

Folk Teleology Drives Persistence Judgments

David Rose, Jonathan Schaffer, and Kevin Tobia
draft of August 2nd, 2017

Persistence judgments are ordinary judgments about whether an object survives a change, or perishes. For instance, if a house fire only superficially damages the kitchen, people judge that the house survived. But if the fire burnt the house to the ground instead, people judge that the house did not survive but was instead destroyed. We are interested in what drives these judgments, in part because objects are so central to our conception of the world, and our persistence judgments get to the very heart of the folk notion of an object.

We aim to connect two research programs. The first of these programs stems from Knobe (2003; see Knobe 2010 for an overview), and has found normative effects on a wide range of intuitions. Recent work by Tobia (2015) and De Freitas *et al.* (2017) shows a normative effect on persistence judgments, at least for a range of artifacts and social objects including universities, rock bands, and research papers. The second of these programs stems from Kelemen (1999), and has found that the folk conception of the world is heavily laden with teleology. Recent work by Rose (2015) and Rose & Schaffer (2017) shows teleological effects on persistence as well as mereological judgments, for rocks, rowboats, and even people shaking hands.

Since both normative and teleological effects have been documented for persistence judgments, the question naturally arises as to how these effects are related. Are they independent? Does either mediate the other? Along these lines, De Freitas *et al.* (2017: 397) speculate—but do not test—that the normative effect they observe might be driven by teleological considerations. We take up this question.

We find that, in accord with the Kelemen-style “promiscuous teleology” view and the De Freitas *et al.* speculation, it is folk teleology that drives persistence judgments, across cases including those that De Freitas *et al.* consider. Across our experiments, any effect of normativity is *screened off* by teleology. So we aim to connect the normativity and teleology research programs. And we claim to put teleology in the driver’s seat, while at the same time shedding further light on our folk notion of an object.

1. Review: Normative and Teleological Effects on Persistence Judgments

1.1 Normative effects

We begin with a summary of the normative effects on persistence judgments, documented by De Freitas *et al.* (2017; see also Tobia 2015), and fitting into the wider Knobe-inspired research program of documenting normative effects across a wide range of judgments. Various normative effects have been documented for judgments about intentional action (Knobe 2003), choice (Petit & Knobe 2009), and knowledge (Beebe & Buckwalter 2010), as well as intuitions about causation (Alicke, Rose & Bloom 2011; Hitchcock & Knobe 2009; Rose 2017), *inter alia*. This seems like a surprising but deep feature of human cognition.

Against this backdrop, De Freitas *et al.* (2017: 384) ran studies to “explore whether valence—that is, whether valuing certain traits as good versus bad—similarly influences persistence judgments.” We focus on their first line of studies (2017: §5), which were a group of five studies demonstrating “the basic effect.” What De Freitas *et al.* find is that—in a range of cases concerning a science paper, a university, a conference, a nation, and a band, normative improvements led to significantly greater judgments of persistence than paired cases with normative deteriorations. For instance, they presented a case where a university in Nazi-era Germany was divided between academic and propaganda functions, and compared changes in which the university focused on academic functions (improvement), with changes in which the university focused on propaganda functions (deterioration). Participants were significantly more likely to say that the university

persisted through improvement than through deterioration. De Freitas *et al.* (2017: 388) depict their basic effect as follows (higher scores indicate stronger judgments of persistence):




Figure 1: De Freitas et al.’s Effect of Deterioration/Improvement on Persistence

Our own studies (§2) corroborate this normative effect on persistence judgments.

1.2 Teleological effects

We turn to a second effect on persistence judgments, documented by Rose (2015; see also Rose & Schaffer 2017), and fitting into the wider Kelemen-inspired research program of documenting teleological effects across a wide range of judgments. The human inclination to teleological thinking begins in childhood, with children maintaining that lions are for “going to the zoo,” that clouds are “for raining” (Bloom 2007: 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: 1444–45). This inclination persists into adulthood, with even college-educated adults endorsing statements such as “the sun radiates heat because warmth nurtures life,” “fungi grow in forests to help decomposition,” and “lightening occurs to release electricity” (Kelemen & Rosset 2009). This tendency to endorse teleological explanations is amplified when background causal beliefs are prevented from intruding, such as when college-educated adults are put in speeded tasks (Kelemen *et al.* 2013). This effect similarly appears among professional physicists and professionals in the humanities (Kelemen *et al.* 2013). Moreover, people with Alzheimer’s disease—who display deficits in causal beliefs—naturally default to teleological thinking (Lombrozo *et al.* 2007). Here is a second surprising but deep feature of human cognition.

Specific evidence of teleological effects on object cognition surfaces in Rips (1989). He reports that participants, considering an object that looks like a lampshade, will judge that it is really an umbrella when told that it was originally designed to protect people from rain. Rose & Schaffer (2017: 247–8) find that people tend to say that a collection of parts forms a whole when those parts serve a collective purpose: “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 backdrop, Rose (2015: 101) began with the following hypothesis:

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.

He tested this hypothesis across several studies, of which we shall focus on his rock cases. Rose’s rock cases involved a rock whose function was to provide minerals for certain micro-organisms. He allowed two sorts of physical changes for the rock: suffering a minor dent, and suffering complete pulverization. And he allowed that such changes might preserve or destroy the function of providing minerals for the micro-organisms, leading to a 2x2 grid of cases. One might have thought that rocks can survive being dented but not being pulverized. But Rose found a strong effect of preservation/destruction of function on persistence judgments, and did not find any effect of denting/pulverizing (nor did he find an interaction). He (2015: 112) depicts these results as follows (higher scores indicate stronger judgments of persistence):




Figure 2: Rose’s Effect of Denting/Pulverizing and Loss/Preservation of Function on Persistence

Our own studies (§2) corroborate this teleological effect on persistence judgments.

1.3 Our question: what is the relationship between normative and teleological effects on persistence judgments?

So far we have presented two effects on persistence judgments: normative effects (§1.1) and teleological effects (§1.2). The question naturally arises as to whether these effects are related, and if so how. Perhaps these are simply independent effects, but a first clue that they are related comes from De Freitas *et al.*'s (2017: 390–1) follow up studies showing that essence judgments significantly mediate normative effects.¹

¹ A wide range of studies show that people tend to weigh “superficial” features such as color less than “deep” and potentially unobservable features when assessing object identity (e.g., Blok *et al.* 2001; Hall *et al.* 2003;

Participants were asked both to make persistence judgments about an object following either a normative improvement or deterioration, and also to judge whether the object “after the changes no longer reflects the true essence of the original.” It was found that essence judgments (“Essence”) significantly mediated the relationship between normativity (“Valence”) and persistence (“Identity”):




Figure 3: De Freitas et al.'s Mediation Model

De Freitas *et al.* (2017: 396–7) offer, as a “plausible explanation” of these results, the idea that “our very understanding of an entity’s identity is that it consists of those traits that we value as good.” They then speculate—without testing—that their normative/essence effect might be teleologically driven, so as to be “most likely to arise in cases where there is believed to be some purpose or teleology of a particular entity.” They take this to suggest that normative effects on persistence judgments arise “only for entities that are seen as having a deeper purpose in this relevant sense.”

We take up this speculation as an invitation for empirical research, and aim to demonstrate a connection between the normative and teleological effects on persistence judgments. For all that has been shown, these effects might be causally independent, or either might mediate the other, or any more complicated causal interaction might be involved. So we ask, *how if at all are these effects connected?*

Our results indicate that teleology is a direct cause of persistence judgments but normativity is not, and that any effect of normativity is screened off by teleology. (We also think—in agreement with Kelemen but perhaps in disagreement with De Freitas *et al.*—that all objects are seen as having a purpose in the relevant sense. So we expect these effects to arise across the board.) This answer is compatible with a number of background views on the general role of normative and teleological considerations in cognition. But one view—which we take as a working hypothesis—is the following:

- People tend to have normatively laden conceptions of the purposes of objects,² and thereby tend to see normative improvements in an object as more “on track” with its purpose, while tending to see normative deterioration as more “off track.”

Blok *et al.* 2005; Newman *et al.* 2014). This is understood in terms of the deeper features being viewed as “essential.”

² The idea that people tend to have normatively laden conceptions of the purposes of objects fits the idea that folk teleology is tied into a *folk theism*, on which all objects are viewed as part of the divine plan. There is also the idea that folk teleology may stem from a natural *Gaia hypothesis* of a living earth or cosmos. We are not sure if that fits normatively laden purposes as neatly. See Kelemen *et al.* 2013 for discussion of these options. This idea also coheres with specific work on the “good true self” from Newman *et al.* 2015 (see also Strohminger *et al.* 2017). Also—as De Freitas (*personal correspondence*) points out—many of Kelemen’s examples are positively valenced, such as “The sun radiates heat because warmth nurtures life,” and “Earthworms tunnel underground to aerate the soil.”

- Persistence judgments are directly driven by the extent to which a given change keeps the object “on track” with its purpose.

1.4 *Who cares?*

We hope that our discussion will be of interest to both psychologists and philosophers, and to anyone interested in *folk metaphysics*. For psychologists, our results bear directly on the question of what drives folk judgments of object persistence, and they thereby bear on the central folk concept of *object* (or perhaps domain-specific concepts, like *social object* and *artifact*—we are neutral on whether there is one general concept in play, or many specific concepts). Moreover, our results bear on the underlying relations between normative and teleological effects on cognition generally, and thereby connect the broader Knobe and Kelemen research programs.

For philosophers, our results bear on the program of “descriptive metaphysics” (Strawson 1959), and on the program of “commonsense metaphysics,” which aims to show—in the words of Thomasson (2007: 3)—“how, reflectively, we can make sense of our unreflective common sense worldview.” Philosophers in both programs make many—and sometimes conflicting—armchair claims about folk persistence judgments (see Rose 2015: 98–100 for examples), without any empirical tests. We see this as an opportunity to provide empirical input into this project.

There is also a deeper methodological question for philosophers: What if any role should “intuitions” and “folk judgments” play in metaphysics (or elsewhere)? Our results are compatible with virtually any view, and we ourselves are divided on the matter. Two of us (following Rose & Schaffer 2017: 261–64)—hold that, to the extent that our folk judgments are infused with teleological considerations, they are thereby infused with benighted superstition, and should be treated as *debunked*. On this view, our results liberate the question of when objects really persist from any demanded conformity to folk intuitions. On this matter, the folk deserve to be ignored.

2. Studies: Connecting Normative and Teleological Effects

In order to test the relationship between normative and teleological effects on persistence judgments, we began with a sample of the basic cases from De Freitas *et al.* (2017), namely their university, band, and science paper cases. This ensured continuity with their results (which we were also able to corroborate). It also reduced researcher degrees of freedom, since their cases were designed to test for normative but not for teleological effects. We also added in a fourth case with a gardening tool. This allowed us to extend the scope of the findings, and overall to consider normative and teleological effects on persistence judgments concerning two social objects (university and band), alongside two artifactual objects (science paper and gardening tool).

For each of these four cases, we ran two types of studies. The first is based on the De Freitas *et al.* template, and looks at correlations between a wide range of variables. These studies allow us to replicate and extend the De Freitas *et al.* results, and consider a wide range of interactions. But ultimately we are interested in causal hypotheses and causal models, and so our second type of study uses a pared-down set of variables, in order to directly construct tractable causal models that pit *Norm* versus *Telos* vis-à-vis persistence judgments. (We think that both sorts of studies have independent interest, though from the perspective of our conclusion we regard the second sort of study with the causal models as most central.)

Our first type of study works by varying three factors: *Start*, *Norm*, and *Telos*. The *Start* variable manipulates whether the object is characterized as good or bad at the outset. The *Norm* variable manipulates whether the change is an improvement or a deterioration. And the *Telos* variable manipulates whether the change preserves or destroys the object’s original purpose. For instance, in the university case, we manipulate

whether the Nazi-era university begins with a focus on academic subjects or on propaganda (*Start*), whether the university then shifts further in an academic or propagandistic direction (*Norm*), and whether the university retains a focus on education (*Telos*).

We then ask participants to make a persistence judgment as to whether or not the original object still exists after the changes. We do this by introducing two characters disagreeing on the issue, with one character claiming that the original object no longer exists, and the other claiming that the original object still exists. Participants are asked to rate which character they agree with more. Participants are also asked whether the changed object still fits the purpose of the original object, how they would morally evaluate the changed object, and whether the changed object still reflects the true essence of the original. We thus aim to measure judgments of persistence, teleology, normativity, and essence, so as to find the correlations.

Our second type of study uses a pared-down set of variables, in order to directly construct tractable causal models. For these we fixed a setting for *Start* (we ran two of these fixed at *Start*=good, and two fixed at *Start*=bad) and we omitted the essence probe. This enabled us to directly pit *Norm* versus *Telos* vis-à-vis persistence judgments, and get to the key issue of what causes what.

2.1 Study 1: university

2.1.1 Version 1A: persistence, teleology, normativity, and essence measures

400 participants were recruited from Amazon Mechanical Turk ($M_{age}=37$, 38% female). Participants were randomly assigned to one of eight conditions in a $2(Start: \text{good, bad}) \times 2(Norm: \text{better, worse}) \times 2(Telos: \text{preserved, destroyed})$ design. We used the following cases (variations indented):

During the Nazi regime, some educational institutions taught a mixture of courses on traditional academic subjects (science, literature, etc.) and courses in Nazi ideology (often with strong anti-Semitic messages). But the Iserlohn Institute was different.

[*Start*=good] Even though it taught a mix of these two kinds of courses, everyone who enrolled could tell that the real essence of the institution was its focus on academic subjects like science and literature. The material they taught on Nazi ideology was just a thin veneer over this more essential part of the curriculum.

[*Start*=bad] Even though it taught a mix of these two kinds of courses, everyone who enrolled could tell that the real essence of the institution was its focus on Nazi ideology and anti-Semitism. The material they taught on traditional academic subjects was just a thin veneer over this more essential part of the curriculum.

Then, after a number of years, there was a sudden administrative change. The rector of the institute was replaced by a new rector who decided to shake things up in certain ways.

[*Telos*=preserved, *Norm*=better] Specifically, the new rector decided to completely eliminate all courses on Nazi ideology and anti-Semitism. Instead, from that day onwards the institute always taught courses in just traditional academic subjects.

[*Telos*=preserved, *Norm*=worse] Specifically, the new rector decided to completely eliminate all courses on traditional academic subjects (science, literature, etc.). Instead, from that day onwards the institute always taught courses in just Nazi ideology and anti-Semitism.

[*Telos*=destroyed, *Norm*=better] Specifically, the new rector decided to completely eliminate all courses and any efforts at education at all, and transform the institute into a Jewish aid organization that would do what it could to hinder the Nazi genocide. So from that day onwards the institute stopped teaching and focused on helping Jews.

[*Telos*=destroyed, *Norm*=worse] Specifically, the new rector decided to completely eliminate all courses and any efforts at education at all, and transform the institute into a weapons research center that would research chemical weapons to help the Nazi war effort. So from that day onwards the institute stopped teaching and focused on weapons research.

After reading one of these eight cases, participants were asked the following (presented in a fixed order):

Identity: Imagine that Alex and Thomas are discussing these changes. Alex says that, because of these changes, the old institute no longer exists. He thinks it has been replaced by something new and different. But Thomas disagrees and says that the old institute still exists despite the changes. He thinks that the old institute has survived in a modified form.

Who do you agree with more, Alex or Thomas? (1 = I agree with Alex, 4 = I equally agree with both persons, 7 = I agree with Thomas).

Purpose: To what extent would you say that the new rector's decision for Iserlohn's focus fits the institutions true purpose or departs from the institutions true purpose (1=it definitely departs from its true purpose 7=it definitely fits with its true purpose)

Normativity: How would you morally evaluate Iserlohn after the new rector's decision about the direction to take the institute in? (1=very bad, 7=very good)

Essence: Iserlohn Institute after the changes no longer reflected the true essence of the original Iserlohn Institute (1=completely disagree, 7=completely agree)

Finally, participants were given two comprehension questions:

Comprehension1: The original Iserlohn Institute was an educational institution mainly focused on: (1) traditional academic subjects (2) Nazi ideology and anti-Semitism

Comprehension2: Alex thinks that because of the changes that the original Iserlohn Institute underwent the old institute: (1) no longer exists (2) still exists, though in a modified form

107 participants were removed for failing one or more of the comprehension questions. Data were then analyzed from the remaining 293 participants. We present the results.

Identity: There was a main effect of *Start* (good: $M=3.14$, $SE=.151$; bad: $M=3.85$, $SE=.150$), $F(1, 285)=11.133$, $p<.01$, *Norm* (better: $M=3.83$, $SE=.151$; worse: $M=3.15$, $SE=.150$), $F(1, 285)=10.126$, $p<.01$, and *Telos* (preserved: $M=4.24$, $SE=.155$; destroyed: $M=2.75$, $SE=.146$), $F(1, 285)=48.549$, $p<.001$. While *Start* and *Norm* had small-sized effects on identity judgments ($\eta_p^2=.038$ and $.034$ respectively), *Telos* produced a large sized effect ($\eta_p^2=.146$).³ These main effects were qualified by a small three-way interaction, $F(1, 285)=5.117$, $p<.05$, $\eta_p^2=.018$.

Purpose: There was a main effect of *Start* (good: $M=3.21$, $SE=.139$; bad: $M=3.84$, $SE=.138$), $F(1, 285)=10.191$, $p<.01$, *Norm* (better: $M=3.76$, $SE=.139$; worse: $M=3.30$, $SE=.138$), $F(1, 285)=5.566$, $p<.05$, and *Telos* (preserved: $M=4.30$, $SE=.143$; destroyed: $M=2.76$, $SE=.135$), $F(1, 293)=61.462$, $p<.001$. While *Start* and *Outcome* had small-sized effects on purpose judgments ($\eta_p^2=.035$ and $.019$ respectively), *Telos* produced a

³ For η_p^2 we follow Ellis (2010) in interpreting 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.

large-sized effect ($\eta_p^2=.177$). These main effects were qualified by a small three-way interaction, $F(1, 285)=7.912, p<.01, \eta_p^2=.027$.

Normativity: There was no main effect of *Start* (good: $M=3.80, SE=.124$; bad: $M=3.86, SE=.123$), $F(1, 285)=.123, p>.05$, or *Telos* (preserved: $M=3.82, SE=.127$; destroyed: $M=3.84, SE=.120$), $F(1, 285)=.009, p>.05$, on normative judgments. Only *Norm* (better: $M=5.58, SE=.124$; worse: $M=2.09, SE=.123$), $F(1, 285)=399.442, p<.001, \eta_p^2=.584$ had an effect, though this main effect was qualified by a small three-way interaction, $F(1, 293)=5.472, p<.05, \eta_p^2=.019$.

Essence: There was a main effect of *Start* (good: $M=4.90, SE=.140$; bad: $M=4.22, SE=.139$), $F(1, 285)=12.036, p<.01$, *Norm* (better: $M=4.30, SE=.140$; worse: $M=4.30, SE=.139$), $F(1, 285)=7.167, p<.01$, and *Telos* (preserved: $M=3.83, SE=.143$; destroyed: $M=5.29, SE=.135$), $F(1, 285)=55.274, p<.001$. While *Start* and *Norm* had small-sized effects on essence judgments ($\eta_p^2=.041$ and $.025$ respectively), *Telos* produced a large-sized effect ($\eta_p^2=.162$). These main effects were qualified by a small three-way interaction, $F(1, 285)=11.997, p<.01, \eta_p^2=.040$.

These results may be visualized as follows (higher scores indicate stronger judgments of persistence):




Figure 4: Effect of Norm on Each DV




Figure 5: Effect of Telos on Each DV

Overall, *Start*, *Norm*, and *Telos* each affected identity judgments, but while *Start* and *Norm* produced a small effect, *Telos* produced a large-sized effect. We found a similar pattern for Purpose and Essence. For Normativity, only *Norm* produced a main effect. The crucial question, to which we now turn, is whether Purpose or Normativity causes Identity.

2.1.2 Version 1B: causal model pitting teleology versus normativity

250 people were recruited from Amazon Mechanical Turk ($M_{age}=39$, 41% female). Each participant read a *Start*=good version of Study 1A which varied *Norm* and *Telos*, resulting in a $2(Norm: \text{better, worse}) \times 2(Telos: \text{preserved, destroyed})$ design. Participants were then asked each question (presented in random order) from Study 1A except for the essence question, omitted so we could pit teleology directly against normativity. 47 people were removed for failing one or more comprehension questions. Data were then analyzed from the remaining 203 participants. We present the results.

A multiple regression model with Purpose, Normativity, *Norm*, *Telos* and a *Norm* by *Telos* interaction revealed that a full model was able to account for 60% of the variance in identity judgments, $F(5, 197)=59.143, p<.001, R^2=.600$. However, the only variables that had significant effects in the full model were *Telos* ($\beta=-.221, t=-2.774, p<.01$) and Purpose ($\beta=.611, t=7.914, p<.001$).

We then conducted a causal search on the data, using Greedy Equivalence Search (GES),⁴ which returned the following model:

⁴ 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

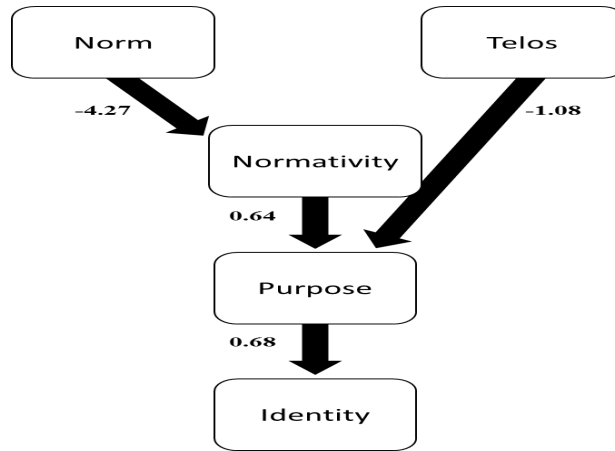


Figure 6: Causal Model of Identity Judgments

This model fits the data well $\chi^2(6)=12.097$, $p>.05$, $BIC=-19.782$. This model positions Normativity as an indirect cause of Identity, but only via Purpose. Importantly, this model recognizes only Purpose as a direct cause of Identity. It is here that we see a first clear sign that teleology and not normativity is what directly causes identity judgments.

2.2 Study 2: band

2.2.1 Version 2A: persistence, teleology, normativity, and essence measures

Four hundred participants were recruited from Amazon Mechanical Turk ($M_{age}=34$, 33% female). Participants were randomly assigned to one of eight conditions in a $2(Start: good, bad) \times 2(Norm: better, worse) \times 2(Telos: preserved, destroyed)$ design. We used the following cases (variations indented):

There are many bands that make songs that are intentionally deeply moving and meaningful and also songs that are intentionally superficial and commercial. But the band Breath String is different.

[*Start=good*] Even though it makes a mix of these kinds of songs, everyone can tell that the real essence of the band is its focus on making deeply moving and meaningful songs. The superficial and commercial songs are just a thin veneer over the more essential parts of the band.

[*Start=bad*] Even though it makes a mix of these kinds of songs, everyone can tell that the real essence of the band is its focus on making superficial and commercial songs. The deeply moving and meaningful songs are just a thin veneer over the more essential parts of the band.

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). Chickering (2002) shows 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 & Nichols 2013; Rose 2017; and Turri *et al.* 2016.

Over the years, some of the original band members left, and some new members joined, resulting in some differences.

[*Telos*=preserved, *Norm*=better] Specifically, the new members decided to completely stop making any superficial and commercial songs. Instead, the band now makes only deeply moving and meaningful songs.

[*Telos*=preserved, *Norm*=worse] Specifically, the new members decided to completely stop making any deeply moving and meaningful songs. Instead, the band now makes only superficial and commercial songs.

[*Telos*=destroyed, *Norm*=better] Specifically, the new members decided to completely stop making songs altogether. Instead, the band now just covers children's songs at free charity concerts for the community, and focuses only on raising money for charity.

[*Telos*=destroyed, *Norm*=worse] Specifically, the new members decided to completely stop making songs altogether. Instead, the band now just covers classic rock songs at private parties for drug-dealers, and focuses only on raising money for drugs.

After reading one of the cases participants were given the same questions (suitably adjusted) as those used in Study 1A. 153 participants were removed due to failing one or more of the comprehension questions. Data were then analyzed from the remaining 247 participants. We present the results.

Identity: There was no main effect of *Start* (good: $M=3.12$, $SE=.143$; bad: $M=3.17$, $SE=.150$), $F(1, 239)=.068$, $p=.794$, but there was a main effect of *Norm* (better: $M=3.43$, $SE=.146$; worse: $M=2.87$, $SE=.146$), $F(1, 239)=7.297$, $p<.01$, and *Telos* (preserved: $M=4.00$, $SE=.147$; destroyed: $M=2.29$, $SE=.146$), $F(1, 239)=68.345$, $p<.001$. *Norm* produced a small-sized effect on identity judgments ($\eta_p^2=.030$) and *Telos* produced a large sized effect ($\eta_p^2=.222$). These main effects were qualified by a small three-way interaction, $F(1, 239)=13.974$, $p<.001$, $\eta_p^2=.055$.

Purpose: There was no main effect of *Start* (good: $M=3.29$, $SE=.132$; bad: $M=3.35$, $SE=.139$), $F(1, 239)=.089$, $p=.765$, but there was a main effect of *Norm* (better: $M=3.61$, $SE=.135$; worse: $M=3.03$, $SE=.135$), $F(1, 239)=9.313$, $p<.01$, and *Telos* (preserved: $M=4.09$, $SE=.136$; destroyed: $M=2.55$, $SE=.134$), $F(1, 239)=65.221$, $p<.001$. While *Norm* had a small-sized effect on purpose judgments ($\eta_p^2=.038$), *Telos* produced a large-sized effect ($\eta_p^2=.214$). These main effects were qualified by a medium-sized three-way interaction, $F(1, 239)=23.723$, $p<.001$, $\eta_p^2=.090$.

Normativity: There was no main effect of *Start* (good: $M=4.19$, $SE=.119$; bad: $M=3.91$, $SE=.125$), $F(1, 239)=2.682$, $p=1.03$, but there was a main effect of *Telos* (preserved: $M=4.46$, $SE=.122$; destroyed: $M=3.63$, $SE=.121$), $F(1, 239)=22.998$, $p<.001$, $\eta_p^2=.088$, and of *Norm* (better: $M=5.14$, $SE=.122$; worse: $M=2.96$, $SE=.122$), $F(1, 239)=160.078$, $p<.001$, $\eta_p^2=.401$ on normative judgments. These effects were qualified by a small two-way interaction between *Telos* and *Norm*, $F(1, 239)=14.884$, $p<.001$, $\eta_p^2=.059$.

Essence: There was no main effect of *Start* (good: $M=4.89$, $SE=.129$; bad: $M=4.84$, $SE=.135$), $F(1, 239)=.086$, $p=.769$, but there was a main effect of *Norm* (better: $M=4.50$, $SE=.132$; worse: $M=5.23$, $SE=.132$), $F(1, 239)=15.323$, $p<.001$, and of *Telos* (preserved: $M=4.23$, $SE=.133$; destroyed: $M=5.50$, $SE=.132$), $F(1, 239)=46.500$, $p<.001$. While *Norm* had a medium-sized effect on essence judgments (η_p^2

=.060), *Telos* produced a large-sized effect ($\eta_p^2=.163$). These main effects were qualified by a medium-sized three-way interaction, $F(1, 239)=16.382, p<.001, \eta_p^2=.064$.

These results may be visualized as follows (higher scores indicate stronger judgments of persistence):




Figure 7: Effect of Norm on Each DV




Figure 8: Effect of Telos on Each DV

In contrast to the university cases in Study 1A (§2.1.1), we found that *Start* didn't produce a main effect on Identity. *Norm* and *Telos* however produced a main effect on Identity though the effect of *Telos* was much larger. We found a similar pattern for Purpose and Essence. For Normativity, only *Telos* and *Norm* produced main effects but *Norm* produced a much larger effect. The crucial question again is whether Purpose or Normativity causes identity judgments.

2.2.2 Version 2B: causal model pitting teleology versus normativity

250 people were recruited from Amazon Mechanical Turk ($M_{age}=35$, 32% female). Each participant read a *Start*=bad version of Study 2A which varied *Norm* and *Telos*, resulting in a $2(Norm: \text{better, worse}) \times 2(Telos: \text{preserved, destroyed})$ design. Participants were then asked each question (presented in random order) from Study 1A except for the essence question (again omitted so we could pit teleology directly against normativity). 59 people were removed for failing one or more comprehension question. Data was then analyzed from the remaining 191 participants.

A multiple regression model with Purpose, Normativity, *Norm*, *Telos* and a *Telos* by *Norm* interaction revealed that a full model was able to account for 37.5% of the variance in identity judgments, $F(5, 185)=22.241, p<.001, R^2=.375$. However, the only variables that had significant effects in the full model were Normativity ($\beta=.222, t=2.595, p<.01$) and Purpose ($\beta=.491, t=7.445, p<.001$).

We then conducted a causal search on the data, using Greedy Equivalence Search (GES), which returned the following model:




Figure 9: Causal Model of Identity Judgments

This model fits the data well $\chi^2(4)=8.577, p>.05, BIC=-12.432$. This model does not position Normativity as even an indirect cause of Identity (actually it positions Identity as a direct cause of Normativity).⁵ But the important point is that, again, only Purpose comes out as a direct cause of Identity. This provides a second clear sign that teleology and not normativity is what directly causes identity judgments.

2.3 Study 3: science paper

2.3.1 Version 3A: persistence, teleology, normativity, and essence measures

400 participants were drawn from Amazon Mechanical Turk ($M_{age}=35, 39\%$ female). Participants were randomly assigned to one of eight conditions in a $2(Start: good, bad) \times 2(Norm: better, worse) \times 2(Telos: preserved, destroyed)$ design. We used the following cases (variations indented):

These days, many science papers contain both well-supported points that follow naturally from the analyses, as well as points that make very big claims that aren't well supported. But a physicist's draft paper entitled Atom Dynamics is different.

[Start=good] Even though it contains a mix of these kinds of points, everyone who reads it can tell that the real essence of the paper is its focus on advancing well-supported points that follow naturally from the analyses. The material it included that made very big claims that are not well-supported was just a thin veneer over the more essential parts of the paper.

[Start=bad] Even though it contains a mix of these kinds of points, everyone who reads it can tell that the real essence of the paper is its focus on making very big claims that aren't well supported. The material it included that made well-supported points that follow naturally from the analyses was just a thin veneer over the more essential parts of the paper.

The physicist's collaborator then spends many more hours editing the paper, removing some parts and adding some new parts.

⁵ We constructed a directed acyclic graph with all edges the same as in the model returned by GES except we reversed the edge from Normativity to Identity. This model was rejected as a poor fit of the data ($p<.05$).

[*Telos*=preserved, *Norm*=better] Specifically, the collaborator decided to completely eliminate all the material that made very big claims that were not well-supported. Instead, the paper included just well-supported points that follow naturally from the analyses.

[*Telos*=preserved, *Norm*=worse] Specifically, the collaborator decided to completely eliminate all the material that made well-supported points that follow naturally from the analyses. Instead, the paper included just very big claims that were not well-supported.

[*Telos*=destroyed, *Norm*=better] Specifically, the collaborator decided to completely eliminate all the scientific material and transform the paper into a literary and artistic work celebrating diversity. So from that day onwards the paper stopped focusing on science and started focusing on celebrating diversity.

[*Telos*=preserved, *Norm*=worse] Specifically, the collaborator decided to completely eliminate all the scientific material and transform the paper into a literary and artistic work celebrating racism. So from that day onwards the paper stopped focusing on science and started focusing on celebrating racism.

After reading one of the cases participants were given the same questions (suitably adjusted) as those used in Study 1A. 163 participants were removed due to failing one or more of the comprehension questions. Data were analyzed from the remaining 237 participants.

Identity: There was no main effect of *Start* (good: $M=3.19$, $SE=.181$; bad: $M=3.10$, $SE=.185$), $F(1, 229)=.124$, $p=.794$, or *Norm* (better: $M=3.29$, $SE=.182$; worse: $M=3.01$, $SE=.184$), $F(1, 229)=1.155$, $p=.284$, but there was a main effect of *Telos* (preserved: $M=3.94$, $SE=.195$; destroyed: $M=2.35$, $SE=.171$), $F(1, 229)=37.960$, $p<.001$. *Telos* produced a large-sized effect ($\eta_p^2=.142$). The main effect of *Telos* was qualified by a small three-way interaction, $F(1, 239)=7.494$, $p<.01$, $\eta_p^2=.032$.

Purpose: There was no main effect of *Start* (good: $M=2.93$, $SE=.194$; bad: $M=3.08$, $SE=.158$), $F(1, 229)=.424$, $p=.516$, but there was a main effect of *Norm* (better: $M=3.42$, $SE=.156$; worse: $M=2.59$, $SE=.157$), $F(1, 229)=14.163$, $p<.01$, and *Telos* (preserved: $M=3.89$, $SE=.166$; destroyed: $M=2.12$, $SE=.146$), $F(1, 229)=63.734$, $p<.001$. While *Norm* had a small-sized effect on purpose judgments ($\eta_p^2=.058$), *Telos* produced a large-sized effect ($\eta_p^2=.218$). These main effects were qualified by a three-way interaction, $F(1, 239)=28.829$, $p<.001$, $\eta_p^2=.112$.

Normativity: There was a main effect of *Start* (good: $M=3.57$, $SE=.137$; bad: $M=3.18$, $SE=.140$), $F(1, 229)=3.998$, $p<.05$, $\eta_p^2=.017$, a main effect of *Telos* (preserved: $M=4.02$, $SE=.147$; destroyed: $M=2.74$, $SE=.129$), $F(1, 229)=42.544$, $p<.001$, $\eta_p^2=.157$ and *Norm* (better: $M=4.27$, $SE=.138$; worse: $M=2.49$, $SE=.139$), $F(1, 229)=82.158$, $p<.001$, $\eta_p^2=.264$ on normative judgments. These effects were qualified by a small three-way interaction between *Telos* and *Norm*, $F(1, 229)=7.981$, $p<.01$, $\eta_p^2=.034$.

Essence: There was no main effect of *Start* (good: $M=4.92$, $SE=.139$; bad: $M=5.18$, $SE=.143$), $F(1, 229)=1.172$, $p=.191$, but there was a main effect of *Norm* (better: $M=4.67$, $SE=.140$; worse: $M=5.43$, $SE=.142$), $F(1, 229)=14.611$, $p<.001$, and *Telos* (preserved: $M=4.12$, $SE=.150$; destroyed: $M=5.99$, $SE=.132$), $F(1, 229)=88.391$, $p<.001$. While *Norm* had a medium-sized effect on essence judgments ($\eta_p^2=.060$), *Telos* produced a large sized effect ($\eta_p^2=.278$). These main effects were qualified by a three-way interaction, $F(1, 229)=22.365$, $p<.001$, $\eta_p^2=.089$.

These results may be visualized as follows (higher scores indicate stronger judgments of persistence):




Figure 10: Effect of Norm on Each DV




Figure 11: Effect of Telos on Each DV

In this case—our first with an artifact rather than a social object—we now find that only *Telos* produced a main effect on identity judgments. *Norm* and *Telos* produced main effects on Purpose and Essence, though the effect of *Telos* was larger. Each factor affected Normativity in this case, though here the effect of *Norm* was largest. Again, our main question is whether Purpose or Normativity directly causes identity judgments.

2.3.2 Version 3B: causal model pitting teleology versus normativity

250 people were recruited from Amazon Mechanical Turk ($M_{age}=31$, 38% female). Each participant read the *Start=good* version of Study 3A which varied *Norm* and *Telos*, resulting in a 2(*Norm*: better, worse) x 2(*Telos*: preserved, destroyed) design. Participants were then asked each question (presented in random order) from Study 1A except for the essence question (again omitted to pit teleology directly against normativity). 82 people were removed for failing one or more comprehension question. Data was analyzed from the remaining 168 participants.

A multiple regression model with Purpose, Normativity, *Norm*, *Telos* and a *Telos* by *Norm* interaction revealed that a full model was able to account for 61% of the variance in identity judgments, $F(5, 163)=50.529, p<.001, R^2=.608$. Tellingly, the only variable that had a significant effect in the full model was Purpose ($\beta=.696, t=6.681, p<.001$).

We then conducted a causal search on the data, using Greedy Equivalence Search (GES), which returned the following model:




Figure 12: Causal Model of Identity Judgments

This model fits the data well $\chi^2(5)=1.645, p>.05, BIC=-24.004$. This model (like the model in §2.1.2, but unlike that in §2.2.2) positions Normativity as an indirect cause of Identity, but only via Purpose. Importantly, as in all our models, only Purpose is a direct cause of Identity. Here is the third sign that Purpose and not Normativity is what drives identity judgments.

2.4 Study 4: gardening tool

2.4.1 Version 4A: persistence, teleology, normativity, and essence measures

400 participants were recruited from Amazon Mechanical Turk ($M_{age}=30, 38\%$ female). Participants were randomly assigned to one of eight conditions in a $2(Start: \text{good, bad}) \times 2(Norm: \text{better, worse}) \times 2(Telos: \text{preserved, destroyed})$ design. Here are the cases:

Many gardening tools contain both high quality parts aimed at doing the job well, as well as low quality parts that are aimed at making the manufacturer a profit. Louis built a tool for trimming hedges, named “Snippy,” that was different.

[*Start*=good] Even though Snippy contains a mix of high and low quality parts, everyone can tell that the real essence of the tool is its high quality parts. The low quality parts are just a thin veneer over the more essential parts of the tool.

[*Start*=bad] Even though Snippy contains a mix of high and low quality parts, everyone can tell that the real essence of the tool is its low quality parts. The high quality parts are just a thin veneer over the more essential parts of the tool.

One day Louis sells Snippy to his neighbor, who replaces some of Snippy’s parts.

[*Telos*=preserved, *Norm*=better] Specifically, the neighbor completely replaced all of the low quality parts with high quality parts. Now, the tool includes just high quality parts and the neighbor uses the tool to trim his hedges.

[*Telos*=preserved, *Norm*=worse] Specifically, the neighbor completely replaced all of the high quality parts with low quality parts. Now, the tool includes just low quality parts and the neighbor uses the tool to trim his hedges.

[*Telos*=destroyed, *Norm*=better] Specifically, the neighbor completely removed the parts needed for doing any gardening work. The neighbor instead uses the tool at the center of an art project he is building, which is a large sculpture celebrating diversity.

[*Telos*=destroyed, *Norm*=worse] Specifically, the neighbor completely removed the parts needed for doing any gardening work. The neighbor instead uses the tool at the center of an art project he is building, which is a large sculpture celebrating racism.

After reading one of the cases participants were given the same questions (again, suitably adjusted) as those used in Study 1A. 125 participants were removed due to failing one or more of the comprehension questions. Data were analyzed from the remaining 275 participants.

Identity: There was no main effect of *Start* (good: $M=3.77$, $SE=.159$; bad: $M=4.05$, $SE=.171$), $F(1, 267)=1.456$, $p=.229$, or *Norm* (better: $M=4.02$, $SE=.168$; worse: $M=3.81$, $SE=.163$), $F(1, 267)=.820$, $p=.366$, but there was a main effect of *Telos* (preserved: $M=4.35$, $SE=.163$; destroyed: $M=3.47$, $SE=.167$), $F(1, 267)=14.021$, $p<.001$. *Telos* produced a small sized effect ($\eta_p^2=.050$). The main effect of *Telos* was qualified by a small three-way interaction, $F(1, 267)=10.791$, $p<.01$, $\eta_p^2=.039$.

Purpose: There was no main effect of *Start* (good: $M=3.55$, $SE=.131$; bad: $M=3.55$, $SE=.141$), $F(1, 267)=.000$, $p=.998$, but there was a main effect of *Norm* (better: $M=3.78$, $SE=.138$; worse: $M=3.33$, $SE=.134$), $F(1, 267)=5.655$, $p<.05$, and *Telos* (preserved: $M=5.13$, $SE=.134$; destroyed: $M=1.98$, $SE=.138$), $F(1, 267)=266.666$, $p<.001$. While *Norm* had a small-sized effect on purpose judgments ($\eta_p^2=.021$), *Telos* produced a large sized effect ($\eta_p^2=.500$). These main effects were qualified by a small three-way interaction, $F(1, 267)=10.373$, $p<.01$, $\eta_p^2=.037$.

Normativity: There was no main effect of *Start* (good: $M=4.41$, $SE=.108$; bad: $M=4.35$, $SE=.116$), $F(1, 267)=.112$, $p=.738$, but there was a main effect of *Telos* (preserved: $M=4.75$, $SE=.111$; destroyed: $M=4.01$, $SE=.114$), $F(1, 267)=21.363$, $p<.001$, $\eta_p^2=.074$, and *Norm* (better: $M=4.88$, $SE=.114$; worse: $M=3.88$, $SE=.111$), $F(1, 267)=40.326$, $p<.001$, $\eta_p^2=.131$ on normative judgments. These effects were qualified by a small two-way interaction between *Telos* and *Norm*, $F(1, 267)=7.111$, $p<.01$, $\eta_p^2=.026$.

Essence: There was no main effect of *Start* (good: $M=4.58$, $SE=.131$; bad: $M=4.37$, $SE=.142$), $F(1, 267)=1.157$, $p=.283$ or *Norm* (better: $M=4.31$, $SE=.138$; worse: $M=4.64$, $SE=.134$), $F(1, 267)=3.169$, $p=.076$. There was a large-sized main effect of *Telos* (preserved: $M=3.52$, $SE=.135$; destroyed: $M=5.43$, $SE=.138$), $F(1, 267)=98.568$, $p<.001$, $\eta_p^2=.270$. This main effect was qualified by a small three-way interaction, $F(1, 267)=9.718$, $p<.01$, $\eta_p^2=.035$.

These results may be visualized as follows (higher scores indicate stronger judgments of persistence):

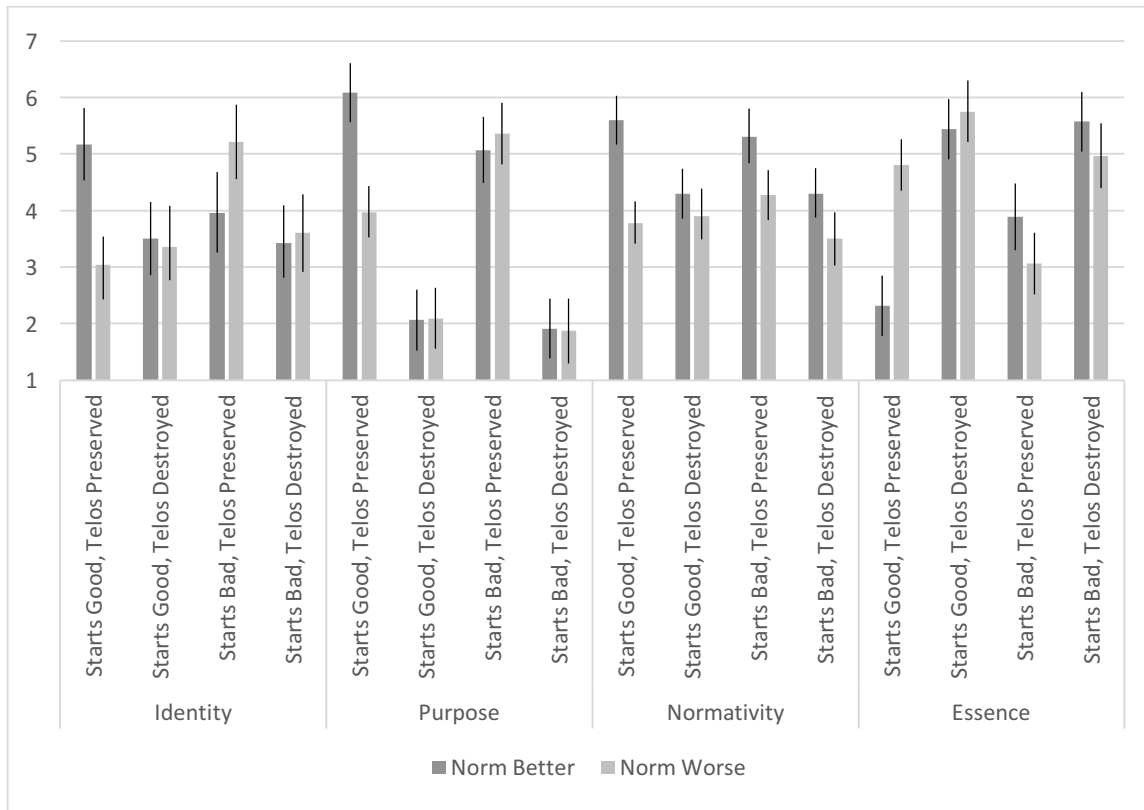


Figure 13: Effect of Norm on Each DV




Figure 14: Effect of Telos on Each DV

As with our other case involving an artifact (Study 3's science paper cases), here again we find that only *Telos* produced a main effect on identity judgments. We also found that only *Telos* produced a main effect on essence judgments. *Telos* and *Norm* produced main effects on Purpose and Normativity, though the effect of *Telos* on Purpose was much larger while the effect of *Norm* on Normativity was much larger. Again, the crucial question is whether Purpose or Normativity causes identity judgments.

2.4.2 Version 4B: causal model pitting teleology versus normativity

250 people were recruited from Amazon Mechanical Turk ($M_{age}=39$, 43% female). Each participant read the *Start=bad* version of Study 4A which varied *Norm* and *Telos*, resulting in a $2(Norm: \text{better, worse}) \times 2(Telos: \text{preserved, destroyed})$ design. Participants were then asked each question (presented in random order) from Study 1A except for the essence question (again we could pit teleology directly against normativity). 58 people were removed for failing one or more comprehension question. Data was then analyzed from the remaining 192 participants.

A multiple regression model with Purpose, Normativity, *Norm*, *Telos* and a *Telos* by *Norm* interaction revealed that a full model was able to account for 24% of the variance in identity judgments, $F(5, 187)=11.803, p<.001, R^2=.240$. Tellingly, the only variable that had a significant effect in the full model was Purpose ($\beta=.474, t=5.485, p<.001$).

We then conducted a causal search on the data, using Greedy Equivalence Search (GES), which returned the following model:




Figure 15: Causal Model of Identity Judgments

This model fits the data well $\chi^2(6)=1.682, p>.05, BIC=-29.894$. This model (like that in §2.2.2, and unlike those in §2.1.2 and §2.3.2) does not position Normativity as even an indirect cause of Identity (actually it positions Normativity and Identity as correlates of a common cause in Purpose). Importantly, as in all our models, only Purpose is a direct cause of Identity. So we see yet another clear and consistent sign that Purpose and not Normativity is what directly causes identity judgments.

3. Discussion of Results

We asked, how if at all are the normative and teleological effects on persistence judgments connected? Our results indicate that teleology is a direct cause of persistence judgments but normativity is not, and that any effect of normativity is screened off by teleology.

We see three main components to our results. First, in our cases with social objects (the university and band studies: §§2.1–2.2), we found that *Norm*—our variable manipulating whether the change was normatively an improvement or a deterioration—did produce a main effect on Identity—our measure of persistence judgments. But this effect was still much smaller than the effect of *Telos*—our variable manipulating whether the object preserved its purpose. Secondly, in our cases with artifacts (the science paper and gardening tool cases: §§2.3–2.4), we found that only *Telos* produced a main effect on identity judgments. Thirdly—and by our lights most convincingly—we found that, when we constructed causal models of all four cases, only Purpose—our measure of teleological shift—directly caused Identity. Any effect observed from *Norm* or Normativity was mediated through Purpose. Thus, our results indicate that folk teleology drives persistence judgments.

We take these results to connect Knobe’s program on normative effects with Kelemen’s program on teleological effects. At least with respect to persistence judgments (where both effects are observed), we offer a causal model of how these effects are connected. We take these results to shed light on our folk notion of an object (or perhaps more domain-specific notions such as *social object* and *artifact*). Just as our judgments about what something is (sortal) are determined by its purpose, so our judgments of whether a plurality forms a whole (composition) are determined by whether the plurality shares a purpose, and so our judgments about whether something persists is determined by whether its purpose is preserved.

We conclude by flagging three follow-up questions, as invitations for further research. A first follow-up question concerns scope. We have considered social objects (universities and bands) and artifacts (science papers and gardening tools). But we have not considered whether the teleological effect on persistence judgments extends beyond these domains. So it remains open whether this effect extends to organisms, persons, or natural non-biological objects like rocks or clouds. We speculate that the effect does so extend, since Kelemen (1999) shows that promiscuous teleological thinking extends widely, and Rose (2015) extends the effect to rocks. We also have not considered whether a normative effect might be observed in these domains as well, and if so, the extent to which it would be screened-off by a teleological effect.

A second follow-up question concerns essentialism, and the connection between teleological and essence effects. We speculate that the connection is that the essences of objects are typically conceived of teleologically. This fits our findings, especially how manipulation on *Telos* consistently affected Essence—our measure of essence judgements—and consistently produced a larger effect than any other variable.

A third follow-up question concerns how generally to view the relation between normative and teleological factors. We speculate that teleology is in the driver’s seat, and that normativity comes into play due to a background tendency to impute normatively-laden purposes to objects, and thereby view normative improvement as more “on track” than deterioration. One interesting sort of case to examine would be a case where people tend to impute an evil purpose to an object (perhaps a weapon, or a malevolent spirit, could serve in this role). We speculate that—if a case could be found where people did impute an evil purpose—any effect of normativity on persistence judgments would likely be reversed.⁶

References

Alicke, M., Rose, D., & Bloom, D. (2011). Causation, Norm Violation and Culpable Control. *Journal of Philosophy*, 108, 670–96.

⁶ Thanks to Julian De Freitas, and especially to Joshua Knobe, who helped initiate this project.

- Beebe, J. R. & Buckwalter, W. (2010). The Epistemic Side-Effect Effect. *Mind & Language*, 25, 474–98.
- Blok, S., Newman, G., Behr, J., & Rips, L. J. (2001). Inferences about personal identity. In J. D. Moore & K. Stenning (Eds.), *Proceedings of the Twenty-Third Annual Conference of the Cognitive Science Society* (pp. 80–85). Mahwah, NJ: Erlbaum.
- Blok, S. V., Newman, G., & Rips, L. J. (2005). Individuals and their concepts. In W. K. Ahn, R. L. Goldstone, B. C. Love, A. B. Markman, & P. Wolff (Eds.), *Categorization inside and outside the lab* (pp. 127–49). Washington, D.C.: American Psychological Association.
- Bloom, P. (2007). Religion is Natural. *Developmental Science*, 10, 147–51.
- Chickering, D. (2002). Optimal Structure Identification with Greedy Search. *Journal of Machine Learning Research*, 3, 507–54.
- De Freitas, J., Tobia, K., Newman, J. E., & Knobe, J. (2017). Normative judgments and individual essence. *Cognitive Science*, 41, 382–402.
- Ellis, P. (2010). *The Essential Guide to Effect Sizes: Statistical Power, Meta-Analysis and the Interpretation of Research Results*. Cambridge: Cambridge University Press.
- Hall, D. G., Waxman, S. R., Bredart, S., & Nicolay, A. C. (2003). Preschoolers' use of form class cues to learn descriptive proper names. *Child Development*, 74, 1547–60.
- Hitchcock, C. & Knobe, J. (2009). Cause and Norm. *Journal of Philosophy*, 11, 587–612.
- Kelemen, D. (1999). Why are Rocks Pointy? Children's Preference for Teleological Explanations of the Natural World. *Developmental Psychology*, 35, 1440–52.
- Kelemen, D., & Rosset, E. (2009). The Human Function Compunction: Teleological Explanation in Adults. *Cognition*, 111, 138–43.
- 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, 1074–83.
- Knobe, J. (2003). Intentional Action and Side Effects in Ordinary Language. *Analysis*, 63, 190–93.
- Knobe, J. (2010). Person as Scientist, Person as Moralist. *Behavioral and Brain Sciences*. 33, 315–65.
- 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.
- Meek, C. (1997). *Graphical Models: Selecting Causal and Statistical Models*. PhD Thesis, Carnegie Mellon University.
- Newman, G. E., Bartels, D. M., & Smith, R. K. (2014). Are artworks more like people than artifacts? Individual concepts and their extensions. *Topics in Cognitive Science*, 6, 647–62.
- Newman, G. E., De Freitas, J., & Knobe, J. (2015). Beliefs About the True Self Explain Asymmetries Based on Moral Judgment. *Cognitive Science*, 39, 96–125.
- Pettit, D. & Knobe, J. (2009). The Pervasive Impact of Moral Judgment. *Mind & Language*, 24, 586–604.
- Rips, L. (1989). Similarity, Typicality and Categorization. In S. Vosniadou & A. Ortony (Eds.), *Similarity and Analogical Reasoning* (pp. 21–59). Cambridge: Cambridge University Press.
- Rose, D. (2015). Persistence Through Function Preservation. *Synthese*, 192, 97–146.

- Rose, D. (2017). Folk Intuitions of Actual Causation: A Two-Pronged Debunking Explanation. *Philosophical Studies*, 174, 1323–61.
- Rose, D., & Nichols, S. (2013). The Lesson of Bypassing. *Review of Philosophy and Psychology*, 4, 599–619.
- Rose, D. and Schaffer, J. (2017). Folk Mereology is Teleological. *Noûs*, 51, 238–70
- Rose, D., Livengood, J., Sytsma, J., & Machery, E. (2011). Deep Trouble for the Deep Self. *Philosophical Psychology*, 25, 629–46.
- Strawson, P. F. (1959). *Individuals*. London: Methuen.
- Strohmingner, N., Knobe, J., & Newman, G. (2017). The True Self: A Psychological Concept Distinct from the Self. *Perspectives on Psychological Science*, 12, 551–60.
- Thomasson, A. (2007). *Ordinary Objects*. Oxford: Oxford University Press.
- Tobia, K. (2015). Personal identity and the Phineas Gage effect. *Analysis*, 75, 396–405.
- Turri, J., Buckwalter, W. and Rose, D. (2016), Actionability Judgments Cause Knowledge Judgments. *Thought: A Journal of Philosophy*, 5, 212–22.