

Folk Teleology Drives Persistence Judgments

David Rose, Jonathan Schaffer, and Kevin Tobia
draft of August 9th, 2018

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 tend to judge that the house survived. But if the fire burnt the house to the ground instead, people tend to 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 heart of our 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 our naïve 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 that teleology is in the driver’s seat, in ways that shed 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, leading Knobe (2010: 316) to the “person-as-moralist” perspective on which “[M]oral considerations actually figure in the competencies people use to make sense of human beings and their actions.”

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):

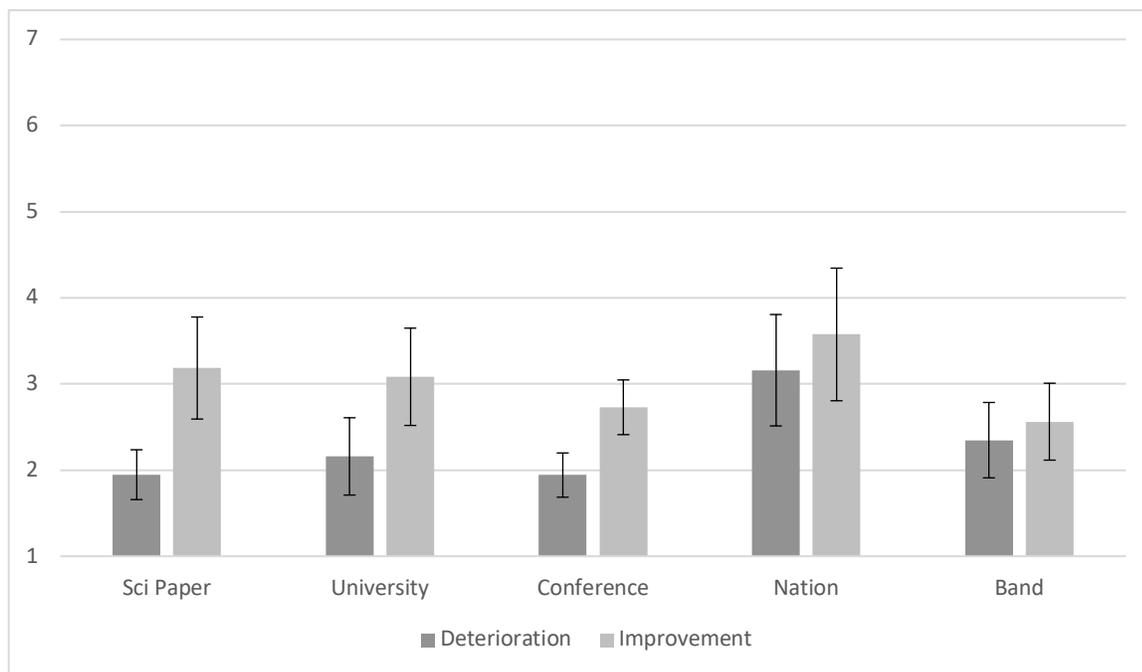


Figure 1: De Freitas et al.’s Effect of Deterioration/Improvement on Persistence

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. As Dawkins (1995: 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.”

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):

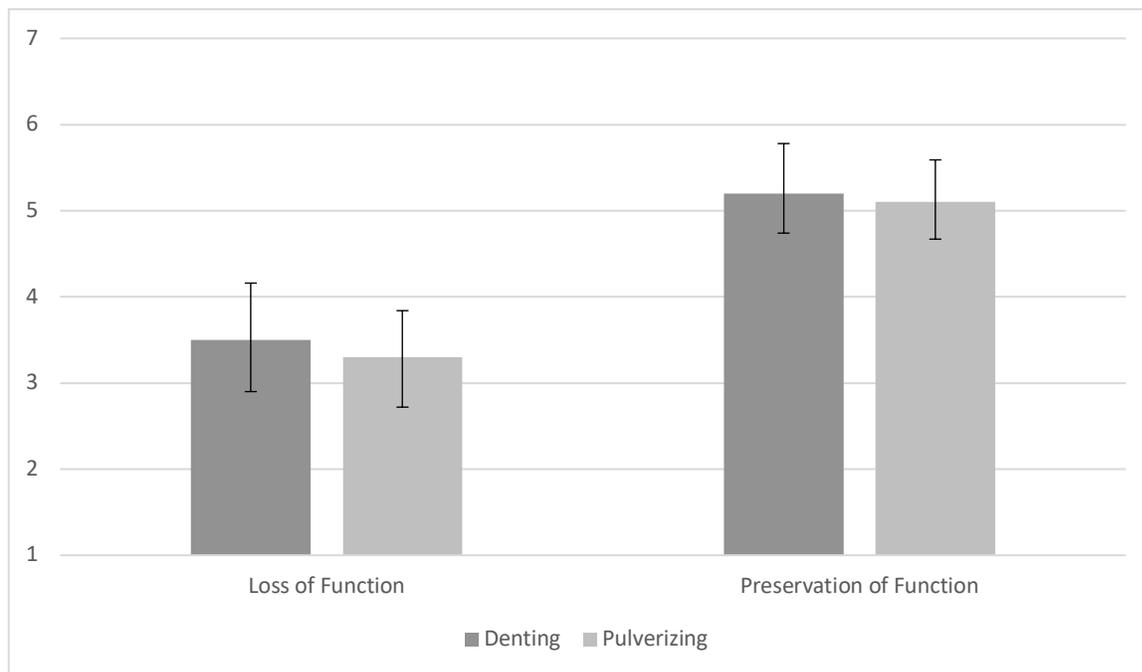


Figure 2: Rose’s Effect of Denting/Pulverizing and Loss/Preservation of Function on Persistence

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). It is important to see that these are distinct and separable. For instance, if a university in Nazi-era Germany changes to become a Jewish aid organization aimed to halt the Nazi genocide, that looks like a normative improvement but a teleological shift “off track” from the mission of a university. Or if

a rock band changes from producing commercial songs for the radio, to playing cover songs at parties for drug dealers, that looks like a normative deterioration while remaining “on track” with the mission of a band.¹

Given that normative and teleological effects on persistence judgments are distinct and separable, 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.² Participants were asked both to make persistence judgments about an object following either a normative improvement or deterioration, and also to make essence judgments as to 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”):

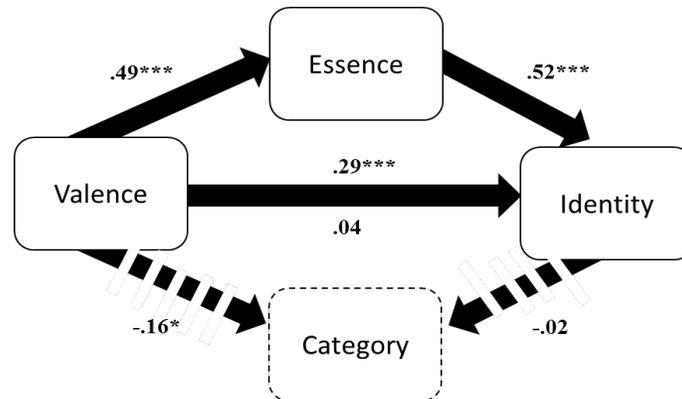


Figure 3: De Freitas et al.'s Mediation Model

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 we 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

¹ We say that a change is “on track” if it fits the apparent purpose of the object, and “off track” if it departs from the apparent purpose of the object. For instance, if the apparent purpose of a university is to educate students, then a change such as doubling student enrollments is still “on track” (the university can still educate students), while a change such as eliminating student enrollments is “off track” (the university can no longer educate students).

² 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; Blok *et al.* 2005; Newman *et al.* 2014). This is understood in terms of the deeper features being viewed as “essential.” More precisely, as De Freitas *et al.* (2017: 382–3) clarify, essence judgments are understood as judgments as to the range of changes through which an object can and cannot persist.

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,³ and thereby tend to see normative improvements in an object as more “on track,” in the sense of fitting its purpose, while tending to see normative deterioration as more “off track,” in the sense of departing from its purpose.
- 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

³ 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.”

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 sought to construct tractable causal models in order to pit normative versus teleological effects on persistence judgments. While we found some interesting variation across cases, our central finding—present in all cases—is that any normative effects on persistence judgments are mediated by an intermediate judgment as to whether the change fits with or departs from the purpose of the thing. Holding this judgment of purpose fixed screens off any effect of normativity.

2.1 Study 1: *university*

250 people were recruited from Amazon Mechanical Turk ($M_{\text{age}}=39$, 41% female). Participants were randomly assigned to one of four conditions in a $2(\text{Norm: better, worse}) \times 2(\text{Telos: preserved, destroyed})$ design, using 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. 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.

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 four cases, participants were asked the following (presented in random 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 institution's true purpose or departs from the institution's 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)

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

47 participants were removed for failing one or more of the comprehension questions. Data were then analyzed from the remaining 203 participants.

A multiple regression model with Purpose, Normativity, *Norm*, *Telos* and a *Norm* x *Telos* interaction revealed that a full model was able to account for 60% of the variance in identity judgments, $F(5, 197)=59.143, p<.001, R^2=.600$. The only variables that had significant effects in the full model were *Telos* ($\beta=-.221, t=-2.774, p<.01$) and Purpose ($\beta=.611, t=7.914, p<.001$). Follow-up t-tests are provided in Table 1. As Table 1 indicates, there was a medium-sized effect of *Telos* on Purpose and a very large-sized effect of *Norm* on Normativity.

	Telos Preserved	Telos Destroyed	t-value	p-value	Cohen's d⁴
Purpose	3.77 (2.51)	2.56 (1.68)	4.09	<.001	.57 (medium)
Normativity	3.68 (2.59)	3.48 (2.34)	.586	.558	.08 (n/a)
Identity	3.79 (2.19)	2.62 (1.56)	4.42	<.001	.62 (medium)
	Norm Better	Norm Worse	t-value	p-value	Cohen's d
Purpose	4.52 (2.14)	1.96 (1.43)	10.16	<.001	1.41 (large)
Normativity	5.91 (1.39)	1.64 (1.06)	24.86	<.001	3.45 (very large)
Identity	4.11 (2.04)	2.39 (1.52)	6.89	<.001	.96 (large)

Table 1: T-tests for University

We then conducted a causal search on the data, using Greedy Equivalence Search (GES),⁵ which returned the following model:

⁴ Cohen (1988) offers the following benchmarks for interpreting the magnitude of effect sizes: $d \leq .20$ is small-sized; $d \leq .50$ is medium-sized; and $d \leq .80$ is large-sized. Rosenthal (1996) adds that $d \leq 1.30$ is very large-sized. This is what we will follow in interpreting the magnitude of effect sizes.

⁵ 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 that would

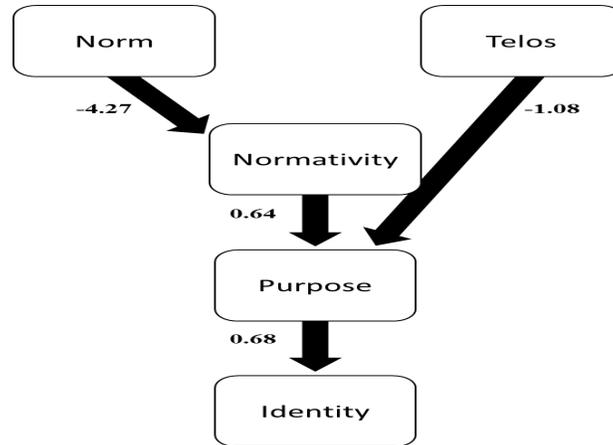


Figure 4: Causal Model of Identity Judgments for University

This model fits the data well $X^2(6)=12.097, p>.05, BIC=-19.782$. This model positions Normativity as an indirect cause of Identity, but only via Purpose (a dependent variable measuring the extent to which people say that the change fit with or departed from the “true purpose” of the thing). 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

We wanted to look at a wider sample of the basic cases from De Freitas *et al.* (2017), so we took up their band case. We also wondered if the fact that the Iserlohn Institute was portrayed as having started off normatively “good” (with an academic focus) might have mattered, so we wanted to look at a case that started off in a “bad” state by having the band start off as focused on superficial commercial songs. Accordingly, 250 people were recruited from Amazon Mechanical Turk ($M_{age}=35, 32\%$ female). Participants were randomly assigned to one of four conditions in a $2(Norm: \text{better, worse}) \times 2(Telos: \text{preserved, destroyed})$ design, using the following cases:

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. 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.

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.

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.

[*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 1 (§2.1), again in random order. 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 *Norm* × *Telos* interaction revealed that a full model was able to account for 37.5% of the variance in identity judgments, $F(5, 185)=22.241, p<.001, R^2=.375$. The only variables that had significant effects in the full model were Normativity ($\beta=.222, t=2.595, p<.01$) and Purpose ($\beta=.491, t=7.445, p<.001$). Table 2 includes results from follow-up t-tests. As Table 2 indicates, there was a medium-sized effect of *Telos* on Purpose and a large-sized effect of *Norm* on Normativity.

	Telos Preserved	Telos Destroyed	t-value	p-value	Cohen’s d
Purpose	4.27 (1.80)	3.22 (1.81)	3.99	<.001	.58 (medium)
Normativity	4.60 (1.43)	3.75 (1.99)	3.33	<.01	.49 (small)
Identity	4.05 (1.97)	2.94 (1.90)	3.93	<.001	.57 (medium)
	Norm Better	Norm Worse	t-value	p-value	Cohen’s d
Purpose	3.21 (1.67)	4.19 (1.96)	-3.73	<.001	.54 (medium)
Normativity	5.27 (1.20)	2.95 (1.56)	11.55	<.001	1.18 (large)
Identity	3.20 (1.92)	3.69 (2.07)	-1.68	.10	.25 (small)

Table 2: T-tests for Band

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

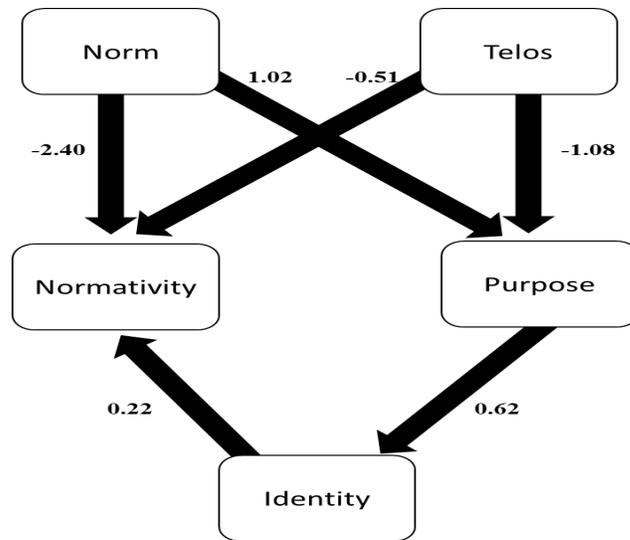


Figure 5: Causal Model of Identity Judgments for Band

This model fits the data well $X^2(4)=8.577, p>.05, BIC=-12.432$. This model does not position Normativity as even an indirect cause of Identity (actually it positions Identity as a direct cause of Normativity).⁶ 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

We wanted to look at one more of the basic cases from De Freitas *et al.* (2017). We also wondered if the fact that universities and bands are social groups might have mattered, so we took up their science paper case (and alternated back to starting off in a “good” state). Accordingly, 250 people were recruited from Amazon Mechanical Turk ($M_{age}=31, 38\%$ female). Participants were randomly assigned to one of four conditions in a $2(Norm: \text{better, worse}) \times 2(Telos: \text{preserved, destroyed})$ design, using the following cases:

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. 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.

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

[*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.

⁶ 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*=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 1 (§2.1), again in random order. 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 ($\beta=.696, t=6.681, p<.001$). T-tests are in Table 3, indicating a large-sized effect of *Telos* on Purpose and a large-sized effect of *Norm* on Normativity.

	Telos Preserved	Telos Destroyed	t-value	p-value	Cohen's d
Purpose	3.61 (2.38)	1.87 (1.40)	5.78	<.001	.92 (large)
Normativity	3.70 (2.17)	2.47 (1.55)	4.25	<.001	.56 (medium)
Identity	3.64 (2.20)	2.20 (1.66)	4.81	<.001	.74 (medium)
	Norm Better	Norm Worse	t-value	p-value	Cohen's d
Purpose	3.70 (2.35)	1.98 (1.60)	5.67	<.001	.36 (small)
Normativity	4.19 (1.96)	2.22 (1.51)	7.37	<.001	1.12 (large)
Identity	3.62 (2.23)	2.37 (1.77)	4.07	<.001	.62 (medium)

Table 3: T-tests for Science Paper

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

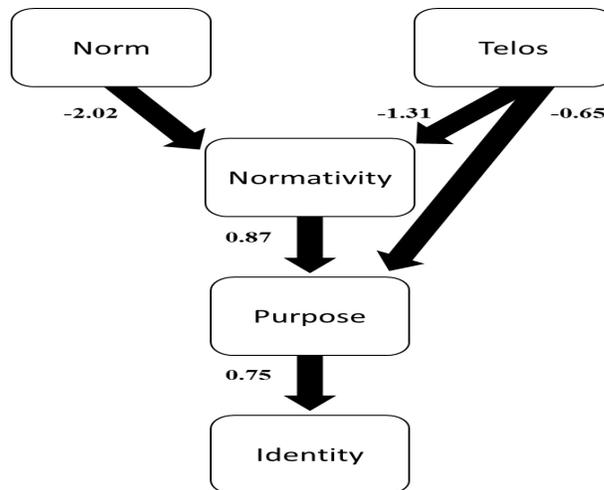


Figure 6: Causal Model of Identity Judgments for Science Paper

This model fits the data well $X^2(5)=1.645, p>.05, BIC=-24.004$. This model (like the model in §2.1, but unlike that in §2.2) positions Normativity as an indirect cause of Identity, but still 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

We wanted to construct a further case of our own looking at a further artifactual object but alternating back to starting off in a “bad” state. Accordingly, 250 people were recruited from Amazon Mechanical Turk ($M_{age}=39, 43\%$ female). Participants were randomly assigned to one of four conditions in a $2(Norm: \text{better, worse}) \times 2(Telos: \text{preserved, destroyed})$ design, using the following 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. 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 (suitably adjusted) as those used in Study 1 (§2.1), again in random order. 58 people were removed for failing one or more comprehension question. Data was 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, $F(5, 187)=11.803, p<.001, R^2=.240$. The only variable that had a significant effect in the full model was Purpose ($\beta=.474, t=5.485, p<.001$). Table 4 includes t-tests, indicating a very large-sized effect of *Telos* on Purpose and a medium-sized effect of *Norm* on Normativity.

	Telos Preserved	Telos Destroyed	t-value	p-value	Cohen's d
Purpose	4.72 (1.65)	1.99 (1.47)	12.19	<.001	1.75 (very large)
Normativity	4.33 (1.24)	4.08 (1.24)	1.43	.154	.20 (small)
Identity	4.68 (1.77)	3.45 (1.75)	4.86	<.001	.70 (medium)
	Norm Better	Norm Worse	t-value	p-value	Cohen's d
Purpose	3.17 (2.13)	3.37 (2.01)	-.67	.504	.10 (n/a)
Normativity	4.63 (1.24)	3.76 (1.08)	5.17	<.001	.75 (medium)
Identity	3.95 (1.86)	4.09 (1.86)	-.54	.588	.08 (n/a)

Table 4: T-tests for Gardening Tool

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

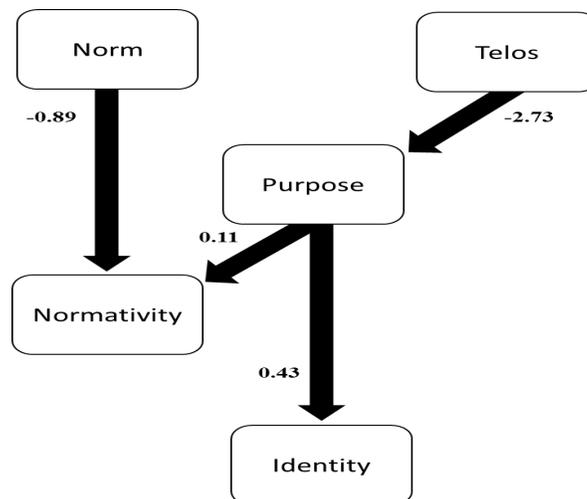


Figure 7: Causal Model of Identity Judgments for Gardening Tool

This model fits the data well $X^2(6)=1.682, p>.05, BIC=-29.894$. This model (like that in §2.2, and unlike those in §2.1 and §2.3) 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 the normative and teleological effects on persistence judgments are related, if at all? Our results indicate that teleology is a direct cause of persistence judgments but normativity is not. Across four studies that involved very different things—from social objects (the university and band studies: §§2.1–2.2) to artifacts (the science paper and gardening tool cases: §§2.3–2.4)—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 (persistence) is determined by whether its purpose is preserved.

Obviously our results are subject to further interpretation, and our conclusions open to revision in light of further empirical research. We note two aspects of our studies as potential places for re-interpretation. First, our probe for *Purpose* uses the wording of ‘true purpose’ and one might worry that ‘true’ is being read in a normatively loaded way, so that we are not eliciting “pure” purpose judgments.⁷ While this is possible, we see no reason yet to think this is happening. Indeed our data already suggests that it is not. Although *Norm* affects Purpose in University, Band and Science Paper—where the magnitude of the effect ranges from small to large—it does not have a significant effect in Gardening Tool. This suggests that, although judgments of purpose are often affected by perceived normativity, this is not always the case—and it is certainly not an effect that occurs wherever “true purpose” is invoked. We welcome further research on this matter, and we note that this worry could be explored by varying the wording of the probes.

Secondly, our manipulation of *Telos* might seem to be of greater magnitude than our manipulation of *Norm*. If, for instance, in our University case, the manipulation of *Norm*—a focus on either traditional academic subjects or Nazi ideology—led to something not as overall dissimilar from the original than the manipulation of *Telos*—a focus on either providing Jewish aid or developing chemical weapons for the Nazi’s, then this might on its own explain why differences in Purpose are viewed as more disruptive of identity. So one might worry that there is a confound with sheer magnitude of dissimilarity.⁸ One option to address this concern would be to run pre-tests for similarity judgments, and revise the vignettes to use matched changes. For example, pre-testing would ensure that the Norm change resulted in a university that participants rate “as dissimilar” from the original as the university produced from the Telos change. But this is worrisome to us, since we expect that similarity judgments will themselves be connected to persistence judgments, so that a prime driver of similarity judgments will be *having the same purpose*. (This is of course an empirical matter.)

But actually we think that our data already contain the raw materials needed to address this concern about dissimilarity. Consider our results concerning the magnitude of the effect of *Norm* on Normativity and *Telos* on Purpose reported in Tables 1-4. These Tables show two important facts. First, there was an effect of *Norm* on Normativity in all of these cases. Secondly, in two cases (University and Band) the effect of *Norm* on

⁷ We thank an anonymous referee for helpful discussion.

⁸ We thank a second anonymous referee for raising this concern, and for encouraging us to report more of our statistical results (which we appeal to in response to this concern).

Normativity was larger than the effect of *Telos* on Purpose; in Science Paper, the effect of *Norm* on Normativity and of *Telos* on Purpose are of the same magnitude; and in Gardening Tool, the effect of *Norm* on Normativity is smaller than the effect of *Telos* on Purpose. So in most of these cases, our normativity manipulation had a larger direct effect than our purpose manipulation. Yet, in all these cases only Purpose directly caused Identity. So regardless of the magnitude of the *Norm* manipulation effect (which was typically greater than or equal to the magnitude of the *Telos* manipulation effect), Purpose is still the only factor playing a direct role in judgments of Identity. Whatever downstream effects these differences in the magnitude of the manipulations are having on Identity, it doesn't appear to be one of unfairly advantaging Purpose.

Indeed there is a second way in which our data can already address this concern about dissimilarity. Recall that, in University, the effect of Norm on Identity is larger than the effect of *Telos* on Identity, while in Science paper the effect of both Norm and *Telos* on Identity is of the same magnitude. But throughout these studies, Purpose is the only factor that plays a direct causal role in Identity. We take this to suggest that the amount of dissimilarity associated with *Telos* can be varied, without varying the core result of Purpose being the only factor that plays a direct causal role in Identity. We welcome further research on this matter as well, and we note that this worry could be explored by measuring judgments of dissimilarity and constructing larger causal models integrating these judgments as well.

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 are currently conducting follow-up studies on the hypothesis that the effect does so extend, since Kelemen (1999) shows that promiscuous teleological thinking extends widely, and Rose (2015) shows that the effect extends to rocks. We also hypothesize that an indirect normative effect will be observed in these domains as well, 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 a range of further studies we have conducted, which we aim to publish separately. This also fits a broadly "Aristotelian" conception of essences.)

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 be reversed.⁹

References

- Alicke, M., Rose, D., & Bloom, D. (2011). Causation, Norm Violation and Culpable Control. *Journal of Philosophy*, 108, 670–96.
- Beebe, J. R. & Buckwalter, W. (2010), The Epistemic Side-Effect Effect. *Mind & Language*, 25, 474–98.

⁹ Thanks to Julian De Freitas, and especially to Joshua Knobe, who helped initiate this project. Thanks also to two anonymous referees for *Synthese* for insightful comments.

- 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). Hillsdale, 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.
- Cohen, J. (1988). *Statistical Power Analysis for the Behavioral Sciences*, 2nd Edition. Hillsdale, NJ: Erlbaum.
- Dawkins, Richard 1995. *River out of Eden: A Darwinian View of Life*. New York, NY: Basic Books.
- 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.
- Rosenthal, J. A. (1996). Qualitative Descriptors of Strength of Association and Effect Size. *Journal of Social Service Research*, 21, 37–59.
- 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*, 5, 212–22.