, indicated who they agreed with on the same seven point scale that was used in the first study (see Appendix B).

As with first study, the results continued to suggest that the folk operate with a teleological view of persistence in that they tend to intuit that an object persist when it preserves its function.[5] The results are shown in the following images:

[5] A total of 310 participants were recruited from Amazon's Mechanical Turk. Participants were randomly assigned to one of nine conditions in a 3 (Function: None, Replacement Preserves Function, Original Parts Preserves Function) x 3 (Object Type: Rowboat, Organism, Rock) design. After reading the case, participants rated the extent to which they agreed with either Andy or Suzy. They were then taken to a separate screen where they were asked comprehension questions (see appendix for details). 13 people were excluded from the data analysis for missing one or more comprehension questions, leaving a total of 297 responses. The results indicated that there was a large-sized effect of Function, $F(2, 288)=34.475$, $p<.001$, $\eta p^2=.193$; no effect of Object Type, $F(2, 288)=2.592$, $p>.05$; and no interaction between Function and Type of Object, $F(4, 288)=.088$, $p>.05$.

Rowboat: Planned pairwise comparisons in the Rowboat condition revealed a large-sized significant difference between Replacement Preserves Function (M=2.78, SD=2.01) and Original Parts Preserves Function (M=4.85, SD=1.91), $t(64)=4.27$, $p<.001$, $d=1.05$; a medium-sized significant difference between No Function (M=3.74, SD=2.16) and Original Parts Preserves Function, $t(65)=2.50$, $p<.05$, $d=.55$; and a small-sized marginally significant difference between No Function and Replacement Preserves Function, $t(65)=1.85$, $p=.07$, $d=.46$.

Organism: Planned pairwise comparisons in the Organism condition revealed a large-sized significant difference between Replacement Preserves Function (M=3.15, SD=1.82) and Original Parts Preserves Function (M=5.57, SD=1.65), $t(66)=5.75$, $p<.001$, $d=1.39$; a medium-sized significant difference between No Function (M=4.37, SD=2.06) and Original Parts Preserves Function, $t(68)=2.69$, $p<.01$, $d=.64$; and a medium-sized significant difference between No Function and Replacement Preserves Function, $t(66)=2.58$, $p<.05$, $d=.63$.

Rock: Planned pairwise comparisons in the Rock condition revealed a large-sized significant difference between Replacement Preserves Function (M=3.10, SD=1.78) and Original Parts Preserves Function (M=5.39, SD=1.66), $t(59)=5.17$, $p<.001$, $d=1.32$; a medium-sized significant difference between No Function (M=4.27, SD=1.89) and Original Parts Preserves Function, $t(62)=2.49$, $p<.05$, $d=.62$; and a medium-sized significant difference between No Function and Replacement Preserves Function, $t(61)=2.52$, $p<.05$, $d=.64$.

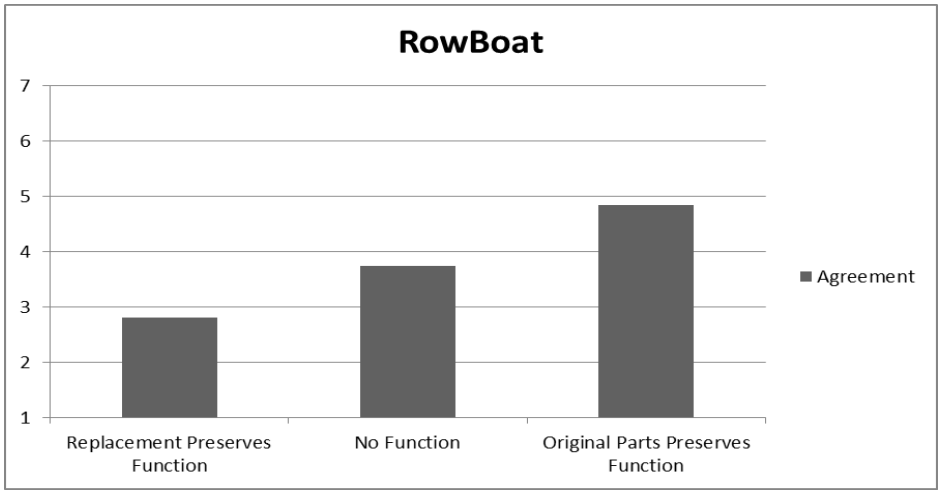

Figure 16: Rowboat Modified

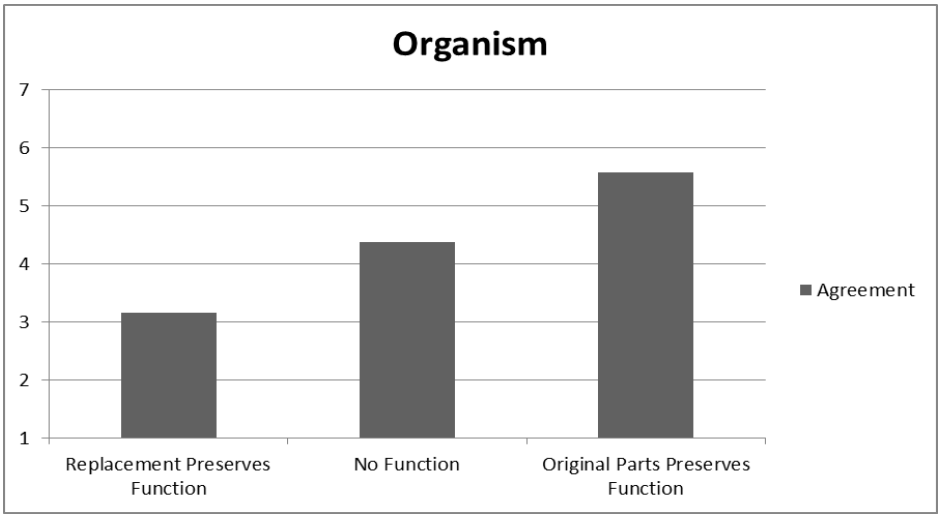

Figure 17: Organism Modified

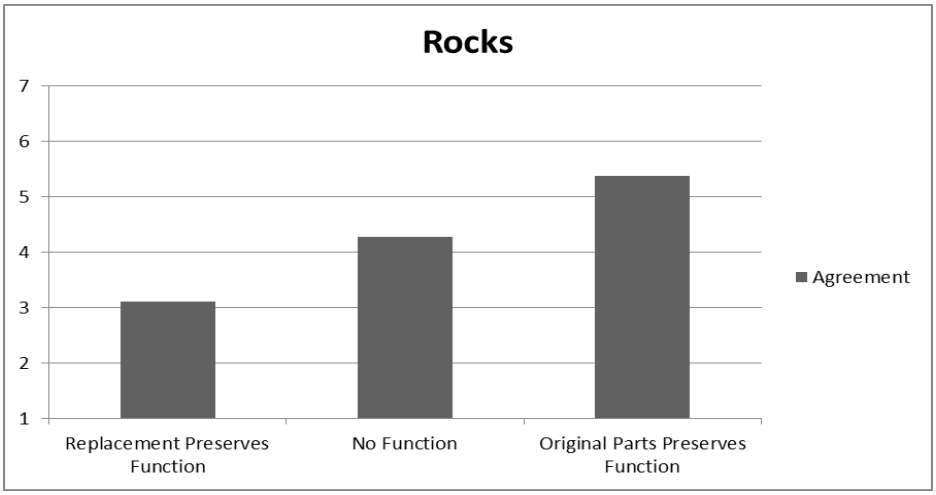

Figure 18: Rocks Modified

Despite the explicit attempt to remove the "not being a jerk element" (e.g., switching from borrowing to painting a rowboat in the rowboat case, having John say that he does not care what the organism or rock is for in the organism and rock cases), and despite switching the function away from being for a person, participants continued to show the same pattern of teleologically driven intuitions as they did in the first study.

The above result suggest that the folk operate with a teleological view of persistence in that they tend to judge that an object persists when it preserves its function. Though the results from the above cases appear to support the view that the folk operate with a teleological view of persistence, it would be good to find out whether the results are robust and hold up in cases that do not involve gradual part replacement. Furthermore, I suggested that the folk are promiscuous teleologists in that teleological considerations extend to judgments involving non-living natural objects like rocks. But the results in support of this, so far, have only been restricted to cases involving gradual part replacement. To continue to explore the hypothesis that the folk are promiscuous

teleologists and that this extends to their general view of how objects persist, I'll focus on cases involving only rocks in what follows.

I'll begin by looking at cases where a rock is smashed into three pieces and either preserves or loses its function.

2.4.3. Rock Smashed Into Three Pieces

I created two cases. In both, John is hiking and finds a rock which is glowing because it hosts special microorganisms. He takes the rock home to study and notices that the rock begins to fade as the microorganisms begin dying. He realizes that the microorganisms feed on minerals in the rock's interior and can't access them. So, he hits the rock with a hammer, breaking it into three pieces so that the microorganisms can access the minerals. In one version, the microorganisms continue to die and it stops glowing while in the other the microorganisms can access the minerals and it glows brighter than before. After reading the cases, participants were asked "Has the rock John found survived being hit by the hammer, or has it been destroyed?" and made ratings on a 7-pt scale anchored with 1=It has definitely survived and 7=It has definitely been destroyed.[6]

The results continued to show a pattern of teleologically driven intuitions, with participants tending to say the rock was destroyed when it lost its function and tending to say that the rock survived when it preserved its function.[7] Graphically:

---

[6] For full cases, see Appendix B.

[7] A total of 95 participants were recruited from Amazon's Mechanical Turk. Participants were randomly assigned to one of two conditions (Function: Lost, Preserved). After reading the case, participants rated the extent to which they thought the rock survived or was destroyed. They were then taken to a separate screen where they were asked two comprehension questions (displayed on separate screens):

(1) John's experiment worked. (Yes/No)

(2) John hit a rock with a hammer and broke it into three pieces. (Yes/No)

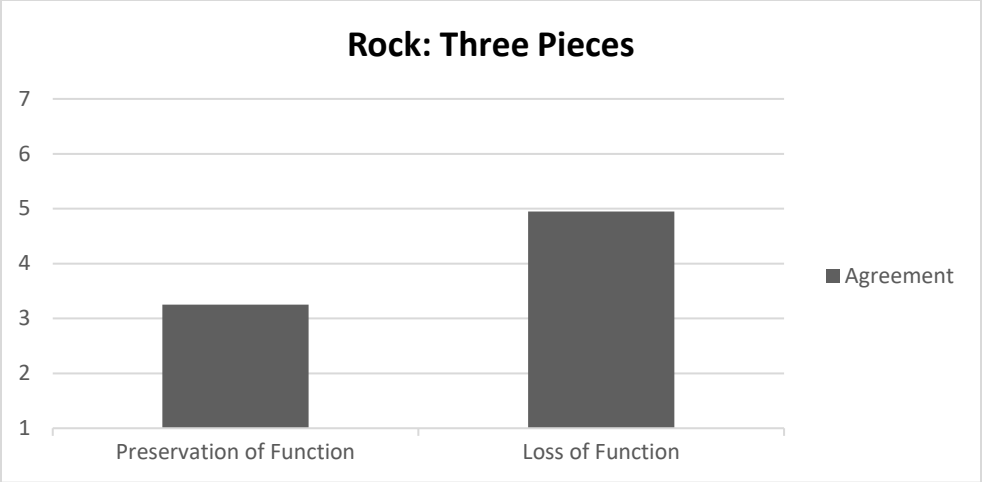

Figure 19: Rock Three Pieces

Even when a rock is smashed into three pieces, participants tend to say that it was destroyed when it loses its function and tend to say that it survives when it preserves its function. Thus, even in non-gradual part replacement cases, the results continue to support the hypothesis that folk intuitions about material object persistence are infused with teleological considerations.

Perhaps the effect would not continue to hold up under conditions where the rock undergoes more radical alterations, such as being pulverized. Or perhaps the effect would not continue to show up when the alterations to the rock are only minor, such as being dented. I'll explore both of these in the next study.

2.4.4. Denting and Pulverizing

7 people were excluded from the data analysis for missing one or more comprehension questions, leaving a total of 88 responses. The results indicated that there was a large-sized effect of whether the rock Lost (M=4.93, SD=1.92) or Preserved (M=3.27, SD=1.84) its function, t(86)=4.13, p<.001, d=.88.

To explore whether teleological considerations would continue to influence ordinary judgments of persistence, even when the object undergoes only minor alterations, I gave participants a case similar to the one is Section 2.4.3 where the rock stopped glowing and the microorganisms died after the rock was hit with a hammer. The main difference was that the rock was only dented when hit with the hammer. This was contrasted with a control case where participants were only told that John hit a rock with a hammer and dented it.

And to explore whether teleological considerations would continue to influence ordinary judgments of persistence, even when the object undergoes radical alteration, I gave people a case that was similar to the one in Section 2.4.3 where the rock continues to glow and the microorganisms survive after the rock is hit with the hammer. The only difference was that participants were told that John smashed the rock when he hit it with the hammer. This case was contrasted with a control case where participants were only told that John smashed a rock with a hammer. In all cases, participants made ratings on the same scale as used in the previous study (Section 2.4.3).[8]

For the cases which involved the rock being dented, the results indicated that while people tended to agree that the rock survived being dented in the control case, when the rock lost its function as a result of being dented, participants tended to say that the rock had been destroyed.[9]

---

[8] For full cases, see Appendix B.

[9] A total of 95 participants were recruited from Amazon's Mechanical Turk. Participants were randomly assigned to one of two conditions (Control, Loss of Function). After reading the case, participants rated the extent to which they thought the rock survived or was destroyed. They were then taken to a separate screen where they were asked two comprehension questions (each displayed on a separate screen):

(1) John's experiment worked. (Yes/No)

(2) John hit a rock with a hammer. (Yes/No)

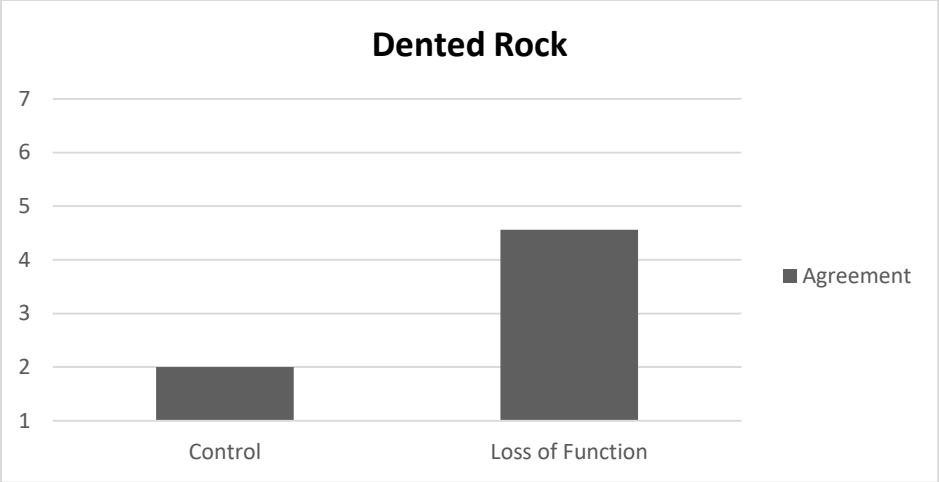

Figure 20: Dented Rock

In the pulverized rock version, when the rock was shattered but preserved its function, people tended to agree that the rock survived the smashing in comparison to the control case where participants tended to say that the rock was destroyed.[10]

6 people were excluded from the data analysis for missing one or more comprehension questions, leaving a total of 89 responses. The results indicated that there was a large-sized significant difference between the Control (M=2.06, SD=1.54) and Loss of Function Cases (M=4.56, SD=2.10), t(87)=6.40, p<.001, d=1.36.

[10] A total of 95 participants were recruited from Amazon's Mechanical Turk. Participants were randomly assigned to one of two conditions (Control, Pulverized). After reading the case, participants rated the extent to which they thought the rock survived or was destroyed. They were then taken to a separate screen where they were asked two comprehension questions (each displayed on a separate screen):

(1) John's experiment worked. (Yes/No)

(2) John hit a rock with a hammer. (Yes/No)

4 people were excluded from the data analysis for missing one or more comprehension questions, leaving a total of 91 responses. The results indicated that there was a large-sized significant difference between the Control (M=5.58, SD=1.30) and Pulverized Cases (M=3.73, SD=2.24), t(89)=4.76, p<.001, d=1.01.

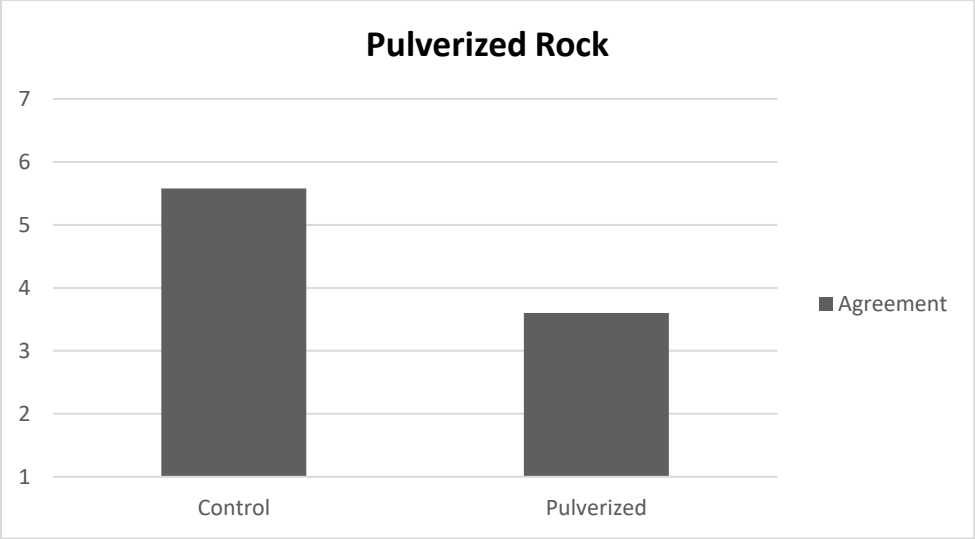

Figure 21: Pulverized Rock

The above results suggest that when a rock undergoes even minor alterations i.e., denting, if it is described as losing its function, people tend to agree that it was destroyed in comparison to a control case. And similarly, in comparison to a control case, when a rock undergoes radical alteration i.e., is pulverized, people tend to say that it survives when it preserves its function.

Perhaps one issue with the previous study is that the extent of the destruction is too ambiguous. Perhaps if pains were taken to make explicit that the rock was radically altered—e.g., smashed to dust—people would be unwilling to say that it survives, despite any preservation of function. Aside from this, the above study used only control cases for comparison and so did not directly manipulate loss/preservation of function in denting or loss/preservation of function in pulverizing. So, this leaves open whether there might be some effect of denting or pulverizing. I'll take all this up in the next and final study.

2.4.5. Denting and Pulverizing with Loss and Preservation of Function

The final set of cases involved John, who discovered a new kind of rock which he named "Zenyte". He noticed tiny worms in the crevices of the rock and gave it to his biologist friend, Frank. Frank notices that the rock transmits chemicals which provide essential nutrients for the worms. The worms begin dying and so he decides to hit the rock with a hammer.

In the denting versions, participants were told that Frank dented the rock with the hammer. They were then either told that the rock did not transmit the unique chemical and the worms continued to die or that the rock began transmitting the chemicals and the worms stopped dying. In the pulverizing versions, participants were told that Frank pounded the rock until it was broke into pieces the size of dust. They were then either told that the rock did not transmit the unique chemical and the worms continued to die or that the rock began transmitting the chemical and the worms stopped dying.[11]

In all cases, participants were told that Frank and John disagreed over whether "Zenyte" was destroyed. Frank says it was not destroyed while John says it was destroyed. In each version of the story, participants were asked to rate the extent to which they agreed with either John or Frank on a 7-pt scale anchored with 1=John is right and 7=Frank is right.

The results indicated a strong effect of whether the rock lost or preserved its function. There was no effect of denting or pulverizing and no interaction between denting or pulverizing and loss or preservation of function.[12] Graphically:

---

[11] For full cases, see Appendix B.

[12] A total of 180 participants were recruited from Amazon's Mechanical Turk. Participants were randomly assigned to one of four conditions in a 2 (Function: Lost, Preserved) x 2 (Damage Type: Denting, Pulverizing) design. After reading the case, participants rated the extent to which they agreed with either John or Frank. They were then taken to a separate screen where they were asked two comprehension questions (each displayed on a separate screen):

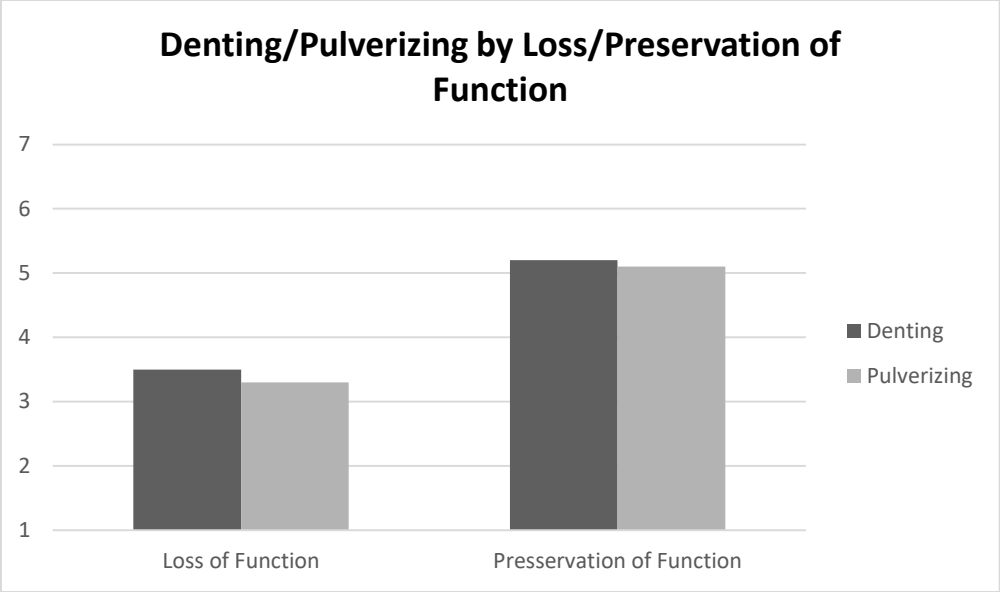

Figure 22: Denting and Pulverizing with Loss and Preservation of Function

Thus, even when making explicit that the rock has undergone radical alteration—being broken into pieces the size of dust—one continues to see a pattern of teleologically driven intuitions about persistence. Indeed, whether the alteration was radical or minor had no effect whatsoever: only teleological considerations affected ordinary intuitions about material object persistence.

Taken together, the above results indicate that the folk tend to judge that a material object survives alterations when it preserves its function and that a material object is destroyed

---

(1) John is a geologist. (Yes/No)
(2) Frank broke the rock into tiny pieces. (Yes/No)
12 people were excluded from the data analysis for missing one or more comprehension questions, leaving a total of 168 responses. The results indicated that there was a large-sized effect of Function, $F(1, 164)=40.842$, $p<.001$, $\eta p^2=.199$; no effect of Damage Type, $F(1, 164)=.467$, $p>.05$; and no interaction between Function and Damage Type, $F(1, 164)=.049$, $p>.05$. Planned pairwise comparisons revealed a large-sized significant difference between Loss ($M=3.53$, $SD=2.09$) and Preservation of Function ($M=5.26$, $SD=1.73$) in the Denting Condition, $t(80)=4.03$, $p<.001$, $d=.91$ and a large-sized significant difference between Loss($M=3.28$, $SD=1.84$) and Preservation of Function ($M=5.13$, $SD=1.51$) in the Pulverizing Condition, $t(84)=5.09$, $p<.001$, $d=1.09$.

when it loses its function. And this persists despite whether the alterations are radical or minor. Given the results and the background psychological discussion of promiscuous teleology, I claim empirical support for the hypothesis that the folk operate with a teleological view of material object persistence. More specifically, I hold that the folk operate with something like the following implicit view of material object persistence:

> *Persistence Through Function Preservation*: Material objects persist through alterations when their functions are preserved.

It may be objected that the results only support the view that teleology is only *one* factor in folk judgments about material object persistence.[13] Perhaps there are additional factors that feature in folk judgments of persistence. I acknowledge that there may be other factors involved in folk judgments of persistence. I would reiterate that my results, as with any empirical work, are defeasible and that I'm open to further work showing that further factors feature in folk judgments of material object persistence. That said, I would only point out that some candidate further factors investigated here—type of object (Sections 2.4.1 and 2.4.2) and extent of alterations (Section 2.4.5)—did not affect folk judgments of persistence.[14] So, while there may be further factors at play in folk

[13] Relatedly, one might wonder whether the above results even support the view that an objects function or purpose affects folk judgments of persistence. Instead, it may be that the preservation or loss of *any* salient feature or property of an object serves as a criterion in folk judgments of persistence. There is, however, considerable reason to doubt that this is the case given the background psychological literature on the role of teleology in our conception of objects. Nonetheless, I looked at vignettes where the rock changed/preserved owners, just to check that people weren't blindly picking up on whatever the vignette discussed as a criterion for persistence. As expected there was no effect of ownership on persistence judgments (see Appendix B).

[14] Connected to this, Rose and Schaffer (forthcoming) found no effect of contact or fusion, of familiarity or labeling or of quantifier restrictions on folk judgments of mereological composition. But they found a large, robust and persistent effect of function or purpose on folk judgments of mereological composition with the folk judging that composition occurred when the plurality had a purpose. Also, after completing this manuscript, Josh Knobe brought to my attention recent work by Nina Strominger and Shaun Nichols (2014) on the persistence of persons as well as a manuscript of his own investigating the role of moral valence in persistence judgments.

judgments of persistence, I take it that teleological considerations at least play a significant role in folk judgments of persistence. Setting this aside, my goal is to debunk folk intuitions on the grounds that they are rooted in primitive, superstitious thinking (Section 2.5). And for that purpose, whether there are additional factors makes no difference to my ultimate conclusion. Thus, I'll set the issue of whether there are further factors featuring in folk judgments of persistence aside in what follows.

2.5. Meeting the Challenge from Folk Belief

Taken together, the results from Section III suggest that the folk view of material object persistence is teleological in that the folk tend to intuit that an object survives alterations when the object preserves its function and is destroyed when an object loses its function. The folk are not selective teleologists; rather, they are promiscuous teleologists in that teleological considerations infuse judgments about persistence for artifacts, organisms and non-living natural objects (Section 2.4.1 and 2.4.2). Furthermore, teleological considerations continue to strongly influence judgments about material object persistence even when alterations to a candidate object are minor or severe (Sections 2.4.3-2.4.5).

Given empirical support for the claim that the folk view of material object persistence is teleological, I now want to consider what, if any, methodological consequences can be drawn. I'll argue that, in the specific case of material object persistence, the folk deserve to be ignored since the folk view is tied into a benighted teleological view of nature. That is, I'll argue that there is a debunking explanation for folk intuitions of material object persistence.

I'll begin by situating the discussion within a background of what Dan Korman (2009) calls the challenge from folk belief (Section 2.5.1). Having presented this challenge, I then show how it can be met in the specific case of material object persistence. To do so, I'll argue that the teleological view that the folk operate with is an illegitimate, benighted form of teleology (Section 2.4.2) and that there is a targeted debunking explanation for folk intuition about material object persistence (Section 2.5.3).

2.5.1 The Challenge from Folk Belief

Metaphysical theories are often judged by their fit with common sense and resultant metaphysical theories are often given more credence when they align more closely with common sense. Indeed, the prescriptive metaphysician who eschews allegiance with common sense accrues an explanatory burden; the prescriptive metaphysician is often charged with the task of explaining just how and where common sense has gone wrong. As Dan Korman (2009) puts it, "[V]irtually 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")." (p. 242).

For the prescriptive metaphysician, providing a plausible account of the conflict is no simple task. Indeed, Paul (2012) tells us that metaphysical theories which "sacrifice central commonsense tenants only rarely convince" (p. 22); Hirsch (2002) claims that "revisionists standardly delude themselves into thinking that they can plausibly explain why people make the mistakes they allege" (p. 117); and Korman (2009) tells us that "despite all that [revisionists] have said on the topic of folk belief—the "scorecard" of

costs and benefits should reflect the fact that they have not met (and probably cannot meet) the challenge from folk belief." (p. 243).

The task for the prescriptive metaphysician then is to meet *the challenge from folk belief*: the prescriptive metaphysician should 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.

An explanation of the alleged error will be globally self-defeating if the alleged source of error is "so thoroughgoing, [that] the most likely source [of the error] would seem to be…a general inability to form true beliefs about the world" since it would then be "unrealistically optimistic for the eliminativists to put any credence in the belief forming mechanisms or lines of reasoning that led them to accept eliminativism" (p. 244). And an explanation of the alleged source of error will be locally self-defeating "to the extent that it undercuts the justification for some specific claim(s) that one has made" (p. 244). For instance, if one locates the alleged source of error in mistaken intuitions and so embraces a general skepticism about intuitive judgments, then one cannot also rely on intuitions to support the premises in an argument for some revisionary or eliminative view. Finally, I would add that, in addition to providing an account that is neither globally nor locally self-defeating, the resulting account must meet a further constraint: (3) the resulting explanation should not be *ad hoc*. In short, the resulting explanation should cohere well with general and independently established psychological claims.

Despite the pessimism over the prospects of meeting the challenge from folk belief, I think the challenge can be met. In particular, the challenge can be met by providing a *targeted* debunking explanation of the relevant folk beliefs which is aided and guided by

work in cognitive science. In this way, one can provide an account which is neither locally self-defeating, globally self-defeating, nor ad hoc and thereby meet the challenge from folk belief. In the specific case of material object persistence, I show that the challenge from folk belief can be met. I'll begin by arguing that the folk operate with a benighted teleological view of nature. Having done so, I will argue that there is a targeted debunking explanation for folk intuitions about material object persistence.

2.5.2. Benighted Teleology

It is widely agreed that teleology is a vestige of an outmoded, pre-enlightenment, Aristotelian perspective on the natural world. It is incompatible with the modern scientific image; introduces occult forces; is mentalistic; is incompatible with mechanistic causation; and is empirically untestable (e.g., Mayr, 1998; Allen and Bekoff, 1994). As such, "the inquisition of Final Causes is barren, and like a virgin consecrated to God produces nothing" (Bacon, 1996, pg., 365). So, insofar as the prescriptive metaphysician views herself as allied to the sciences, she ought to reject a teleological perspective on the natural world and, as such, the folk view of material object persistence, since it is encrusted with the muck and funk of a benighted teleological perspective. As Paul (2012) puts it "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). Given the conflict between the teleological commitments of the folk in the case of material object persistence and the (presumptive) commitments of the prescriptive metaphysician

who views herself as allied with the sciences, in the specific case of material object persistence, I hold that the prescriptive metaphysician should ignore the folk.

One might hold, however, that the folk do not operate with a benighted teleological view. Taking cue from disputes among philosophers of biology, the guiding view is that there are scientifically legitimate forms of teleology and there are scientifically illegitimate forms of teleology. An instance of the former would be teleonaturalism, while an instance of the latter would be teleomentalism. Perhaps the metaphysician who bears allegiance to common sense would hold that folk teleology best fits with teleonaturalism and as such is entirely compatible with a scientific perspective.

Teleomentalists regard "the teleology of psychological intentions, goals, and purposes as the primary model for understanding teleology in biology" (Allen and Bekoff, 1994, p. 13). Whether teleomentalism is taken literally or metaphorically, it is typically regarded as eliminable (Allen, 2009; Allen and Bekoff, 1994). Those who reject teleomentalism—teleonaturalists—"seek naturalistic truth conditions for teleological claims in biology that do not refer to the intentions, goals, or purposes of psychological agents" and so attempt to "reduce teleological language to forms of description and explanation that are found in other parts of science" (Allen and Bekoff, 1994, pp. 13-14).[15] Given that teleonaturalism is typically taken to be a scientifically legitimate form of teleology, perhaps folk teleology is similarly legitimate.

[15] Examples of teleonaturalist accounts can be found in e.g., Millikan (1989) and Cummins (1975).

Though there may be scientifically legitimate teleological notions, folk teleology does not operate with any such notions. Rather folk teleological thinking uses primitive and superstitious notions. Folk teleology best fits the crude superstition of teleomentalism.

A classical demonstration of our adult tendency toward teleomentalism is found in Heider and Simmel (1944), who made a simple movie in which geometrical figures – circles, squares, triangles – moved in certain systematic ways. When shown this movie, people instinctively describe the figures as if they have goals and desires. And more recent research has found that this effect persists even with unbounded figures, such as moving dots and swarms of tiny squares (Bloom & Veres, 1999).

Indeed, Guthrie (1993) presents a range of experiments showing that people attribute purpose and design to a striking range of real-world entities such as cities, clouds, earthquakes, fire, hurricanes, the moon, mountains, plants, rain, the sun, rivers, rocks, trees, volcanoes, water, and wind. Other work on adult judgments of mental states suggests that adults have a tendency to attribute mental states (e.g., feeling pain, being happy) to plants (Arico et al. 2011).

We have what Pascal Boyer (2001) has called a 'hypertrophy of social cognition': a willingness to attribute purpose, agency and design, even when it is inappropriate to do so. We are "hypersensitive to signs of purpose, design and agency, so much so that we see purpose where all that really exists is artifice or accident" (Bloom, 2007, p. 150). Moreover, as Bloom (2007, p. 150) notes:

> We have a bias to attribute an agent when we see nonrandom structure. This is the impetus for the argument for design – the intuition that the design that is apparent in the natural and biological world is evidence for a designer…When we see complex structure, we see it as the product of beliefs and goals and desires. We chew over the

> natural word with our social mode of understanding, and it is difficult to make sense of it in any other way.

As experts who have overcome the naïve folk theory, it can perhaps seem somewhat surprising that the folk teleology best fits with teleomentalism (see "the curse of knowledge" in Section 4.3). In the biological domain, for instance, it can seem incredible that the folk view of biological functions would be rooted in teleomentalism. But as decades of research in scientific education suggest, teleological thinking is one of the primary obstacles in students' path to acquiring an adequate understanding of natural selection (see Galli and Meinardi, 2011 and Kelemen, 2012 for an overview). For instance, students tend to think that a "personified "Mother Nature" responded to animals 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 & 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 & Wood-Robinson, 1985; Demastes, Settlage, & Good, 1995; Evans et al., 2010; Jensen & Finley,1995; Kampourakis and Zogza, 2008). Summing up a range of this work, Kelemen (2012) suggests that people's teleological views are "embedded within a framework of intuitions characterizing Nature as a designing agent" (p. 6).

In support of this, Kelemen (2012) reports on work conducted with Rottman and Seston (2013). In this study participants filled out the Conceptual Inventory of Natural Selection (CINS) and were independently asked a range of questions aimed at assessing religious, scientific and quasi- scientific beliefs such as "I believe Nature is driven to preserve things" . Surprisingly, the results showed that:

> [U]ndergraduates' mean level of agreement…with the scientifically unwarranted statement "I believe Nature is driven to preserve living things" was relatively high (59%) as was their mean agreement with highly correlated statements such as "I believe the Earth is alive" (64%); "I believe that Nature is a powerful being" (73%); "The Earth is driven to provide optimal conditions for Life" (62%). In general then, these students who strongly endorsed natural selection as an explanation of both human (M=82%) and non-human origins (M=81%), had a marked tendency to view the Earth as a powerful, protective, controlling being. More importantly, this agentive view of Nature was found to be highly correlated with students' rather high tendency…to endorse inaccurate…answer options on the CINS... (Kelemen, 2012, pgs 6-7)

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 & Rosset, 2009).

Taken together, the best evidence suggests that telomentalism is the more accurate characterization of folk teleology. This interpretation coheres with a wide swath of research in cognitive science and science education and thus deserves more credence on that basis. That said, one might think that some of the studies I've already presented suggest that the folk view fits with teleonaturalism. For instance, one might think that the rowboat cases (Sections 2.4.1 and 2.4.2) are importantly different from the rock cases (Sections 2.4.3-2.4.5). Specifically, in the rowboat case a person makes it to serve a certain purpose but in the rock case this does not happen, the rock has a sort of "natural purpose". Perhaps this fits a teleonaturalist construal of folk teleology.

To find out whether a teleonaturalist construal of folk teleology best explains the pattern of results in the rock cases, I decided to rerun one of the rock cases, but with a twist. Participants received Rock Three Pieces Cases (see Section 2.4.3) and we're randomly assigned to the Loss and Preservation of Function conditions. The same probe and control questions used in Section 3.3 were used here. Importantly, participants were also asked to rate the extent to which they agreed or disagreed, on a 7-pt scale (anchored with 1=strongly disagree and 7=strongly agree), with the following Gaia Belief probe (taken from Kelemen, Rottman and Seston, 2013):

Gaia Belief Probe: I believe Nature is driven to preserve things.

People's endorsement of quasi-religious Gaia beliefs has been shown to significantly predict endorsement of scientifically illegitimate, teleological explanations for biological and non-living natural phenomenon (Kelemen et al, 2013). Extending this: if Gaia beliefs play a role in people's persistence judgments then we should expect endorsement of Gaia beliefs to predict persistence judgments. Such a finding would support the role of *teleomentalism* in persistence judgments. Alternatively, if Gaia beliefs play *no* role in predicting persistence judgments, then this would provide some evidence that teleomentalism does not play a role in persistence judgments.

First, the pattern of results reported in Section 2.4.3 was replicated, with participants tending to agree that the rock was destroyed when it lost its function (M=5.60, SD=1.54) after being smashed into three pieces and tending to agree that it survived when it preserved its function (M=3.45, SD=2.05) after being smashed into three pieces, F(1,

99)=34.130, p<.001.[16] Second, a multiple regression model, with both Condition and Gaia Belief[17] as predictors of persistence judgments revealed that both factors significantly predicted persistence judgments.[18]

Variable	Beta	t-value	p-value
Condition	-.522	-6.207	.001
Gaia Belief	-.242	-2.883	.005

Table 8: Gaia Belief and Condition as Predictors of Persistence

To get a clearer picture of the relationships among these factors, I ran a causal search over the data using Greedy Equivalence Search (GES).[19] The model returned:[20]

[16] A total of 110 participants were recruited from Amazon's Mechanical Turk. Participants were randomly assigned to one of two conditions (Function: Lost, Preserved). After reading the case, participants rated the extent to which they thought the rock survived or was destroyed. After this they answered the same two comprehension question used above in Section 3.3 and were given the Gaia Belief Probe. The presentation of each of the items—the two comprehension questions and Gaia Belief Probe—was randomized. 10 people were excluded from the data analysis for missing one or more comprehension questions, leaving a total of 100 responses.

[17] Overall, Gaia belief endorsement was fairly high, M=4.8, SD=1.96. The median was 5 and the mode was 7.

[18] Both Condition and Gaia Belief had a large-sized effect on persistence judgments, R^2=.317. R^2 indicates the amount of variance in the dependent variable explained by the linear model. Following Ellis (2010), values greater than or equal to .26 are large, greater than or equal to .13 but less than .26 are medium, and values greater than or equal to .02 but less than .13 are small.

[19] 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; Rose and Nichols, forthcoming.)

[20] This model is a good fit of the data, χ^2 (1)=.3083, p=.5815, BIC=-4.3013

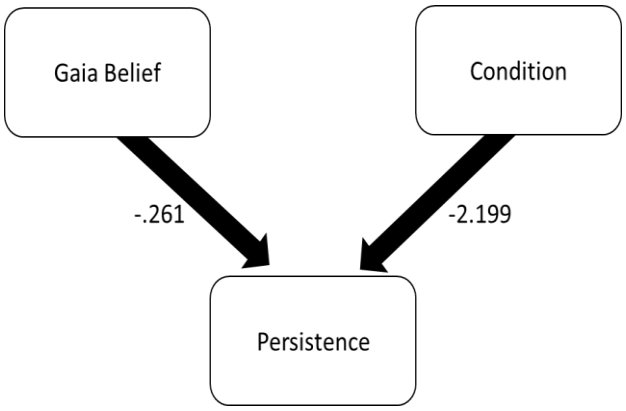

Figure 23: Gaia Belief and Persistence Causal Model

This model shows that persistence is a *collider*:[21] both variables—Gaia Belief and Condition—independently cause persistence judgments. This makes perfect sense. For instance, we would not expect Condition (i.e., whether the rock preserves or loses its function) to cause Gaia Beliefs. Rather, we would expect the quasi-religious Gaia Belief to be held independent of any information about whether an actual object loses or preserves its function. In other words, we would not expect Gaia Beliefs to *mediate* the effect of Condition on persistence judgments.[22] We would, however, expect that if Gaia Beliefs are playing a role in persistence judgments, that they would play a *direct causal role* in persistence judgments. And indeed, this is just what was found: Gaia Beliefs play a direct causal role in people's function-based persistence judgments. Given that Gaia Beliefs did not fail to predict persistence judgments, a teleonaturalist construal of the folk understanding of teleology seems untenable. The evidence suggests that the proper construal of folk teleology is in terms of teleomentalism.

[21] For more, see Alicke, Rose and Bloom, 2011
[22] Put differently: we would not expect changes in whether an actual object preserves or loses its function *to cause* differences in the extent to which somebody endorsed the Gaia Belief.

Given the current result and its coherence with a wide swath of independent psychological results suggesting that folk teleology best fits teleomentalism, I hold that the overall pattern of teleologically driven intuitions observed in cases of material object persistence best fits teleomentalism. As such, the folk operate with a benighted view of teleology. Folk teleology is thus unfit for real metaphysics.

2.5.3 Targeted debunking

Concerning the folk view of material object persistence, I hold that the folk suffer from an *undermining defeater* (e.g., Pollock, 1987). Just as the force of testimony is undermined if it is discovered that the testimony is based on a lunatic view on the topic, so too the force of intuition is undermined if it is discovered that the intuitions are based on a hopeless theory of the topic. As such, I hold that there is a *debunking explanation* for folk intuitions about material object persistence.[23]

Work in folk biology suggests that we tend to judge that something is alive when it exhibits motion. Piaget (1929/1960) famously uncovered an interesting tendency among preschool aged children to attribute life to bicycles, cats, clouds, the sun, snails and water and a tendency to deny that plants are alive. The reason for attributing life to some artifacts, organisms and non-living natural objects but denying that plants are alive is because these varied objects tend to exhibit motion while plants do not (e.g., Richards and Seigler, 1986). And though adults do not tend to judge whether something is alive on the basis of whether it exhibits motion, recent work suggests that these childhood

[23] I follow Kahane (2011) in associating undermining defeaters and debunking explanations: "Debunking arguments are arguments that show the causal origins of a belief to be an undermining defeater." (p. 106). For a connected strategy for debunking—via process debunking—see e.g., Nichols, 2014; Rose, 2014; and Rose and Nichols, forthcoming.

tendencies are never fully outgrown but rather are masked into adulthood. Work by Goldberg and Thompson-Schill (2009)—placing adults in speeded up conditions to prevent their background beliefs from intruding—found that college aged students and even biology professors tended to classify something (rivers, cars, etc) as alive when it displayed motion and deny that it was alive when it failed to display motion (e.g., plants). They conclude that these "biases within biological knowledge appear to reflect developmental ''roots'' that cannot be completely overwritten, or replaced, with the acquisition of more advanced knowledge" (p. 435).

Assuming that intuitions about whether something is alive could be shown to trace from the conceptual connection between life and motion, our intuitions about whether something is alive would be debunked and so should be accorded no weight in disputes about biological classification. Likewise if it is discovered that we intuit that material objects survive through alterations when their functions are preserved and are destroyed when their functions are destroyed, intuitions about material object persistence would equally be debunked and should be accorded no weight in disputes about material object persistence.

Given that the folk view of material object persistence is teleological, I claim that there is a debunking explanation for folk intuitions about material object persistence. But, in offering a debunking explanation for folk intuitions about material object persistence, one must provide an account which is neither globally nor locally self-defeating (Section 2.5.1). Thus, one must show that the view of the prescriptive metaphysician is not likewise infused with teleological considerations.

Research suggests that formal scientific training plays some role in conferring the benefit of overcoming the naïve teleological perspective on the world (Casler and Kelemen, 2008). But, other factors may also play a role. In a study by Kelemen, Rottman and Seston (2013) a range of interesting findings emerged. Participants were current college students, non-student laypersons from the community, professional physical scientists and professionals in various humanities departments.[24] All participants were given various statements that were either scientifically legitimate (e.g., conception occurs because sperm and egg fuse together) or scientifically, illegitimate teleological statements (e.g., the sun radiates heat because warmth nurtures life). Participants were randomly assigned to one of two conditions: a speeded and an unspeeded condition. The purpose of this was to see whether limitations on cognitive resources affected participants' willingness to accept the scientifically illegitimate teleological statements. It turns out that it did: regardless of one's background, participants were more willing to accept teleological statements in speeded conditions. Interestingly, college students (Speeded=56%, Unspeeded=45%) and non-student laypersons (Speeeded=53%, Unspeeded=40%) displayed a strong overall tendency to endorse teleological statements. In contrast, professionals in the sciences (Speeded=29%, Unspeeded=15%) and humanities (Speeded=32%, Unspeeded=21%) displayed a much weaker overall tendency to endorse teleological statements. Thus, college students and non-student laypersons, showed a strong, overall tendency toward teleological explanations, regardless of whether they were making more reflective judgments or not. In contrast, professionals in the

[24] Scientists were from chemistry, physics and geoscience departments and had held a PhD for an average of eight years. Humanities professionals were drawn from classics, English and history departments and had held a PhD for an average of seven years. Professionals from the sciences and humanities were drawn from Columbia, Boston University, Brown, Harvard, Yale and MIT.

sciences and humanities, though having a stronger tendency to endorse teleological explanations when their cognitive resources were limited, showed only a slight tendency to endorse teleological explanations when making more reflective judgments. So, the *reflective judgments* of the experts appear to substantially abate—to the point of near extinction—the tendency toward teleological explanation.[25]

Moreover, Kelemen et al. found that endorsement of the quasi-religious Gaia belief that "Nature is driven to preserve things" did not differ between professionals in the sciences (M=2.2) and humanities (M=2.5).[26] And while college students (M=3.7) and non-student laypersons (M=3.9) did not differ in the extent to which they endorsed the Gaia belief, college students and non-student laypersons showed a much stronger, overall tendency to endorse the Gaia belief in comparison to professionals in the sciences and humanities. So, overall professionals in the sciences and humanities showed both a substantially *weaker* tendency to endorse teleological explanations and the quasi-religious Gaia belief that "Nature is driven to preserve things".

But what accounts for this decreased tendency to endorse both teleological explanations and the quasi-religious Gaia belief among professional scientists and professionals in the humanities? A natural hypothesis is that professionals in the sciences and humanities have significantly more background scientific knowledge. As it turned out, however, Kelemen *et. al.* found significant differences between these groups in background

---

[25] Indeed, one might take low frequency of teleological endorsement among professionals in the unspeeded conditions to reflect mere noise. Though it's not entirely clear what the threshold for mere noise is, one reasonable standard is to treat frequencies significantly below or non-significantly different from 20% to reflect mere noise (see e.g., Murray, Sytsma and Livengood, 2013). If this is right, then the tendency among the professionals to endorse teleological explanations on the basis of reflective judgments might be treated as reflecting nothing more than mere noise. But notice the same could not be said about their unreflective judgments.

[26] Ratings were made of a 5-pt scale with 1=strongly disagree and 5=strongly agree

scientific knowledge. Though professional scientists displayed significantly more background knowledge—as measured by the Conceptual Inventory of Natural Selection and the Geoscience Concept Inventory—professionals in the humanities did *poorly*. Indeed, professionals in the humanities displayed no more scientific knowledge than college students or non-student laypersons. Though it seems that formal scientific training may make *some* difference in the move away from teleological explanations and Gaia beliefs, another factor which may be quite important is the development of more sophisticated conceptual, analytical skills which comes along with extended educational training and experience. And so perhaps this accounts for both the overall decreased tendency among professionals in the sciences and humanities to endorse teleological explanations and Gaia beliefs.

Taken together, these results seem to fit the pattern we would expect if philosophers are indeed experts. Indeed, those who endorse the expertise defense often charge that the intuitions of the folk cannot be trusted since they cannot be relied on to understand the relevant concepts, draw salient distinctions and so forth. Only the intuitions of the philosopher, with her extensive experience in analyzing concepts, drawing salient distinctions and so forth, should be trusted (see e.g., Ludwig, 2007; Rose and Schaffer, 2014; Williamson, 2011). Given that professionals in the sciences and humanities in the Kelemen *et al* studies were not inclined to reflectively endorse scientifically illegitimate teleological explanations it seems that, in the case of persistence, we should expect philosophers to tend *not* to reflectively endorse teleological explanations of material object persistence. And, indeed, this is exactly what we see: metaphysicians do not themselves offer teleological views of material object persistence. Only the folk, who are

caught up in the muck and funk of teleology, display strong teleological tendencies in determining whether an object that undergoes alterations persists.

But, even though expertise confers the benefit of overcoming naïve folk theory, it also comes at the cost of losing track of how the folk think. While the folk might be cursed with a benighted teleological view of nature, the experts are saddled with the curse of knowledge (Camerer, Loewenstein and Weber, 1989): expertise in a given domain biases the experts to project their own sophisticated views onto the folk and so comes with the cost of losing track of how the folk think (Hinds, 1999). Indeed, in the case of material object persistence, it seems that we see the mark of expertise: metaphysicians do not offer teleological views of material object persistence in their explicit theorizing and they also seem to be blinded to the prospects that the folk view of material object persistence is teleological. So, insofar as the metaphysician enjoys in expertise—and so has overcome the naïve folk theory—it seems that a *targeted* debunking explanation for folk intuitions about material object persistence is on offer: the folk, given their lack of expertise, operate with a benighted teleological view of material object persistence while the metaphysician, in virtue of her expertise, has overcome the naïve folk theory and so enjoys in a more enlightened perspective. Thus, I clam that there is a targeted debunking explanation for folk intuitions about material object persistence.

In meeting the challenge from folk belief, three conditions need to be met. The account must not be (1) globally self-defeating, (2) locally self-defeating or (3) ad hoc. Thus, to meet the challenge from folk belief one must offer a *targeted* debunking explanation,

aided by cognitive science.[27] First, I provided empirical evidence of what the folk view is, finding that it is teleological. Second, I situated folk teleology within the psychological background of promiscuous teleology, drawing on the best current empirical evidence which suggests that teleomentalism is the best characterization of folk teleology. I then argued that the views of the metaphysician are not likewise teleologically infused by drawing on empirical work and invoking an expertise defense in the specific case of material object persistence. Thus (1) and (2) are met since the metaphysicians display the marks of expertise and so their reflective judgments appear to issue from an enlightened perspective. And (3) is met since the account, taken together, coheres well with independent psychological claims. The targeted debunking explanation goes through; the challenge from folk belief is met. Taken together, there is a case specific reason for holding that the folk view of material object persistence does not deserve to be taken seriously in disputes over material object persistence.[28] On this matter, *the folk deserve to be ignored.*

2.6. Conclusion

Metaphysicians have wanted a view of material object persistence that fits with common sense. Yet there is disagreement over what the folk view of persistence is and no empirical discipline to the dispute. I suggested that, in measuring metaphysical theories against their fit with common sense, empirically discerning what the folk view is can help

[27] Here I join Goldman (2007) and Paul (2010) in thinking that cognitive science can be useful to metaphysics.

[28] In offering a targeted debunking explanation via an expertise defense for material object persistence, I am endorsing what Rose and Schaffer (2014) call *the nuanced view:* "the prospects for the expertise defense must be evaluated on a case-by-case basis, in empirically disciplined ways" (p. 32; also see Rose and Danks, 2013)

in deciding whether it deserves to be taken seriously. My hypothesis was that the folk view of material object persistence is teleological. I motivated this by situating it within a background discussion of promiscuous teleology (Section 2.3). Given that we tend to determine whether something is by determining whether there is something it is for, I suggested that this basic pattern would extend to the folk view of material object persistence: whether something persists is given by whether it continues to serve its function. And, in a range of studies (Section 2.4), I provided support for the hypothesis that the folk view of material object persistence is teleological.

I then went on, in Section 2.5, to draw out the methodological consequences by situating the discussion within a background of the challenge from folk belief. I then argued for a targeted debunking explanation of folk intuitions of material object persistence by arguing that the folk view is tied into a benighted view of nature (teleomentalism) and by offering an expertise defense in the specific case of persistence.

Taken together, I hold that discussion over material object persistence should be liberated from any demanded conformity with folk intuitions: the prescriptive metaphysician should not be compelled to square her account with the verdicts of the folk. Thus, in the dispute over material object persistence, the folk view should not be taken seriously; there is a case specific reason for holding that metaphysical theories of persistence should not be measured against their fit with common sense.

In liberating the prescriptive metaphysician from any demanded conformity with folk intuitions, one might wonder what metaphysicians can appeal to in deciding between

competing theories of material object persistence.[29] Given that expert metaphysicians disagree over material object persistence and given that the folk operate with a benighted teleological view of material object persistence and, as such, should be ignored in the dispute over material object persistence, what's left? Perhaps appealing to consistency or coherence might help in deciding between competing metaphysical theories of material object persistence. But it appears that some competing metaphysical theories of material object persistence—for instance, three dimensionalism, four dimensionalism and the standard sortal based account—are neither internally inconsistent nor internally incoherent. Perhaps there are still virtues of simplicity, elegance and coherence with wider theory that may help favor certain approaches. Overall, I suspect that liberating metaphysical theories of material object persistence from conformity with folk intuitions tilts the scales in favor of the more elegant four dimensionalist approach. But this is obviously a matter which falls beyond the scope of the current discussion. I'm only defending a descriptive claim about when the folk think that material object persistence occurs and drawing out its methodological implications. I am not defending any claim about when material objects, in fact, persist, nor am I defending any metametaphysical claim about the prospects for metaphysical knowledge.[30]

---

[29] I would like to thank an anonymous referee for raising this issue.
[30] I would like to thank David Danks, Josh Knobe, Shaun Nichols, Laurie Paul, Jonathan Schaffer, Stephen Stich, John Turri and an anonymous referee for helpful comments on earlier versions of this paper.

Chapter 3

Folk Mereology is Teleological (Expanded Version)

3.1. Introduction

When does mereological composition occur? For instance, if a paper plate is positioned on a table between a plastic knife and a metal fork, does this scattered plurality of diverse objects make up a single composite object (a 'table setting') or not? Or if two people shake hands, does this connected plurality of similar objects make up a single composite object (shaped like a sculpture of two people shaking hands) or not? In general, when does a collection of things form a whole? Many metaphysicians have wanted a view of composition that respects folk intuitions, and have charged leading views with failing on this score. For instance, Hirsch (2002, p. 60) declares that 'the linguistic evidence indicates that fluent speakers of English do not speak the mereologist's language.' And Markosian (1998, p. 211) sets out from the claim that 'no one has yet defended a view… consistent with standard, pre-philosophical intuitions about the universe's composite objects.'

Yet there is widespread disagreement among metaphysicians as to what the folk intuit about mereological composition and why they do so, and no empirical discipline to the debate. We see this situation as an opportunity to put the tools of experimental philosophy to constructive use. Accordingly we aim to discover when the folk tend to think that composition occurs, and why they do so. So our question is: *when do the folk*

think that mereological composition occurs? That is, what is *folk mereology*, against which metaphysical accounts of real mereology might be measured?

Our question—beyond whatever intrinsic interest it might possess—should be of interest to anyone interested in the psychological question of how humans conceptualize the world, and in the connected project of descriptive metaphysics. Whether our question is also relevant to prescriptive metaphysics is controversial. For those who take conformity with folk intuitions to be at least one desideratum of theory choice in prescriptive metaphysics, our question bears obvious relevance. But even those who would dismiss folk intuitions as irrelevant to real metaphysics (either because they deny that *intuitions* should play any role, or because they deny that the intuitions of *the folk* should play any role) may still want to know what they would dismiss. Indeed, it seems to us that understanding folk mereology is a precondition to considering whether it deserves to be taken seriously.

Our own view is twofold. First, we hold that folk mereology is *teleological*, in that the folk tend to think that composition occurs in restricted circumstances, in which the question of whether the plurality has a purpose plays a major role. So, for instance, we predict that people will tend to say that composition has occurred with the knife, fork, and plate (since they collectively serve as a table setting) but not with the two people shaking hands (unless they are accorded a collective function). This view seems not even to be considered in the contemporary discussion, though it coheres with a wide swath of current psychological work on object concepts. Secondly, we regard such a folk theory as tied into a benighted teleological view of nature, and thus fit for *debunking*. As such we think that understanding folk mereology should actually lead us to liberate the discussion

of when composition really occurs from any demanded conformity with folk intuitions. On this matter, the folk deserve to be ignored.

Overview: In Section 3.2 we review the existing discussion about when composition occurs, with an eye to claims that have been made about what the folk think. In Section 3.3 we review current psychological work on object cognition and promiscuous teleology, with an eye to documenting the extent to which the folk worldview is teleological. In Sections 3.4-3.5 we use the methods of experimental philosophy to extract a teleological account of folk mereology, articulate some of the details, and connect this to current psychological research. Finally in Section 3.6 we address methodological issues about the role of folk intuitions in real metaphysics, having (we hope) achieved sufficient understanding of folk mereology to see why it should be ignored.

3.2. Composition and Intuition

When does mereological composition occur? The metaphysical debate often centers on claims about what common sense would say. Yet there is disagreement about what common sense would say and why it would say so, and no empirical discipline to the debate.

3.2.1 The Special Composition Question

When does mereological composition occur? That is, under what conditions does some plurality of individuals *X*s compose some one individual *y*? This question is what van Inwagen (1990, p. 30) calls the ‘special composition question,’ and it has set the agenda for one of the main debates in metaphysics spanning the last two decades.

Strictly speaking, we—following van Inwagen—are focused on a question that is more restricted in two respects. First, we are only focused on when composition occurs *for material objects*. So in what follows all quantifiers should be read as restricted to material objects (unless the local context makes obvious otherwise). Perhaps mereological composition can occur among events or among abstract objects or among entities of some other sort, or even across categories of objects. Such is not our concern.[1] Secondly, we are only interested in when *many* make one. It is standard to use 'plurality' in a way that is actually number *neutral*, allowing for pluralities of one or more individuals. With 'degenerate pluralities' of just a single individual, it is—at least on one usage of terms—trivial that composition occurs (everything composes itself). It is smoother to state the views in ways that ignore degenerate pluralities, and so we follow suit. Bringing this together, a more explicit formulation of our question is: under what conditions does some non-degenerate plurality of material objects *X*s compose some one material object *y*?

The literature offers a wide variety of proposed answers to the special composition question. Perhaps the most standard answer is *always, under any condition whatsoever*. This is the answer of the Universalist (built into the classical mereology of Lesniewski and Goodman [c.f. Simons 1987, pgs. 37-41], and defended by Lewis, 1991). A second sort of answer to this question is *never, under no condition whatsoever*. This is the answer of the Nihilist (explored by Rosen & Dorr, 2002 and recently advocated by Sider,

[1] There is debate as to whether there is a single category-neutral relation of composition, or perhaps a range of analogous category-restricted relations (c.f. van Inwagen 1990, pgs. 18-20). On the one hand we do use 'part' for relations not just among material objects, but also among spatiotemporal regions, events ('the inning is part of the baseball game'), and abstracta ('the chapter is part of the book,' 'the hypotenuse is part of the triangle'). On the other hand it is not obvious that all of these are literal applications of one and the same neutral notion. We also use temporal notions in describing abstracta ('the sequence converges rapidly'), but no one takes this to indicate that we need a single neutral account of time that equally covers abstracta.

2013). Various intermediate answers are considered as well, such as *sometimes, when the plurality is in contact,* and *sometimes, when the activities of the plurality constitute a life* (the former is van Inwagen's first 'representative answer,' and the latter is his final considered position).

Without further discussion, we will simply tabulate various answers in the current literature. This table is not intended to be exhaustive, but just to illustrate the main options under consideration:[2]

Universalism (Lesniewski, Goodman, Lewis)	Composition always occurs
Nihilism (Rosen & Dorr, Sider)	Composition never occurs
Contact (van Inwagen's first illustrative view)	Composition occurs when the plurality is in contact
Fastening (van Inwagen's second considered view)	Composition occurs when the plurality is fastened together
Vitalism (van Inwagen's final considered view)	Composition occurs when the activities of the plurality constitute a life
Emergentism (Merricks)	Composition occurs when the plurality exhibits novel and irreducible collective powers

[2] Serialism and Brutalism differ from the preceding seven views on the table over whether a general and informative answer to the special composition question can be given, though this difference plays no role in our discussion.

Regionalism (Markosian's latter view)	Composition occurs when the fusion of the regions occupied by the plurality is occupied by an individual
Serialism (Sanford)	Composition occurs if the plurality is so and this condition is met, or if the plurality is such and that condition is met, or if the plurality is thus and the other condition is met, or …
Brutalism (Markosian's earlier view)	Composition occurs when it does, as a brute matter of fact

Table 9: Answers to the Special Composition Question

3.2.2 The Role of Folk Intuitions

Virtually every single view on the table as to when composition occurs has been charged with violating common sense, though the charges are usually disputed. In this vein Markosian (2014) charges the main views with generating 'wildly counterintuitive consequences' and so evincing 'mereological madness.' In the interests of brevity we will focus on the role of intuitions in the debate over *Universalism*.

It is often said that *Universalism* posits bizarre fusions which common sense would reject. Here is a representative quote, from Hirsch (2002, p. 60):

> I understand perfectly well what it means to talk (in plain English) about such things as cars, bees, human beings, books, and the Eiffel Tower, or even to talk about such marginal things as noses and car-hoods. But it's crazy to say (in plain English) that there exists something composed of my nose and the Eiffel Tower.

Indeed Markosian (1998, p. 228) considers this sort of charge to be ‘a fatal objection’ to Universalism, and Kriegel (2011, p. 198) tells us what the folk will think: ‘Commonsense shuns [arbitrary] fusions, but… mereological universalists… embrace them. The folk’s intuitive verdicts will be *against* ‘there is a fusion of this table and the moon’…’

But the friends of *Universalism* usually reject this charge. As Korman (2008, p. 320; c.f. van Inwagen 1990, p. 75) observes: ‘[U]niversalists typically take the view to be entirely compatible with what the folk say in ordinary discourse about material objects.’ In this vein, Lewis (1991, p. 80), discussing the fusion of the front of a trout with the back of a turkey, explains away hesitation to affirm existence via quantifier domain restriction:

> Only if you speak with your quantifiers wide open must you affirm the trout-turkey’s existence. If, like most of us all the time and all of us most of the time, you quantify subject to restrictions, then you can leave it out.

And in a slightly different vein, Thomasson—who (2007, p. 3) explicitly aims to show ‘how, reflectively, we can make sense of our unreflective common sense worldview’—maintains (2007, p. 183) that the folk simply have no view one way or another on arbitrary sums, never having considered them:

> Certainly it is true that common sense does not recognize the existence of gollyswoggles, mereological sums, and the like. Nor, of course, does it *deny* their existence—There are no terms in ordinary English for these things, and common sense understandably does not consider such things at all, since given our current range of practices, such entities would be quite irrelevant and uninteresting.

Indeed she (2007, p. 184) then speculates that if ‘we explained to ‘normal’ people’ a term for some arbitrary sum, and then asked them if there is such a thing, they ‘would certainly accept that there is.’

So we find a debate over what the folk think. Is *Universalism* a form of 'mereological madness' or 'entirely compatible with what the folk say'? We see little hope in settling this debate without empirical work.

We could easily spend many more pages documenting further appeals to folk intuitions in the debate as to when composition occurs, but will instead just flag some main points:

- virtually every answer to the special composition question (not just *Universalism*) has been charged with violating folk intuitions,[3]
- metaphysicians have moreover constructed psychological theories about why the folk have the intuitions they allegedly have,[4] and
- teleological notions are almost entirely absent from the current debate, including when the metaphysicians are speculating about what intuitions the folk have and why they have them.

Just on the last point: teleological notions do not come up when metaphysicians are presenting their own theories.[5] This is not so surprising: orthodoxy has it that teleological notions are a vestige of an obsolete conception of nature—'part of a superseded, pre-

[3] For the interested reader: On *Nihilism*, Bennett (2009, p. 44), Schaffer (2009, p. 358), and Kriegel (2011, p. 198) are among the many who charge that it violates commonsense, though Rosen & Dorr (2002, p. 158) claim that commonsense is on reflection neutral. On *Contact* and *Fastening*, van Inwagen (1990, p. 34) and Markosian (1998, p. 223) claim these views to be close to commonsensical, but van Inwagen (1990, p. 35; pgs. 57-59; cf. Markosian 1998, p. 224) offers an intuitive counterexample (his handshake case). On *Vitalism*, Hirsch (2002, p. 67) says that it is hard 'to keep a polite straight face' at the idea that there are apple trees but no apples, while van Inwagen (1990, p. 103) claims no conflict with 'Universal Belief.'

[4] For instance, Lewis (1986, p. 211) claims that the folk focus on qualities, locations, and causal connections, Hoffman and Rosenkrantz's (1997, p. 73) speak of a 'key commonsense intuition' in terms of the parts being joined together so as to be pushed or pulled together, and Simons (1987, p. 303) speaks of a kind of intuitive 'integrity of internal connectedness' of wholes which he understands in terms of ontological dependence.

[5] The literature is vast and of course there are exceptions to the general rule. For instance, Rea (1998, p. 354) speaks of having a possible purpose as sufficient for composition. And—as we discovered after writing this paper—Bowers (*manuscript*) defends a teleological view of composition more or less exactly in line with the view we attribute to the folk.

scientific muddle about how the world works' (Hawthorne & Nolan 2006, p. 267; c.f. Jenkins & Nolan, 2008)—and thus unfit for real metaphysics. What is more surprising is that teleological notions do not come up even when metaphysicians are speculating about what the folk think. It almost seems as if the metaphysicians regard teleology as being so muddled that they cannot charitably imagine the folk wallowing in it.

3.2.3 Aristotelian Roots

While teleological notions are almost entirely absent from the current debate, we would be remiss not to mention their roots in a broadly Aristotelian worldview.[6]

For Aristotle, substances are composites of matter and form, where the form plays the role of unifying the parts of a composite substance. The form unifies the parts of a composite substance by lending them a unified purpose: the parts become one because they act for the sake of a common end. This is the perspective that Aristotle seems to take when defending natural teleology in *Physics* (Aristotle 1984a, p. 340; *Phys*199a.30-33), saying: '[S]ince nature is twofold, the matter and the form, of which the latter is the end, and since all the rest is for the sake of the end, the form must be the cause in the sense of that for the sake of which.' And this perspective is perhaps most explicit when Aristotle is explaining the many senses in which things are called one in *Metaphysics*, in the course of which he (1984b, p. 1605; *Meta*1016b.12-16) specifies the sense in which something is called one as 'a whole,' which he glosses in terms of having 'one form,' and illustrates with the example of the shoe: '[I]f we saw the parts of a shoe put together anyhow we

[6] It is probably impossible to say anything about Aristotle without incurring scholarly controversy. For present purposes we only claim 'broadly Aristotelian' roots. For a detailed discussion of Aristotle on mereology, see Koslicki, 2008 (ch. 6.).

should not call them one…; we do this only if they are put together so as to be a shoe and have thereby some one form.'

We are not alone in finding a teleological account of composition in Aristotle. Indeed, we find Charles's (2001, p. 100) discussion of Aristotle on unity especially clear on this point:

> [I]t is the presence of a final cause which makes the relevant planks and bricks into a house. *Being a house*, on this view, cannot be reduced to being a mereological sum of bricks and planks. There is more to its unity than that of the sum of its components and their physical interrelations; for merely to specify these is to ignore the goal whose attainment is required if there is to be a house… Houses are the result of the operation of the final cause as a principle which organizes the relevant type of matter…

So we see in Aristotle a teleologically-laden view on when composition occurs, and thus want to credit Aristotle with delivering an account that proves to be insightful for folk metaphysics (even if that rules it out for real metaphysics).

3.3. Psychological context

Empirical work on folk mereology need not begin *ex nihilo*. For while there is no psychological work that specifically considers the special composition question, there is a rich body of psychological work surrounding our object concepts, which ought to guide informed inquiry into folk mereology. A major theme emerging in this psychological work is the extent to which the folk worldview is *teleological*. We have a deep-seated tendency to view not just animals but all of nature—every rock and cloud—as infused with agency and purpose. As Dawkins (1995, p. 96) observes: 'We humans have purpose on the brain. We find it hard to look at anything without wondering what it is 'for,' what the motive for it is, or the purpose behind it.'

It is fairly well-established that when it comes to *object categorization*—saying what sort of thing something is—people tend to associate what something is with what it is for. Against this psychological background, our thesis can be understood as the thesis that people also tend to associate *whether* something is with *whether* there is something it is for.

3.3.1 Selective Teleology: Artifacts and Organisms

It is widely accepted that we are at least 'selectively teleological,' in that our artifact and organism concepts are infused with notions of purpose and function. Most of the relevant psychological work concerns principles of *object categorization* (how we determine what a thing is). With both artifacts and organisms, we tend to identify what a thing is with what it is for.

So starting with artifact concepts, as German & Johnson (2002, pgs. 279-80) note, it is well established that people take 'the *design stance,* in which an entity's properties, behavior, and existence is explained in terms of its having been designed to serve a particular purpose.' Likewise Bloom (1996, p. 3), reviewing earlier work by Rips and by Keil, notes: 'This has suggested to many scholars that the psychological 'core' of artifact concepts is that their members share a common intended function.' To illustrate, Rips (1989) found that adults judge that an object that looks like a lampshade is actually an umbrella when they are told that it was originally designed to protect people from rain.

With artifact concepts, there is a robust debate over the developmental details. For instance, Kelemen (1999a) argues that children take the design stance as early as four, Matan & Carey (2001) argue for a more complicated developmental process in which the design stance is not fully in place until six years of age, and German & Johnson (2002)

argue for a different developmental process in which the ability to use an integrated adult-level design stance continues to develop in the period between five to seven years of age. We are only interested in adult competence, which all sides agree is based on the design stance, and we remain neutral on how and when this competence develops.

Turning to the organism concepts of 'folk biology,' it is well established that we strongly tend towards teleological thinking in this domain as well. For instance Keil (1995, p. 245) writes:

> Historically there have been many arguments for a 'design' stance, which can include teleological interpretations and tool construction and use… Notions of functional architecture are among the most cognitively compelling ways of approaching the biological world and much of the artificial world as well, …

Likewise Atran (1998, pgs. 550-51) speaks of the folk idea of a 'biological essence' as 'an intrinsic… teleological agent, which physically… causes the biologically relevant parts and properties of a generic species to function and cohere 'for the sake of' the generic species itself.'

Again there is robust debate concerning 'folk biology,' as to the extent to which biological cognition is domain-specific (perhaps subserved by a partially encapsulated 'folk biology module') or an application of more general-purpose cognition. And there is debate concerning the details of how we naturally organize the biological world hierarchically.[7] Again we remain neutral on these debates, maintaining only the minimal and uncontroversial claim that organism concepts are teleologically infused.

3.3.2 Promiscuous teleology

[7] See the response pieces to Atran (1998) for a useful overview of the main positions.

So far we have discussed object categorization with artifact and organism concepts. One might think, in accord with the approach that Kelemen (1999a, p. 243) calls 'selective teleology,' that teleological thinking is largely restricted to these domains. But, as Kelemen (1999a, p. 244) argues, the more psychologically plausible view is that of 'promiscuous teleology,' according to which teleological thinking is not restricted:

> [T]he teleological stance derives from children's understanding of agency and intentional object-directed behavior and may never become entirely autonomous from the intentional domain… [D]ue to these origins the teleological stance is applied broadly rather than selectively early in development: Infants may start out generally assuming that objects exist to be used by agents in some way and subsequently, in lieu of alternative explanations, develop the teleological belief that virtually all sorts of living and non-living entities are intentionally caused for a purpose. Children may only begin to revise and restrict this belief once they begin to assimilate more formal scientific ideas.

Indeed Kelemen (1999a, p. 245) goes on to note a historical tendency to view all of nature as an artifact:

> [T]hroughout history, non-living natural objects have also been considered in such [teleological] terms… The earth, its climates, landforms, water sources, and elements, were seen as intentionally designed to create a habitat for, and meet the needs of, people. In other words, natural objects of all kinds—particularly those fulfilling a significant function in people's lives—were candidates for construal as quasi-artifacts.

Moreover Kelemen & DiYanni (2005; c.f. Kelemen, 2004) report a strong tendency among children—both from religious and nonreligious backgrounds—to an 'intuitive theism' in which nature is viewed as an artifact of a creator, as well as a significant correlation between viewing something teleologically and regarding it as created. So one should expect teleological thinking to extend through to our general conception of an object. As Bloom (2007, p. 150)—in an article entitled 'Religion is Natural'—summarizes:

> One of the most interesting discoveries in the developmental psychology of religion is

> that the bias towards creationism appears to be cognitively natural. Four-year-olds insist that everything has a purpose, including things like lions ('to go to the zoo') and clouds ('for raining'). When asked to explain why a bunch of rocks are pointy, adults prefer a physical explanation, while children choose functional answers, such as 'so that animals could scratch on them when they get itchy.'

Thus it seems that the four-year old view has it that 'everything has a purpose,' not just lamps and lions, but even rocks and clouds.

We take the main unresolved issue concerning 'promiscuous teleology' to be the extent to which it extends beyond children to adults, with even adults retaining the mindset of 'everything has a purpose.' While children show a strong tendency to prefer teleological explanations to mechanistic explanations across the board, adults in contemporary Western cultures tend to resist teleological accounts when considering inanimate natural things like rock piles. For instance, Kelemen (1999b, pgs. 1443-44) asked both children and adults why a certain rock was pointy, and found that children tend to resist a mechanistic explanation in terms of bits of stuff being piled up, and instead prefer the following (bizarre) teleological explanation: 'so that animals wouldn't sit on them and smash them.' She found a strong preference for this style of teleological explanation among first and second graders, which persisted (albeit in diminished form) even among fourth graders, but was finally reversed with adults. Extending this research, Lombrozo & Carey (2006, p. 180) found that 'adults accepted teleological explanations selectively,' summarizing (2006, p. 184):

> We found that adults accept teleological explanations when two conditions obtain: (a) the function invoked in the explanation played a causal role in bringing about what is being explained and (b) the process by which the function played a causal role seems general, in the sense that it conforms to a predictable pattern.

That said, there is recent converging evidence that our childhood tendencies to teleological thinking persist through adulthood, being merely occasionally masked. Thus Lombrozo, Kelemen & Zaitchik—investigating the recurrence of teleological explanation in patients with Alzheimer's Disease—claim (2007, pgs. 999-1000) that 'an underlying tendency to construe the world in terms of functions persists throughout life' and represents an 'explanatory default.' They (2007, p. 1004) conclude:

> [T]he preference for teleology is never outgrown. Rather, the preference persists throughout life, reemerging when causal beliefs that might otherwise constrain it are limited or compromised. In short, these findings provide evidence for a basic human preference to understand the world in terms of purpose. When faced with an object that supports a plausible function, humans make an immediate but defeasible inference to design, and assume a teleological explanation is warranted.

In a similar vein, Kelemen & Rosset—speeding up adults to prevent their background beliefs from intruding—elicited explicitly teleological judgments even in scientific contexts, concluding (2009, p. 143): '[T]he bottom line implied by the current findings remains that, like children, college-educated adults display scientifically unwarranted teleological explanations with ease.' And perhaps the most interesting and recent results, due to Kelemen, Rottman & Seston (2013, p. 1079), involved research on an expert population of physical scientists, with the finding that: 'even physical scientists, despite their extensive scientific training, routine adoption of physical-causal explanations, and anti-teleological norms, default to scientifically inaccurate teleological explanations when their cognitive resources are limited.'[8]

[8] It may be useful to invoke the image of a 'dual processing system,' on which the human mind is viewed as having both 'Type 1' automated, encapsulated, and intuitive animalistic systems, as well as 'Type 2' deliberative, general, and reflective systems layered over the Type 1 systems by evolution. Then we might say that teleological thinking is the product of Type 1 systems. To the extent that college-educated adults can sometimes avoid teleological thinking, it is only through the effortful employment of a trained Type 2 cognitive mechanism. In this vein Kelemen, Rottman & Seston (2013, p. 1075) characterize their view as 'akin to dual-processing models that characterize early developing intuitions as heuristics that can be

3.3.3 Folk Teleology as Unscientific

In the psychological literature, the folk tendency to teleological explanation is generally recognized as *an error*, and indeed an obstacle students face in properly understanding processes such as natural selection (c.f. Galli & Meinardi, 2011, Kelemen, 2012). In this vein, Kelemen (2012, p. 68; c.f. Kampourakis & Zogza, 2008, Moore *et al.*, 2002, Gregory, 2009) notes that students tend to think that a 'personified 'Mother Nature' or 'Evolution' responded to the functional needs of animals by generating or conferring the functional part with a view to preserving the animal's survival,' such as by stretching a giraffe's neck so it could reach its food. She (2012, p. 71) goes on to explain:

> 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.

The dismissal of teleological explanation from real science goes back at least to Bacon and the emergence of modern science from medieval Aristotelianism. As Bacon himself (1996, p. 365) memorably declares: 'For the inquisition of Final Causes is barren, and like a virgin consecrated to God produced nothing.'

That said, we recognize that some philosophers—especially in response to the recurrent

---

increasingly overridden later in development by effortful processing, but which can nevertheless persistently reemerge in cases when intuitions are favored or forced.'

use of teleological notions within biology—have sought to legitimize a revised kind of 'naturalistically acceptable' teleological thinking (Cummins, 1975, Millikan, 1989). Indeed one reaction we have encountered to our work (in certain philosophical quarters) is to ask whether folk teleology is so wrong. So—to pick up on the terminology of Allen & Bekoff (1995, pgs. 13-14)—one may distinguish the crude and unscientific idea of *teleomentalism* (or 'teleology-heavy') which regards 'the teleology of psychological intentions, goals, and purposes as the primary model for understanding teleology in biology,' from a revisionary notion of *teleonaturalism* ('teleology-lite') through which one may 'seek naturalistic truth conditions for teleological claims in biology that do not refer to the intentions, goals, or purposes of psychological agents' and so attempt to 'reduce teleological language to forms of description and explanation that are found in other parts of science.'

We grant—if only for the sake of the argument—that there may be some revised and scientifically legitimate teleonaturalist notions which philosophers of biology might usefully identify. Our point is simply that the folk are not teleonaturalists but full-blown teleomentalists, indulging in the heavy mentalistic projection of agency onto the entirety of nature. A classical demonstration of our adult tendency toward teleomentalism is found in Heider & Simmel (1944), who made a simple movie in which various geometrical figures—circles, squares, triangles—moved in certain systematic ways. When shown this movie, people instinctively describe the figures as if they have goals and desires. This effect persists even with unbounded figures, such as moving dots and swarms of tiny squares (Bloom & Veres, 1999). Other work suggests that this tendency to attribute agency and purpose extends to real world entities including cities, clouds,

earthquakes, fire, hurricanes, the moon, mountains, plants, rain, the sun, rivers, rocks, trees, volcanoes, water, and wind (Guthrie 1993).

(In what follows we will continue to use the term ‘teleology,’ but we are throughout referring to the illegitimate superstition of teleomentalism, and not to any potentially respectable but revisionary version of teleonaturalism that philosophers of biology might identify.)

Putting this together: Background work in psychology suggests that adults indulge in promiscuous teleological thinking. In particular, we tend to classify objects by their purposes. There is a natural connection between object categorization and the special composition question, even beyond the fact that both involve object cognition (and for that reason alone might already be thought to follow common principles). For one can think of the special composition question as asking *whether a thing is* (that is, asking whether there is an individual *y* that fuses the plurality of individuals *X*s). And there is a natural connection between how we determine *what a thing is* (object categorization) and how we determine *whether a thing is* (folk mereology). Indeed, one can plausibly equate the question of whether a thing is with the question of whether there is anything that it is (that is, one can equate existence with falling under some category or other). When the plurality of individuals collectively falls under an object category, then (and only then) should we be expected to say that the plurality has a fusion. When the plurality is *for something* then it *is something*. But when the plurality lacks a purpose—when it is *for nothing*—then it *is nothing*. Against this background of promiscuous teleological thinking even by adults, our hypothesis that folk mereology is teleological should be unsurprising. Our hypothesis deserves high initial credence.

3.4. Folk Mereology is Teleological

So far we have displayed a range of armchair disputes about folk mereology (Section 3.2), and documented some background psychological work on folk teleology (Section 3.3). We now describe a range of studies we conducted which support the idea that folk mereology is teleological, in the sense that people tend to intuit that a plurality of objects has a fusion partly on the basis of considering when that plurality serves a purpose.

We did face a methodological difficulty worth flagging from the start. Simply asking people the special composition question directly ('Under what conditions does some plurality of individuals *X*s compose some one individual *y*?') will presumably elicit something between blank incomprehension and pure noise. Attempting to teach the relevant concepts first comes with no guarantee of success (and may generate bias, and will generate cognitive load). So instead we chose to design a variety of surveys describing various particular cases and then posing composition questions in various ways (e.g., do two mice glued together 'compose a new object'? do two ropes tied together 'create a single, unified object' or still leave one with 'two, distinct pieces of rope'?) We hoped that our results might prove *robust* across these diverse vignettes and probes, thereby buttressing the conclusion that we have been uniformly successful in eliciting the intended mereological concepts. Our hopes turn out to be realized. But still, we would flag the concern that some of our questions may have been read in an unintended way.

As with all empirical work, our results are subject to potential confounds and diverse interpretations, and defeasible in the light of future inquiry. We think that any open-minded person who considers all of our studies together will agree that they point overall

towards a teleological view of folk mereology (especially given the psychological background of promiscuous teleological thinking). But each specific study may be questioned, and of course future results may always point in other directions, or towards a more specific form of a teleological view. Overall we hope to have provided the first but not the last word on folk mereology.

3.4.1 Handshake Cases

We begin with studies we ran based on van Inwagen's (1990, p. 35; pgs. 57-59) famous 'handshake case.' Van Inwagen—after hypothesizing that the folk theory is something like *Contact* or *Fastenation*—asks one to imagine two people shaking hands (thereby coming into contact), or even gluing their hands together (so as to become fastened together), and then to consider whether 'a new thing' has 'at that moment come into existence.' If the folk theory were *Contact* or *Fastenation* one should think 'yes.' But if the folk theory were *Nihilism, Vitalism,* or *Emergentism* one should think: 'no, there is no sum, nothing has coming into existence.' While if the folk theory were based on *Universalism* one should think: 'no, there always was a sum—nothing *new* has come into existence.' (If the folk theory were based on *Regionalism*, *Serialism,* or *Brutalism* then no prediction can be made one way or another.)

We had two goals in mind. First we wanted to test whether van Inwagen was right to predict that most people would say 'no' to the question of whether a new thing has come into existence. Secondly we wanted to see if we could *flip intuitions* by manipulating function. So we first tested van Inwagen's initial question, as to whether two people shaking hands and thereby coming into contact thereby create a new thing. We set up a

vignette involving a handshake, and then described a disagreement between two characters—Andy and Liz—as to whether a new larger object was thereby created:

> *Handshake*
>
> Sally and Tom are leaders of rival political factions, and have recently decided to lay aside their differences. They have worked out all the details, signed all the official papers, and will now seal their deal with a public and historic handshake.
>
> Later that day, Andy and Liz—who were both present for the historic handshake—have a disagreement over whether Sally and Tom created a new object when they shook hands. Andy says that simply coming into contact with someone or something is not enough to create some new thing, and claims that Sally and Tom did not create a new larger object in the moment when they shook hands.
>
> Liz, however, disagrees. She thinks that when Sally and Tom came into contact in this case, they thereby created a new larger thing, made from Sally and Tom together.

Participants were asked the extent to which they agreed with either Andy or Liz (they were given a seven point scale with 1 marked 'Andy is right,' 4 marked 'Neither is right,' and 7 marked 'Liz is right'). Just as van Inwagen predicted, participants tended to side with Andy (M=2.48, SD=1.84).

We then made minimal adjustments to the vignette so that the larger object formed served a purpose, by adding that a sculptor would use it as a model for a sculpture of two people shaking hands:

Handshake with function

Sally and Tom are leaders of rival political factions, and have recently decided to lay aside their differences. They have worked out all the details, signed all the official papers, and will now seal their deal with a public and historic handshake. To commemorate this historic event, a sculptor has been commissioned to sculpt the handshake. Sally and Tom together, while they are shaking hands, will be providing a model for the sculpture, which will be dubbed 'Unity.'

Later that day, Andy and Liz—who were both present for the historic handshake—have a disagreement over whether Sally and Tom created a new object when they shook hands. Andy says that simply coming into contact with someone or something is not enough to create some new thing, and claims that Sally and Tom did not create a new larger object in the moment when they shook hands.

Liz, however, disagrees. She thinks that when Sally and Tom came into contact in this case, they thereby created a new larger thing, made from Sally and Tom together, which served to provide a model for the sculpture.

Adding a function produced a large-sized effect on intuitions. Participants in *Handshake with function* tended to side with Liz (M=4.86, SD=1.60), agreeing that 'when Sally and Tom came into contact in this case, they thereby created a new larger thing, made from Sally and Tom together, which served to provide a model for the sculpture.' Here is a visual depiction of the effect of function on intuitions:[9]

[9] A total of 62 participants were recruited through Amazon's Mechanical Turk and randomly assigned to either *Handshake* or *Handshake with function*. After reading the case and indicating who they agreed with (Liz or Andy), participants were then taken to a separate page where they answered three comprehension questions:

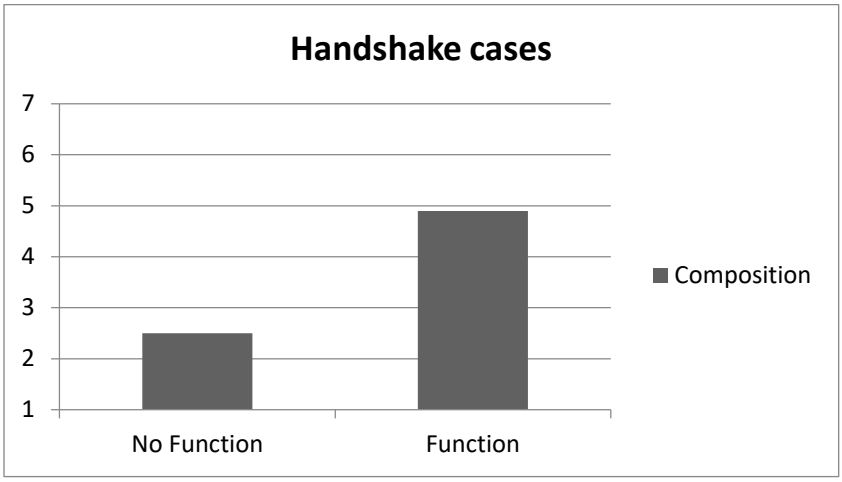

Figure 24: Handshake Cases

While our primary claim is just that adding a function produced a large-sized effect on intuitions, we also note that the overall effect in this case was to flip intuitions. In both *Handshake* and *Handshake with function* the responses were significantly different ($p<.01$ in both cases) from the midpoint (4.0). So we were able to take a paradigm case in

1. Sally and Tom are leaders of rival political factions (yes/no)
2. Andy thinks that the handshake did not create a larger object (yes/no)
3. Liz thinks that the handshake did create a larger object (yes/no)

Four participants missed one or more comprehension questions in *Handshake* while three participants missed one or more comprehension questions in *Handshake with function*. After excluding these participants, a total of 55 responses were analyzed using a one-way analysis of variance. We found that whether or not the larger object served a purpose (No Function: M=2.48, SD=1.84; Function: M=4.86, SD=1.60) produced a statistically significant large-sized effect of people's composition judgments $F(1, 55)=26.824$, $p<.001$, $\eta p^2=.383$. Moreover, responses in both the No Function and Function cases were significantly different from the midpoint: No Function $t(28)=-4.43$, $p<.001$; Function, $t(27)=2.82$, $p<.01$.

Throughout we report effect sizes. For the studies (such as the Handshake cases) which included a scaled, dependent variable, we report effect sizes using partial Eta squared (ηp^2), which is the amount of variance in the dependent variable explained by a given independent variable plus its associated error variance. For studies with a binary dependent variable, we report effect sizes using Cramer's V, which is a nonparametric correlation coefficient that indicates the strength of the relationship between nominal variables. Both of these measures deliver a value between 0 and 1. We follow Ellis (2010) for interpreting the magnitude of the effect sizes. So for Cramer's V we interpret values greater than or equal to .5 as *large*, greater than or equal to .3 but less than .5 as *medium*, and greater than or equal to .1 but less than .3 as *small*. And for ηp^2 we interpret values greater than or equal to .14 as *large*, greater than or equal to .06 but less than .14 as *medium*, and greater than or equal to .01 but less than .06 as *small*.

the literature of when composition fails to occur, and flip intuitions about composition by manipulating function. None of the extant answers to the special composition question (Section 3.2.1) predict this pattern. So, to the extent that metaphysicians are beholden to folk intuitions, we have just provided an empirically substantiated counterexample to virtually every extant account.

With these Handshake cases we see a first sign of an underlying pattern, in which intuitions about whether or not composition occurs are significantly impacted by whether or not there is a function for the candidate larger thing to serve. Our remaining studies further illustrate and clarify this pattern.

3.4.2 Gollywag cases

In another series of studies we wanted to consider an unfamiliar type of artifact ('gollywags') and consider the effect of function along multiple dimensions. We wanted to see whether there was an effect of function, and also look for any interactions with contact or with fastening. So we began with:

> *Gollywags with fusion*
>
> Acme Inc. is a large research company. Two Acme Inc. researchers, Jones and Smith, have recently discovered a new thing, a 'gollywag.' Nobody has ever seen or heard of such a thing so Jones and Smith were quite surprised to stumble upon this new thing.
>
> One day, Jones takes two gollywags and superglues them together. He thinks that he has created a new object.
>
> Later that same day, Smith and Jones have a disagreement over whether the arrangement of gollywags composes a new object. Smith claims that the

arrangement of gollywags does *not* compose a new object at all. He claims that simply supergluing some stuff together with some other stuff is *not* sufficient to compose a new object.

Jones, however, disagrees. He claims that simply supergluing some stuff together with some other stuff is sufficient to compose a new object. Concerning the specific case that they are considering, Jones concludes that the arrangement of gollywags does compose a new object.

We also had a *Gollywags with contact* case, just like *Gollywags with fusion* except that Jones, instead of supergluing the gollywags together, simply placed them into contact with one another.

We compared *Gollywags with fusion* with a case that added a function for the superglued gollywags:

Gollywags with fusion with function

Acme Inc. is a large research company. Two Acme Inc. researchers, Jones and Smith, have recently discovered a new thing, a 'gollywag.' Nobody has ever seen or heard of such a thing so Jones and Smith were quite surprised to stumble upon this new thing.

Jones is a very hard worker and spends long hours in his office, spending most of his time sitting at his desk. On most days, his back becomes very sore from sitting for so long at his desk.

One day, Jones is working at his desk as he usually does. His back starts to become very sore from sitting for so long. He decides that he will take two of the

> gollywags and superglue them together, making what he thinks to be a new object, what he calls the 'Gollywag-Supporter.'
>
> He places the Gollywag-Supporter on his desk chair and continues to work. At the end of the day, his back is not sore at all.
>
> The next day, Smith and Jones have a disagreement over whether the arrangement of gollywags composes a new object. Smith claims that the arrangement of gollywags does not compose a new object at all. He claims that simply super-gluing some stuff together with some other stuff is not sufficient to compose a new object.
>
> Jones, however, disagrees. He claims that simply supergluing some stuff together with some other stuff is sufficient to compose a new object. Concerning the specific case that they are considering, Jones concludes that the arrangement of gollywags does compose a new object, namely a Gollywag-Supporter.

We also had a *Gollywags with contact with function case*, just like *Gollywags with fusion and function* except that Jones, instead of supergluing the gollywags together, simply placed them into contact with one another.

We found, in line with our Handshake cases (Section 3.4.1) but not in line with any of the extant answers to the special composition question, that function continued to have a significant (in this case medium-sized) impact on judgments about composition. We also found no effect of contact versus fusion (thus undermining the idea that causal joining plays a significant role in our intuitions), and no interaction between the presence or absence of function and the presence of contact or fusion, as may be visualized in:[10]

[10] 121 participants were recruited from Amazon's Mechanical Turk, and randomly assigned to one of four conditions. After reading the case, participants were asked the extent to which they agreed with either Jones

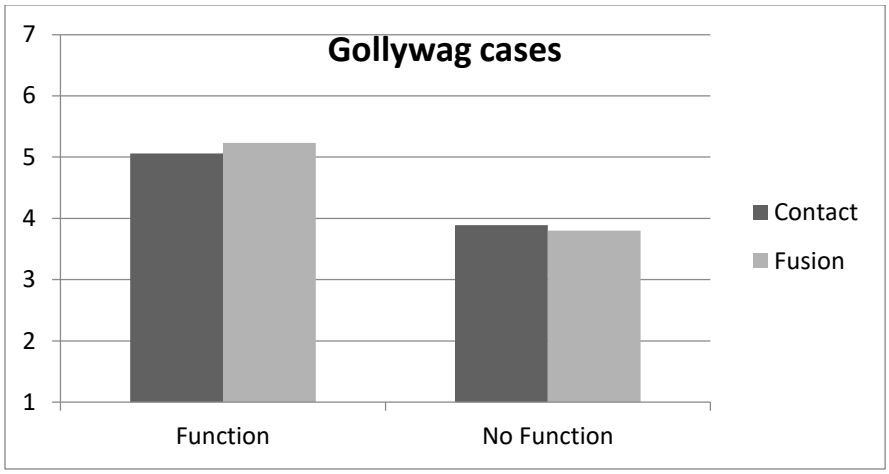

Figure 25: Gollywag Cases

With the Gollywag cases we thus see a continuation of the underlying pattern of teleologically influenced intuitions. Further, we see no evidence that other factors, in particular the kind of causal connectedness or joining that fastening adds to contact, play any role in folk intuitions about when composition occurs, either alone or in interaction with function.

---

or Smith (they were given a seven point scale with 1 marked 'Smith is right,' 4 marked 'Neither is right,' and 7 marked 'Jones is right'). Participants were then given, on a separate page, three comprehension questions:

1. Jones and Smith were talking about what is required to create a new thing (yes/no)
2. When Jones said that supergluing two things together [putting two things into contact] is sufficient for creating a new thing, what he was saying was that if you superglue any two things together [put any two things into contact with one another], then that is enough to create a new thing (yes/no)
3. Suppose someone said that supergluing two things together [putting two things into contact] is sufficient for creating a new thing. If that person were to take a wheel and a piece of metal and then superglue them together [put them into contact with one another], then that person would think that a new thing was created (yes/no)

Six people were excluded from the data analysis for failing at least one of the comprehension questions. Analyzing the remaining 115 responses using a two-way analysis of variance, we found that whether the gollywags had a function or not produced a statistically significant medium-sized effect on people's judgments $F(1, 115)=12.492$, $p=.001$, $\eta p^2=.101$, with people tending to disagree that the gollywags composed an object when the gollywags had no function ($M=3.85$, $SD=1.94$) and agreeing that the gollywags composed an object when they had a function ($M=5.15$, $SD=1.99$). There was no effect of the relation type (contact, fusion) on people's composition judgments $F(1, 115)=.016$, $p=.901$ and there was no interaction between function and relation type on people's judgments $F(1, 115)=.329$, $p=.567$. Since there was only an effect of function, we looked at whether responses in the No Function and Function cases differed from the midpoint. We found that responses in the Function cases were significantly different from the midpoint, $t(60)=4.48$, $p<.001$ but that responses in the No Function cases were not, $t(55)=-.557$, $p=.580$.

We also note that this is the one case of ours in which we did not see a flip in intuitions, insofar as we did not detect a statistically significant difference between responses in the No Function cases and the midpoint (marked as 'Neither is right'). The presence or absence of function is continuing to produce an effect on intuitions, as per our primary claim. Moreover, responses in the Function cases were significantly above midpoint (in the direction of 'Jones is right'). But our data in this one case is consistent with the thought that people tend to be neutral, undecided, or just confused about what to say in some cases when no function is specified.

3.4.3 Mouse cases

With our Gollywag cases we looked at composition judgments for artifacts. We wanted to see if the same pattern of intuitions extended for biological organisms, and we also wanted to explore potential effects of familiarity, and of having a name to label the larger thing. So we began with familiar organisms—mice—and looked at the effect of function between:

> *Mice with fusion*
>
> Acme Inc. is a large research company. Two Acme Inc. researchers, Jones and Smith, are experimenting with mice.
>
> One day, Jones takes two mice and superglues them together. He thinks that he has created a new object.
>
> Later that same day, Smith and Jones, have a disagreement over whether the arrangement of mice composes a new object. Smith claims that the arrangement of mice does not compose a new object at all. He claims that simply supergluing some together with some other stuff is not sufficient to compose a new object.

Jones, however, disagrees. He claims that simply supergluing some stuff together with some other stuff is sufficient to compose a new object. Concerning the specific case that they are considering, Jones concludes that the arrangement of mice does compose a new object.

And:

Mice with fusion and function

Acme Inc. is a large research company. Two Acme Inc. researchers, Jones and Smith, are experimenting with mice.

The FBI has commissioned Jones and Smith and given them the task of determining whether mice can be used to sniff out explosives. Typically, dogs are used to sniff out explosives but in some cases dogs are way too large to enter certain types of spaces. So, the FBI wants to see if a smaller creature can sniff out explosives with the same degree of accuracy as a dog.

For many months, Jones and Smith have been running the mice through various types of mazes, trying to determine how quickly and accurately the mice can find explosives. But they are having little luck: the mice are much slower and much less accurate than dogs in finding explosives.

One day, Jones takes two mice and superglues them together. He runs a wide range of experiments and finds that when the two mice are superglued together they are both very fast and highly accurate at detecting explosives. As a matter of fact, when two mice are superglued together, they are both faster and more accurate than dogs at detecting explosives. Jones thus thinks that he has created a new object, the ‘Mini-Bomb Detector.’

> The next day, Smith and Jones have a disagreement over whether the arrangement of mice composes a new object. Smith claims that the arrangement of mice does not compose a new object at all. He claims that simply super-gluing some stuff together with some other stuff is not sufficient to compose a new object.
>
> Jones, however, disagrees. He claims that simply supergluing some stuff together with some other stuff is sufficient to compose a new object. Concerning the specific case that they are considering, Jones concludes that the arrangement of mice does compose a new object, namely a Mini-Bomb Detector.

Alongside all of this, we also wanted to explore potential effects of familiarity, and of having a name to label the thing. So we looked at counterpart cases in which we replaced the mice with an unfamiliar sort of organism, for which we again used 'gollywag.' And we also looked at further counterpart cases in which we deleted the label 'mini-bomb detector,' to check if the presence of the label was influencing intuitions.[11]

Our results confirmed and extended the results of our Handshake and Gollywag (artifact) cases (Sections 3.4.1-3.4.2). We continued to find a significant medium-sized effect of function on judgments about composition, and we found no other effects (either alone or in interaction). Whether the organisms were familiar (mice) or unfamiliar (gollywags), and whether they were jointly labeled (a 'mini-bomb detector') or left unlabeled made no detectable difference to intuitions. All that seemed to drive intuitions was whether or not the candidate larger thing served a purpose. Since only function had any effect on

[11] We thus had a 2 (Function: Yes, No) x 2 (Label: Yes, No) x 2 (Familiarity: Gollywag, Mouse) design, resulting in a total of eight conditions.

composition, we've aggregated responses across all remaining conditions in the following visualization:[12]

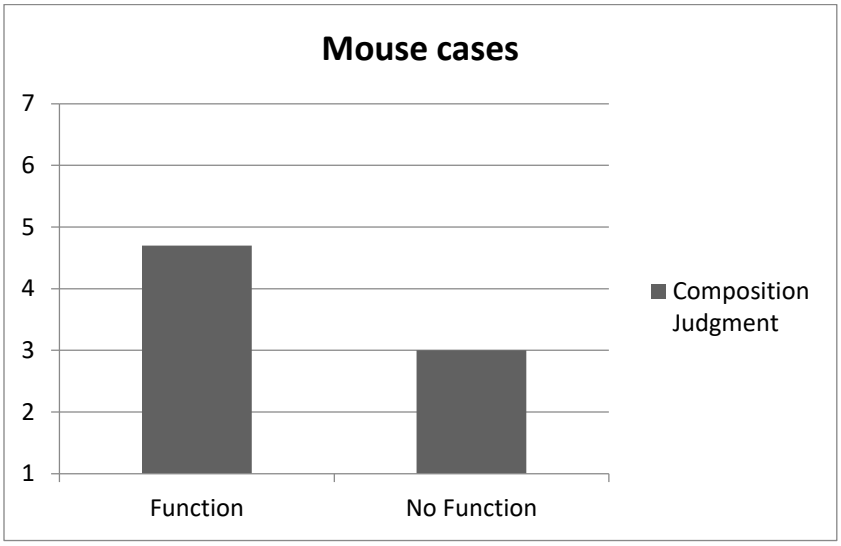

Figure 26: Mouse Cases

[12] 212 participants were recruited from Amazon's Mechanical Turk, and randomly assigned to one of eight conditions. After reading the case, participants were asked the extent to which they agreed with either Smith or Jones (they were given a seven point scale with 1 marked 'Smith is right,' 4 marked 'Neither is right,' and 7 marked 'Jones is right'). Participants were then given, on a separate page, five comprehension questions:

1. Jones said that supergluing the [gollywags/mice] together did create a new object (yes/no)
2. Smith said that supergluing the [gollywags/mice] together did not create a new object (yes/no)
3. Suppose someone said that supergluing two things together is sufficient for creating a new thing. If that person were to take a wheel and a piece of metal and then superglue them together, then that person would think that a new thing was created (yes/no)
4. When Jones said that supergluing two things together is sufficient for creating a new thing, what he was saying was that if you superglue any two things together, then that is enough to create a new thing (yes/no)
5. Jones and Smith were talking about what is required to create a new thing (yes/no)

Eight people were excluded from the data analysis for failing at least one of the comprehension questions. Analyzing the remaining 204 responses using a three-way analysis of variance, we found that having a function or not produced a statistically significant medium-sized effect on whether people judged that the arrangement of mice or gollywags composed a new object $F(1, 204)=30.115$, $p<.001$, $\eta p^2=.133$: when the mice or gollywags served a function people tended to judge that they composed a new object ($M=4.70$, $SD=1.15$) but when they did not serve a function people tended to deny this ($M=3.00$, $SD=1.73$). Neither Label $F(1, 204)=.033$, $p=.856$, Familiarity $F(1, 204)=.219$, $p=.640$, nor any two- or three-way interactions were found (all p's >.05). Since there was only an effect of function, we looked at whether responses in the No Function and Function cases differed from the midpoint and found that responses in both the No Function and Function cases were significantly different from the midpoint: No Function, $t(118)=-5.32$, $p<.001$; Function, $t(86)=2.82$, $p<.01$.

With the Mouse cases we thus see a continuation of the underlying pattern of teleologically influenced intuitions. We moreover see that the effect of function was to again flip intuitions (i.e., responses in both the Function and No Function cases were significantly different from the midpoint). Further, we see no evidence that other factors including familiarity and labeling play any role in folk intuitions about when composition occurs.

3.4.4 Avalanche Cases

Our cases so far have been limited in two main respects. First, we have mainly focused on artifacts and organisms, for which it is well established (even given just 'selective teleology': Section 3.3.1) that people tend to think teleologically. We have not yet looked at natural non-biological things such as rocks (where one needs 'promiscious telelology' [Section 3.3.2] to predict an effect). Secondly, we have only looked at larger things that are at least spatiotemporally connected, and sometimes even fastened together. We have not yet looked at scattered things, so left open whether there still might be some effect of scattering versus contact, or some interaction between scattering and function. So we turned to cases involving a scattered collection of rocks, to see if attributing a function (having the rocks designed to serve as a rock garden) would have an effect, and to continue checking for effects of labeling. We also wanted to start getting 'inside' the notion of a function by comparing the case in which a thing was designed from the start to serve a function, with the case in which a thing has already come to be (seemingly by accident) and is then accorded a function after the fact.

So in order to look at scattered collections of rocks, we began with:

Avalanche (base)

Jones lives on the side of a mountain. He has just been awoken by a series of loud crashes from a small avalanche on the mountain. Jones wakes up and looks outside, and sees a bunch of rocks strewn across his lawn from the avalanche.

We wanted to vary whether or not the rocks were given a label, whether or not they were accorded a function as a rock garden, and whether or not Jones then rearranged the rocks so that they were actually designed to serve that function. This resulted in six different extensions of *Avalanche (base*):

Avalanche

[*Avalanche (base)* plus] Though he is surprised, he just goes back to bed.

Avalanche with label

[*Avalanche (base)* plus] He thinks to himself, 'Looks like I have mountain man's rock garden!' He goes on thinking to himself, 'What a useless mess—looks like I'll have to clean all this up in the morning.'

Avalanche with accorded function

[*Avalanche (base)* plus] Even though avalanches are usually quite annoying for Jones, he decides at that moment that the rocks are actually strewn across his lawn in such a way that they will make his front lawn beautiful. He thus thinks that the arrangement of rocks from the avalanche compose a new object, namely an object that would make his front lawn beautiful.

Avalanche with accorded function and label

[*Avalanche (base)* plus] Even though avalanches are usually quite annoying for Jones, he decides at that moment that the rocks are actually strewn across his lawn in such a way that they will do perfectly for a beautiful rock garden. He thus

thinks that the arrangement of rocks from the avalanche compose a new object, namely a rock garden.

Avalanche with designed function

[*Avalanche (base)* plus] Even though avalanches are usually quite annoying for Jones, he decides that in the morning he'll use them to make his front lawn beautiful.

The next day, he arranges all of the rocks in such a way that he thinks he has created something that makes his front lawn beautiful. He thus thinks that the arrangement of rocks compose a new object, namely an object that makes his front lawn beautiful.

Avalanche with designed function and label

[*Avalanche (base)* plus] Even though avalanches are usually quite annoying for Jones, he decides that in the morning he'll use them to make a rock garden.

The next day, he arranges all of the rocks in such a way that he thinks he has created a rock garden. He thus thinks that the arrangement of rocks compose a new object, namely a rock garden.

Our results on these six cases confirmed and extended our previous results. We continued to find a significant (now large-sized) effect of function on judgments about composition, and we found no effect of naming and no interaction between naming and function. We also found a significant difference between accorded function and designed function, as may overall be visualized in:[13]

[13] 173 participants were recruited from Amazon's Mechanical Turk, and randomly assigned to one of six conditions. After reading the case participants were asked to indicated their agreement, on a 7-point scale (anchored with 1=completely disagree and 7=completely agree), with the following statement: 'Rather than being a bunch of scattered objects that do not in any way compose some one thing, the arrangement of

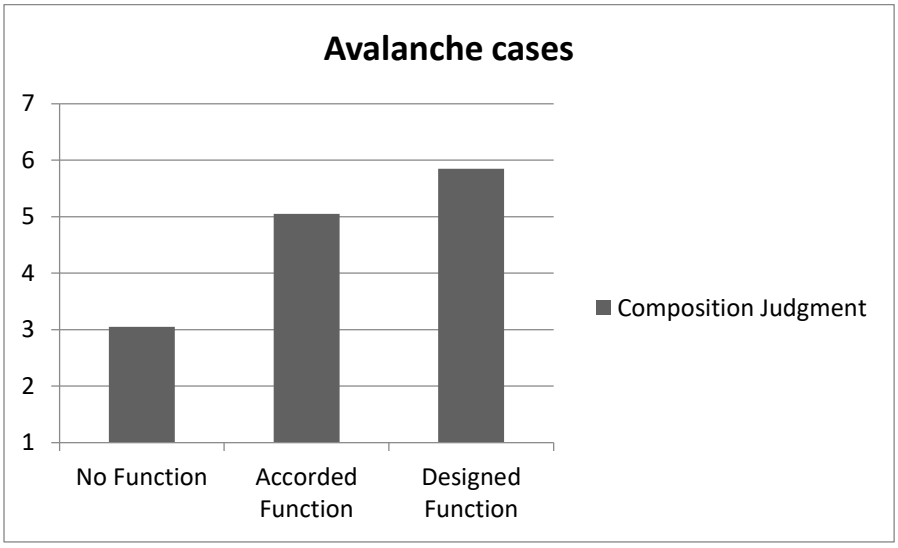

Figure 27: Avalanche Cases

Thus with the Avalanche cases we see a continuation of the underlying pattern of teleologically influenced intuitions, *extended even to cases with scattered rocks*. And we start to see inside this pattern, by seeing that merely according a function is already sufficient to influence intuitions, but that having the thing designed all along to serve the function produces an even stronger effect. Putting this all together, and considering the totality of our studies, it seems that the folk operate with a restricted and teleologically-

rocks actually compose something.' Participants were then given, on a separate page, two comprehension questions:

1. Jones lives on the side of a mountain (yes/no)
2. Because of the avalanche, rocks were strewn across Jones' lawn (yes/no)

Four people were excluded from the data analysis for failing at least one of the comprehension questions. Analyzing the remaining 169 responses using a two-way analysis of variance, we found a statistically significant large-sized effect of function on people's composition judgments $F(1, 169)=46.341$, $p=.000$, $\eta p^2=.362$. Bonferonni post-hoc tests showed that composition judgments across each level of function were significantly different from the others: No Function Conditions ($M=3.05$, $SD=1.34$) were significantly different from both the Accorded Function ($M=5.05$, $SD=1.77$, $p<.001$) and Designed Function Conditions ($M=5.84$, $SD=1.52$, $p<.001$), and the Accorded Function Conditions were significantly different from Designed Function Conditions ($p<.05$). Moreover, responses in each of the conditions were significantly different from the midpoint: No Function, $t(56)=-3.73$, $p<.001$; Accorded Function, $t(50)=4.06$, $p<.001$; Designed Function, $t(63)=14.34$, $p<.001$. There was no effect of Naming $F(1, 169)=2.185$, $p=.141$ nor was there an interaction between Naming and Function $F(2, 169)=.533$, $p=.588$.

laden view of when composition occurs. Through a range of very different cases we find a significant effect of function on judgments of when composition occurs, as per our main thesis. Indeed, in all cases save that of our Gollywag cases (Section 3.4.2) we saw an effect that we would describe as 'flipping intuitions,' from significantly below to significantly above the midpoint.

3.4.5 Restricted Composition or Restricted Domain?

The friend of *Universalism* who (like Lewis and Thomasson) thinks that her theory does not conflict with folk intuitions has one more trick up her sleeve. Instead of saying that the folk operate with a restricted and teleologically-laden view of when composition occurs, she might say that the folk operate with an unrestricted view of when composition occurs, coupled with a teleological restriction of the domain of quantification. That is, she can treat the folk as upholding *Universalism* alongside:

> *Teleologically Restricted Domain*: In normal contexts, quantifiers are restricted to things that have a purpose

There is something to be said for this move, insofar as it will be vague whether something has a purpose, and insofar as Teleologically Restricted Domain relocates this vagueness from the metaphysics to the semantics (where vagueness most plausibly belongs). If this move could be shown to work, we would regard this as a very interesting result in its own right!

With this trick in mind, we returned to our Handshake cases (Section 3.4.1) and had one of the characters (Andy) speak explicitly on behalf of *Universalism*, saying the sort of things one might reasonably say to ensure that the domain of quantification is sufficiently open. This yielded:

Handshake without restrictions

Sally and Tom are leaders of rival political factions, and have recently decided to lay aside their differences. They have worked out all the details, signed all the official papers, and will now seal their deal with a public and historic handshake. Later that day, Andy and Liz—who were both present for the historic handshake—have a disagreement over whether Sally and Tom created a new object when they shook hands.

Andy says that Sally and Tom did compose a single object. Indeed, he says that Sally and Tom always made a single, larger object, not just during the handshake, but before and after too. He says that in addition to the two people, Sally and Tom, there is, was, and will also be this larger object made of the two of them together. He says, 'This larger object, which I will call 'the Sally-Tom hybrid' has two parts: a Sally part, and a Tom part. It weighs just as much as Sally and Tom put together. Before the handshake, its two parts were scattered, during the handshake its two parts came into contact briefly, and after the handshake its two parts scattered. I know that we have no usual name for this larger object, and I know that it is not the sort of thing we usually chat about, but it is there all the same.'

Liz, however, disagrees. She says that, both before Sally and Tom shook hands, during their handshake, and after, there were just two people present. There was never any third, larger object made from Sally and Tom together, either before, during, or after the handshake. She clarifies: 'There is no such thing of any sort as

> your alleged ‘Sally-Tom hybrid.’ It's not just that we don't usually chat about this thing. There is no such thing.’

Despite Andy’s explicit description of the sort of fusion he has in mind, and despite his explicit moves to open up the quantifiers, our participants still tended to agree with Liz that ‘It's not just that we don't usually chat about this thing. There is no such thing.’[14]

Accordingly we see little prospect in sustaining Lewis’s idea that our hesitation to accept the existence of arbitrary sums is merely due to contextual domain restrictions, and little prospect in sustaining Thomasson’s (2007, p. 184) speculation that if ‘we explained to ‘normal’ people’ a term for some arbitrary sum, and then asked them if there is such a thing, they ‘would certainly accept that there is.’ Instead we find empirical substantiation for Korman’s (2008) skepticism that folk intuitions can be explained by *Universalism* plus domain restrictions. Of course it is possible that—despite Andy’s explicit description of the sort of sum he has in mind, and his explicit moves to open up the quantifiers—he still failed to open up the quantifiers sufficiently. But pending real evidence for this it must be considered doubtful. (*Compare*: In paradigm cases of quantifier domain restriction such as when one says ‘All the beer is in the fridge,’ it is very easy for an interlocutor to open up the quantifiers such as by saying ‘If we need more beer there is more at the grocery story on the corner.’) And given the background psychological

---

[14] A total of 33 participants were recruited through Amazon’s Mechanical Turk. After reading the case and indicating who they agreed with (Liz or Andy: again, using the same 7-point scale as in Handshake above), participants were then taken to a separate page where they answered three comprehension questions:

1. Sally and Tom are leaders of rival political factions (yes/no)
2. Andy thinks that the handshake did not create a larger object (yes/no)
3. Liz thinks that the handshake did create a larger object (yes/no)

Four participants were excluded for missing one or more comprehension questions, which left a total of 29 responses. Overall, participants showed strong agreement with Liz ($M=1.88$, $SD=1.48$), with responses being significantly below the midpoint $t(29)=-7.16$, $p<.001$.

evidence that we are natural teleologists (Section 3.3) this view must be considered more doubtful still. It's not just that we usually choose to focus on things with a purpose; it's rather that we naturally equate *being a thing* with *having a purpose.*

3.4.6 Teleologically Restricted Composition

Putting this all together, and considering the totality of our studies in light of the background psychological evidence about the role of teleology in folk metaphysics, we find empirical support for the idea that teleology plays a significant role in folk judgments about composition:

> *Teleologically Restricted Composition*: Composition occurs in restricted circumstances, in which the question of whether the plurality has a purpose plays a significant role

We would emphasize that we are offering Teleologically Restricted Composition as an account of the folk theory, and not as a metaphysical claim. We sometimes hear: 'Surely the folk accept that there are composites like rocks and clouds which serve no purpose.' But it should not be presumed that the folk think that rocks and clouds serve no purpose. Indeed, as we have already noted (Section 3.3.2), the folk tend to regard *everything* as having a purpose, including every rock and cloud ('clouds are for rain').

We would further emphasize that Teleologically Restricted Composition is not an exhaustive claim about what factors drive folk intuitions. We claim that teleology is *one factor* driving folk intuitions, not that it is *the only* such factor. That would require the strong claim:

> *Purely Teleologically Restricted Composition*: Composition occurs if and only if the plurality has a purpose

We leave open whether Purely Teleologically Restricted Composition holds, and will only operate with the more modest claim that Teleologically Restricted Composition holds.

To establish Purely Teleologically Restricted Composition would be to prove the negative existential that no other factors play a role in the folk theory, which obviously we cannot prove. Indeed there may be good empirical reason to doubt Purely Teleologically Restricted Composition, stemming from the role that *gestalt principles* play in visual cognition.[15] As Goldman (1993, p. 108) usefully summarizes:

> We do not readily consider something as a physical body if it lacks cohesion (a pile of leaves), lacks bounds (a drop of water in a pool), or lacks continuity (a row of flashing lights). These may be considered collections of objects or parts of objects, but they are not unitary and independent objects for us.

Indeed it may be the case that the visual system operates with an implicit 'theory' of composition that differs from that used by other cognitive systems. (Though we would expect any such influence to be operative only when the plurality is visually presented, and so probably not triggered in philosophical intuitions given the usual narrative presentation of cases.)

That said, our studies did not turn up any effect of gestalt factors like contact or fusion (Section 3.4.2), or of other factors like familiarity or labeling (Sections 3.4.3.-3.4.4) or quantifier restrictions (Section 3.4.5). While inferences from a null result on a particular study to a 'no effect' conclusion are notoriously fraught (though see Machery, 2012), these null results at least provide some support for the idea that contact, fusion, familiarity, labeling, and quantifier restrictions are not playing a significant role in folk

[15] We thank Jonathan Weinberg for this point.

intuitions about when composition occurs, at least in narrative presentations of cases. This is especially relevant since philosophers have speculated that contact and/or fusion may be the main drivers of folk intuition (Sections 3.2.2-3.2.3). Thus our studies do at least provide some support for the idea that the most plausible additional factors other than teleology are not in fact driving folk intuitions, at least for the usual narratively presented cases. (Perhaps contact, fusion, and other gestalt-related factors might have played a role in judgments about visual stimuli.)

In any case, our ultimate purpose is to *debunk* folk intuitions on grounds that they are laden with primitive teleological thinking (Section 3.6.2), and for that purpose *Teleologically Restricted Composition* is sufficient. So in that sense the empirical question of whether there are further factors driving folk mereology—interesting as it may be—does not make a difference to our ultimate conclusion. Accordingly we will leave the prospect of further factors aside in what remains, as an invitation for follow-up empirical work. (Though we briefly return to the issue in Section 3.6.2 when we consider the prospects for a 'cleaned up' folk theory.)

3.5. Inside Folk Teleology

So far we have documented a range of studies—our Handshake, Gollywag, Mouse, and Avalanche cases—which all converge on the conclusion that the folk judgments fit a teleologically-laden theory of composition, namely *Teleologically Restricted Composition*. We have also started to get a bit more precise about the folk notion of 'purpose,' by seeing (Section 3.4.4) that being accorded a function after the fact matters, but actually being designed from the start to serve a function matters even more. We want to get further inside folk teleology and explore a range of further distinctions, including

whether the thing can in fact successfully perform the function it was accorded or designed for (or whether it in fact fails to serve the function), and what happens in cases where a thing takes in multiple functions (perhaps it was designed to serve one function, but is then accorded another), as well as interactions between these factors.

3.5.1 Rope Cases

One aspect of the folk notion of purpose we wanted to explore further was the extent to which success matters. The existing psychological literature on object categorization focuses almost exclusively on successful cases, where the object in fact works as intended. We also wanted to consider cases in which something fails to serve the purpose it was designed for, as well as take the opportunity to vary our probe further. Does a thing which fails to serve the purpose it was designed for still count as having a purpose (in the sense relevant to folk mereology), or not? So we sought to compare a case with a failed design and a case with a successful design, via the following paired vignettes:

> *Ropes failure*
>
> Jones has an old, rusty water heater lying in his backyard. He thinks the water heater is ugly and so decides to move it out of his yard. So he ties a piece of rope around the water heater and pulls on it. But it does not move at all. Jones then thinks to himself, 'The rope must be too short. I need something longer to get more leverage.' He thus grabs another piece of rope and ties it in a knot around the other piece of rope that is already around the water heater. Jones then pulls on it but still the water heater does not move at all.
>
> *Ropes success*

> Jones has an old, rusty water heater lying in his backyard. He thinks the water heater is ugly and so decides to move it out of his yard. So he ties a piece of rope around the water heater and pulls on it. But it does not move at all. Jones then thinks to himself, 'The rope must be too short. I need something longer to get more leverage.' He thus grabs another piece of rope and ties it in a knot around the other piece of rope that is already around the water heater. Jones then pulls on it and is easily able to drag the water heater out of his yard.

We asked people whether they thought that Jones 'created a single, unified object' or whether 'When Jones tied the pieces of rope together, he did not create a single, unified object out of the rope. Rather, he simply had two, distinct pieces of rope.'

We found that success matters in a significant small-sized way to people's intuitions. In *Ropes failure* only 36% of our participants chose the option of 'a single, unified object,' while in *Ropes success* 63% of our participants chose the option of 'a single, unified object.' Graphically:[16]

[16] 75 participants were recruited from Amazon's Mechanical Turk, and randomly assigned to one of two conditions. After reading the case, participants were given two response options:

a. When Jones tied the pieces of rope together, he created a single, unified object out of the rope
b. When Jones tied the pieces of rope together, he did not create a single, unified object out of the rope. Rather, he simply had two, distinct pieces of rope

After answering, they were then asked, on a separate page, two comprehension questions:

1. Jones was trying to move a water heater (yes/no)
2. Jones was able to move the water heater (yes/no)

Seven people were excluded from the data analysis for failing at least one of the comprehension questions. Analyzing the remaining 68 responses, we found that whether or not the rope successfully fulfilled its function had a statistically significant small-sized effect on whether people said that the rope was a single, unified rope or two separate pieces of rope, $X^2(68)=4.769$, $p=.029$, Cramer's $V=.265$.

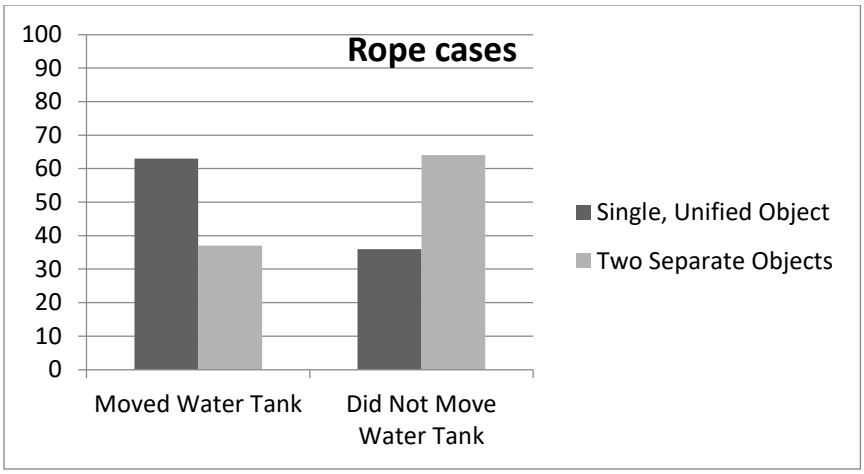

Figure 28: Rope Cases

We thus see further confirmation for the view that the folk operate with a restricted and teleologically-laden view of when composition occurs, alongside further and new clarification of the folk notion of function. Comparing these results with our earlier results on Avalanche cases (Section 3.4.4), we can contrast two different dimensions along which a thing might be said to have a purpose (in the sense relevant to the folk). First, the thing might be originally designed to serve a given function, or it might merely be accorded a function after the fact. Secondly, the thing might succeed in serving whatever function it was either designed for or accorded, or it might fail to serve this function. It seems (from the Avalanche cases) that accordance versus design is not crucial for serving a purpose (though design seems to bolster intuitions somewhat). And it seems (from the Rope cases) that success tends to be required.[17]

3.5.2 Rock garden cases

[17] When we say that success tends to be required, we mean to remain neutral on whether the thing in question might survive a later failure to serve its function. Contrast (i) a botched watch that never worked, with (ii) a broken watch that initially worked but later broke. We suspect that most people would deny that the botched watch is really a watch (it's just a collection of cogs and gears that doesn't really do anything), but may well uphold that the broken watch is still a watch if it can easily be fixed so as to resume its successful performance. See Rose (2015) for empirical evidence that folk judgments of persistence through time are tied to preservation of function.

We wanted to look further into the distinction between according something a function after the fact, and designing something with a function. In the Avalanche cases (Section 3.4.4) we saw that designed function played a stronger role than accorded function in influencing intuitions of composition. So we wondered whether intuitions triggered by accorded function were robust, and could survive explicit discussion of the idea that the thing was not in fact designed to serve that function. That is, we wondered whether designed function (or lack thereof) *trumped* accorded function when both were made explicit. We also wanted to continue looking at intuitions in the context of a scattered plurality.

So we returned to our Avalanche cases, again comparing cases in which the scattered plurality of rocks is just accorded the function of being a rock garden, with cases in which the scattered plurality of rocks is designed in advance to be a rock garden. But we also introduced a new character, Smith, whom we used to explicitly voice the lack of a designed function. We thus ran the following two cases:

> *Rock garden with accorded function*
>
> Jones was taking a walk in the forest one day. After walking for several hours, he came upon a field. In the field he found a bunch of rocks of all different shapes and sizes. Though nobody had ever encountered this before, Jones decided that the rocks were not merely a bunch of scattered objects but that they actually composed a new object: a rock garden.
>
> Later that day, Jones was discussing the rocks that he discovered with his neighbor Smith. Jones told Smith that even though nobody has ever encountered the rocks before, he thought that the rocks composed a new object, a rock garden.

Smith, however, disagreed. He claimed that since the rocks were not out there for the purposes of creating a new object—a rock garden—the arrangement of rocks did not make a new object. Thus Smith claimed that the arrangement of rocks did not compose a new object, namely a rock garden.

Rock garden with designed function

Jones was taking a walk in the forest one day. After walking for several hours, he came upon a field. In the field he found a bunch of rocks of all different shapes and sizes. Though nobody had ever encountered this before, Jones decided that he wanted to make a rock garden. So he rearranged the rocks in such a way that he thought the new arrangement of rocks composed a new object: a rock garden.

Later that day, Jones was discussing the rocks that he discovered with his neighbor Smith. Jones told Smith that he took the rocks, rearranged them, and thereby composed a new object, a rock garden.

Smith, however, disagreed. He claimed that simply rearranging a bunch of rocks with the purpose of creating a new object—a rock garden—was not enough to actually make a new object. Thus Smith claimed that the rearrangement of rocks did not compose a new object, namely a rock garden.

In *Rock garden with accorded function*, Jones accords the function of being a rock garden to a scattered plurality of rocks and Smith challenges Jones, claiming that the arrangement of rocks does not compose a rock garden since they were *not designed* with the purpose of being a rock garden. If designed function takes priority over accorded function, then we should expect participants in this case to largely deny that the arrangement of rocks composes a rock garden. In contrast, in *Rock garden with designed*

function, Jones arranges the rocks with the purpose of designing a rock garden, but Smith claims that arranging the rocks with the purpose of creating a rock garden was not enough to actually make a rock garden. If designed function plays a crucial role in composition intuitions, then we should expect that even in the face of Smith's downplaying the importance of designed function, designed function will still continue to guide composition judgments.

We found that in *Rock garden with accorded function* our participants tended to agree with Smith's claim that the rocks did not compose a new object, while in *Rock garden with designed function* our participants tended to agree with Jones's claim that the rocks did compose a new object. This can be seen via:[18]

[18] 68 participants were recruited through Amazon's Mechanical Turk, and randomly assigned to one of two conditions. After reading the case and indicating who they agreed with, participants were then given, on a separate page, two comprehension questions:

1. Jones and Smith were talking about what is required to create a new thing (yes/no)
2. Smith said that the rocks did not compose a rock garden (yes/no)

Seven people missed one or more comprehension questions. Analyzing the remaining 61 responses using a one-way analysis of variance, we found a statistically significant large-sized effect of accorded (M=3.33, SD=2.08) versus designed (M=5.22, SD=1.78) function on people's composition judgments $F(1, 61)=14.506$, $p<.001$, $\eta p^2=.197$, with people being more willing to say that the arrangement of rocks composed a rock garden in *Rock garden with designed function*. Moreover, we found that responses were marginally different from the midpoint when the function was accorded, $t(30)=-1.75$, $p=.09$, and significantly different from the midpoint when the function was designed, $t(31)=3.83$, $p<.01$.

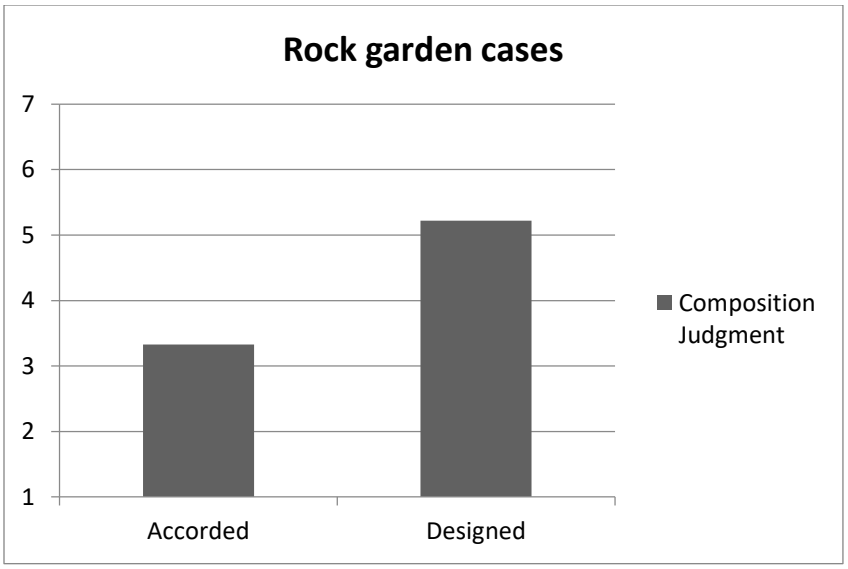

Figure 29: Rock Garden Cases

We saw in *Avalanche with accorded function* and *Avalanche with designed function* (Section 3.4.4) that both accorded and designed function played an important role in people's composition judgments, though designed function appeared to play a more significant role. In *Rock garden with accorded function* we saw that when designed function is emphasized, accorded function ceases to guide composition judgments. But even when designed function is downplayed in *Rock garden with designed function*, designed function still continues to be pivotal in people's judgments. Thus it seems that intuitions triggered by accorded function are not robust. Designed function (or lack thereof) trumps accorded function when both are made explicit. Putting our results in Rope cases and Rock cases together, we see that designed function tends to trump accorded function, and that success tends to be required.

3.5.3 Umbrella Cases

In a final series of studies, we wanted to explore the interaction between accorded versus designed function, and success versus failure. We also wanted to elicit judgments of

function from our participants so that we could have some direct measures of function, and also see how judgments of function (what something is for) would align with categorization judgments (what something is).

So we used stories about things that were originally designed to serve one function but were then later accorded a distinct function, varying whether or not the thing succeeded in serving the purpose it was originally designed for. Our first pair of cases—inspired by Rips, 1989—involved a thing that was originally designed to be an umbrella, but later re-purposed as a lampshade:

> *Bad umbrella good lampshade*
>
> In the early 1400s people typically stayed indoors when it rained since they had nothing that would protect them from the rain. One day Jones decided that he would make something that could protect people from the rain. So, he took a piece of material, some screws and wood, and put them all together. He called it an 'umbrella.' To Jones' surprise, it did a terrible job at protecting people from rain: they still got just as wet as they would have if they did not use anything at all.
>
> Four hundred years later, electricity was invented. Some people had lights in their house. But they were very bright since there was nothing to dim them.
>
> Smith lived in the same house that Jones did over four hundred years ago. One day Smith was rummaging around in his attic and stumbled across one of Jones' old boxes. On the box was the label 'umbrellas.' Smith opened up the box and took out one of the objects. He then took it and placed it over his lamp. He called

it a ‘lampshade.’ To Smith’s surprise, it worked perfectly for dimming the lights: when placed over a light, it was dimmed to just the right level of luminance.

One afternoon, Smith and Frank were having a disagreement. Frank claimed that the object was not a lampshade. Rather, he claimed that it was an umbrella. Smith, however, disagreed. He claimed that the object was a lampshade and not an umbrella.

In *Bad umbrella good lampshade* we find an object designed to be an umbrella but bad at it, and yet very good at serving the unexpected purpose of being a lampshade. We wanted to contrast this with a case in which the object did succeed as an umbrella:

Good umbrella good lampshade

In the early 1400s people typically stayed indoors when it rained since they had nothing that would protect them from the rain. One day Jones decided that he would make something that could protect people from the rain. So, he took a piece of material, some screws and wood, and put them all together. He called it an ‘umbrella.’ To Jones’ surprise, it worked perfectly at protecting people from rain: when outside in the rain, they stayed completely dry.

Four hundred years later, electricity was invented. Some people had lights in their house. But they were very bright since there was nothing to dim them.

Smith lived in the same house that Jones did over four hundred years ago. One day Smith was rummaging around in his attic and stumbled across one of Jones’ old boxes. On the box was the label ‘umbrellas.’ Smith opened up the box and took out one of the objects. He then took it and placed it over his lamp. He called

> it a 'lampshade.' To Smith's surprise, it worked perfectly for dimming the lights: when placed over a light, it was dimmed to just the right level of luminance.
>
> One afternoon, Smith and Frank were having a disagreement. Frank claimed that the object was not a lampshade. Rather, he claimed that it was an umbrella. Smith, however, disagreed. He claimed that the object was a lampshade and not an umbrella.

We found that intuitions flipped—in a connected way—both on the question of what the arrangement of material, wood, and screws composes (umbrella or lampshade), and on the question of what the object is for (protecting people from rain or dimming light). In *Bad umbrella good lampshade,* 68% of our participants said both that the object was a lampshade and that its function was to dim light, while in *Good umbrella good lampshade,* 73% said both that the object was an umbrella and that its function was to protect people from rain. Visually:[19]

[19] 57 participants were recruited through Amazon's Mechanical Turk, and randomly assigned to one of two conditions. After reading the case, participants were then presented with two questions (which were randomized):

1. The arrangement of material, screws and wood composes a: (lampshade/umbrella)
2. What best describes the function of the object made of material, wood and screws? (dimming light/protecting people from rain)

After answering these questions, participants were then taken to a separate page and presented with two comprehension questions:

1. Smith thought the object was a lampshade (yes/no)
2. Frank thought the object was an umbrella (yes/no)

Six people missed one or more comprehension questions. We analyzed the remaining 51 responses and found a large effect of condition on composition judgments, $X^2(51)=16.525$, $p<.001$, Cramer's $V=.569$, and a medium-sized effect of condition on function judgments $X^2(51)=12.244$, $p=.001$, Cramer's $V=.490$. Within *Bad umbrella good lampshade*, there was a strong relationship between composition and function judgments $X^2=8.383$, $p=.004$, Cramer's $V=.579$, while within *Good umbrella good lampshade*, there was a very strong relationship between composition and function judgments $X^2=15.954$, $p<.001$, Cramer's $V=.783$.

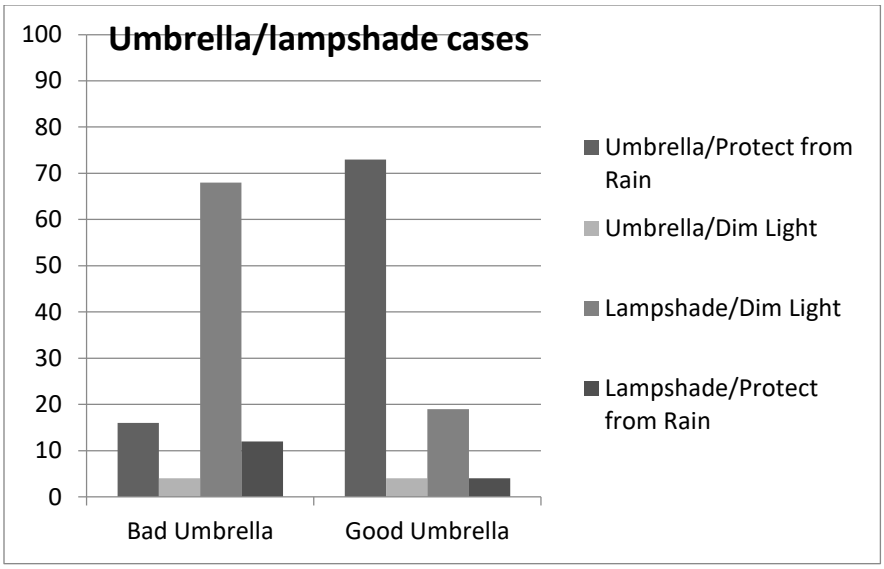

Figure 30: Umbrella Cases

It seems that both designed function and success are playing a role in our connected intuitions about what a thing is and what a thing is for. Accorded function plus success seems to override designed function when they pull apart in *Bad umbrella good lampshade*. Yet designed function seems to be playing at least a tie-breaker role in *Good umbrella good lampshade*: when the thing has multiple successful functions, people attribute to it the function it was designed to produce over the function it was later accorded.

3.5.4. Tog Cases

To extend this basic finding we considered two final variant cases in which the unintended function was not even for human beings but for an unfamiliar sort of organism ('togs'). These cases were:

Bad umbrella good tog-helper

In the early 1400s people typically stayed indoors when it rained since they had nothing that would protect them from the rain. One day Jones decided that he

would make something that could protect people from the rain. So, he took a piece of material, some screws and wood, and put them all together. He called it an ‘umbrella.’ To Jones’ surprise, it did a terrible job at protecting people from rain: they still got just as wet as they would have if they did not use anything at all. Jones was so frustrated that he just threw the object in his backyard.

In those days, a rare and dying species of animal—togs—had incredible difficulty catching prey. Togs were very small creatures that were almost completely blind. To hunt, togs had to completely rely on their ability to track sounds. But, unlike bats, togs were terrible at hearing sounds as they echoed off of objects. The object that Jones threw in his backyard, however, was perfect for echoing sounds at a frequency that only togs could hear. Since the object relayed the perfect frequency of sound for togs and no other creature could register the frequency of sound, togs were able to both hunt down prey and avoid danger with incredible efficiency.

One afternoon, Smith and Frank were having a disagreement. Frank claimed that the object was not an umbrella at all. Rather he claimed that it was a ‘tog echolocation device.’ Smith, however, disagreed. He claimed that the object was an umbrella and not a tog echolocation device at all.

Good umbrella good tog-helper

In the early 1400s people typically stayed indoors when it rained since they had nothing that would protect them from the rain. One day Jones decided that he would make something that could protect people from the rain. So, he took a piece of material, some screws and wood, and put them all together. He called it an ‘umbrella.’ To Jones’ surprise, it worked perfectly at protecting people from

rain: when outside in the rain, they stayed completely dry. Jones was so excited that he accidentally left the object in his backyard.

In those days, a rare and dying species of animal—togs—had incredible difficulty catching prey. Togs were very small creatures that were almost completely blind. To hunt, togs had to completely rely on their ability to track sounds. But, unlike bats, togs were terrible at hearing sounds as they echoed off of objects. The object that Jones threw in his backyard, however, was perfect for echoing sounds at a frequency that only togs could hear. Since the object relayed the perfect frequency of sound for togs and no other creature could register the frequency of sound, togs were able to both hunt down prey and avoid danger with incredible efficiency.

One afternoon, Smith and Frank were having a disagreement. Frank claimed that the object was not an umbrella at all. Rather he claimed that it was a 'tog echolocation device.' Smith, however, disagreed. He claimed that the object was an umbrella and not a tog echolocation device at all.

We found the same connection between intuitions of what a thing is, and intuitions of what it is for. And we found the same pattern of interaction between designed function and success. That is, in *Bad umbrella good tog helper* our participants (70%) tended to say both that the object was a tog echolocation device, and that it was for echolocation. While in *Good umbrella good tog helper* our participants (80%) tended to say both that the object was an umbrella, and that it was for protecting people from the rain. In an image:[20]

[20] 59 participants were recruited through Amazon's Mechanical Turk, and randomly assigned to one of two conditions. After reading the case, participants were then presented with two questions (which were randomized):

1. The arrangement of material, screws and wood composes a: (umbrella/tog echolocation device)

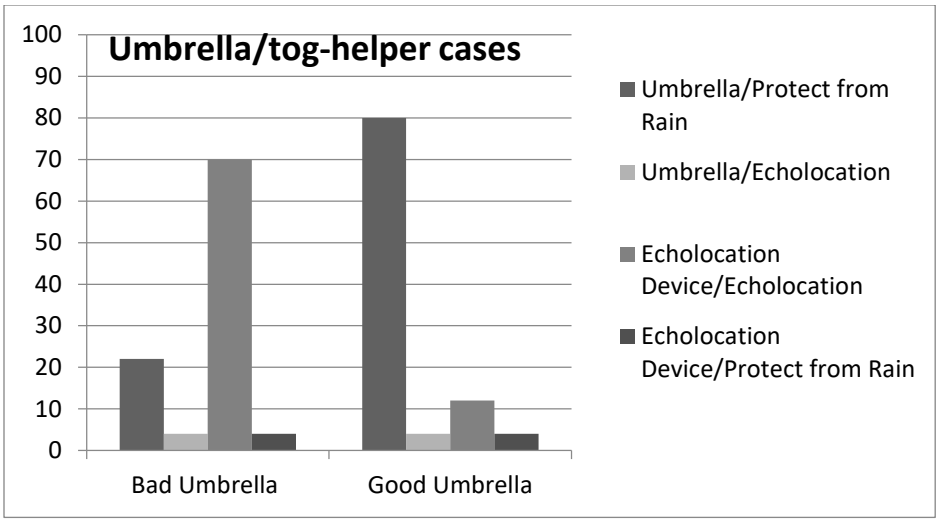

Figure 31: Tog Cases

Together our Rope, Rock, and Umbrella cases suggest a fairly complicated interaction between accorded versus designed function, and success versus failure. When something is both designed to serve a certain purpose and succeeds at that purpose, then it seems that people will tend to lock onto that purpose in saying what a thing is as well as what it is for. But when the thing is accorded a different purpose after the fact, and only succeeds at that accorded purpose, then it seems that people may lock onto that accorded purpose in saying what a thing is as well as what it is for, at least so long as the lack of a (successful) designed function is not sufficiently emphasized.

2. What best describes the function of the object made of material, wood and screws? (protecting people from rain/echolocation)

After answering these questions, participants were then taken to a separate page and presented with two comprehension questions:

1. Smith thought the object was an umbrella (yes/no)
2. Frank thought the object was a tog echolocation device (yes/no)

Seven people missed one or more comprehension questions. We analyzed the remaining 52 responses and found a large effect of condition on composition judgments, $X^2(52)=20.77$, $p<.001$, Cramer's $V=.624$, and a large effect of condition on function judgments, $X^2(52)=15.205$, $p<.001$, Cramer's $V=.541$. Within *Bad umbrella good tog-helper*, there was a very strong relationship between composition and function judgments, $X^2=16.636$, $p<.001$, Cramer's $V=.807$, and within *Good umbrella good tog-helper* there was a very strong relationship between composition and function judgments, $X^2=17.590$, $p<.001$, Cramer's $V=.739$.

But for the purposes of continuing to confirm that folk mereology is teleologically restricted, we would just emphasize how closely object categorization judgments (what the thing is) and function judgments (what the thing is for) line up. It is natural to think of teleologically driven folk mereology as something of the limit case of our conceptual equation of what a thing is and what it is for. When the plurality is *for something* then it *is something*. But when the plurality lacks a purpose—when it is *for nothing*—then it *is nothing*.

3.6. Implications

We began by documenting various answers to the special composition question and the charges of 'mereological madness' that have played a major role in the metaphysical debate (Section 3.2). We then offered an empirically driven case for regarding folk mereology as teleological (Sections 3.3-3.5). We conclude by discussing some of the metaphysical and methodological implications of our research.

3.6.1 Psychological implications

Our results are broadly in line with the perspective that sees folk teleology as not merely selectively applied to the domains of artifacts and organisms, but promiscuously applied to all of nature (Section 3.3.2). Indeed we take our results—and in particular our continued finding of a teleological effect on composition judgments even for our avalanche cases (Section 3.4.4)—to broadly confirm the emerging 'promiscuous teleology' perspective. Our results suggest that people are intuiting teleologically, even in cases involving scattered bunches of rocks.

Our results moreover advance the discussion of the folk notion of function. There is wide agreement among cognitive scientists that object function is primarily determined by

intended (/designed) function and that intended function largely guides object categorization judgments. For example, Rips (1989) found that adults judge that an object that looks like a lampshade is actually an umbrella when they are told that it was originally designed to protect people from rain. Matan & Carey (2001) and Hall (1995) found that adults judge category membership on the basis of original design despite the objects' being successfully used for some other purpose. And Kelemen (1999a) gave children and adults cases where an object, either intentionally or accidentally, succeeds in fulfilling some other purpose and found that both children and adults tend to ignore successful use, favoring categorization judgments in line with original intended design. These findings cohere well with our findings targeting composition judgments (Section 3.5), on which function is primarily determined by design plus success.

But these studies have almost exclusively focused on cases in which the designed function was also a success (the object worked as intended). Cases in which the object fails to work as intended have not really been explored. Our Umbrella cases (Section 3.5.3) begin the exploration. Our results confirm the 'standard view' that designed function—when successful—runs the show. But our results also suggest the somewhat surprising coda that when the designed function fails there is then room for successful accorded function to enter the picture. This is not predicted on any existing story. Folk teleology may turn out to be more complicated than has hitherto been supposed.

3.6.2 Metaphysical implications

If we are right that folk mereology is teleological, what follows for the special composition question? That is, what if any morals should the metaphysician draw from our discussion? First and foremost, we hope to have made an empirically substantiated

case that answers to the special composition question should not be beholden to folk intuitions. Accordingly we claim to have *liberated* the prescriptive metaphysician from the various charges of 'mereological madness' in the literature. Whatever problems Universalism or Nihilism or any of the other extant views about composition might have, failure to fit folk intuitions is not among them. (Not because these views fit folk intuitions, but because folk intuitions in this domain are tied into a benighted teleological view of nature.)

Folk mereology looks to us in many ways akin to folk physics. Both involve fascinating conceptual machinery worthy of extended psychological study, but neither can make any serious claim to fitting reality. Objecting to a philosophical theory of mereology on grounds that it violates folk intuitions about composition looks to us roughly on par with objecting to Newtonian mechanics on grounds that it violates folk intuitions about persistent curvilinear motion (Kaiser, McCloskey, and Proffitt, 1986). In both cases the theories do conflict with folk intuitions, but the folk intuitions deserve dismissal.[21]

We imagine three likely follow-up reactions to our claim that folk mereology is teleological.[22] One sort of follow-up reaction—which certain sorts of opponents of experimental philosophy might be drawn to—would be to say 'let us then ignore the ignorant folk and turn to the real experts, namely the metaphysicians.' We do not think it

[21] At this point we join Goldman (2007), Paul (2010), and Schaffer (*forthcoming*) in thinking that cognitive science can prove relevant to real metaphysics, if only to teach us which intuitions should be ignored. As Goldman (2007, p. 2) explains: 'Cognitive organs or mechanisms play a critical role in the causal production of appearances, including metaphysical appearances (whatever exactly we take that to connote). In considering whether such metaphysical appearances should be accepted at face value or, alternatively, should be superseded through some sort of metaphysical reflection, it obviously makes sense to be as informed as possible about how these mechanisms of cognition work. That is why cognitive science is relevant.'

[22] Our discussion of reactions is not intended to be exhaustive. For instance, we will not consider the reaction of attempting to revive teleological thinking. See Hawthorne & Nolan (2006) for further discussion, as well as Bowers (*manuscript*) for a more optimistic attempt at revival.

is obvious that the metaphysicians have earned the mantle of ‘experts’ on the matter. But setting that aside, while we do not have empirical evidence as to the distribution of views among metaphysicians, it certainly seems from the current literature as if there is little agreement among the metaphysicians, not just with respect to theory but with respect to intuitions reported. We see intuitions reported swinging all the way from Universalism to Nihilism and stopping at many places between (Section 3.2.2). The folk may be benighted, but the ‘experts’ seem divided.

A second sort of follow-up reaction—which some partisans in the metaphysical dispute might be drawn to—would be to appeal to ‘corrected’ intuitions in a revised argument for mereological sanity. Perhaps if we could just wipe the teleological muck off from the folk theory we would find clean intuitions by which to judge theories. We are skeptical, in part because we can imagine different partisans in the metaphysical dispute spinning this move in different ways. We can imagine the friend of *Universalism* saying: ‘The folk only ever reject composition when they think the result has no purpose; wipe off the teleological muck and the folk would no longer ever reject composition.’ But we can equally imagine the friend of *Nihilism* saying: ‘The folk only ever accept composition when they think the result has a purpose; wipe off the teleological muck and the folk would no longer ever accept composition.’ And we can just as equally imagine partisans of various restrictive views saying: ‘The folk clearly hold a restricted theory and we should try to honor that commitment; wipe off the teleological muck and the folk would only need a more respectable restriction.’ The folk concept can be cleaned up in many ways. We doubt there is a determinate fact as to ‘the’ corrected folk theory.

(It may be that further empirical work will reveal that the folk operate with a plurality of factors, or perhaps even recruit different factors for different tasks). It may then be possible to consider a more determinate cleaned-up folk theory operating with just those other factors. But for present purposes we can only say that we have not uncovered any positive evidence for other factors (including factors such as contact and fusion which had been thought operative), and we cannot guess what further factors could potentially be involved. Overall we suggest that those attracted to correcting the folk theory do the empirical work to determine if there are any residual and respectable factors remaining in the folk theory, which might still be retained. We would welcome such follow-up empirical work.)

A third sort of follow-up reaction—which some skeptics about prescriptive metaphysics might be drawn to—would be to say that the special composition question has been revealed to be hopeless, in that we lack the epistemic wherewithal to select the right answer.[23] As Rosen & Dorr (2002, pgs. 154-156) clarify, each proposal is analytically consistent, and all the proposals seem to be empirically equivalent. If intuitions cannot help either, what remains? (We expect that this is why intuitions loom so large in the current debate: Section 3.2.1.)

Perhaps the skeptic is right. Though we think that there are still virtues of simplicity, elegance, and coherence with wider theory in play that may still help favor certain approaches. Overall we suspect that liberating the special composition question from folk

[23] See Bennett (2009) and Kriegel (2013) for defenses of skeptical views on the composition debate. Both Bennett and Kriegel accept that there is a meaningful special composition question, and merely recommend withholding belief as to which is the right answer. Though of course there is also room for a skeptical view that denies that the special composition question is even meaningful.

intuitions helps tilt the overall balance somewhat in favor of the elegant but 'radical' extreme views of *Universalism* and *Nihilism,* over the more complex but 'conservative' intermediate positions. But this is a matter that obviously falls beyond the scope of the current discussion.

3.6.3 Methodological implications

We take our results to support a *targeted debunking* of folk intuitions about when composition occurs. Epistemically speaking, we hold that any line of argument that is premised on folk intuitions about composition suffers from an *undermining defeater* (cf. Pollock, 1987). Just as the force of testimony is undermined if it is discovered that the testimony is based on a lunatic view on the topic, so the force of intuition is undermined it is discovered that the intuitions are based on a hopeless theory of the topic. (As is standard with defeaters, there can also be *reinstaters*. For instance the lunatic views of the testifier may happen by luck to agree with sane views in the case at hand. But any such reinstatement must be shown.)

In other words, we want to say that there is a *debunking explanation* for folk intuitions about when composition occurs.[24] Imagine that it is discovered that we intuit that something is alive when we think the thing has a spirit. Our intuitions about when something is alive would thereby be debunked, and should be accorded no weight in disputes about biological classification. Likewise if it is discovered that we intuit that some things fuse when we think they have a collective purpose, our intuitions would thereby be equally debunked, and would thereby equally deserve no weight in disputes

[24] We follow Kahane (2011, p. 106) in associating undermining defeaters and debunking explanations: 'Debunking arguments are arguments that show the causal origins of a belief to be an undermining defeater.' See Mason (2010) for a useful overview of the structure of debunking arguments.

about mereological composition. (As is standard with debunking explanation, there can be *restorers*. For instance it might be found that the 'spirited' view of life is connected with a tendency to attribute being spirited on the basis of reliable indicators of life, such as self-movement. But any such restoration must be shown.)

There are of course difficult background issues about where to draw the line with debunking explanations. On the one hand, if we demand that our folk theories be *perfect* to be trusted, then we will be left with a blanket skepticism about intuitions. Perhaps such a skepticism is defensible; we only do not wish to premise our argument on such a radical view. On the other hand, if we allow our folk theories to be *terrible* and still trust them, then we are left with a naive credulity about intuitions. Theorizing about this issue is obviously beyond the scope of the current discussion. But for present purposes our claim is that there are clear cases of undermining defeaters and debunking explanations (such as our hypothetical spiritual folk biology), and that teleological folk mereology is one such clear case.

We note that our attempt at debunking is *targeted*, insofar as we have tried to identify a very specific flaw in our folk worldview. Korman (2009, p. 244)) raises the general worry that attempts at debunking folk intuitions may either be *globally self-defeating*, undermining all reasoning whatsoever, or *locally self-defeating*, undermining specific premises in the very argument used for debunking. Insofar as our reasoning targets specific features of our naïve view of the natural world, it does not undermine reasoning generally, nor do we ever appeal to naive teleological reasoning ourselves. In this way we hope to have illustrated a stable and targeted strategy for debunking folk intuitions.

3.7. Conclusion

To summarize: Many metaphysicians have wanted a view of mereological composition that fits with folk intuitions, and have charged leading views with failing to do so, while failing themselves to agree as to what the folk intuit or why. So we have tried to put the tools of experimental philosophy to constructive use to break this impasse. We have found something that, though unconsidered by any of the metaphysicians, coheres very well with recent psychological work: folk intuitions are based on a crude teleologically-laden conception of when composition occurs. The folk tend to connect composition to purpose. And we have suggested, in conclusion, that this finding should lead us to liberate the metaphysics of composition from any demand of fitting with folk intuitions. Folk mereology is teleological, and hence unenlightened.[25]

[25] Thanks to Jason Bowers, Andrew Higgins, Joshua Knobe, Dan Korman, Uriah Kriegel, Jonathan Livengood, Ned Markosian, David McElhoes, Shaun Nichols, Laurie Paul, Angel Pinillos, John Turri, Jonathan Weinberg, audiences at Buffalo and Arizona State, and an anonymous *Noûs* referee.

Chapter 4

Cognitive Science for the Revisionary Metaphysician

4.1 Introduction

Many philosophers insist that the revisionary metaphysician—i.e., the metaphysician who offers a metaphysical theory which conflicts with folk intuitions—bears a special burden to explain why certain folk intuitions are mistaken. But though it is widely agreed that the revisionary metaphysician incurs such an explanatory burden, many philosophers think that the revisionist is unlikely to be able to successfully discharge this explanatory burden (e.g., Paul, 2012, p. 22; Hirsch, 2002, p. 117; and Korman, 2009). My plan is to offer some resources to the revisionary metaphysician to meet this explanatory burden. Specifically, my proposal is that evidence from cognitive science can help the revisionary metaphysician discharge the explanatory burden of providing a plausible account of how the folk have gone wrong. In taking this up, I'm not going to engage in a discussion of the abstract prospects for how cognitive science *might* end up helping discharge this explanatory burden. Instead, my strategy will be to show how cognitive science can *actually* help the revisionary metaphysician discharge this explanatory burden by taking up a concrete illustration. In particular, I will take up metaphysical disputes over composition and persistence and various charges of violating folk intuitions. I'll then discuss a range of empirical evidence which suggests that the folk operate with a benighted teleological view of composition and persistence and go on to argue that there is a debunking explanation for folk intuitions of composition and persistence. Given an

empirically informed understanding of the folk view of composition and persistence, the revisionary metaphysician has the resources to offer a plausible debunking explanation of folk intuitions. In this way, I'll take myself to have illustrated one key way in which evidence from cognitive science can help the revisionary metaphysician discharge the explanatory burden of providing a plausible account of how the folk have gone wrong.

Though my main goal is to illustrate how cognitive science can contribute to metaphysics by helping discharge the explanatory burden put to the revisionary metaphysician, I have another goal. Recently, Dan Korman and Chad Carmichael (forthcoming) have provided a range of objections against the particular debunking explanation I'll be focusing on. Two objections are particularly troubling. The first is that the teleological view of composition that the folk operate with may well be legitimate. If this is right, then it will undercut the attempt to debunk folk intuitions and leave the explanatory burden faced by the revisionary metaphysician untouched. The second is that, in the studies I'll be discussing, the folk don't report their intuitions about the relevant cases. Instead, they merely provide answers. If this is right, then the relevant folk intuitions don't deserve to be debunked. The studies tell us nothing about what the folk intuit. Since the debunking argument has as its focus folk intuitions and since the studies allegedly don't tell us anything about what the folk actually intuit, the debunking argument fails. If either of these two objections is right, the illustration of how cognitive science can help the revisionary metaphysician discharge the explanatory burden will be an utter failure. So my other goal will be to address these two objections.

The Plan: I will begin, in Section 4.2, with some stage setting, discussing how I'll be understanding intuitions and their targets. I'll also set out one reason why conflict with

folk intuitions gives rise to an explanatory burden and briefly set out a discussion of metaphysical disputes over composition and persistence and the explanatory burden that arises for the revisionary metaphysician in connection with these disputes. In Section 4.3, I'll illustrate how cognitive science can help the revisionary metaphysician discharge the explanatory burden in the case of composition and persistence. I'll then respond to two main objections in Section 4.4.

4.2. Stage Setting

I want to begin with a bit of stage setting before getting on to the illustration of how cognitive science can help the revisionary metaphysician discharge the explanatory burden (Section 4.3). To this end, I want to briefly set out:

- How I'll be understanding intuitions
- How I'll be understanding the target of intuitions
- A rationale for why folk intuitions might be taken seriously and the explanatory burden that arises when a theory conflicts with folk intuitions
- The main metaphysical disputes I will focus on (i.e., composition and persistence)

Having accomplished this, Section II will set out the kind of teleological view I take the folk to operate with (i.e., promiscuous teleomentalism), some of the evidence suggesting that teleological considerations play a key role in folk intuition of composition and persistence, and the debunking argument for folk intuitions of composition and persistence.

4.2.1. Intuitions

What are intuitions? There's an important debate over how intuitions are best characterized with a wide range of proposals on offer (e.g., Bealer 1998; Cappelen 2012; Deutsch 2015; Devitt 2015; Pust 2016; Sosa 1998; Weinberg 2014; Williamson 2007). That debate is a swamp, which I won't wade into here. So, I'm simply going to set out the way I'll be understanding intuitions.

I'll be adopting an inclusive account of intuitions (Williamson, 2007; Stich, forthcoming) and will take intuitions to be spontaneous judgments—which immediately arise without any awareness of their origin and without having gone through any conscious process of reflection or reasoning—that some person, object, event, etc. described in a scenario has (or lacks) some interesting or important property or relation (Stich, forthcoming). Understanding intuitions in this way reflects how they are understood within dual-processing models of cognition (e.g., Epstein, 1994; Sloman, 1996; Chaiken and Trope, 1999; Kahneman and Frederick, 2002, Stanovich and West, 2000) and is analogous to how they are understood in linguistics (Stich, forthcoming).

Linguists typically gather data by presenting native speakers with sentences and asking participants, for instance, whether the target sentence is grammatical. Often the target sentence will elicit an intuitive judgment about whether it is grammatical, that is an intuition about grammaticality that arises spontaneously without the subject being aware of its origin or of having gone through any conscious process of reflection or reasoning. "Colorless green ideas sleep furiously" is a familiar example. We have the intuition that it is grammatical (though meaningless): the judgment that it is grammatical arises spontaneously and immediately without our engaging in any conscious reflection or reasoning. Similarly, philosophers present scenarios or "thought experiments" intended

to elicit intuitive judgments from their audience. To take one example, in Trolley cases, a scenario is presented in which a protagonist in the scenario is faced with a decision to flip a switch to divert a trolley onto a sidetrack. If the switch is not flipped, the trolley will kill five people on the main track; if it is flipped, the trolley will kill only one person who is on the sidetrack. The audience is then asked whether it would be morally permissible to flip the switch. Many people offer a judgment which arises spontaneously and immediately without being aware of its origin or of having gone through any conscious process of reflection or reasoning (e.g., Hauser, Young and Cushman, 2008). In short, they intuit whether or not it is morally permissible to flip the switch.

4.2.2. The Target of Intuitions

What is the target of intuitions? For present purposes, I take it that intuitions can have one of two targets. The first target is *mentalist* where intuitions have as their target in-the-head psychological entities, most notably concepts; the second target is *extra-mentalist*, where intuitions have as their target outside-the-head non-psychological entities (Goldman and Pust, 1998).

Which view one takes on the target of intuitions depends, in part, on what philosophical project is being undertaken. If one's project is to engage in conceptual analysis, such as an analysis of the concept of causation, knowledge and the like, then insofar as the theorist takes intuitions to play a role in conceptual analysis, intuitions will be targeted at uncovering the features of the relevant concept, its content or its extension. But if one's project is to understand the outside-the-head phenomena themselves, for instance, the nature of causation, the nature of knowledge and the like, then intuitions will have as their target outside-the-head non-psychological entities, the phenomena themselves.

There's some controversy over whether the appropriate target for intuitions is mentalist or extra-mentalist. Goldman (2007) takes it that the only legitimate target for intuitions is mentalist since it is deeply puzzling how intuitions could provide evidence about the extra-mentalist phenomena themselves (e.g., knowledge, causation, etc.). So intuitions, insofar as they play a legitimate role in philosophical theorizing, do so largely in the course of conceptual analysis. But as Stich and Tobia (forthcoming) note, "[F]or most of human history, perception posed a comparable mystery, and the appeal to intuition in mathematics still does" (p. 11). So while we may lack an account of how intuitions could be linked to extra-mentalist phenomena, this shouldn't lead to skepticism that they can be legitimately used in this way. And as Sosa (2007) points out, intuitions are often used in this way. That is, philosophers who rely on intuitions often take the content of the intuition to be about the world, and not merely about what is in one's head, and so on this view the content of the intuition is taken to be true (provided the intuition hasn't been triggered in error or is due to a bias, etc.).

In the metaphysical disputes I'll be discussing below, it is not always clear what target for intuitions the theorist has in mind. For present purposes, I'll simply assume that they take the target of intuitions to be extra-mentalist. There are two reasons for this. First, given the mereological disputes I'll be discussing, it is more natural to understand the appeal to folk intuitions as being useful insofar as they serve as a guide as to what there is in the world. For instance, mereological nihilism—the view that composition never occurs—holds that there are no composite objects. It's hard to make sense of this view if it is born out of an analysis of our concept of composition. Second, conceptual analysis

has fallen out of favor for quite some time[1] and so it seems reasonable to take contemporary metaphysicians, unclear as they are as to what they take the targets of intuitions to be, to assume that intuitions are targeted at the extra-mentalist phenomena themselves.

4.2.3. Taking Folk Intuitions Seriously and The Explanatory Burden

Why might folk intuitions be taken seriously? And why think an explanatory burden arises for those who put forward theories that conflict with folk intuitions? A number of different rationales could be offered for why folk intuitions might be taken seriously. But since my main goal is to illustrate how cognitive science can help the revisionary metaphysician discharge the explanatory burden of providing a plausible explanation of how the folk have gone wrong, I'm only going to focus on one rationale. Lewis (1986) tells us that "Common sense is a settled body of theory— unsystematic folk theory— which at any rate we do believe; and I presume that we are reasonable to believe it. (Most of it.)" (p. 134). So if it is taken that we are reasonable to believe common sense, then a theory that conflicts badly with it would appear to have some explaining to do. Specifically, a theory which conflicts with it should give us reasons for thinking that the relevant aspect of the folk view is wrong and tell us just how it is wrong. Indeed, as Korman notes "[V]irtually 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-

[1] Though see e.g., Jackson, 1998 and Kriegel, forthcoming

philosophers ("the folk")" (Korman, 2009, p. 242). That is, the revisionary metaphysician, in offering a metaphysical theory which departs from common sense or folk intuitions accrues an explanatory burden in that she is required to provide a plausible account of how the folk have gone wrong.[2]

4.2.4. Mereology and The Explanatory Burden

To illustrate how cognitive science can help the revisionary metaphysician discharge the explanatory burden, I'm going to focus on two issues connected with mereology. Mereology is roughly concerned with parthood relations, the relations of part to whole and the relations of part to part within a whole. Specifically, I'm going to focus on disputes over the question of when mereological composition occurs—when do the parts make a whole?—and a connected question over how a whole—an object—persists through alterations to its parts.

Metaphysical theories of composition and persistence are often charged with violating folk intuitions. In short: a number of metaphysical theories of composition and persistence are revisionary in that they allegedly conflict with folk intuitions. With respect to composition, one example of a revisionary theory is universalism. The universalist thinks that composition occurs under any circumstance whatsoever. It is completely unrestricted. It always occurs. On this view, there exists, for instance, a fusion of my nose and a doorknob, a fusion of a trout and a turkey, and even a fusion of the moon and a piece of cheese. Perhaps unsurprisingly, universalism is typically charged with violating folk intuitions since presumably the folk do not recognize the kind

[2] This isn't to say metaphysicians are the only philosophers confronted with this kind of explanatory burden. For instance it also arises in ethics (Mackie, 1977; see Rose and Nichols, 2016 for an overview).

of fusions that the universalist recognizes. Universalism is thus typically regarded as a revisionary metaphysical theory.[3] And indeed one of the main objections against universalism is that it violates folk intuitions, though the charges are typically disputed.[4] Thus the task put to the universalist is to provide some plausible explanation of how it is that the folk have gone wrong.

So too with persistence. Again to take just one example, Burke's (1994) sortal based account of persistence is typically charged with violating folk intuitions. On Burke's view, when a sculptor takes a piece of copper and fashions it into a statue, the piece of copper is destroyed, and in its place there comes to exist a new piece of copper which is in turn identical to the resultant statue. One main objection against Burke's sortal based account is that it violates folk intuitions. Nobody, not even children, thinks that an object is destroyed upon merely assuming a certain shape (Lowe, 1995). Burke however disputes these charges.[5] Insofar as this account does conflict with folk intuitions, it counts as a revisionary metaphysical theory. Given this, the metaphysician who embraces such a view of persistence is faced with the task providing some plausible explanation of how it is that the folk have gone wrong.

In general, if a metaphysical theory is revisionary, then the theorist will often be confronted with the explanatory burden. But if this is the case, how might the revisionary metaphysician discharge the explanatory burden? One way would be to argue that the

[3] I'm focusing on universalism simply for illustrative purposes. A number of other views have also been charged with violating folk intuitions and so count as revisionary metaphysical theories. For an overview see Rose and Schaffer, forthcoming.
[4] See Rose and Schaffer, forthcoming
[5] Here again, I've only focused on Burke's view for illustrative purposes. For more see Rose (2015) and Sidelle (2002).

theory doesn't conflict with the relevant folk intuitions.[6] On this way of proceeding, the alleged conflict between the relevant folk intuitions and the deliverances of the theory would only be apparent, not genuine. Another way would be for the revisionist to embrace the conflict between the relevant folk intuitions and the deliverances of the theory as being genuine and seek out an empirically informed debunking explanation of the relevant folk intuitions. This is my preferred strategy. So, moving on to this, I want to now discuss the particular view of composition and persistence—folk mereology—I take the folk to be operating with and some of the evidence supporting this before getting on to the debunking explanation.

4.3. Promiscuous Teleomentalism, Folk Mereology and Debunking

4.3.1. Promiscuous Teleomentalism

My view is that the folk heavily indulge in promiscuous teleomentalist thinking. There are two parts to this view. The first is that the folk are promiscuous teleologists in that teleological considerations inform their understanding of artifacts, organisms and even non-living natural things like rocks. Second, the specific teleological view the folk adopt is teleomentalism where teleology is understood psychologically, in terms of intentions, goals and purposes.[7] In other words, the folk view reality as a whole as being infused

[6] See Rose and Schaffer (forthcoming) and Rose (2015) for documentation of a number of philosophers claiming that their view doesn't really conflict with folk intuitions.

[7] This contrasts with teleonaturalism where teleological claims are to be understood in ways that do not refer to the intentions, goals, or purposes of psychological agents' (see Allen and Bekoff, 1994, p. 13 for a discussion of teleonaturalism and teleomentalism). Though teleonaturalism may be a respectable form of teleology, I would note that this is a stipulated/revisionary idea introduced to legitimize certain uses of teleology in biology, and was never intended to correspond to the actual concept the folk operate with.

with agency and purpose and thus the folk indulge in promiscuous teleological thinking because they are promiscuous telelomentalists.

Why think that the folk are promiscuous teleomentalists? There are two relevant lines of psychological research. One concerns promiscuous teleology; the other concerns teleomentalism underwriting promiscuous teleological thinking. I'll briefly discuss each line of evidence (for further details see Rose, 2015 and Rose and Schaffer, forthcoming).

First, regarding promiscuous teleology, it has been found, for instance, that children insist that lions are for "going to the zoo", that clouds are "for raining" (Bloom, 2007, p. 150), that "mountains exist to give animals a place to climb" and that rocks are pointy "so that animals won't sit on them and smash them" (Kelemen, 1999, pgs. 1444-1445). Promiscuous teleological thinking isn't merely confined to children. Instead, it extends into adulthood. For instance, Lombrozo et al (2007) found that adults with Alzheimer's disease, who tend to suffer from deficits in background causal beliefs, naturally default to accepting promiscuous teleological explanations. Even those without cognitive deficits, such as college educated adults, accept promiscuous teleological explanations such as "the sun radiates heat because warmth nurtures life", "fungi grow in forests to help decomposition", and "lightning occurs to release electricity" with ease. In fact, even professional physical scientists naturally accept promiscuous teleological explanations when their cognitive resources are limited, such as when in a speeded task. This research suggests that promiscuous teleological thinking emerges early in childhood and is retained into adulthood. And though it may be masked in certain kinds of cases (e.g., as with professional physical scientist explicitly considering the acceptability of

promiscuous teleological explanations), it is retained nonetheless. It represents a deep-seated, natural default perspective on the world.

Second, concerning teleomentalism, the evidence suggests that the folk take an agentive perspective on reality as a whole. Importantly, teleomentalism isn't tied into any particular background religious views. In this way, teleomentalism is open to all. For instance, Kelemen and DiYanni (2005) report a strong tendency among children—both from religious and nonreligious backgrounds—to an 'intuitive theism' in which nature is viewed as an artifact of a creator, as well as a significant correlation between viewing something teleologically and regarding it as created. Moreover, in recent research by Rottman et al (2016), they found that even Chinese people, despite being in a largely atheistic culture, succumb to promiscuous teleological thinking. Though background religious views don't appear to be heavily tied into promiscuous teleological thinking, endorsement of background Gaia beliefs (Kelemen et al., 2013; see also Kelemen, 2012)—beliefs that nature is a living, powerful, goal directed being—has been shown to significantly predict peoples tendency to engage in promiscuous teleological thinking. On top of all this, decades of research in science education has found that an agentive perspective on reality is one of the main obstacles students confront—religious and non-religious alike—in acquiring a proper understanding of natural selection (see Kelemen, 2012 for an overview). We have what Pascal Boyer (2001) has called a "hypertrophy of social cognition"— a willingness to attribute purpose, agency and design, even when it is inappropriate to do so—and are "hypersensitive to signs of purpose, design and agency, so much so that we see purpose where all that really exists is artifice or accident"

(Bloom, 2007, p. 150). In short: the evidence suggests that promiscuous teleological thinking is underwritten by teleomentalism.

So research in psychology suggest that the folk are promiscuous teleomentalists. The main question now is whether promiscuous teleomentalism plays a role folk mereology. I'll take this up next.

4.3.2. Folk Mereology

I won't rehearse the full details of the work on folk mereology (see Rose, 2015 and Rose and Schaffer, forthcoming for further details). Instead, I will discuss some of the highlights to set out some support for the view that promiscuous teleomentalist thinking underwrites folk intuitions of composition and persistence. First, composition.

Beginning with artifacts, people were given a case featuring an unfamiliar artifact, a "gollywag". They were told about two individuals, Smith and Jones, experimenting with the gollywags. In one case, the no purpose case, participants were told that Jones superglues some of the gollywags together. In the purpose case, people were again told that Jones superglued the gollywags together. But they were also told that Jones had a sore back, placed the superglued gollywags on his chair, sat on them for the rest of the day and no longer had a sore back. In both cases a disagreement ensues between Smith and Jones with Smith saying that the gollywags do not compose an object and Jones saying they do. Whether the gollywags had a purpose strongly affected people's intuitions about whether composition had occurred: people were significantly more inclined to think that composition had occurred when the gollywags had a purpose.

It turns out that teleological considerations aren't merely confined to intuitions about composition when considering artifacts. They also play a role in intuitions about composition when considering biological organisms and non-living natural things like rocks. Concerning biological organisms, people were again told that Smith and Jones are experimenting, but this time with mice. In one case Jones superglues some mice together; in another case he superglues them together, runs them through a maze, and finds that they are very successful at detecting bombs. In the former case, the superglued mice had no purpose; in the latter case, they did. Again, there was an impact of teleology on intuitions about composition with people being significantly more inclined to think that the superglued mice composed an object when they had a purpose. For the cases involving non-living natural things, rocks, people were told about an individual who lives on the side of a mountain. One evening an avalanche occurs and rocks are scattered across the individual's front lawn. What was varied was whether the protagonist thought that the scattered plurality of rocks had no purpose, thought they had a purpose, or arranged them for a purpose. The key finding was that when the scattered plurality of rocks had no purpose people denied that composition had occurred but when the protagonist thought the arrangement of rocks had a purpose or arranged them for a purpose people thought that composition had occurred.

Promiscuous teleology also plays a significant role in intuitions about persistence. In one study people were told about an individual, John, who is hiking and spots a glowing rock. The rock houses a special sort of microorganism which feeds off minerals in the rocks interior. But the microorganisms can't access the minerals deep in the rocks interior. The microorganisms begin dying and the rock begins fading. So John tries an experiment.

He smashes the rock into three pieces with a hammer. In one case, the microorganisms continue dying and the thing fades to black. In the other case, the microorganisms can access the minerals and the thing resumes glowing. In the former case people thought the rock did not survive the smashing; in the latter case people thought the rock survived the smashing. So even when turning to persistence, we continue to see a pattern of promiscuous teleological intuitions.

Turning now to teleomentalism, one of the main reasons for thinking that promiscuous teleology is underwritten by teleomentalism comes from the background research in psychology. But there is some more direct evidence. In particular the same case concerning persistence (discussed above) was rerun. But this time people's background Gaia beliefs were assessed.[8] Utilizing a causal modeling procedure, it was found that Gaia beliefs caused people's teleologically laden intuitions about persistence. This suggest that the tendency toward promiscuous teleology is underwritten by teleomentalism. Taken together, the evidence from Rose (2015) and Rose and Schaffer (forthcoming) suggests that folk intuitions of composition and persistence are tied into promiscuous teleomentalism.[9]

---

[8] Kelemen et al (2013) utilize the following probes as a measure of belief that reality as a whole is infused with agency (i.e., Gaia beliefs):
(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
In the study under consideration only (1) was used (see Rose, 2015, p. 118). See Dink and Rips (forthcoming) for a critical discussion of whether these measures, especially (3), tap in to background beliefs about agentive forces affecting the natural world.

[9] I would also note that some recent work on folk intuitions about causation has also found that Gaia beliefs cause people's promiscuous teleological causal intuitions (Rose, 2016). I would also add, though not pursue in any detail here, that it may be that folk mereology is underwritten by causal intuitions. The idea here is that people understand composition and persistence causally and given that causal intuitions are tied into promiscuous teleomentalism, it may be that these promiscuous teleomentalist tendencies run through causal intuitions to affect both composition and persistence intuitions.

4.3.3. Debunking Folk Intuitions and Meeting the Explanatory Burden

We saw that the revisionary metaphysician, in offering a metaphysical theory that conflicts with folk intuitions, is often faced with an explanatory burden in that she is required to provide a plausible explanation of how the folk could have gone wrong. And one case where we see the explanatory burden put forth is in disputes over mereology, specifically in discussions of composition and persistence. While one typical strategy for meeting the explanatory burden is to deny that there is any conflict between the deliverances of the theory and folk intuitions, my suggestion was that the revisionary metaphysician might do better to confront the explanatory burden head on and seek out an empirically informed debunking explanation of the relevant folk intuitions. Having now discussed some empirical evidence that the folk operate with a promiscuous teleomenatlist view of composition and persistence, we're now in a position to set out how the revisionary metaphysician might debunk these aspects of folk mereology and thereby discharge the explanatory burden.

I take it that one constraint on a metaphysical theory is that it shouldn't conflict with what our best science says about the world. This isn't to say that every metaphysical question is answerable to our best science. Nor is it to say that metaphysical theorizing is held entirely hostage by our best science (Ladyman and Ross, 2007). Instead I am only making the more modest claim that metaphysics, insofar as it is appropriate to the issue at hand, should be constrained, in part, by what our best science says about the world. In this way, I view science and metaphysics as enjoying a joint partnership, mutually informing one another (Paul, 2012). That said, there is an obvious respect in which promiscuous teleomentalism conflicts with what our best science says about the world.

Teleology has, for the most part, been purged from the natural sciences.[10] Indeed, the rejection of a teleological perspective on all of the natural world traces at least as far back to the emergence of modern science from medieval Aristotelianism. It's arguably due, at least in part, to the rejection of a teleological perspective on all of the natural world that science has made great strides over the past several hundred years. Not to put too fine a point on it, the rejection of an agentive, teleological perspective on the natural world is also one reason why modern day intelligent design "science" is widely rejected. If scientists were to accept some version of intelligent design along with the agentive, teleological perspective on all of the natural world that comes along with it, science would slip into the dark ages. In short: it is widely agreed that a teleological perspective whereby all of nature—every rock and cloud—is infused with agency and purpose is part of a "superseded, pre-scientific muddle about how the world works" (Hawthorne and Nolan, 2006; see also e.g., Mayr, 1998; Allen and Bekoff, 1994).

So for the metaphysician who views herself as allied to the sciences, I take it that she ought to reject a teleological perspective on the natural world and as such the folk view of composition and persistence since it is encrusted with the muck and funk of an outmoded teleological perspective. Indeed, as Laurie Paul (2012) notes: "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). Given the conflict between the

[10] Teleological notions play a role in biology and there is a dispute over whether appeals to teleology in biology might be naturalistically respectable (see Allen, 2009 and Allen and Bekoff, 1994 for overviews). But whether or not there is some revised, naturalistically respectable view of teleology to be had in biology, it's agreed that promiscuous teleomentalism is not one such view.

teleological commitments of the folk in the cases of composition and persistence and the (presumptive) commitments of the revisionary metaphysician who views herself as allied with the sciences, in the specific cases of composition and persistence, there looks to be a basis to debunk folk intuitions.

Debunking can be fleshed out in different ways depending on one's background epistemological view (see e.g., Schaffer, forthcoming for an overview). To take just one example, Nichols (2014) explores one option for debunking—process debunking—which is based on identifying epistemically defective processes generating certain intuitions and is naturally viewed as being tied into reliabilism. But whatever one's preferred epistemological view for supplying debunking arguments, debunking should be aided by a metaphysical assessment (Schaffer, forthcoming). And I would add, as I've already noted, that certain metaphysical assessments should be informed and guided by our best science. In the case of folk mereology, our best science helps issue a clear metaphysical assessment: teleological infused intuitions fail to fit reality. On this basis, they deserve to be debunked.

Operating with a fairly neutral background epistemology, I view the situation as follows: conformity with folk intuitions confers prima facie justification on a theory, perhaps because, as Lewis (1986) points out, it's often reasonable to believe most of the deliverances of common sense. But once we have empirically assessed the relevant folk intuitions, in this case, two aspects of folk mereology—composition and persistence—we see that they are tied to a promiscuous teleomentalist outlook on reality. As such, they fail to fit reality. The assessment that promiscuous teleomentalism—and thereby folk mereological intuitions—fail to fit reality is a metaphysical assessment, one which is

aided by what our best science says about the world. Given this, and thinking of folk intuitions as testifying as to what the relevant extra-mentalist phenomena is like, the intuitional testimony lacks ultima facie justification once we learn from cognitive science that the folk have a promiscuous teleomentalist view of the issue. In other words the evidence from cognitive science, coupled with a metaphysical assessment aided by our best science, helps supply an undermining defeater (Pollock, 1987) for folk intuitions of composition and persistence, which substantially strips these intuitions of their evidential credentials with respect to the extra-mentalist phenomena at hand, thereby cutting them off from achieving ultima facie justification.

Concerning the explanatory burden faced by the revisionary metaphysician in the case of composition and persistence, the revisionary metaphysician has the resources to discharge the explanatory burden. Armed with the results delivered from cognitive science and a background commitment to metaphysical theorizing being constrained, in part, by what our best science has to say about the world, she has an explanation as to how the folk have gone wrong in the case of composition and persistence. She has the resources to supply an undermining defeater for folk intuitions about composition and persistence, thereby debunking the relevant folk intuitions and meeting the explanatory burden. In this way, evidence from cognitive science can help the revisionary metaphysician discharge the explanatory burden of providing a plausible account of how the folk have gone wrong.

4.4. Objections

Having set out the specific illustration of how cognitive science can help the revisionary metaphysician discharge the explanatory burden, I now want to consider two main

objections. If either of these is right, it will undercut the debunking explanation set out above.

4.1.1. Legitimate Teleology

The first main objection is that the teleological view the folk operate with may well be entirely legitimate (Korman and Carmichael, forthcoming). Folk intuitions of composition don't appear to be best explained by claiming that the folk are working with anything like the crude, superstitious view embodied in promiscuous teleomentalism. Rather there is something much more innocuous at work. Folk intuitions of composition, far from being tied into promiscuous teleomentalism, are tied into an assessment of creative intentions. On this view, the intentions of a creator play a significant role in folk intuitions of composition. For instance, if a sculptor was fashioning some copper into a statue with the intention of creating a statue, it would be entirely appropriate to use the creators intention—to create a statue—as a guide to whether the arrangement of copper composes a statue. And as Korman and Carmichael (forthcoming) note, in all of the cases concerning composition in Rose and Schaffer (forthcoming)—cases involving artifacts, biological organisms and even non-living natural things like rocks—the key difference between the purpose and no purpose versions of the cases lies in a difference between whether the relevant agent has or lacks the relevant creative intention.

As I understand the objection, at least two considerations must be in play in each of Rose and Schaffer's composition cases. The first is that there should be an agent and that this agent should either have or lack an intention to create some further thing. The second is that in every case where participants are considering whether an agent has an intention to create something, they are treating the thing as an artifact. A few brief remarks on these

points. When considering artifacts, consulting an agent's creative intentions is arguably relevant in considering the question of whether the arrangement of some parts compose some further object. But creative intentions would be irrelevant when considering composition if there is either no agent or if the would-be composite is a non-artifact. Since most would presumably agree that the creative intentions of an agent are relevant when considering composition for artifacts, it must be the case that creative intentions play a role in cases involving non-artifacts because people are construing the would-be composite as an artifact.

Concerning the cases of composition, Korman and Carmichael are right that in each case, whether it involved an artifact, organism or non-living natural thing, there was an agent with (or without) an intention to create some further thing.[11] Given this, it may well be that participants are construing the would-be composite as an artifact.

I think the creative intentions account fails to provide a satisfactory explanation of folk mereological intuitions. First, it is not even clear that it can provide a general account of folk intuitions of composition. Given that Korman and Carmichael only put it to work in attempting to provide an alternative explanation of the Rose and Schaffer results, it is unclear how the creative intentions account would extend to other cases of composition. How do the folk make composition judgments when no (actual) agent is involved in the case? What about when the would-be composite isn't plausibly construed as an artifact?

[11] One exception to this is the Avalance Accorded Function case from Rose and Schaffer (forthcoming) since in this case the relevant agent didn't plausibly intend to *create* a rock garden from the scattered arrangement of rocks. He only accorded the plurality of rocks scattered across his lawn the purpose of being a rock garden. Korman and Carmichael maintain though that the agent at least has an intention that the scattered plurality of rocks be a rock garden (see Endnote V).

What about when the agent's intention are unsuccessful? We would need a lot more detail regarding their account to answer these and related questions.

Perhaps the creative intentions account isn't aimed at providing a general account of folk intuitions of composition. In that case we would need some alternative account of folk intuitions of composition for cases in which the creative intentions account is ill suited. But the view that folk intuitions of composition are tied into promiscuous teleomentalism is aimed at providing a general account of folk intuitions of composition. It thus provides a simpler, more general account of folk intuitions of composition.

Second, the creative intentions account doesn't provide a plausible explanation of folk intuitions about persistence. In those case, there was a human agent involved but the relevant agent lacked creative intentions. For instance, in the rock smashing cases, John smashes the glowing rock—which is fading as the microorganisms can't access minerals deep in the rocks interior—into three pieces. In one case (see Rose, 2015), the microorganisms all begin dying and the thing fades to black; in the other case, the microorganisms can access the minerals after the smashing and the thing resumes glowing. It's doubtful that in either of these cases John had an intention to create anything. Even assuming that John did possess some creative intention, it seems that in both cases John presumably possesses either the intention to create a glowing rock, a better situation for the microorganisms or both. Yet, people make different persistence judgments in this case. If creative intentions play a role in persistence judgments as well, then given that they are held fixed across these cases, we shouldn't see any difference in

persistence judgments. But we do. So there's good reason to think that the creative intentions account fails to fit the pattern of findings from this case.[12]

Perhaps the creative intentions account isn't aimed at providing a more general account of folk mereological intuitions. Korman and Carmichael only focused on folk intuitions about composition in the Rose and Schaffer cases so perhaps they think that there must be some alternative account for why the folk have teleologically laden intuitions when considering persistence. On my proposal, however, there is a single view underwriting both aspects of folk mereology i.e., composition and persistence. That is, promiscuous teleomentalism underwrites folk mereological intuitions and so provides a unifying explanation of the processes underwriting folk intuitions of composition and persistence. Given this and given that it provides a smooth, simple fit to the data, we're owed some explanation of what might be going on in the persistence cases since the creative intentions account is inadequate. But Korman and Carmichael offer no such proposal.

Korman and Carmichael might point out that nonetheless even in the persistence cases involving rocks there is an agent involved and so plausibly the agent's intentions, whatever they are, may still be guiding folk judgments. Moreover, they might note, the folk are nonetheless considering the relevant things in terms of artifacts. So agent intentions are playing some yet to be specified role and moreover in each case, people treat the relevant thing as an artifact. Against this, I would point out two things. First, it is implausible that people treat the glowing rock as an artifact. Nobody created the glowing rock. John found it by the side of a trail. To insist that the folk treat the rock as

[12] I would flag that I'm assuming creative intentions are held fixed across these cases. But they may not be. Given that a probe assessing creative intentions wasn't used for this study, it's an open empirical question whether people actually view the creative intentions similarly in the cases.

an artifact would be to saddle the folk with the kind of "uncharitable" view that Korman and Carmichael want to avoid. Second, I doubt that the creative intentions of a human agent are playing a role in the persistence cases because Gaia beliefs cause teleologically laden persistence intuitions (see Rose, 2015). So an (actual) agent's creative intentions are not playing a role in underwriting teleologically laden persistence intuitions. Instead, and in line with my hypothesis, promiscuous teleomentalism looks to be underwriting these intuitions. Given that, and it's coherence with background work in psychology, promiscuous teleomentalism also plausibly underwrites folk intuitions of composition. In short: I take it that promiscuous teleomentalism drives, plays a general role in, folk mereological intuitions and that the folk don't treat every collection of things as an artifact. Given this, I maintain the claim that the folk operate with a benighted view of composition and persistence. Folk mereology is unfit for real metaphysics.

4.4.2. Answers, Not Intuitions

The second objection I want to consider is that the folk, in responding to the prompts in the experiments, aren't reporting their intuitions. Instead they are merely reporting their answers.[13] If this is right, then folk mereological *intuitions* wouldn't be debunked (Korman & Carmichael, forthcoming).

This kind of objection tends to be broad in scope.[14] If right, it wouldn't simply threaten the work on folk mereology. Instead it would extend to every claim anyone ever makes about folk intuitions, including any claim philosophers make about folk intuitions. No

[13] Of course, reporting an intuition is to report a kind of answer. Korman and Carmichael are using "reporting an answer"" in the sense of "not reporting an intuition".
[14] Though Korman and Carmichael only raise it for the Rose and Schaffer studies, it is easily extended to any study (see e.g., Bengson, 2013).

one would be in a position to speak to the content of folk intuitions. As a general criticism, I think there are good reasons to be suspicious of such a claim. I also think that, in the specific case of folk mereology, there is good reason to doubt that the folk are reporting answers and not intuitions.

Korman and Carmichael claim that, in contrast to reporting intuitions, giving answers involves reporting considered judgments after talking one's self out of her intuitive reactions (p. 7). This suggests that reporting an answer, as opposed to an intuition, involves something like the following: Upon having an intuition, a subject would, for instance, have to bring some background considerations to bear on the task, leading her to suppress reporting the intuition and instead report an answer which is independent of the intuition. This is a somewhat complicated procedure which would surely require a good deal of cognitive effort to execute. And indeed, there is some reason for thinking that the folk have a difficult time executing such a process. Take the Cognitive Reflectivity Test (CRT) developed by Frederick (2005). Here's one test item from the CRT:

A bat and a ball cost $1.10 in total. The bat costs $1.00 more than the ball. How much does the ball cost?

This question has an incorrect but intuitive answer, $.10. To arrive at the correct answer, $.05, one needs to suppress the intuitive judgment, bring background knowledge of algebra to bear on the case and report the answer delivered from the algebraic computation. Reporting an answer as opposed to an intuition requires a good deal of cognitive effort. And it turns out that in general, the majority of people fail the CRT (i.e., they report the intuitive but incorrect answer for one or more items). Indeed, even those who score high on the CRT (i.e., give the correct answer to each item) are not immune

from giving intuitive responses. For instance, one study found that 38% of those scoring high on the CRT commit the conjunction fallacy, that some thing possessing two properties is more probable than that thing possessing a single property (Ochssler, Roider and Schmitz, 2009). Given that the majority of people fail the CRT and given that some people who score high on the CRT even naturally default to reporting intuitions in some cases, there is good reason to think that, as a default, in any given study probing folk intuitions, people are reporting intuitions and not answers.

In general, I think the claim that people are reporting answers and not intuitions should be viewed with suspicion. There is good reason for thinking that people do indeed tend to report intuitions and that it takes a good deal of cognitive effort to do otherwise. And in the specific case of folk mereology, I also think that the claim that the folk are reporting answers and not intuitions fails.

One reason it fails is because when one looks at work on folk teleology in psychology, one finds that at a very young age children naturally offer and accept promiscuous teleological explanations, for instance, thinking that "mountains exist to give people a place to climb". It's doubtful that children aren't reporting intuitions in this case. It is implausible that they are going through something like the kind of complicated cognitive procedure suggested above in order to provide answers as opposed to intuitions. And given that adults display teleologically laden responses which reflect these childhood tendencies, there is good reason for thinking that they are indeed reporting intuitions, not answers. Moreover, even though some groups of individuals tend to avoid explicit teleological thinking, such as professional physical scientists, these same individuals naturally default to a teleological thinking when their cognitive resources are limited,

such as when in a speeded task. This suggest that teleological thinking is an intuitive default, suppressed only by engaging in a good deal of cognitive effort. Given that folk mereological intuitions reflect this default aspect of intuitive teleological thinking, this suggest that their judgments reflect their intuitions and not merely their answers.

4.5. Conclusion

The revisionary metaphysician—in offering a theory that conflicts with folk intuitions—is typically confronted with an explanatory burden in that she is required to provide a plausible explanation of how the folk have gone wrong. My suggestion was that evidence from cognitive science can help the revisionary metaphysician discharge this explanatory burden. I illustrated this by taking up metaphysical disputes over mereological composition and persistence. I then discussed evidence from cognitive science which suggests that the folk are promiscuous teleomentalists and that promiscuous teleomentalist thinking underwrites folk intuitions of composition and persistence. On the basis of this, I went on to argue that there is a debunking explanation for folk intuitions. I then responded to two main objections to the debunking argument I set out, finding both objections to be insufficient to undercut the debunking argument put forth. Thus, I uphold the view that folk mereological intuitions concerning composition and persistence deserve to be debunked. In this way, I take myself to have illustrated one key role cognitive science can play in metaphysics; namely by helping the revisionary metaphysician discharge the explanatory burden of providing a plausible explanation of how the folk have gone wrong.[15]

[15] Thanks to Chad Carmichael, Dan Korman, Shaun Nichols, Jonathan Schaffer, Steve Stich and John Turri for helpful comments on earlier versions of this paper.

Appendix A

Folk Intuitions of Actual Causation

A.1. Study 1

A. 1.1 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.

A.1.2. 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.

A.1.3. 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.

A.1.4. 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.

A.2. Study 2

A.2.1. Cases A.1.3 and A.1.4 (Above)

A.3. Study 3

A.3.1. 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.

A.3.2. 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.

A.3.3. Andy, 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. 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.

A.3.4. Andy, 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. 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.

A.3.5. 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.

A.3.6. 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.

A.3.7. 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.

A.3.8. 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.

Appendix B

Persistence Through Function Preservation

B.1. Study 1

B.1.1. Rowboat Cases

No Function: John is an accomplished woodworker and sailor, whose lifelong hobby is building rowboats by hand. He built his first rowboat—which he named "Drifter"—thirty years ago. Over the years there has been wear and tear, and it turns out that every single one of the original planks has needed to be replaced.

John—never one to throw anything out—has stored all of the original planks in his shed over the years. Last month John—realizing that he had accumulated enough old planks in his shed for a whole rowboat—took out his old plans for "Drifter" and assembled these planks exactly according to his old plans. John now has two rowboats of the same design: the rowboat originally built thirty years ago that has none of its original parts, and the rowboat just built one month ago with all of the original parts from "Drifter".

John has promised two of his friends—Suzy and Andy—that they can borrow Drifter for an outing. But Suzy and Andy aren't sure which of the two rowboats is actually Drifter. Andy thinks that the rowboat just built a month ago is actually Drifter, since it has exactly the same parts, arranged in exactly the same way as Drifter originally had. But Suzy thinks that the rowboat built thirty years ago with all of its parts completely replaced is actually Drifter, since, even though it has all new parts, this was just the result of normal maintenance.

Replacement Preserves Function: John is an accomplished woodworker and sailor, whose lifelong hobby is building rowboats by hand. He built his first rowboat—which he named "Drifter"—thirty years ago. It was an excellent rowboat: it was very stable, never let water in and sailed smoothly. Over the years there has been wear and tear, and so John, to keep it in perfect working order, has replaced various parts. And, it turns out that, over the years, every single one of the original planks has needed to be replaced.

John—never one to throw anything out—stored all of the original planks in his shed over the years. Last month John—realizing that he had accumulated enough old planks in his shed for a whole rowboat—took out his old plans for "Drifter" and assembled these planks exactly according to his old plans. This rowboat, however, worked terribly: it was very rickety, always let water in and sank after just a few minutes in the water. So, John now has two rowboats of the same design: the rowboat originally built thirty years ago that has none of its original parts and works perfectly as a rowboat, and the rowboat just built one month ago with all of the original parts from "Drifter", which works terribly as a rowboat.

John has promised two of his friends—Suzy and Andy—that they can borrow Drifter for an outing. But Suzy and Andy aren't sure which of the two rowboats is actually Drifter. Andy thinks that the rowboat just built a month ago is actually Drifter, since, even though it works terribly as a rowboat, it has exactly the same parts, arranged in exactly the same way as Drifter originally had. But Suzy disagrees. She thinks that the rowboat built thirty years ago with all of its parts completely replaced is actually Drifter, since even though it has all new parts, this was just the result of the normal maintenance required to keep the rowboat in perfect working order.

Original Parts Preserves Function: John is an accomplished woodworker and sailor, whose lifelong hobby is building rowboats by hand. He built his first rowboat—which he named "Drifter"—thirty years ago. It was an excellent rowboat: it was very stable, never let water in and sailed smoothly. But John was always thinking of ways to try and make the rowboat even better. And over the years, every single one of the original planks ended up being replaced. The end result, however, was not an improved rowboat. Rather the rowboat ended up being terrible: it was very rickety, always let water in and sank after just a few minutes in the water.

Fortunately, John—never one to throw anything out—stored all of the original planks in his shed over the years. Last month John—realizing that he had accumulated enough old planks in his shed for a whole rowboat—took out his old plans for "Drifter" and assembled these planks exactly according to his old plans. This rowboat worked perfectly: it was very stable, never let in water and sailed smoothly across the water. So, John now has two rowboats of the same design: the rowboat originally built thirty years ago that has none of its original parts and works terribly as a rowboat, and the rowboat just built one month ago with all of the original parts from "Drifter", which works perfectly as a rowboat.

John has promised two of his friends—Suzy and Andy—that they can borrow Drifter for an outing. But Suzy and Andy aren't sure which of the two rowboats is actually Drifter. Andy thinks that the rowboat just built a month ago is actually Drifter, since it has exactly the same parts, arranged in exactly the same way as Drifter originally had and works perfectly. But Suzy thinks that the rowboat built thirty years ago with all of its

parts completely replaced is actually Drifter, since, even though it has all new parts and is a terrible rowboat, this was just the result of John's attempt to improve it.

Comprehension Checks:

(1) The rowboat built a month ago is made from all the original parts of "Drifter". (Yes/No)

(2) The rowboat built thirty years ago has all new parts. (Yes/No)

(3) Suzy thinks that the rowboat built thirty years ago with all of its parts completely replaced is actually Drifter. (Yes/No)

(4) Andy thinks that the rowboat just built a month ago is actually Drifter. (Yes/No)

B.1.2. Organism Cases

No Function: John is an accomplished biochemist who has devoted his life to studying organisms. When he first began his career thirty years ago, he discovered a new organism. He named the organism "Gollywag" and immediately logged the exact details of the organism. Over the years, John has conducted many experiments on the organism. Each time he conducts an experiment on the organism, he cuts off a part of it, which he uses for testing, and replaces it with the same type of part from another organism of the same type. Indeed, John has conducted so many experiments over the years that every single piece of the organism has been replaced.

After testing though, John always stores each part of the organism from the original organism. Last month John—realizing that he had accumulated enough parts of the organism for a whole organism—took out his logbook and assembled the parts exactly

according to his notes. John now has two organisms of exactly the same design: the organism discovered thirty years ago with all of its parts completely replaced and the organism assembled just a month ago with all of the original parts from “Gollywag”.

John has just hired two interns—Andy and Suzy—to work in his lab. He tells them that their first assignment will be to perform a series of experiments on “Gollywag”. But, Suzy and Andy aren't sure which of the two organisms is actually Gollywag. Andy thinks that the organism assembled just a month ago is actually Gollywag, since it has exactly the same parts, arranged in exactly the same way as Gollywag originally had. But Suzy thinks that the organism discovered thirty years ago with all of its parts completely replaced is actually Gollywag, since, even though it has all new parts, this was just the result of years of experimentation.

Replacement Preserves Function: John is an accomplished biochemist who has devoted his life to studying organisms. When he first began his career thirty years ago, he discovered a new organism. He named the organism “Gollywag” and immediately logged the exact details of the organism.

John has always suffered from Eczema, a condition where, despite the use of lotions and so forth, his hands are itchy, dry and red. But, after first handling the organism, he noticed that his Eczema disappeared: his hands were no longer itchy, dry or red. So, he decided to start conducting experiments on the organism.

Over the years, John has conducted many experiments on the organism. Each time he conducts an experiment on the organism, he cuts off a part of it, which he uses for testing, and replaces it with the same type of part from another organism of the same type.

Indeed, John has conducted so many experiments over the years that every single part of the organism has been replaced.

After testing though, John always stores each part from the original organism. Last month John—realizing that he had accumulated enough parts of the organism for a whole organism—took out his logbook and assembled the parts exactly according to his notes. When John handled this organism, however, he noticed his that his Eczema is severely aggravated: his hands were severely itchy, dry and red. So, John now has two organisms of exactly the same design: the organism discovered thirty years ago with all of its parts completely replaced which makes his Eczema completely disappear and the organism assembled just a month ago with all of the original parts from "Gollywag", which severely aggravates his Eczema.

John has just hired two interns—Andy and Suzy—to work in his lab. He tells them that their first assignment will be to perform a series of experiments on "Gollywag". But, Suzy and Andy aren't sure which of the two organisms is actually Gollywag. Andy thinks that the organism assembled just a month ago is actually Gollywag, since even though it severely aggravated John's Eczema, it has exactly the same parts, arranged in exactly the same way as Gollywag originally had. But Suzy thinks that the organism discovered thirty years ago with all of its parts completely replaced is actually Gollywag, since it makes John's Eczema disappear and even though it has all new parts, this was just the result of years of experimentation.

Original Parts Preserves Function: John is an accomplished biochemist who has devoted his life to studying organisms. When he first began his career thirty years ago, he

discovered a new organism. He named the organism “Gollywag” and immediately logged the exact details of the organism.

John has always suffered from Eczema, a condition where, despite the use of lotions and so forth, his hands are itchy, dry and red. But, after first handling the organism, he noticed that his Eczema disappeared: his hands were no longer itchy, dry or red. So, he decided to start conducting experiments on the organism.

Over the years, John has conducted many experiments on the organism. Each time he conducts an experiment on the organism, he cuts off a part of it, which he uses for testing, and replaces it with the same type of part from another organism of the same type. Indeed, John has conducted so many experiments over the years that every single part of the organism has been replaced. In the end, however, John noticed that when he handled the organism that, instead of relieving his itchy, dry, red skin that it was actually severely aggravated: his hands were dryer, itchier and redder than they had ever been.

Fortunately, after testing, John always stores each part from the original organism. Last month John—realizing that he had accumulated enough parts of the organism for a whole organism—took out his logbook and assembled the parts exactly according to his notes. When John handled this organism, he noticed his Eczema completely disappeared: his hands were no longer itchy, dry or red. So, John now has two organisms of exactly the same design: the organism discovered thirty years ago with all of its parts completely replaced which severely aggravates his Eczema and the organism assembled just a month ago with all of the original parts from “Gollywag”, which makes his Eczema completely disappear.

John has just hired two interns—Andy and Suzy—to work in his lab. He tells them that their first assignment will be to perform a series of experiments on “Gollywag”. But, Suzy and Andy aren't sure which of the two organisms is actually Gollywag. Andy thinks that the organism assembled just a month ago is actually Gollywag, since it has exactly the same parts, arranged in exactly the same way as Gollywag originally had and makes John’s Eczema completely disappear. But Suzy thinks that the organism discovered thirty years ago with all of its parts completely replaced is actually Gollywag, since, even though it has all new parts and severely aggravates John’s Eczema, this was just the result of years of experimentation.

Comprehension Checks:

(1) The organism assembled a month ago is made from all the original parts of "Gollywag".(Yes/No)

(2) The organism discovered thirty years ago has all new parts. (Yes/No)

(3) Suzy thinks that the organism discovered thirty years ago with all of its parts completely replaced is actually Gollywag. (Yes/No)

(4) Andy thinks that the organism assembled just a month ago is actually Gollywag. (Yes/No)

B.1.3. Rock Cases

No Function: John is an accomplished geochemist who has devoted his life to studying rocks. When he first began his career thirty years ago, he discovered a rock, made out of

an unknown mineral. He named the rock "Zenyte" and immediately logged the exact details of the rock. Over the years, John has conducted many experiments on the rock. Each time he conducts an experiment on the rock, he breaks off a piece of it, which he uses for testing, and replaces it with the same type of mineral. Indeed, John has conducted so many experiments over the years that every single piece of the rock has been replaced.

After testing though, John always stores each piece of mineral from the original rock. Last month John—realizing that he had accumulated enough pieces of the mineral for a whole rock—took out his logbook and assembled the minerals exactly according to his notes. John now has two rocks of exactly the same design: the rock discovered thirty years ago with all of its minerals completely replaced and the rock assembled just a month ago with all of the original minerals from "Zenyte".

John has just hired two interns—Andy and Suzy—to work in his lab. He tells them that their first assignment will be to perform a series of experiments on "Zenyte". But, Suzy and Andy aren't sure which of the two rocks is actually Zenyte. Andy thinks that the rock assembled just a month ago is actually Zenyte, since it has exactly the same parts, arranged in exactly the same way as Zenyte originally had. But Suzy thinks that the rock discovered thirty years ago with all of its parts completely replaced is actually Zenyte, since, even though it has all new parts, this was just the result of years of experimentation.

Original Parts Preserves Function: John is an accomplished geochemist who has devoted his life to studying rocks. When he first began his career thirty years ago, he

discovered a rock, made out of an unknown mineral. He named the rock “Zenyte” and immediately logged the exact details of the rock.

John has always suffered from Eczema, a condition where, despite the use of lotions and so forth, his hands are itchy, dry and red. But, after first handling the rock, he noticed that his Eczema disappeared: his hands were no longer itchy, dry or red. So, he decided to start conducting experiments on the rock.

Over the years, John has conducted many experiments on the rock. Each time he conducts an experiment on the rock, he breaks off a piece of it, which he uses for testing, and replaces it with the same type of mineral. Indeed, John has conducted so many experiments over the years that every single piece of the rock has been replaced. In the end, however, John noticed that when he handled the rock that, instead of relieving his itchy, dry, red skin that it was actually severely aggravated: his hands were dryer, itchier and redder than they had ever been.

Fortunately, after testing, John always stores each piece of mineral from the original rock. Last month John—realizing that he had accumulated enough pieces of the mineral for a whole rock—took out his logbook and assembled the minerals exactly according to his notes. When John handled this rock, he noticed his Eczema completely disappeared: his hands were no longer itchy, dry or red. So, John now has two rocks of exactly the same design: the rock discovered thirty years ago with all of its minerals completely replaced which severely aggravates his Eczema and the rock assembled just a month ago with all of the original minerals from “Zenyte”, which makes his Eczema completely disappear.

John has just hired two interns—Andy and Suzy—to work in his lab. He tells them that their first assignment will be to perform a series of experiments on "Zenyte". But, Suzy and Andy aren't sure which of the two rocks is actually Zenyte. Andy thinks that the rock assembled just a month ago is actually Zenyte, since it has exactly the same parts, arranged in exactly the same way as Zenyte originally had and makes John's Eczema completely disappear. But Suzy thinks that the rock discovered thirty years ago with all of its parts completely replaced is actually Zenyte, since, even though it has all new parts and severely aggravates John's Eczema, this was just the result of years of experimentation.

Replacement Preserves Function: John is an accomplished geochemist who has devoted his life to studying rocks. When he first began his career thirty years ago, he discovered a rock, made out of an unknown mineral. He named the rock "Zenyte" and immediately logged the exact details of the rock.

John has always suffered from Eczema, a condition where, despite the use of lotions and so forth, his hands are itchy, dry and red. But, after first handling the rock, he noticed that his Eczema disappeared: his hands were no longer itchy, dry or red. So, he decided to start conducting experiments on the rock.

Over the years, John has conducted many experiments on the rock. Each time he conducts an experiment on the rock, he breaks off a piece of it, which he uses for testing, and replaces it with the same type of mineral. Indeed, John has conducted so many experiments over the years that every single piece of the rock has been replaced.

After testing though, John always stores each piece of mineral from the original rock. Last month John—realizing that he had accumulated enough pieces of the mineral for a whole rock—took out his logbook and assembled the minerals exactly according to his notes. When John handled this rock, however, he noticed his that his Eczema is severely aggravated: his hands were severely itchy, dry and red. So, John now has two rocks of exactly the same design: the rock discovered thirty years ago with all of its minerals completely replaced which makes his Eczema completely disappear and the rock assembled just a month ago with all of the original minerals from “Zenyte”, which severely aggravates his Eczema.

John has just hired two interns—Andy and Suzy—to work in his lab. He tells them that their first assignment will be to perform a series of experiments on “Zenyte”. But, Suzy and Andy aren't sure which of the two rocks is actually Zenyte. Andy thinks that the rock assembled just a month ago is actually Zenyte, since even though it severely aggravated John’s Eczema, it has exactly the same parts, arranged in exactly the same way as Zenyte originally had. But Suzy thinks that the rock discovered thirty years ago with all of its parts completely replaced is actually Zenyte, since it makes John’s Eczema disappear and even though it has all new parts, this was just the result of years of experimentation.

Comprehension Checks:

(1) The rock assembled a month ago is made from all the original parts of "Zenyte".(Yes/No)

(2) The rock discovered thirty years ago has all new parts. (Yes/No)

(3) Suzy thinks that the rock discovered thirty years ago with all of its parts completely replaced is actually Zenyte. (Yes/No)

(4) Andy thinks that the rock assembled just a month ago is actually Zenyte. (Yes/No)

B.2. Study 2

B.2.1. Rowboat Cases

No Function: John is an accomplished woodworker and sailor, whose lifelong hobby is building rowboats by hand. He built his first rowboat—which he named "Drifter"—thirty years ago. Over the years there has been wear and tear, and it turns out that every single one of the original planks has needed to be replaced.

John—never one to throw anything out—has stored all of the original planks in his shed over the years. Last month John—realizing that he had accumulated enough old planks in his shed for a whole rowboat—took out his old plans for "Drifter" and assembled these planks exactly according to his old plans. John now has two rowboats of the same design: the rowboat originally built thirty years ago that has none of its original parts, and the rowboat just built one month ago with all of the original parts from "Drifter".

Two of John's friends—Suzy and Andy—are taking a painting class at the local university. Both of them decide that they would like to paint a picture of Drifter. They call John and he tells them that while he is an extremely private individual who never lets anyone mess with his stuff, he will make one exception for them. But he tells them that they can only paint Drifter and nothing else. He goes on to tell them though that he is heading out for vacation and will have no phone or internet access while he is away. But he tells them that they are free to stop by while he is away and paint Drifter.

So, one afternoon, Suzy and Andy head over to John's place so that they can begin their painting. But Suzy and Andy aren't sure which of the two rowboats is actually Drifter. Andy thinks that the rowboat just built a month ago is actually Drifter, since it has exactly the same parts, arranged in exactly the same way as Drifter originally had. But Suzy thinks that the rowboat built thirty years ago with all of its parts completely replaced is actually Drifter, since, even though it has all new parts, this was just the result of normal maintenance.

Replacement Preserves Function: John is an accomplished woodworker and sailor, whose lifelong hobby is building rowboats by hand. He built his first rowboat—which he named "Drifter"—thirty years ago. It was an excellent rowboat: it was very stable, never let water in and sailed smoothly. Over the years there has been wear and tear, and so John, to keep it in perfect working order, has replaced various parts. And, it turns out that, over the years, every single one of the original planks has needed to be replaced.

John—never one to throw anything out—stored all of the original planks in his shed over the years. Last month John—realizing that he had accumulated enough old planks in his shed for a whole rowboat—took out his old plans for "Drifter" and assembled these planks exactly according to his old plans. This rowboat, however, worked terribly: it was very rickety, always let water in and sank after just a few minutes in the water. So, John now has two rowboats of the same design: the rowboat originally built thirty years ago that has none of its original parts and works perfectly as a rowboat, and the rowboat just built one month ago with all of the original parts from "Drifter", which works terribly as a rowboat.

Two of John's friends—Suzy and Andy—are taking a painting class at the local university. Both of them decide that they would like to paint a picture of Drifter. They call John and he tells them that while he is an extremely private individual who never lets anyone mess with his stuff, he will make one exception for them. But he tells them that they can only paint Drifter and nothing else. He goes on to tell them though that he is heading out for vacation and will have no phone or internet access while he is away. But he tells them that they are free to stop by while he is away and paint Drifter.

So, one afternoon, Suzy and Andy head over to John's place so that they can begin their painting. But Suzy and Andy aren't sure which of the two rowboats is actually Drifter. Andy thinks that the rowboat just built a month ago is actually Drifter, since, even though it works terribly as a rowboat, it has exactly the same parts, arranged in exactly the same way as Drifter originally had. But Suzy disagrees. She thinks that the rowboat built thirty years ago with all of its parts completely replaced is actually Drifter, since even though it has all new parts, this was just the result of the normal maintenance required to keep the rowboat in perfect working order.

Original Parts Preserves Function: John is an accomplished woodworker and sailor, whose lifelong hobby is building rowboats by hand. He built his first rowboat—which he named "Drifter"—thirty years ago. It was an excellent rowboat: it was very stable, never let water in and sailed smoothly. But John was always thinking of ways to try and make the rowboat even better. And over the years, every single one of the original planks ended up being replaced. The end result, however, was not an improved rowboat. Rather the rowboat ended up being terrible: it was very rickety, always let water in and sank after just a few minutes in the water.

Fortunately, John—never one to throw anything out—stored all of the original planks in his shed over the years. Last month John—realizing that he had accumulated enough old planks in his shed for a whole rowboat—took out his old plans for "Drifter" and assembled these planks exactly according to his old plans. This rowboat worked perfectly: it was very stable, never let in water and sailed smoothly across the water. So, John now has two rowboats of the same design: the rowboat originally built thirty years ago that has none of its original parts and works terribly as a rowboat, and the rowboat just built one month ago with all of the original parts from "Drifter", which works perfectly as a rowboat.

Two of John's friends—Suzy and Andy—are taking a painting class at the local university. Both of them decide that they would like to paint a picture of Drifter. They call John and he tells them that while he is an extremely private individual who never lets anyone mess with his stuff, he will make one exception for them. But he tells them that they can only paint Drifter and nothing else. He goes on to tell them though that he is heading out for vacation and will have no phone or internet access while he is away. But he tells them that they are free to stop by while he is away and paint Drifter.

So, one afternoon, Suzy and Andy head over to John's place so that they can begin their painting. But Suzy and Andy aren't sure which of the two rowboats is actually Drifter. Andy thinks that the rowboat just built a month ago is actually Drifter, since it has exactly the same parts, arranged in exactly the same way as Drifter originally had and works perfectly. But Suzy thinks that the rowboat built thirty years ago with all of its parts completely replaced is actually Drifter, since, even though it has all new parts and is a terrible rowboat, this was just the result of John's attempt to improve it.

Comprehension Checks:

(1) The rowboat built a month ago is made from all the original parts of "Drifter". (Yes/No)

(2) The rowboat built thirty years ago has all new parts. (Yes/No)

(3) Suzy thinks that the rowboat built thirty years ago with all of its parts completely replaced is actually Drifter. (Yes/No)

(4) Andy thinks that the rowboat just built a month ago is actually Drifter. (Yes/No)

B.2.2. Organism Cases

No Function: John is an accomplished microbiologist who has devoted his life to studying microorganisms. When he first began his career thirty years ago, he discovered a new microorganism. He named the microorganism "Gollywag" and immediately logged the exact details of the microorganism. Over the years, John has conducted many experiments on the microorganism. Each time he conducts an experiment on the microorganism, he cuts off a part of it, which he uses for testing, and replaces it with the same type of part from another microorganism of the same type. Indeed, John has conducted so many experiments over the years that every single piece of the microorganism has been replaced.

After testing though, John always stores each part of the microorganism from the original microorganism. Last month John—realizing that he had accumulated enough parts of the microorganism for a whole microorganism—took out his logbook and assembled the parts exactly according to his notes. John now has two microorganisms of exactly the

same design: the microorganism discovered thirty years ago with all of its parts completely replaced and the microorganism assembled just a month ago with all of the original parts from “Gollywag”.

John has just hired two interns—Andy and Suzy—to work in his lab. He tells them that their first assignment will be to perform a series of experiments on “Gollywag”. But, Suzy and Andy aren't sure which of the two microorganisms is actually Gollywag. Andy thinks that the microorganism assembled just a month ago is actually Gollywag, since it has exactly the same parts, arranged in exactly the same way as Gollywag originally had. But Suzy thinks that the microorganism discovered thirty years ago with all of its parts completely replaced is actually Gollywag, since, even though it has all new parts, this was just the result of years of experimentation.

Replacement Preserves Function: John is an accomplished microbiologist who has devoted his life to studying microorganisms. When he first began his career thirty years ago, he discovered a new microorganism. He named the microorganism “Gollywag” and immediately logged the exact details of the microorganism.

When John began examining the microorganism, he noticed it emitted a unique sequence of chemicals. Curious as to why the microorganism emitted this particular sequence of chemicals, John showed the microorganism to one of his friends, Frank, who is a biochemist. Frank kept the microorganism for several days and after careful examination he reported back to John. Frank told John that the microorganism is very delicate and can only survive if its body temperature stays between 60 and 65 degrees. Frank goes on to tell John that the way it maintains its body temperature is by having other microorganisms come into contact with it. And the only way to get the other

microorganisms in contact with it is by emitting this unique sequence of chemicals. So, Frank tells John that the unique sequence of chemicals is actually for signaling to the others that they need to come into contact with it to help maintain its body temperature.

But John, being a microbiologist, isn't interested in the chemicals emitted by the microorganism. Rather he is interested in investigating and studying the various parts of the microorganism "Gollywag". So, he starts conducting experiments on the microorganism.

Over the years, John has conducted many experiments on the microorganism. Each time he conducts an experiment on the microorganism, he cuts off a part of it, which he uses for testing, and replaces it with the same type of part from another microorganism of the same type. Indeed, John has conducted so many experiments over the years that every single part of the microorganism has been replaced.

After testing though, John always stores each part from the original microorganism. Last month John—realizing that he had accumulated enough parts of the microorganism for a whole microorganism—took out his logbook and assembled the parts exactly according to his notes. When John handled this microorganism, however, he noticed his that it no longer emitted the unique sequence of chemicals. So, John now has two microorganisms of exactly the same design: the microorganism discovered thirty years ago with all of its parts completely replaced and which emits the unique sequence of chemicals and the microorganism assembled just a month ago with all of the original parts from "Gollywag", which no longer emits the unique sequence of chemicals.

John has just hired two interns—Andy and Suzy—to work in his lab. He tells them that their first assignment will be to perform a series of experiments on "Gollywag". But, Suzy and Andy aren't sure which of the two microorganisms is actually Gollywag. Andy thinks that the microorganism assembled just a month ago is actually Gollywag, since even though it does not emit the unique sequence of chemicals, it has exactly the same parts, arranged in exactly the same way as Gollywag originally had. But Suzy thinks that the microorganism discovered thirty years ago with all of its parts completely replaced is actually Gollywag, since it emits the unique sequence of chemicals, and even though it has all new parts, this was just the result of years of experimentation.

Original Parts Preserves Function: John is an accomplished microbiologist who has devoted his life to studying microorganisms. When he first began his career thirty years ago, he discovered a new microorganism. He named the microorganism "Gollywag" and immediately logged the exact details of the microorganism.

When John began examining the microorganism, he noticed it emitted a unique sequence of chemicals. Curious as to why the microorganism emitted this particular sequence of chemicals, John showed the microorganism to one of his friends, Frank, who is a biochemist. Frank kept the microorganism for several days and after careful examination he reported back to John. Frank told John that the microorganism is very delicate and can only survive if its body temperature stays between 60 and 65 degrees. Frank goes on to tell John that the way it maintains its body temperature is by having other microorganisms come into contact with it. And the only way to get the other microorganisms in contact with it is by emitting this unique sequence of chemicals. So,

Frank tells John that the unique sequence of chemicals is actually for signaling to the others that they need to come into contact with it to help maintain its body temperature.

But John, being a microbiologist, isn't interested in the chemicals emitted by the microorganism. Rather he is interested in investigating and studying the various parts of the microorganism "Gollywag". So, he starts conducting experiments on the microorganism.

Over the years, John has conducted many experiments on the microorganism. Each time he conducts an experiment on the microorganism, he cuts off a part of it, which he uses for testing, and replaces it with the same type of part from another microorganism of the same type. Indeed, John has conducted so many experiments over the years that every single part of the microorganism has been replaced. In the end, however, John noticed that when he examined the microorganism it no longer emitted the unique sequence of chemicals.

After testing though, John always stores each part from the original microorganism. Last month John—realizing that he had accumulated enough parts of the microorganism for a whole microorganism—took out his logbook and assembled the parts exactly according to his notes. When John examined this microrganism, he noticed that it emitted the unique sequence of chemicals. So, John now has two organisms of exactly the same design: the microorganism discovered thirty years ago with all of its parts completely replaced which no longer emits the unique sequence of chemicals and the microorganism assembled just a month ago with all of the original parts from "Gollywag", which emits the unique sequence of chemicals.

John has just hired two interns—Andy and Suzy—to work in his lab. He tells them that their first assignment will be to perform a series of experiments on “Gollywag”. But, Suzy and Andy aren't sure which of the two microorganisms is actually Gollywag. Andy thinks that the microorganism assembled just a month ago is actually Gollywag, since it has exactly the same parts, arranged in exactly the same way as Gollywag originally had and emits the unique sequence of chemicals. But Suzy thinks that the microorganism discovered thirty years ago with all of its parts completely replaced is actually Gollywag, since, even though it has all new parts and no longer emits the unique sequence of chemicals, this was just the result of years of experimentation.

Comprehension Checks:

(1) The microorganism assembled a month ago is made from all the original parts of "Gollywag".(Yes/No)

(2) The microorganism discovered thirty years ago has all new parts. (Yes/No)

(3) Suzy thinks that the microorganism discovered thirty years ago with all of its parts completely replaced is actually Gollywag. (Yes/No)

(4) Andy thinks that the microorganism assembled just a month ago is actually Gollywag. (Yes/No)

B.2.3. Rock Cases

No Function: John is an accomplished geochemist who has devoted his life to studying rocks. When he first began his career thirty years ago, he discovered a rock, made out of an unknown mineral. And interestingly, due to this mineral, the rock took on a

distinctive, hollowed-out seashell shape. He named the rock “Zenyte” and immediately logged the exact details of the rock.

Over the years, John has conducted many experiments on the rock. Each time he conducts an experiment on the rock, he breaks off a piece of it, which he uses for testing, and replaces it with the same type of mineral. Indeed, John has conducted so many experiments over the years that every single piece of the rock has been replaced.

After testing though, John always stores each piece of mineral from the original rock. Last month John—realizing that he had accumulated enough pieces of the mineral for a whole rock—took out his logbook and assembled the minerals exactly according to his notes. John now has two rocks of exactly the same design: the rock discovered thirty years ago with all of its minerals completely replaced and the rock assembled just a month ago with all of the original minerals from “Zenyte”.

John has just hired two interns—Andy and Suzy—to work in his lab. He tells them that their first assignment will be to perform a series of experiments on “Zenyte”. But, Suzy and Andy aren't sure which of the two rocks is actually Zenyte. Andy thinks that the rock assembled just a month ago is actually Zenyte, since it has exactly the same parts, arranged in exactly the same way as Zenyte originally had. But Suzy thinks that the rock discovered thirty years ago with all of its parts completely replaced is actually Zenyte, since, even though it has all new parts, this was just the result of years of experimentation.

Replacement Preserves Function: John is an accomplished geochemist who has devoted his life to studying rocks. When he first began his career thirty years ago, he discovered

a rock, made out of an unknown mineral. And interestingly, due to this mineral, the rock took on a distinctive, hollowed-out seashell shape. He named the rock "Zenyte" and immediately logged the exact details of the rock.

When John began examining the rock, he noticed that the rock housed an interesting species of worm. Since he had never seen this particular type of worm before, John showed the rock to one of his friends, Frank, who is a biologist. Frank kept the rock for several days and after careful examination he reported back to John. Frank told John that the worms living in the rock are actually very delicate. Indeed, Frank told John that the rock actually creates the perfect conditions for the worms to breed and flourish in and that there is no other environment in which the worms could survive. So, Frank tells John that the rock is actually for providing a perfectly hospitable environment for the worms.

But John, being a geochemist, isn't interested in the worms. Rather he is interested in investigating and studying the unknown mineral that makes up "Zenyte". So, he starts conducting experiments on the rock.

Over the years, John has conducted many experiments on the rock. Each time he conducts an experiment on the rock, he breaks off a piece of it, which he uses for testing, and replaces it with the same type of mineral. Indeed, John has conducted so many experiments over the years that every single piece of the rock has been replaced.

After testing though, John always stores each piece of mineral from the original rock. Last month John—realizing that he had accumulated enough pieces of the mineral for a whole rock—took out his logbook and assembled the minerals exactly according to his

notes. When John examined this rock, however, he noticed that all the worms died. So, John now has two rocks of exactly the same design: the rock discovered thirty years ago with all of its minerals completely replaced which creates a perfectly hospitable environment for the worms to reproduce and flourish and the rock assembled just a month ago with all of the original minerals from "Zenyte", which kills all the worms.

John has just hired two interns—Andy and Suzy—to work in his lab. He tells them that their first assignment will be to perform a series of experiments on "Zenyte". But, Suzy and Andy aren't sure which of the two rocks is actually Zenyte. Andy thinks that the rock assembled just a month ago is actually Zenyte, since even though it kills the worms, it has exactly the same parts, arranged in exactly the same way as Zenyte originally had. But Suzy thinks that the rock discovered thirty years ago with all of its parts completely replaced is actually Zenyte, since it creates a perfectly hospitable environment for the worms to reproduce and flourish and even though it has all new parts, this was just the result of years of experimentation.

Original Parts Preserves Function: John is an accomplished geochemist who has devoted his life to studying rocks. When he first began his career thirty years ago, he discovered a rock, made out of an unknown mineral. And interestingly, due to this mineral, the rock took on a distinctive, hollowed-out seashell shape. He named the rock "Zenyte" and immediately logged the exact details of the rock.

When John began examining the rock, he noticed that the rock housed an interesting species of worm. Since he had never seen this particular type of worm before, John showed the rock to one of his friends, Frank, who is a biologist. Frank kept the rock for several days and after careful examination he reported back to John. Frank told John that

the worms living in the rock are actually very delicate. Indeed, Frank told John that the rock actually creates the perfect conditions for the worms to breed and flourish in and that there is no other environment in which the worms could survive. So, Frank tells John that the rock is actually for providing a perfectly hospitable environment for the worms.

But John, being a geochemist, isn't interested in the worms. Rather he is interested in investigating and studying the unknown mineral that makes up "Zenyte". So, he starts conducting experiments on the rock.

Over the years, John has conducted many experiments on the rock. Each time he conducts an experiment on the rock, he breaks off a piece of it, which he uses for testing, and replaces it with the same type of mineral. Indeed, John has conducted so many experiments over the years that every single piece of the rock has been replaced. In the end, however, John noticed that the worms could no longer survive in the rock.

After testing though, John always stores each piece of mineral from the original rock. Last month John—realizing that he had accumulated enough pieces of the mineral for a whole rock—took out his logbook and assembled the minerals exactly according to his notes. When John examined this rock, he noticed that the worms reproduced and flourished. So, John now has two rocks of exactly the same design: the rock discovered thirty years ago with all of its minerals completely replaced which kills all the worms and the rock assembled just a month ago with all of the original minerals from "Zenyte", which creates a perfectly hospitable environment for the worms to reproduce and flourish.

John has just hired two interns—Andy and Suzy—to work in his lab. He tells them that their first assignment will be to perform a series of experiments on “Zenyte”. But, Suzy and Andy aren't sure which of the two rocks is actually Zenyte. Andy thinks that the rock assembled just a month ago is actually Zenyte, since it has exactly the same parts, arranged in exactly the same way as Zenyte originally had and creates a hospitable environment for the worms to reproduce and flourish. But Suzy thinks that the rock discovered thirty years ago with all of its parts completely replaced is actually Zenyte, since, even though it has all new parts and kills all the worms, this was just the result of years of experimentation.

Comprehension Checks

(1) The rock assembled a month ago is made from all the original parts of "Zenyte".(Yes/No)

(2) The rock discovered thirty years ago has all new parts. (Yes/No)

(3) Suzy thinks that the rock discovered thirty years ago with all of its parts completely replaced is actually Zenyte. (Yes/No)

(4) Andy thinks that the rock assembled just a month ago is actually Zenyte. (Yes/No)

B.3. Study 3

B.3.1. Loss of Function

John is out hiking and he spots something glowing in a strange way by the side of the trail. It turns out to be a rock, glowing because it serves as a home to a special sort of microorganism. John takes the thing home to study it further. But it begins to fade as the

microorganisms it hosts start dying. John realizes that the microorganisms are feeding on minerals in the rock but can't access the minerals in the interior of the rock. So he tries an experiment: he hits the rock with a hammer, breaking it into three pieces. But the experiment does not work: the microorganisms all begin quickly dying and it stops glowing, fading completely to black.

B.3.2. Preservation of Function

John is out hiking and he spots something glowing in a strange way by the side of the trail. It turns out to be a rock, glowing because it serves as a home to a special sort of microorganism. John takes the thing home to study it further. But it begins to fade as the microorganisms it hosts start dying. John realizes that the microorganisms are feeding on minerals in the rock but can't access the minerals in the interior of the rock. So he tries an experiment: he hits the rock with a hammer, breaking it into three pieces. The experiment works: the microorganisms are then able to access all the minerals and so it resumes glowing even brighter than before.

B.4. Study 4

B.4.1. Denting

John is out hiking and he spots something glowing in a strange way by the side of the trail. It turns out to be a rock, glowing because it serves as a home to a special sort of microorganism. John takes the thing home to study it further. But it begins to fade as the microorganisms it hosts start dying. John realizes that the microorganisms are feeding on minerals in the rock but can't access the minerals in the interior of the rock. So he tries an experiment: he hits the rock with a hammer. As a result the rock is dented. But the

experiment does not work: the microorganisms all begin quickly dying and it stops glowing, fading completely to black.

B.4.2. Denting Control

John is out hiking and he spots something by the side of the trail. It turns out to be a rock. He takes the thing home. Later that evening, John gets bored and he hits it with a hammer. As a result, the rock is dented.

B.4.3. Pulverizing

John is out hiking and he spots something glowing in a strange way by the side of the trail. It turns out to be a rock, glowing because it serves as a home to a special sort of microorganism. John takes the thing home to study it further. But it begins to fade as the microorganisms it hosts start dying. John realizes that the microorganisms are feeding on minerals in the rock but can't access the minerals in the interior of the rock. So he tries an experiment: he smashes the rock into pieces with a hammer. The experiment works: the microorganisms are then able to access all the minerals and so it resumes glowing even brighter than before.

B.4.4. Pulverizing Control

John is out hiking and he spots something by the side of the trail. It turns out to be a rock. He takes the thing home. Later that evening John gets bored and he smashes it into pieces with a hammer.

B.5. Study 5

B.5.1. Denting, Loss of Function

John is a geologist. Recently, he went on an expedition to Antarctica in search of new kinds of rocks. While in Antarctica, John discovered a very strange rock, unlike any rock he had ever seen, and noticed what appeared to be tiny worms in the crevices of the rock. He named the rock “Zenyte”.

John took Zenyte back to Arizona and gave it to a biologist, Frank. Frank kept Zenyte for several days and after careful examination he reported back to John. Frank told John that the worms living in Zenyte are actually very delicate. Indeed, Frank told John that Zenyte actually transmits an incredibly rare combination of chemicals and that this combination of chemicals creates the perfect conditions for the worms to breed and flourish. Frank tells John that there is no other environment in which the worms could survive and that Zenyte is actually for providing a perfectly hospitable environment for the worms.

Frank kept Zenyte so that he could continue studying the worms. But as time went on, Frank noticed that Zenyte was reacting negatively to the lab environment and so was getting denser. As Zenyte became denser the combination of chemicals transmitted by it were slowly decreasing and the worms were beginning to die.

In an attempt to preserve Zentye and prevent the worms from dying, Frank decided that he would hit the rock with a hammer. As a result, the rock is dented and, unfortunately, now the combination of chemicals is not transmitted at all and the worms are all quickly dying.

Later that day, John wants to examine Zenyte and so goes to Frank’s lab. John asks Frank if he can examine Zenyte and Frank takes him over to it. When John looks at it, he

says “this is not Zenyte, you’ve completely destroyed it and it no longer performs the functions that are essential to it: it no longer transmits the combination of chemicals and all the worms it once housed are dead”. But Frank disagrees, saying that though it looks different it is still Zenyte.

B.5.2. Denting, Preservation of Function

John is a geologist. Recently, he went on an expedition to Antarctica in search of new kinds of rocks. While in Antarctica, John discovered a very strange rock, unlike any rock he had ever seen, and noticed what appeared to be tiny worms in the crevices of the rock. He named the rock “Zenyte”.

John took Zenyte back to Arizona and gave it to a biologist, Frank. Frank kept Zenyte for several days and after careful examination he reported back to John. Frank told John that the worms living in Zenyte are actually very delicate. Indeed, Frank told John that Zenyte actually transmits an incredibly rare combination of chemicals and that this combination of chemicals creates the perfect conditions for the worms to breed and flourish. Frank tells John that there is no other environment in which the worms could survive and that Zenyte is actually for providing a perfectly hospitable environment for the worms.

Frank kept Zenyte so that he could continue studying the worms. But as time went on, Frank noticed that Zenyte was reacting negatively to the lab environment and so was getting denser. As Zenyte became denser the combination of chemicals transmitted by it were slowly decreasing and the worms were beginning to die.

In an attempt to preserve Zentye and prevent the worms from dying, Frank decided that he would hit the rock with a hammer. As a result, the rock is dented and now, the combination of chemicals is being perfectly transmitted—just like they were before Zenyte became dense—and the worms are continuing to flourish.

Later that day, John wants to examine Zenyte and so goes to Frank's lab. John asks Frank if he can examine Zenyte and Frank takes him over to it. When John looks at it, he says "this is not Zenyte, you've completely destroyed it". But Frank disagrees, saying that it is still Zenyte. He tells John that though it looks different, it performs all of the functions that are essential to it: it still transmits the rare combination of chemicals and still creates a perfectly hospitable environment for the worms.

B.5.3. Pulverizing, Loss of Function

John is a geologist. Recently, he went on an expedition to Antarctica in search of new kinds of rocks. While in Antarctica, John discovered a very strange rock, unlike any rock he had ever seen, and noticed what appeared to be tiny worms in the crevices of the rock. He named the rock "Zenyte".

John took Zenyte back to Arizona and gave it to a biologist, Frank. Frank kept Zenyte for several days and after careful examination he reported back to John. Frank told John that the worms living in Zenyte are actually very delicate. Indeed, Frank told John that Zenyte actually transmits an incredibly rare combination of chemicals and that this combination of chemicals creates the perfect conditions for the worms to breed and flourish. Frank tells John that there is no other environment in which the worms could

survive and that Zenyte is actually for providing a perfectly hospitable environment for the worms.

Frank kept Zenyte so that he could continue studying the worms. But as time went on, Frank noticed that Zenyte was reacting negatively to the lab environment and so was getting denser. As Zenyte became denser the combination of chemicals transmitted by it were slowly decreasing and the worms were beginning to die.

In an attempt to preserve Zentye and prevent the worms from dying, Frank decided that he would carefully break up the rock. He breaks Zenyte into more fine grained pieces until the pieces are so small—much like dust particles—that he cannot break them down any further. Unfortunately, now the combination of chemicals is not transmitted at all and the worms are all quickly dying.

Later that day, John wants to examine Zenyte and so goes to Frank's lab. John asks Frank if he can examine Zenyte and Frank takes him over to it. When John looks at it, he says "this is not Zenyte, you've completely destroyed it and it no longer performs the functions that are essential to it: it no longer transmits the combination of chemicals and all the worms it once housed are dead". But Frank disagrees, saying that though it looks different it is still Zenyte.

B.5.4. Pulverizing, Preservation of Function

John is a geologist. Recently, he went on an expedition to Antarctica in search of new kinds of rocks. While in Antarctica, John discovered a very strange rock, unlike any rock he had ever seen, and noticed what appeared to be tiny worms in the crevices of the rock. He named the rock "Zenyte".

John took Zenyte back to Arizona and gave it to a biologist, Frank. Frank kept Zenyte for several days and after careful examination he reported back to John. Frank told John that the worms living in Zenyte are actually very delicate. Indeed, Frank told John that Zenyte actually transmits an incredibly rare combination of chemicals and that this combination of chemicals creates the perfect conditions for the worms to breed and flourish. Frank tells John that there is no other environment in which the worms could survive and that Zenyte is actually for providing a perfectly hospitable environment for the worms.

Frank kept Zenyte so that he could continue studying the worms. But as time went on, Frank noticed that Zenyte was reacting negatively to the lab environment and so was getting denser. As Zenyte became denser the combination of chemicals transmitted by it were slowly decreasing and the worms were beginning to die.

In an attempt to preserve Zentye and prevent the worms from dying, Frank decided that he would carefully break up the rock. As he breaks Zenyte into more fine grained pieces he notices that the chemical transmission is slowly being restored to normal levels. So, he continues until the pieces are so small—much like dust particles—that he cannot break them down any further. Now, the combination of chemicals is being perfectly transmitted—just like they were before Zenyte became dense—and the worms are continuing to flourish.

Later that day, John wants to examine Zenyte and so goes to Frank's lab. John asks Frank if he can examine Zenyte and Frank takes him over to it. When John looks at it, he says "this is not Zenyte, you've completely destroyed it ". But Frank disagrees, saying that it is still Zenyte. He tells John that though it looks different, it performs all of the

functions that are essential to it: it still transmits the rare combination of chemicals and still creates a perfectly hospitable environment for the worms.

B.6. Study 6 *(Footnote 13: Ownership and Function)*

B.6.1. Same Owner, Preserves Function

One day while mowing, John discovered a rock in his backyard which was glowing in a strange way. After studying the rock he realized that it was glowing because it serves as a home to some special microorganisms which feed on the minerals in the rock.

After keeping the rock in his house for a year, John decides that he does not want the rock in his house anymore. He considers giving it to his neighbor Frank but instead decides that he'd rather keep it for himself. Since he wants to keep it but doesn't want it in his house, he decides to just throw it in his backyard so he can enjoy it when he is mowing.

The rock now glows even brighter than before and the microorganisms flourish as they continue to feed on the minerals in the rock

B.6.2. Same Owner, Loses Function

One day while mowing, John discovered a rock in his backyard which was glowing in a strange way. After studying the rock he realized that it was glowing because it serves as a home to some special microorganisms which feed on the minerals in the rock.

After keeping the rock in his house for a year, John decides that he does not want the rock in his house anymore. He considers giving it to his neighbor Frank but instead decides that he'd rather keep it for himself. Since he wants to keep it but doesn't want it

in his house, he decides to just throw it in his backyard so he can enjoy it when he is mowing.

The rock, however, stops glowing and the microorganisms it hosts begin dying.

B.6.3. Different Owner, Preserves Function

One day while mowing, John discovered a rock in his backyard which was glowing in a strange way. After studying the rock he realized that it was glowing because it serves as a home to some special microorganisms which feed on the minerals in the rock.

After keeping the rock in his house for a year, John decides that he does not want the rock anymore. He decides to give it to his neighbor Frank. To surprise him, John decides that he'll throw it in Frank's front yard so that he'll see it when he gets the morning paper. So, he throws the rock in Frank's front yard. The rock now glows even brighter than before and the microorganisms flourish as they continue to feed on the minerals in the rock.

The next morning Frank finds the rock when he gets the paper. He picks it up and decides to take it inside his house to display over his fireplace.

B.6.4. Different Owner, Loses Function

One day while mowing, John discovered a rock in his backyard which was glowing in a strange way. After studying the rock he realized that it was glowing because it serves as a home to some special microorganisms which feed on the minerals in the rock.

After keeping the rock in his house for a year, John decides that he does not want the rock anymore. He decides to give it to his neighbor Frank. To surprise him, John

decides that he'll throw it in Frank's front yard so that he'll see it when he gets the morning paper. So, he throws the rock in Frank's front yard. The rock, however, stops glowing and the microorganisms it hosts begin dying.

The next morning Frank finds the rock when he gets the paper. He picks it up and decides to take it inside his house to display over his fireplace.

B.6.5. Probe

Is the rock that [Frank/John] now has really the same rock that John originally found in his backyard? 1=No it is different, 7= Yes, it is the same

B.6.6. Results

There was a large-sized effect of whether the rock lost (M=4.71, SD=2.16) or preserved (M=6.49, SD=.971) its function on persistence judgments $F(1, 137)=40.99$, $p<.001$, $np2=.230$. There was no effect of whether John (M=5.63, SD=1.89) or Frank (M=5.71, SD=1.82) owned the rock, $F(1, 137)=.002$, $p=.967$ and no interaction between ownership and function $F(1, 137)=.054$, $p=.816$

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. Plenum Press.

Arico, A., Fiala, B., Goldberg, R., and Nichols, S. (2011). The Folk Psychology of Consciousness. *Mind & Language*, 26, 327-352.

Aristotle. (1984a). Physics. In J. Barnes (ed.), The complete works of Aristotle, vol.1, 315–446. New Jersey: Princeton University Press.

Aristotle. (1984b). Metaphysics. In J. Barnes (ed.), The complete works of Aristotle, vol. 2, 1552–1728. New Jersey: Princeton University Press.

Atran, S. (1998). Folk Biology and the Anthropology of Science: Cognitive Universals and Cultural Particulars. *Behavioral and Brain Sciences*, 21, 547-609.

Bacon, F. (1996). *Collected Works of Francis Bacon, vol. 4: Translations of the Philosophical Works*, eds. James Spedding, Robert Leslie Ellis, and Douglas Denon Heath. Routledge/Thoemmes Press.

Bealer, G. (1998). Intuition and the autonomy of philosophy. In M. DePaul & W. Ramsey (Eds.), *Rethinking Intuition*. Lanham, MD: Rowman & Littlefield.

Beebee, H. (2004). Causing and Nothingness. In L. A. Paul, E. J. Hall & J. Collins (eds.), *Causation and Counterfactuals*. MIT Press. 291-308.

Bengson, J. (2013). Experimental Attacks on Intuitions and Answers. *Philosophy and Phenomenological Research*, 86, 495–532.

Bennett, K. (2009). Composition, colocation, and metaontology. In D. Chalmers, D. Manley, & R. Wasserman (eds), Metametaphysics: New essays on the foundations of ontology. 38–76. New York: Oxford University Press.

Benovsky, J. (2006). *Persistence through Time, and Across Possible Worlds*. Ontos Verlag.

Bloom, P. (1996). Intention, History, and Artifact Concepts. *Cognition,* 60, 1-29.

Bloom, P. (2007). Religion is Natural. *Developmental Science,* 10, 147-151

Bloom, P., & Veres, C. (1999). The Perceived Intentionality of Groups. *Cognition,* 71, B1-B9.

Bowers, J. (ms). A teleological answer to the special composition question.

Boyer, P. (2001). *Religion Explained*. New York: Basic Books.

Burke, M. (1994). Preserving the Principle of One Object to a Place: A Novel Account of the Relations among Objects, Sorts, Sortals and Persistence Conditions. *Philosophy and Phenomenological Research,* 54, 591–624.

Burke, M (1997). Coinciding Objects: Reply to Lowe and Denkel. *Analysis,* 57, 11–18.

Camerer, C., Loewenstein, G. & Weber, M. (1989). The Curse of Knowledge in Economic Settings: An Experimental Analysis. *Journal of Political Economy,* 97, 1232-1254.

Cappelen, H. (2012). *Philosophy Without Intuitions*. Oxford: Oxford University Press.

Casler, K., & Kelemen, D. (2008). Developmental Continuity in Teleo-Functional Explanation: Reasoning about Nature among Romanian Romani Adults. *Journal of Cognition and Development*, 9, 340-362.

Chaiken, S. & Trope, Y. (1999). *Dual-Process Theories in Social Psychology*. New York: Guilford Press.

Channon, S., Lagnado, D. A., Drury, H., Matheson, E., Fitzpatrick, S., Shieff, C., Mendoza, N. & Maudgil, D. (2010). Causal reasoning and intentionality judgments after frontal brain lesions. *Social Cognition*, 28, 509-522.

Charles, D. (2001). Matter and form: Unity, persistence, and identity. In T. Scaltsas, D. Charles,& M. L. Gill (eds), Unity, identity, and explanation in Aristotle's metaphysics, 75–106. New York: Oxford University Press.

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, 125-30.

Cohen, S. (1984). Justification and Truth. *Philosophical Studies*, 46, 279–295.

Cummins, R. (1975). Functional Analysis. *Journal of Philosophy*, 72, 74-765.

Danks, D., Rose, D., & Machery, E. (2013). Demoralizing Causation. *Philosophical Studies*, 1-27.

Dawkins, R. (1995). *River out of Eden*. New York: Basic Books.

De Freitas, J., Tobia, K., Newman, G., Knobe, J. (manuscript). The Good Ship Theseus: The Effect of Valence on Object Identity Judgments.

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.

Deutsch, M. (2015). *The Myth of the Intuitive*. Cambridge, MA: MIT Press.

Devitt, M. (2105). Relying on intuitions: Where Cappelen and Deutsch go wrong. *Inquiry*, 58: 7-8, 669-699.

Dink, J. & Rips, L. (forthcoming). Folk Teleology and its Implications. In (ed) D. Rose, *Experimental Metaphysics*. Bloomsbury.

Dowe, P. (2000). Physical Causation. Cambridge University Press.

Dowe, P. (2004). Causes are Physically Connected to their Effects: Why Preventers and Omissions are not Causes. In Hitchcock, C. (ed.) *Contemporary Debates in Philosophy of Science.* Blackwell, 189-196.

Driver, J. (2008). Attributions of Causation and Moral Responsibility. In Walter Sinnott-Armstrong (ed.)., Moral Psychology Volume 2. MIT press, 423-439.

Ellis, P. (2010). *The Essential Guide to Effect Sizes: Statistical Power, Meta-Analysis and the Interpretation of Research Results*. Cambridge University Press.

Epstein, S. (1994). Integration of the Cognitive and Psychodynamic Unconscious. *American Psychologist*, 49, 709–24.

Evans, E. M., Spiegel, A.N., Gram, W., Frazier, B.N., Tare, M., Thompson, S., & Diamond, J. (2010). A Conceptual Guide to Natural History Museum Visitors Understanding of Evolution. *Journal of Research in Science Teaching*, 47, 326-353.

Frederick, S. (2005). Cognitive Reflection and Decision Making. *Journal of Economic Perspectives*, 19, 25-42.

Galli, L., & Meinardi, E. (2011). The role of teleological thinking in learning the Darwinian model of evolution. *Evolution Education Outreach*, 4, 145–152.

German, T. & Johnson, S. (2002). Functions and the Origin of the Design Stance. *Journal of Cognition and Development*, 3, 279-300

Goldberg, R. & Thompson-Schill, S. (2009). Developmental "Roots" in Mature Biological Knowledge. *Psychological Science*, 20, 480-487.

Goldman, A. (1979). What Is Justified Belief? In G. Pappas (ed.), Justification and Knowledge, Dordrecht: Reidel.

Goldman, A. (1987). Cognitive science and metaphysics. *Journal of Philosophy*, 84, 537-544.

Goldman, A. (1993). Philosophical applications of cognitive science. Westview Press.

Goldman, A. (2007). A Program for "Naturalizing" Metaphysics, with Application to the Ontology of Events. *The Monist*, 90, 457-479.

Goldman, A. (2007). Philosophical Intuitions: Their Target, Their Source, and Their Epistemic Status, *Grazer Philosophische Studien*, 74, 1-26.

Goldman, A. (2010). Philosophical Naturalism and Intuitional Methodology (Romanell Lecture), *Proceedings and Addresses of the American Philosophical Association*, 115-150.

Goldman, A. (2015). Naturalizing metaphysics with the help of cognitive science. In eds., K Bennett and D. Zimmerman, *Oxford Studies in Metaphysics*, Volume 8, 171-214. Oxford: Oxford University Press.

Goldman, A. and Pust, J. (1998). Philosophical Theory and Intuitional Evidence. In M Depaul and W. Ramsey (Eds.), *Rethinking Intuitions: The Psychology of Intuition and its Role in Philosophical Inquiry*. Rowman and Littlefield, 179-197.

Gonzales Galli, L., & Meinardi, E. (2011). The Role of Teleological Thinking in Learning the Darwinian Model of Evolution. *Evolution Education Outreach*, 4, 145-152.

Gregory, T. R. (2009). Understanding Natural Selection: Essential Concepts and Common Misconceptions. *Evolution: Education and Outreach*, 2, 156-175.

Guthrie, S.E. (1993). *Faces in the Clouds: A New Theory of Religion*. New York: Oxford University Press.

Hall, D. G. (1995). Artifacts and origins. University of British Columbia.

Halpern, J., & Hitchcock, C. (forthcoming). Graded Causation and Defaults. *British Journal for the Philosophy of Science.*

Hauser, M. Young, L., & Cushman, F. (2008). Reviving Rawls linguistic analogy: Operative principles and the causal structure of moral actions. In (ed.) Sinnott-Armstrong, W., *Moral Psychology Volume 2: The Cognitive Science of Morality: Intuition and Diversity*, pp. 107-143.

Hawley, K. (2001). *How Things Persist*. Oxford University Press.

Hawthorne, J. (2008). Three Dimensionalism vs Four Dimensionalism. In Sider, T., Hawthorne, J. & Zimmerman, D. (Eds.), *Contemporary Debates in Metaphysics*: Blackwell.

Hawthorne, J., & Nolan, D. (2006). What would teleological causation be? Metaphysical essays: 265–284. New York: Oxford University Press.

Heider, F., & Simmel, M. (1944). An Experimental Study of Apparent Behavior. *American Journal of Psychology*, 57, 243–259.

Hinds, P. (1999). The Curse of Expertise: The Effects of Expertise and Debiasing Methods on Predictions of Novice Performance. *Journal of Experimental Psychology: Applied*, 5, 202-221.

Hirsch, E. (1982). *The Concept of Identity*. Oxford: Oxford University Press

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.

Hoffman, J., & Rosenkrantz, G. (1997). *Substance: Its nature and existence*. New York: Routledge.

Jackson, F. (1998). *From Metaphysics to Ethics: A Defense of Conceptual Analysis*. Oxford University Press.

Jenkins, C. S., & Nolan, D. (2008). Backwards explanation. *Philosophical Studies*, 140, 103–115.

Jensen, M. S. & Finley, F. N. (1995). Teaching Evolution Using Historical Arguments in a Conceptual Change Strategy. *Science Education*, 79, 147-166.

Kahane, G. (2011). Evolutionary Debunking Arguments. *Nous*, 45, 103-125.

Kahneman, D. & Frederick, S. (2002). Representativeness Revisited: Attribute Substitution in Intuitive Judgment. In (eds.) T. Gilovich, D. Griffin and D. Kahneman , *Heuristics and Biases: The Psychology of Intuitive Judgment.* New York: Cambridge University Press, pp. 49–81.

Kaiser, M., McCloskey, M., & Proffitt, D. (1986). Development of intuitive theories of motion: Curvilinear motion in the absence of external forces. *Developmental Psychology*, 22, 67–71.

Kampourakis, K. & Zogza, V. (2008). Students' Intuitive Explanations of the Causes of Homologies and Adaptations. *Science and Education*, 17, 27–47.

Keil, F. (1995). The Growth of Causal Understandings of Natural Kinds. Causal Cognition: A Multidisciplinary Debate. (Eds.). Sperber, D., Premack, D & Premack, A. J. pp. 234-262. Oxford University Press

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. (2008). Unrestricted composition and restricted quantification. *Philosophical Studies*, 140, 319–334.

Korman, D. (2009). Eliminativism and the Challenge from Folk Belief. *Noûs*, 43, 242-264.

Korman, D. & Carmichael, C. (forthcoming). What do the Folk Think About Composition and Does it Matter? In (ed). D. Rose, *Experimental Metaphysics*. Bloomsbury.

Koslicki, K. (2008). *The structure of objects.* New York: Oxford University Press.

Kriegel, U. (2011). Two defenses of common-sense ontology. *Dialectica*, 65, 177–204.

Kriegel, U. (2013). The epistemological challenge of revisionary metaphysics. *Philosopher's Imprint*, 13, 1–30.

Kriegel, U. (forthcoming). Metaphysics and Conceptual Analysis: Experimental Philosophy's Place Under the Sun. In (ed). D. Rose, *Experimental Metaphysics*. Bloomsbury.

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.

Leslie, S .J. (2013). Essence and natural kinds: When science meets preschooler intuition. *Oxford Studies in Epistemology*, 4, 108-165.

Lewis, D. (1986). *Philosophical Papers, Volume II*. Oxford: Oxford University Press.

Lewis, D. (1986). *On the Plurality of Worlds*. Oxford: Basil Blackwell

Lewis, D. (1991). *Parts of Classes*. Oxford: Basil Blackwell.

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., & Carey, S. (2006). Functional explanation and the function of explanation. *Cognition*, 99, 167–204.

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.

Lowe, E. J. (1995). Coinciding Objects: In Defense of the 'Standard Account'. *Analysis*, 55, 171–178.

Ludwig, K. (2007). The Epistemology of Thought Experiments: First Person Versus Third Person Approaches. *Midwest Studies in Philosophy*, 31, 128-159.

Machery, E. (2012). Power and negative results. *Philosophy of Science*, 79, 808–820.

Machery, E., Mallon, R., Nichols, S., & Stich, S. (2004). Semantics, cross-cultural style. *Cognition*, 92 (3), 1-12.

Mackie, J. L. (1977). *Ethics: Inventing Right and Wrong*. New York: Penguin.

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 and 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.

Markosian, N. (1998). Brutal composition. *Philosophical Studies*, 92, 211–249.

Markosian, N. (2014). A spatial approach to mereology. In S. Kleinschmidt (ed.), *Mereology and location*, 69–90. New York: Oxford University Press.

Mason, K. (2010). Debunking arguments and the genealogy of religion and morality. *Philosophy Compass*, 5, 770–8.

Matan, A., & Carey, S. (2001). Developmental changes within the core of artifact concepts. *Cognition*, 77, 1–26.

Mayr, E. (1998). The Multiple Meanings of Teleological. *History and Philosophy of the Life Sciences,* 20, 35-40.

McDermott, M. (1995). Redundant Causation. *British Journal for the Philosophy of Science*, 46, 523-544.

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. PhD 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 Helen Beebee, Peter Menzies & Christopher Hitchcock (eds.), The Oxford Handbook of Causation. Oxford University Press. 341-367.

Merricks, T. (1994). Endurance and Indiscernibility. *Journal of Philosophy*, 91, 165-184.

Merricks, T. (2001). *Objects and persons.* New York: Oxford University Press.

Michotte. A. (1963). The Perception of Causality [Translated by T. R. Miles and E. Miles]. London: Methuen. (Originally published in 1946).

Millikan, R. G. (1989). In Defense of Proper Functions. *Philosophy of Science*, 56, 288-302.

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.

Murray, D., Sytsma, J. & Livengood, J. (2013). God Knows (But Does God Believe). *Philosophical Studies*, 166, 83-107.

Nado, J. (2016). The intuition deniers. *Philosophical Studies*, 173(3), 781-800.

Nahmias, E., Morris, S., Nadelhoffer, T., & Turner, J. (2006). Is incompatibilism intuitive? *Philosophy and Phenomenological Research,* 73 (1), 28 - 53.

Nichols, S. (2014). Process Debunking and Ethics. *Ethics*, 124, 727-749.

Nichols, S., & Knobe, J. (2007). Moral responsibility and determinism: The cognitive science of folk intuitions. *Noûs*, 41(4), 663–685.

Ochssler, J., Roider, A., & Schmitz, P. (2009). Cognitive Abilities and Behavioral Biases. *Journal of Economic Behavior and Organization*, 72, 147-152.

Paul, L. A. (2002). Logical Parts. *Nous,* 36, 578-596.

Paul, L. (2010). Temporal experience. *Journal of Philosophy*, 107 (7), 333–359.

Paul, L. A. (2010). A New Role for Experimental Work in Metaphysics. *Review of Philosophy and Psychology*, 1, 461-476.

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 University Press.

Parsons, J. (2000). Must a Four Dimensionalist Believe in Temporal Parts? The Monist, 83, 399-418.

Piaget, J. (1960). The Child's Conception of the World (J. Tomlinson & A. Tomlinson, Trans.). Paterson, NJ: Littlefield, Adams & Co. (Original work published 1929)

Pollock, J. (1987). Defeasible Reasoning. *Cognitive Science*, 11, 481-518.

Pust, J. (2016). Intuition, *The Stanford Encyclopedia of Philosophy* (Spring 2016 Edition) Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/spr2016/entries/intuition/>.

Rea, M. (1998). In defense of mereological universalism. *Philosophy and Phenomenological Research*, 58, 347–360.

Richards, D.D., & Siegler, R.S. (1986). Children's Understanding of the Attributes of Life. *Journal of Experimental Child Psychology*, 42, 1–22.

Ripley, D. (2016). Experimental philosophical logic. In (eds.) Justin Sytsma and Wesley Buckwalter, *A Companion to Experimental Philosophy*. Blackwell.

Rips, L. (1989). Similarity, typicality and categorization. In S. Vosniadou & A. Ortony (eds), *Similarity and analogical reasoning*, 21–59. New York: Cambridge University Press.

Rose, D. (2015). Persistence through function preservation. *Synthese*, 192, 97–146.

Rose, D. (forthcoming). Folk Intuitions of Actual Causation: A Two-Pronged Debunking Explanation. *Philosophical Studies*.

Rose, D. (forthcoming). Introduction. In Rose, D. (ed.), *Experimental Metaphysics*. Bloomsbury.

Rose, D. (forthcoming). Cognitive Science for the Revisionary Metaphysician. In Goldman, A., and McLaughling, B. (eds.), *Cognitive Science and Metaphysics*. Oxford University Press.

Rose, D. & Danks, D. (2013). In Defense of a Broad Conception of Experimental Philosophy. *Metaphilosophy*, 44, 512-534.

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.

Rose, D. & Nichols, S. (2016). From Punishment to Moral Objectivity. Unpublished Manuscript, Rutgers University.

Rose, D., & Schaffer, J. (forthcoming). Folk Mereology is Teleological. *Nous.*

Rose, D., & Schaffer, J. (forthcoming). Folk Mereology is Teleological (Expanded Version). In Rose, D. (ed.), *Experimental Metaphysics*. Bloomsbury.

Rosen, G., & Dorr, C. (2002). Composition as a fiction. In R. Gale (ed.), The Blackwell companion to metaphysics, 151–174. Oxford: Basil Blackwell.

Rosset, E. (2008). It's No Accident: Our Bias for Intentional Explanations. *Cognition*, 108, 771-780.

Rottman, J., Zhu, L., Wang, W., Schillaci, R. S., Clark, K. J., & Kelemen, D. (2016). Cultural influences on the teleological stance: Evidence from China. *Religion, Brain & Behavior*.

Sanford, D. (1993). The problem of the many, many composition questions, and naïve mereology. *Noûs*, 27, 219–228.

Sarkissian, H., Park, J., Tien, D., Wright, J., & Knobe, J. (2011). Folk moral relativism. *Mind and Language*, 26(4), 482-505

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. Blackwell. 197-216.

Schaffer, J. (2005). Contrastive Causation. *Philosophical Review*, 114, 327-358.

Schaffer, J. (2009). On what grounds what. In D. Chalmers, D. Manley, & R. Wasserman (eds), *Metametaphysics: New essays on the foundations of ontology*, 347–383. New York: Oxford University Press.

Schaffer, J. (forthcoming). Cognitive science and metaphysics: Partners in debunking. In H. Kornblith & B. McLaughlin (eds), *Goldman and his critics*. Oxford: Wiley-Blackwell.

Sidelle, A. (2002). Is There a True Metaphysics of Material Objects? *Philosophical Issues*, 12, 118-145.

Sider, T. (2001). *Four Dimensionalism.* Oxford University Press.

Sider, T. (2013). Against parthood. In K. Bennett & D. Zimmerman (eds), *Oxford studies in metaphysics* 8, 237–293. New York: Oxford University Press.

Simons, P. (1987). Parts: A study in ontology. New York: Oxford University Press.

Sloman, S. (1996). The Empirical Case for Two Systems of Reasoning. *Psychological Bulletin*, 119, 3–22.

Sosa, E. (1988). Minimal intuition. In (eds.) M. DePaul & W. Ramsey, Rethinking Intuition, Lanham, MD: Rowman & Littlefield.

Sosa, E. (2007). Experimental philosophy and philosophical intuition. Philosophical Studies, 132: 99-107.

Stanovich, K. & West, R. (2000). Individual Differences in Reasoning: Implications for the Rationality Debate? *Behavioral and Brain Sciences*, 22, 645–726.

Stich, S. (2013). Do different groups have different epistemic intuitions? A Reply to Nagel. *Philosophy and Phenomenological Research*, 87, 151-178.

Stich, S. (forthcoming). Knowledge, Intuition and Culture. In (eds.) Joelle Proust & Martin Fortier, *Metacognitive Diversity An Interdisciplinary Approach*. Oxford: Oxford University Press.

Stich, S. & Tobia, K. (forthcoming) Experimental Philosophy's Challenge to the 'Great Tradition'. *Analytica: Revista de Filosofia*.

Strominger, N. & Nichols, S. (2014). The Essential Moral Self. *Cognition*, 131, 159-171.

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.

Sytsma, J., & Machery, E. (2010). Two conceptions of subjective experience. *Philosophical Studies*, 151 (2), 299-327.

Thomson, J. (1983). Parthood and Identity Across Time. *Journal of Philosophy*, 80, 201-220.

Thomasson, A. (2007). *Ordinary objects*. New York Oxford University Press.

Turri, J. (forthcoming). Exceptionalist naturalism: human agency and the causal order. *Quarterly Journal of Experimental Psychology*.

Turri, J., & Blouw, P. (forthcoming). Excuse Validation: A Study in Rule-Breaking. *Philosophical studies.*

Van Inwagen, P. (1990). *Material Beings.* Ithaca: Cornell.

Weinberg, J. (2014). Cappelen between a rock and a hard place. *Philosophical Studies*, 171, 545-553.

Wiggins, D. (1968). On Being in the Same Place at the Same Time. *Philosophical Review*, 77, 90-95

Wiggins, D. (1980). *Sameness and Substance*. Cambridge: Harvard University Press.

Williamson, T. (2007). The Philosophy of Philosophy. Oxford: Blackwell.

Williamson, T. (2011). Philosophical Expertise and the Burden of Proof. *Metaphilosophy*, 42, 215-229