Not so simple! Mechanisms increase preference for complex explanations

Jeffrey C. Zemla¹, Steven Sloman², Christos Bechlivanidis³ and David Lagnado³

¹ Department of Psychology, University of Wisconsin-Madison. 1202 West Johnson St.,
Madison, WI 53706

² Department of Cognitive, Linguistic, and Psychological Sciences, Brown University. 190
Thayer St., Providence, RI 02912

³ Department of Experiment Psychology, University College London. 26 Bedford Way,
London WC1H 0AP, UK

Correspondence concerning this article should be addressed to Jeffrey Zemla, Department
of Psychology, University of Wisconsin-Madison, Madison, WI 53706.

Contact: zemla@wisc.edu

Abstract

Mechanisms play a central role in how we think about causality, yet not all causal explanations describe mechanisms. Across four experiments, we find that people evaluate explanations differently depending on whether or not they include mechanisms. Despite common wisdom suggesting that explanations ought to be simple (appealing to as few causes as necessary to explain an effect), the previous literature is divided over whether people adhere to this simplicity principle. Our findings suggest that the presence of mechanisms in an explanation is one factor that moderates this preference for simplicity. Without mechanisms, people typically exhibit a preference for simple explanations, consistent with probabilistic accounts. This preference is significantly reduced or even reversed when explanations contain mechanisms, suggesting that mechanisms afford a different way of evaluating explanations. Rather than focusing on probability, complex explanations that contain mechanisms may be preferred because they provide a fuller account of the underlying causal network, promoting a greater sense of understanding.

Keywords: explanation, causal networks, simplicity, mechanisms

1 Introduction

In January of 1986, the Space Shuttle Challenger exploded shortly after launch, killing all seven of its crewmembers. In the months that followed, the Rogers Presidential Commission was tasked with explaining what led to this catastrophic failure. The commission sought to answer three questions (Rogers Commission, 1986, p. 40): What faults contributed to the failure? What evidence supports that conclusion? And finally, what was the mechanism of failure?

The questions posed by the Rogers Commission underscore the idea that explanations can achieve multiple goals. In particular, while most explanations involve determining the most likely causes of an event given the available evidence, many explanations also describe the mechanisms that lead from causes to effect. By *mechanism*, we mean an arrangement of entities which interact in an invariant way to generate an output (Glennan, 2010). Much of the literature on explanatory reasoning suggests that we judge explanations in part based on how well an explanation satisfies a set of explanatory virtues (e.g., Lipton, 2004)—such as simplicity, coherence, and breadth—without much consideration for whether an explanation is postulating causes, describing mechanisms, or both.

Some research has hinted that a person's explanatory goals can change which explanatory virtues are given priority. Zemla et al. (2017) suggest that a preference for simplicity in explanation may shift depending on what the goal of an explanation is. Here, we test whether explanations that describe mechanisms are evaluated differently than explanations that do not contain mechanisms.

We hold that mechanisms re-frame the way people think about and evaluate explanations. Mechanism-free explanations are likely to be judged by their simplicity, in part because simple explanations tend to be more probable. We hypothesize that, in contrast, mechanistic explanations provide us with a deeper understanding of the causal network generating an event. This increases the attractiveness of complex explanations because they more fully describe these mechanisms. In four experiments, we test the hypothesis that simple explanations that appeal to fewer causes are attractive when identifying the causes of an event, and that complex explanations that appeal to more causes are more attractive when they include mechanisms.

1.1 The role of mechanisms

1.1.1 Mechanisms versus covariation

Explanations typically serve several purposes. One purpose is to identify which of many possible causes led to an event. Explanations should identify causes that are both true and explanatorily relevant, i.e., they should identify root causes that made a difference to whether the event occurred in the way that it did (Strevens, 2004). In the case of the Challenger shuttle, this was answered plainly: “[...] the loss of the Space Shuttle Challenger was caused by a failure in the joint between the two lower segments of the right Solid Rocket Motor,” and “[...] no other element of the Space Shuttle system contributed to this failure.” (Rogers Commission, 1986, p. 40). The Rogers Commission found that the joint had indeed failed, and judged this to be the singular cause of failure that led to the explosion. In doing so, the commission ruled out other potential causes such as intentional sabotage, which was a genuine concern at the time (Rogers Commission, 1986).

Nonetheless, this explanation is somewhat unsatisfying. *Why* did the joint fail despite working on previous Shuttle flights? *How* did a faulty joint lead to an explosion? The chain of events and necessary conditions that led to an explosion would provide a more complete account of the Challenger explosion.

A mechanistic explanation unpacks a series of mechanisms into a sequence of steps that reveals the enabling conditions describing how an event occurred, rather than what the event is correlated with (Ahn & Bailenson, 1996). In the case of the Challenger shuttle, the unusually cold weather prevented an O-ring seal from expanding to seal the joint; as a result, exhaust leaked from the motor which heated the external fuel tank, causing it to rupture and explode.

Figure 1. (a) One component of explanation involves determining which causes were present (solid circles) or absent (dashed circles). Each of these causes alone could be sufficient to cause a Shuttle failure. (b) Another component of explanation involves exposition of the causal system, including description of the causal mechanisms that generate effects. Although a faulty joint does not necessarily cause Shuttle failure, a faulty joint in conjunction with cold weather set off a chain of events that resulted in Shuttle failure.

Figure 1 contrasts these two types of explanation. Figure 1a depicts a causal model that outlines several potential causes of the Shuttle's failure. One might infer that a faulty joint was the cause of failure based on covariation: this was the first Shuttle flight in which the joint failed, and the first flight in which an explosion occurred. However, a faulty joint in and of itself does not cause an explosion. Figure 1b depicts a causal model highlighting the

mechanisms of failure—the enabling conditions and sequence of events that began with a faulty joint and resulted in an explosion.

1.1.2 Previous Findings

An understanding of the causal mechanism is often critically important for making predictions about the unknown (Johnson & Ahn, 2017), and can be useful in determining whether the same outcome will recur in the presence of a cause (e.g., would the same faulty joint have caused an explosion if the temperature were warmer?). Mechanistic information is generally preferred over covariation information in forming explanations (Ahn, Kalish, Medin, & Gelman, 1995). Johnson and Keil (2017) found this to be particularly true for evaluating specific claims (e.g., smoking caused John’s lung cancer) compared to general claims (e.g., smoking causes cancer).

Explanations with mechanisms are typically judged as better than their counterparts without mechanisms, perhaps because they improve our subjective understanding of a phenomenon (Vasilyeva & Lombrozo, 2015). When participants realize their inability to articulate the specific mechanisms of an explanation, their self-reported level of understanding decreases (Alter, Oppenheimer, & Zemla, 2010; Fernbach et al., 2013; Rozenblit & Keil, 2002).

Knowledge of explanatory mechanisms can lead to apparent violations of probabilistic reasoning. Park and Sloman (2013) found that beliefs about the mechanism of action predicts whether participants will violate the Markov condition of Bayesian networks (screening off). For instance, smoking causes both lung cancer and blood vessel damage. However, participants will judge the likelihood of a patient having blood vessel damage as higher when told that the patient is both a smoker and has lung cancer

compared to participants who are only told that the patient is a smoker, even though lung cancer and blood vessel damage should be conditionally independent given that the patient is a smoker. This finding can be explained if we assume participants bring latent knowledge about mechanisms into play when making probability judgements. For example, one might believe that smoking causes blood vessel damage in part because lung cancer can invade the bloodstream. Johnson and Ahn (2015) found that judgements of causality violate transitivity when the causal relations are not seen as belonging to a single mechanism. For example, although sex causes pregnancy and pregnancy causes nausea, participants did not judge sex and nausea to be causally related. This contrasts with similar causal chains that encompass a single mechanism: Not flossing causes plaque build up, plaque causes cavities, and not flossing causes cavities were all judged as highly causal.

Overall, mechanisms play a special role in explanation. Judgments of explanatory goodness and self-reported levels of understanding are both affected by the presence of mechanisms. People seek out mechanistic information when developing explanations, and knowledge of mechanisms influences the way we think about causality and probability.

1.2 Simplicity in explanation

1.2.1 A case for simplicity

A guiding principle in explanatory reasoning is that explanations ought to be *simple* (Lipton, 2004; Thagard, 1978). Though several measures of simplicity have been proposed, one measure of simplicity that has received attention in the psychological literature is the number of root causes that an explanation appeals to (Pacer & Lombrozo, forthcoming). An

explanation that appeals to few root causes is said to be a *simple* explanation¹. While fewer causes may be preferred due to an innate preference for parsimony (Lombrozo, 2007), it is also often supported on probabilistic grounds. When the presence of a cause is uncertain, we can infer how likely the cause is (via Bayes' rule), taking into account both the base rate probability of the cause and the conditional probability of an effect given that cause. This offers an intuitive rationale for why we should prefer explanations with a minimal number of causes: a conjunction of two independent causes has a lower base rate than any one of those causes individually. If one's goal is to find the set of most likely causes given an effect, a rational solution is to posit the existence of as few causes as necessary (all else being equal).

However, this criterion does not necessarily align with other explanatory goals. When an explanation contains mechanisms, it affords additional ways of evaluating that explanation beyond probability (e.g., the clarity of an explanation, the completeness of an explanation, the informativeness of an explanation, etc.). It is not clear whether positing fewer causes improves the quality of an explanation with respect to these other criteria that may be more salient in mechanistic explanations.

1.2.2 Previous Findings

Several authors have argued persuasively in favor of a simplicity principle for explanations. Thagard (1989) proposes simplicity as a core tenet of a good explanation, and this hypothesis is empirically supported by others (Read & Marcus-Newhall, 1993; van Overwalle, 1998). For example, Read and Marcus-Newhall (1993) asked participants to rate various explanations for why Cheryl has an upset stomach, has gained weight, and has

¹ For the remainder of this article, we use the terms *simple* and *complex* as shorthand to refer to competing explanations that appeal to fewer or more causes relative to each other.

been feeling tired recently. Participants preferred an explanation that appealed to a common cause (Cheryl is pregnant) as opposed to multiple separate causes (Cheryl has a stomach virus, has stopped exercising, and has Mononucleosis). Similarly, Lombrozo (2007) and Bonawitz and Lombrozo (2012) have provided further evidence that both adults and children exhibit a preference for simple, common cause explanations even when competing explanations are equally probable.

Similarly, an extensive literature on discounting (Kelley, 1973) indicates that people frequently discount the role of a cause in producing an effect given a belief that other causes could have produced that effect (see Morris & Larrick, 1995 and Khemlani & Oppenheimer, 2011 for reviews of the discounting literature). For example, one might infer that Sally failed an exam because she did not study enough. But when told that *most* students in the class failed the exam, one might instead come to the conclusion that Sally *did* study, but the exam was unduly hard—even though the two causes are not mutually exclusive. Because a plausible sufficient cause will discount the role of other causes, people often neglect the possibility of explanations that propose multiple sufficient causes. This tendency to discount causes may be normatively justified (Morris & Larrick, 1995; Griffiths, 2001).

Though much of the literature suggests that simple explanations should be preferred, there are some circumstances in which a conjunction of causes appears to make an explanation better (Tversky & Kahneman, 1983; Leddo, Abelson, & Gross, 1984; Khemlani & Johnson-Laird, 2011). Ahn, Kalish, Medin and Gelman (1995) provide one account for this observation. They suggest that when explaining an event, people seek to understand the mechanisms that relate cause and effect. For example, a good explanation

for why a car accident occurred might be that it was raining and the driver is nearsighted, even though this explanation invokes two causes that might be independently sufficient (nearsightedness and poor weather). This may be because these two causes “cohere with respect to a single mechanism” (Ahn & Bailenson, 1996)—i.e., the weather made visibility poor enough so that a nearsighted driver would have difficulty driving. When two causes appear disjointed (e.g., the driver was not wearing his glasses and the car’s brakes were worn), people may be less likely to endorse complex explanations. Though reasons for the conjunction effect in explanation vary (see McClure, 1998 for a review), it seems clear that explanations sometimes reflect a desire to make sense of an effect in its broader context, even if that means positing additional causes that make the explanation less probable.

The addition of mechanisms to an explanation may mitigate the tendency to discount causes. Fugelsang, Stein, Green, and Dunbar (2004) found that while scientists often discount anomalous data when no plausible mechanism exists to explain covariation, they are significantly less likely to do so when made aware of a plausible mechanism to explain that covariation. Other factors may also attenuate our preference for simple explanations. For instance, Johnson, Valenti, and Keil (2017) replicated the simplicity bias demonstrated by Lombrozo (2007), but found that when causes were described as stochastic (rather than deterministic), preference for simple explanations was reduced.

Thus, the prior literature is divided regarding both how and why we might prefer simple or complex explanations. Although there appears to be a tendency to prefer simplicity, this is not universal and the factors that affect this preference are unclear. We propose that one overlooked factor is that cues in an explanation serve to indicate what the explanation is trying to accomplish. When explanations merely posit causes, a simple

explanation is usually preferred on probabilistic grounds. When explanations contain mechanistic details, people may focus less on the likelihood of the causes being present and more on evaluation of the mechanisms themselves. If the mechanisms are sound, a complex explanation may be judged highly in spite of the number of causes.

2 Experiment 1: Disease

2.1 Methods

2.1.1 Participants

150 participants located in the United States were recruited from Amazon's Mechanical Turk. Ages ranged from 21 to 67 (median age 31; 41% female). An attention check question (Oppenheimer, Meyvis, & Davidenko, 2009) was used to screen participants prior to beginning the experiment². Participants who failed the attention check were not allowed to proceed with the experiment.

2.1.2 Materials and Methods

Participants were told: "Over the course of a lifetime, many people acquire a latent (harmless) viral infection called CPX. There are several different ways of contracting CPX. It can be a result of family history, high alcohol consumption, or iron deficiency. These causes differ in their strength and prevalence." They were informed that strength reflects the probability of acquiring CPX for people who have that trait, and that prevalence reflects how many people in the general population have this trait, but do not necessarily have CPX. These factors convey information about the conditional probability of contracting CPX given a trait and the base rate of a particular trait, respectively.

² Participants were prompted to report their age, but upon further reading were asked to ignore the initial prompt and instead enter a specific number in the textbox to continue.

Participants were asked to choose the best explanation for Sam's infection. A *simple* explanation stated that the infection was the result of a single cause (e.g., family history), while a *complex* explanation stated that the infection was the result of a conjunction of the other two causes (e.g., high alcohol consumption and iron deficiency). A third option allowed participants to indicate no preference between the two explanations.

A table indicated the strength and prevalence of each trait in qualitative terms (see Figure 2). One trait was high risk and medium prevalence, and the remaining two traits were medium risk and high prevalence. These combinations were chosen to plausibly equate the posterior probability of the two competing explanations. An individual with two independent medium risk traits is at higher risk than an individual with only one of those traits. The probability of an individual having two high prevalence traits is less than the probability of having only one of those traits. In other words, an individual with two medium risk traits could conceivably be at high risk of contracting the infection and there could conceivably be a medium prevalence of individuals with a conjunction of two high prevalence traits. The traits were counterbalanced across participants so that any given trait could appear as the single high risk/medium prevalence trait. This trait was always presented as the simple explanation.

Trait	Strength	Prevalence
Family history	HIGH RISK	MEDIUM
High alcohol consumption	MEDIUM RISK	HIGH
Iron deficiency	MEDIUM RISK	HIGH

Figure 2. Table of probabilistic information shown to participants in Experiment 1. All traits were counterbalanced.

Participants were randomly assigned to either the mechanism condition or the no mechanism condition. In the no mechanism condition, participants were told: “Sam was just diagnosed with CPX. Which of the following do you think is the best explanation?” A forced choice was presented, e.g.:

- (1) Sam’s CPX was caused by family history
- (2) Sam’s CPX was caused by iron deficiency and high alcohol consumption
- (3) Both of the above explanations are equally good

Participants in the mechanism condition were told: “Sam was just diagnosed with CPX. Two doctors, who are not familiar with Sam's medical history, were asked to explain why Sam might have developed CPX. Which of the following do you think is the best explanation?”

- (1) Sam has a family history of CPX. People with a family history of CPX have a specific gene which increases their susceptibility to CPX.
- (2) Sam drinks too much alcohol. Alcohol lowers the immune system which allows pathogens to attack healthy cells. He also has an iron deficiency, which interferes with his body's normal ability to conduct enzyme reactions.
- (3) Both of the above explanations are equally good.

The ordering of the simple and complex explanations was counterbalanced between participants.

After choosing their preferred explanation, participants answered two additional questions about their subjective interpretation of the risk and prevalence of the individuals described in each explanation³. Each question was presented on a separate page with the

³ A third question, a one-item variant of the cognitive reflection test (Frederick, 2005), was presented in this experiment and in Experiments 2-4. Responses to this question did not predict explanation preference in any of the four experiments.

table in Figure 2 always visible. To probe the participants' evaluation of *prevalence*, participants were asked: "Which is more likely?"

- (1) A randomly selected individual has a family history of CPX
- (2) A randomly selected individual consumes a lot of alcohol and has an iron deficiency
- (3) Both of these are about equally likely

To probe participants' evaluation of *risk*, participants were asked: "Who has the higher risk of contracting CPX?"

- (1) Someone who has a family history of CPX
- (2) Someone who consumes a lot of alcohol and has an iron deficiency
- (3) These two individuals have about the same risk of contracting CPX

2.2 Results

2.2.1 Main results

The majority of participants (55%) expressed no preference between the explanations, regardless of condition (see Figure 3). Subsequent analyses examine the remaining participants to reveal whether or not there is a preference for either the simple or complex explanation.

Figure 3. Shown are the proportion of participants in each condition who preferred the simple explanation, complex explanation, or expressed no preference.

A chi-square test of independence revealed that explanation preference differed by mechanism condition, $\chi^2(1, N = 68) = 5.13, p = .024, \phi = .27$. A chi-square goodness-of-fit test indicates that participants in the no mechanism condition preferred simple explanations more than complex ones, $\chi^2(1, N = 31) = 5.45, p = .02, \phi = .42$. In contrast, participants in the mechanism condition showed a preference for complex explanations, but this result was not significant, $\chi^2(1, N = 37) = 1.32, p = .25, \phi = .19$.

Collapsed across conditions, a chi-square goodness-of-fit test indicated no general preference for complex or simple explanations ($p = .47, \phi = .09$). Preference for simple or complex explanations varied depending on the particular trait invoked in the simple explanation (see Figure 4), indicating that participants likely used prior knowledge when evaluating the explanations. However regardless of the particular causes in each

explanation, the inclusion of mechanisms consistently resulted in an increased preference for complex explanations relative to simple ones.

Figure 4. Explanation preferences separated by which cause was invoked in the simple explanation. While this factor affected explanation preference, a shift towards preference for the complex explanation was seen across all three conditions.

2.2.2 Risk and prevalence

Chi-square tests of independence revealed that participants' subjective assessment of risk ($p = .33$) and prevalence ($p = .85$) did not differ by mechanism condition. That is, although mechanisms altered participants' explanation preference, they did not influence participants' interpretation of the relevant probabilities.

Linear regression was used to determine whether explanation preference was predicted by subjective assessments of risk or prevalence. Responses on each variable were coded on three levels: -1 for favoring the simple explanation (i.e., a person with a

single trait is at higher risk/is more prevalent), 0 for no preference, and 1 for favoring the complex explanation. Overall, explanation preference was predicted by a subject's evaluation of risk information ($\beta = .25, p < .001$) and to a lesser extent (though not significant) by prevalence information ($\beta = .11, p = .11$). This suggests that explanation preferences aligned more with participants' judgments of conditional probability, i.e. how well the effect is predicted given a set of causes (Pacer, Williams, Chen, Lombrozo, & Griffiths, 2013), rather than judgments of prevalence, echoing the familiar neglect of base rates (Kahneman & Tversky, 1973).

Broken down by mechanism condition, risk and prevalence information were better at predicting explanation preference in the no mechanism condition ($R^2 = .23$) compared to the mechanism condition ($R^2 = .07$). Thus, it seems likely that participants in the mechanism condition relied less on probabilistic information presented in the table and more on the description of the mechanisms when choosing their preferred explanation.

2.3 Discussion

In Experiment 1, we discovered a *mechanism bias*, in which the inclusion of mechanistic information had a noticeable effect on explanation preferences. When presented explanations without mechanisms, participants favored simple explanations (i.e., fewer causes), corroborating some previous findings (Lombrozo 2007; Johnson, Valenti & Keil, 2017). This preference reversed when explanations contained mechanisms. Specifically, mechanisms increased preference for complex explanations relative to simple ones. In addition, subjective assessments of risk and prevalence were less predictive of explanation preference for those in the mechanism condition, suggesting that mechanisms may lead people to rely less on assessment of probabilities when judging explanations.

Nonetheless, Experiment 1 has several potential limitations. The stimuli contained causes that participants may have prior experience with. For instance, knowledge that susceptibility to disease is often linked to genetics may have resulted in a preference for explanations that appealed to family history. It is also possible that participants relied on a simple “explanation length” heuristic to decide which explanation was better. Although the complex explanation was longer in both the no mechanism and mechanism conditions, this difference was larger (and perhaps more salient) in the mechanism condition which listed two causes and two mechanisms.

A table of probabilistic information was presented in Experiment 1 in order to understand how risk and prevalence information guide explanation preference. Explanations don’t often appear with such cues, limiting the ecological validity of the experiment. Though the table was identical for both mechanism and no mechanism conditions, we removed the table in subsequent experiments to ensure that the effect of mechanisms on explanation preference is not contingent on salient probabilistic information.

Lastly, the use of a ternary response variable resulted in a relatively small sample size for analysis, as most participants expressed no preference between explanations. In subsequent experiments, we switched to a within-subjects rating scale which asked participants to rate both simple and complex explanation on a continuous sliding scale, allowing participants to express no preference by rating both explanations equally.

3 Experiment 2: Migraines

In Experiment 2, we sought to replicate the mechanism bias while addressing the limitations of Experiment 1.

3.1 Methods

3.1.1 Participants

122 participants located in the United States were recruited from Amazon's Mechanical Turk. Ages ranged from 19 to 64 (median age 33, 51% female). As with Experiment 1, an attention check question was used to screen participants prior to beginning the experiment and participants who failed the attention check were not allowed to proceed with the experiment.

3.1.2 Materials and Methods

Participants were told: "Sally suffers from migraines. She consults with her doctor to discuss three medications that are known to be very effective at treating migraines: Doloxan, Havidol, and Metazine. She is given a prescription for one or more of these medications, and the treatment is effective at reducing her migraines."

Participants were asked to compare two explanations for why Sally's migraines improved. A simple explanation attributed her migraine relief to a single cause (Doloxan), while a complex explanation attributed her relief to two causes (Havidol and Metazine). Each explanation was rated on a sliding scale from one to 100, ranging from "Poor explanation" to "Great explanation" (midpoint "Average explanation").

Participants were randomly assigned to one of two conditions. Participants in the no mechanism condition received explanations that identified the causes of Sally's migraine relief, i.e., "Sally's migraines improved because she took Doloxan" or "Sally's migraines improved because she took Havidol and Metazine". Participants in the mechanism condition received explanations that identified the causes of Sally's migraine relief but also described the mechanisms by which those causes have an effect. To roughly equate the

length of the simple and complex explanations, additional non-mechanistic details were padded to the simple explanations (see Table 1). Each of the three drugs was associated with one of two mechanisms³. The two drugs in the complex explanation were always described as acting on different mechanisms. The pairing of each drug and mechanism was counterbalanced across participants.

Finally, participants were asked to provide justification for their ratings in an open-response format, which were coded by a research assistant.

Table 1. Each drug was associated with one of two mechanisms. Additional, non-mechanistic detail (shown here in italics) was presented for the simple explanation only. The complex explanation appealed to two drugs acting on two different mechanisms, while the simple explanation appealed to a single drug acting on only one mechanism.

[Doloxan/Havidol/Metazine] works by:

- binding to G protein-coupled receptors (GPCR), leading to an increase in GABA production. GABA is known to inhibit pain by preventing neurotransmitters from binding to pain receptors. *GABA (gamma-Aminobutyric acid) is an amino acid and the chief inhibitory neurotransmitter in the mammalian central nervous system. It plays the principal role in reducing neuronal excitability throughout the nervous system.*
- inhibiting cyclooxygenase (COX) enzymes, leading to a decrease in prostaglandin production. Prostaglandins are known to cause pain by direct

³ These mechanisms correspond to the pain relief mechanisms of opiates and NSAIDs, which act upon different biochemical pathways.

action upon nerve endings. *Prostaglandins are a group of physiologically active lipid compounds having diverse hormone-like effects in animals. Prostaglandins have been found in almost every tissue in humans and other animals. They are derived enzymatically from fatty acids.*

3.2 Results

Mean ratings appear in Figure 5. As in Experiment 1, the preference for complex over simple explanations was greater in the presence of a mechanism than in the absence of a mechanism. We conducted a two-way ANOVA with mechanism (mechanism or no mechanism) and explanation type (simple and complex) as factors. A significant interaction indicated that the relative preference for complex explanations increased in the presence of mechanisms, $F(1, 120) = 4.12, p = .045, \eta_G^2 = .008$. A main effect of explanation type indicated that *complex* explanations were rated significantly higher than simple explanations, $F(1, 120) = 17.52, p < .001, M_{\text{simple}} = 55.2, M_{\text{complex}} = 64.2, \eta_G^2 = .034$. Simple main effect analyses reveal that a preference for complex explanations is significant in the mechanism condition, $t(111.1) = 3.58, p < .001, d = .66$, but not in the no mechanism condition ($p = .33, d = .17$).

Figure 5. Mean ratings for both the complex and simple explanations (within-subject) separated by mechanism condition (between-subject).

In addition, explanations with mechanisms were rated higher than those without mechanisms, $F(1, 120) = 15.75, p < .001, M_{\text{mechanism}} = 67.6, M_{\text{nomechanism}} = 52.3, \eta_G^2 = .091$.

Simple explanations were rated higher with mechanisms, $t(119.1) = 2.46, p = .015, d = .44$, as were complex explanations, $t(106.1) = 4.56, p < .001, d = .81$.

Consistent with our hypothesis, participants in the no mechanism condition tended to justify their ratings by appealing to probability (38% of participants, compared to 12% in the mechanism condition). In contrast, participants in the mechanism condition tended to justify their choice by appealing to the explanation's overall informativeness (39%, compared to 5% in the no mechanism condition) and clarity or ease of understanding

(37%, compared to 0% in the no mechanism condition)⁴. Thus, participants seemed to use very different criteria to rate explanations depending on whether or not mechanisms were provided.

3.3 Discussion

Consistent with Experiment 1, Experiment 2 revealed a mechanism bias in which the relative preference for complex explanations over simple ones increased when those explanations contained mechanisms. This result suggests that the mechanism bias is specifically a result of *mechanistic detail*, rather than general detail or explanation length.

In a slight departure from Experiment 1, participants expressed no preference between simple and complex explanations when explanations did not contain mechanisms (rather than a preference for simple explanations). Conversely, participants did show a significant preference for complex explanations when explanations did contain mechanisms—a result that was not significant in Experiment 1 (though in the same direction). Presumably the content of an explanation plays a role in one’s “baseline” preference for an explanation without mechanisms. One possibility is that participants’ baseline preference was shifted towards preferring complex explanations in Experiment 2 because the explanations invoke intentionality (Jacob, 2014). Specifically, while it seems unlikely that an individual would coincidentally have two independent traits causing an infection (Experiment 1), it is much more plausible that an individual would intentionally take more than one medication to reduce a migraine (Experiment 2). This difference could explain why a null result is observed for explanation preference without mechanisms (as

⁴ Note that these response types were coded independently, as participants could refer to multiple criteria to justify their preference. A coding rubric is provided with the dataset; see the Supplementary Material.

opposed to a simplicity preference in Experiment 1), and why a significant preference for complex explanations is observed with mechanisms (as opposed to a null result in Experiment 1). Despite this nuance, the mechanism bias persisted across both experiments.

In the experiments presented so far, complex causes are related to their effect through a common effect structure. One question that naturally follows is whether the mechanism bias generalizes to explanations with common cause structures as well. This is tested in Experiments 3 and 4.

4 Experiment 3: Buffet

In Experiments 1 and 2 we found support for a mechanism bias in explanation preference: a preference for complex explanations is increased relative to simple ones when those explanations contain mechanisms. In Experiments 1 and 2 we focused on a common effect causal structure (Figure 6a), where either a single cause or multiple causes produce a single effect. In Experiment 3 we follow Lombrozo (2007) and test explanations that have a common cause structure (Figure 6b), in which two separate effects are caused by either a single common cause or two independent causes. Recall that for common effect explanations, multiple causes typically increase the conditional probability of the effect, which may contribute to a preference for complex explanations. Common cause scenarios offer a more stringent test of the mechanism bias, because additional causes do not typically increase the conditional probability of the effect. Thus, there is no reason to believe that multiple cause explanations should be preferred on the basis of probability alone.

Figure 6. (a) In Experiments 1 and 2, a single effect is either caused by a one cause or by two causes. (b) In Experiments 3 and 4, two effects are either caused by a single common cause or by two independent causes.

4.1 Methods

4.1.1 Participants

120 participants located in the United States were recruited from Amazon's Mechanical Turk. Ages ranged from 18 to 75 (median 32, 53% female). Prior to beginning the survey, participants were screened using an attention check question identical to Experiments 1 and 2.

4.1.2 Materials and Methods

Participants were told: "Sally recently visited a local buffet for dinner. Later that night she started feeling sick, experiencing two symptoms: dizziness and redness of the face. Sally knows that many of the buffet foods contain nutrients that may have caused these symptoms." All of the cause-effect relations were provided: Niacin causes dizziness and redness of the face; biotin causes redness of the face but does not causes dizziness; and thiamine causes dizziness but does not cause redness of the face.

Participants were asked to compare two explanations for why Sally is experiencing both symptoms. The simple explanation attributed her symptoms to consuming one

nutrient (niacin) that caused both symptoms, while the complex explanation attributed her symptoms to consuming two nutrients (biotin and thiamine) that each caused a single symptom. Each explanation was rated on a sliding scale from one to 100, ranging from “Poor explanation” to “Great explanation” (midpoint “Average explanation”).

Half of the participants were randomly assigned to the mechanism condition while the other half were assigned to the no mechanism condition. Participants in the no mechanism condition received explanations that merely identified the causes of Sally’s symptoms, e.g., “Sally ate food containing niacin.” Participants in the mechanism condition received explanations that identified the causes of Sally’s symptoms but also described the mechanisms by which those causes have an effect, e.g., “Sally ate food containing niacin. Niacin induces dizziness and redness of the face through dermal cells where the activation of G protein receptors increases arachidonic acid, activating prostaglandins.” Each nutrient was associated with a different mechanism, and the pairing of these mechanisms was counterbalanced across participants. See Table 2 for a list of mechanisms⁵. As with Experiment 3, participants justified their ratings in an open-response format which were coded by a research assistant.

Table 2. Each nutrient was associated with a different mechanism, and these mechanisms were counterbalanced across participants so that each mechanism could appear in the simple explanation.

[Thiamine/biotin/niacin] induces [dizziness/redness of the face/dizziness and redness of the face] through:

⁵ These mechanisms were fabricated for the purpose of this experiment. However in an open-ended prompt, no participant expressed doubt over whether the mechanisms were real.

- production of antibodies attached to mast cells, stimulating release of chemical mediators into tissue.
- a circulatory reflex involving autonomic neurons that affect the capillaries in the venous return system.
- dermal cells where the activation of G protein receptors increases arachidonic acid, activating prostaglandins.

4.2 Results

Mean explanation ratings for each condition are shown in Figure 7. We conducted a two-way ANOVA with mechanism (mechanism or no mechanism) and explanation type (simple and complex) as factors. A significant interaction revealed a mechanism bias: relative to simple explanations, complex explanations were rated higher with mechanisms, $F(1, 118) = 5.53, p = .02, \eta_G^2 = .024$. A main effect of explanation type revealed a preference for simple explanations, $F(1, 118) = 31.15, p < .001, M_{\text{simple}} = 77.1, M_{\text{complex}} = 57.8, \eta_G^2 = .124$. Follow-up analyses revealed that simple explanations were preferred over complex ones in both the mechanism, $t(117.4) = 2.46, p = .015, d = .45$, and no mechanism conditions, $t(111.9) = 5.6, p < .001, d = 1.0$.

Figure 7. Mean explanation ratings for complex and simple explanations shown for two groups of participants (mechanism and no mechanism).

There was no significant main effect of mechanism ($p = .22$, $M_{\text{mechanism}} = 69.4$, $M_{\text{nomechanism}} = 65.5$, $\eta_G^2 = .006$). Simple explanations were not rated significantly differently between conditions, ($p = .34$, $d = .18$), but complex explanations were higher with mechanisms compared to without mechanisms, $t(115.5) = 2.38$, $p = .019$, $d = .43$.

While participants in both the mechanism and no mechanism condition justified their responses by appealing to probabilistic information (47% cf. 42%, respectively), only participants in the mechanism condition appealed to informativeness (18%) or clarity/ease of understanding (12%).

4.3 Discussion

The results from Experiment 3 corroborated a key finding in Experiment 1 and 2 that participants display a mechanism bias, preferring complex explanations more relative

to simple ones when those explanations contained mechanisms. The results of Experiment 3 contrast most with Experiment 2 in that participants showed a general preference for simple explanation, rather than complex ones. A plausible reason for this is the difference in causal structures used in each experiment. This main effect of a simplicity preference for common cause structures mirrors previous findings in the literature (Lombrozo, 2007).

All else being equal, a conjunction of independent causes is less likely than a single cause regardless of causal structure and this may foster a preference for simple explanations. However in a common effect structure (Experiments 1 and 2), a conjunction of causes typically increases the conditional probability of the effect given the causes; that is, multiple causes are typically additive (e.g., Pearl, 1988). This is not the case with a common cause structure (Experiments 3 and 4), as no single effect is produced by more than one cause. As such, it is normatively reasonable to assume that simple explanations should be preferred for common cause scenarios, but not necessarily for common effect scenarios. Indeed, this is what we have found.

This also suggests a possible way to attenuate a preference for simple explanations, which we test in Experiment 4. If participants prefer simple explanations in part because they believe it is less likely that Sally consumed two nutrients (biotin and thiamine) rather than one (niacin), this preference should be reduced if it is known that Sally has consumed all three nutrients. This effectively negates any simplicity bias that is due to base rate effects.

5 Experiment 4: Buffet Round 2

In Experiment 4, we sought to replicate the finding that including a mechanism increases preference for complex explanations. Unlike Experiments 1-3, we control for the base rates of causes by making all causes present (i.e., assigning them a probability of one).

5.1 Methods

5.1.1 Participants

120 participants located in the United States were recruited from Amazon's Mechanical Turk. Ages ranged from 18 to 67 (median age 29, 48% female). Participants were screened prior to completing the experiment using an attention check question identical to the previous experiments.

5.1.2 Materials and Methods

The materials and procedure were nearly identical to Experiment 3 with a subtle but important change: Participants were told that Sally consumed foods containing all three nutrients—niacin, biotin, and thiamine. This manipulation essentially eliminates the potential difference in base rates between the two explanations. In addition, the explanations were altered to specify that a particular nutrient caused a symptom (e.g., “The niacin in Sally’s food caused her dizziness”) rather than merely assert that she consumed that nutrient.

5.2 Results

Mean explanation ratings for each condition are shown in Figure 8. We conducted a two-way ANOVA with mechanism (mechanism or no mechanism) and explanation type (simple and complex) as factors. As in previous experiments, the relative preference for complex explanations was greater in the mechanism condition (i.e., participants showed a

mechanism bias), and this difference was nearly significant, $F(1, 118) = 3.63, p = .059, \eta_G^2 = .015$.

As in Experiment 3, we observed a main effect of explanation type, with simple explanations rated higher than complex explanations, $F(1, 118) = 5.20, p = .024, M_{\text{simple}} = 70.9, M_{\text{complex}} = 63.2, \eta_G^2 = .022$. However, as predicted, this difference is noticeably smaller (40% of the original magnitude) than observed in Experiment 3 (an effect size of $\eta_G^2 = .022$ compared to $\eta_G^2 = .124$ in Experiment 3; classified as “small” and “medium” effects by Bakeman, 2005). Follow-up analyses revealed that simple explanations were preferred over complex explanations in the no mechanism condition, $t(112.5) = 2.43, p = .017, d = .45$, but unlike Experiment 3 there was no difference between ratings in the mechanism condition ($p = .67, d = .08$).

Figure 8. Mean ratings for complex and simple explanations are shown for participants in the mechanism and no mechanism conditions.

Explanation ratings were significantly higher when those explanations contained mechanisms, $F(1, 118) = 4.19, p = .043, M_{\text{mechanism}} = 70.4, M_{\text{nomechanism}} = 63.5, \eta_G^2 = .017$. As in Experiment 3, complex explanations were rated significantly higher with mechanisms compared to without mechanisms, $t(94.1) = 2.59, p = .011, d = .48$, while simple explanations were not significantly different between conditions ($p = .92, d = .02$).

5.3 Discussion

These results build on those of Experiments 1-3 showing that participants have a mechanism bias, where a preference for complex relative to simple explanations is increased when those explanations describe the mechanisms that produce the effects. Although participants generally displayed a preference for simple explanations in

Experiment 4, this preference disappeared when explanations contained mechanisms and the scenario controlled for the base rate probability of the causes.

6 Conclusions

Across four experiments, we find that the presence of mechanisms affects one's preference for simple or complex explanations. Overall, participants preferred explanations that included mechanisms (Experiments 2 and 4)⁶. When explanations do not contain mechanisms, people typically prefer simple explanations that appeal to fewer causes (Experiments 1, 3, and 4). This preference is predicted by probabilistic accounts of explanation (e.g., Jefferys & Berger, 1992) as well non-probabilistic accounts that appeal to explanatory ideals (Lombrozo, 2007). The semantic content of an explanation may still affect the subjective probabilities that accentuate or mitigate this preference (Experiment 2).

We also found strong evidence in favor of a mechanism bias when evaluating explanations. Mechanistic explanations reduced (Experiments 2, 3 and 4) or reversed (Experiment 1) the preference for simple explanations relative to complex ones. This is consistent with previous research that has shown participants are less likely to discount a cause when mechanistic information is available (Fugelsang et al., 2004). This effect appears to be driven largely by an increase in preference for complex explanations (Experiments 2, 3, and 4). We believe that mechanistic details divert attention from which causes are most likely, and instead lead to judgments about the mechanisms themselves. Because mechanisms are known to increase subjective understanding (Vasilyeva &

⁶ The use of a forced-choice dependent variable in Experiment 1 precludes evaluation of this assertion. In Experiment 3, participants rated mechanistic explanations higher, but this result was not significant.

Lombrozo, 2015), participants preferred explanations with multiple mechanisms that provide a fuller description of the causal network surrounding an effect. Participants in the mechanism condition used qualitatively different cues to rate explanations, incorporating factors such as informativeness and ease of understanding that played little or no role in rating explanations for participants in the non-mechanistic condition. While mechanistic detail is not the only factor that affects explanation preference, these results provide further insight into how people evaluate explanations and reconcile some previous findings for when complex explanations triumph over simple ones.

Acknowledgements

This project/publication was made possible through the support of a grant from The Varieties of Understanding Project at Fordham University and The John Templeton Foundation. The opinions expressed in this publication are those of the author(s) and do not necessarily reflect the views of The Varieties of Understanding Project, Fordham University, or The John Templeton Foundation. The authors would like to thank Cristina Ballarini for her research assistance.

Supplementary Material

The raw data from all four experiments is publically available at <https://osf.io/9e7q3/>. In addition, to address the “file drawer” problem, we have provided internal write-ups for two unreported experiments (slight variations of the same experiment) conducted after Experiment 1 (but before Experiments 2, 3, and 4) that guided our research hypothesis.

References

- Ahn, W. K., & Bailenson, J. (1996). Causal attribution as a search for underlying mechanisms: An explanation of the conjunction fallacy and the discounting principle. *Cognitive Psychology*, 31(1), 82-123.
- Ahn, W. K., Kalish, C. W., Medin, D. L., & Gelman, S. A. (1995). The role of covariation versus mechanism information in causal attribution. *Cognition*, 54(3), 299-352.
- Alter, A. L., Oppenheimer, D. M., & Zemla, J. C. (2010). Missing the trees for the forest: a construal level account of the illusion of explanatory depth. *Journal of Personality and Social Psychology*, 99(3), 436-451.
- Bonawitz, E. B., & Lombrozo, T. (2012). Occam's rattle: Children's use of simplicity and probability to constrain inference. *Developmental Psychology*, 48(4), 1156-1164.
- Fernbach, P. M., Rogers, T., Fox, C. R., & Sloman, S. A. (2013). Political extremism is supported by an illusion of understanding. *Psychological Science*, 24(6), 939-946.
- Frederick, S. (2005). Cognitive reflection and decision making. *The Journal of Economic Perspectives*, 19(4), 25-42.
- Fugelsang, J. A., Stein, C. B., Green, A. E., & Dunbar, K. N. (2004). Theory and data interactions of the scientific mind: Evidence from the molecular and the cognitive laboratory. *Canadian Journal of Experimental Psychology*, 58, 86-95.
- Glennan, S. (2010). Ephemeral mechanisms and historical explanation. *Erkenntnis*, 72(2), 251-266.
- Griffiths, T. L. (2001). Explaining away and the discounting principle: Generalising a normative theory of attribution. *Unpublished manuscript*. Retrieved from <http://cocosci.berkeley.edu/tom/papers/social.ps>

- Jacob, P., "Intentionality", The Stanford Encyclopedia of Philosophy (2014), Edward N. Zalta (Ed.). Retrieved from <https://plato.stanford.edu/entries/intentionality/>
- Jefferys, W. H., & Berger, J. O. (1992). Ockham's razor and Bayesian analysis. *American Scientist*, 80(1), 64-72.
- Johnson, S. G., & Ahn, W. K. (2015). Causal networks or causal islands? The representation of mechanisms and the transitivity of causal judgment. *Cognitive Science*, 39(7), 1468-1503.
- Johnson, S. G., & Ahn, W. K. (2017). Causal Mechanisms. In M. Waldmaan (Ed.), *The Oxford Handbook of Causal Reasoning*, 127-146. New York, NY: Oxford University Press.
- Johnson, S. G., & Keil, F. C. (2017). Statistical and Mechanistic Information in Evaluating Causal Claims. In Gunzelmannm, G., Howes, A., Tenbrink, T., and Davelaar, E. (Eds.) *Proceedings of the 39th Annual Meeting of the Cognitive Science Society*. (p. 618-623). Austin, TX: Cognitive Science Society.
- Johnson, S. G., Valenti, J. J., & Keil, F. C. (2017). Opponent Uses of Simplicity and Complexity in Causal Explanation. In Gunzelmannm, G., Howes, A., Tenbrink, T., and Davelaar, E. (Eds.) *Proceedings of the 39th Annual Meeting of the Cognitive Science Society*. (p. 606-611). Austin, TX: Cognitive Science Society.
- Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review*, 80(4), 237-251.
- Kelley, H. H. (1973). The processes of causal attribution. *American Psychologist*, 28(2), 107-128.
- Khemlani, S. S., & Johnson-Laird, P. N. (2011). The need to explain. *The Quarterly Journal of Experimental Psychology*, 64(11), 2276-2288.

- Khemlani, S. S., & Oppenheimer, D. M. (2011). When one model casts doubt on another: a levels-of-analysis approach to causal discounting. *Psychological Bulletin*, 137(2), 195-211.
- Leddo, J., Abelson, R. P., & Gross, P. H. (1984). Conjunctive explanations: When two reasons are better than one. *Journal of Personality and Social Psychology*, 47(5), 933.
- Lipton, P. (2004). *Inference to the best explanation*. Oxford, UK: Oxford University Press. Second edition.
- Lombrozo, T. (2007). Simplicity and probability in causal explanation. *Cognitive Psychology*, 55(3), 232-257.
- McClure, J. (1998). Discounting causes of behavior: Are two reasons better than one? *Journal of Personality and Social Psychology*, 74(1), 7-20.
- Morris, M. W., & Larrick, R. P. (1995). When one cause casts doubt on another: A normative analysis of discounting in causal attribution. *Psychological Review*, 102(2), 331-355.
- Oppenheimer, D. M., Meyvis, T., & Davidenko, N. (2009). Instructional manipulation checks: Detecting satisficing to increase statistical power. *Journal of Experimental Social Psychology*, 45(4), 867-872.
- Pacer, M., Lombrozo, T. (forthcoming). Ockham's Razor cuts to the root: Simplicity in causal explanation. Retrieved from <http://www.concepts.dreamhosters.com/wp-content/uploads/2013/11/OckhamsRazorCutsToTheRoot.pdf>
- Pacer, M., Williams, J., Chen, X., Lombrozo, T., & Griffiths, T. (2013). Evaluating computational models of explanation using human judgments. *arXiv preprint arXiv:1309.6855*.

- Park, J., & Sloman, S. A. (2013). Mechanistic beliefs determine adherence to the Markov property in causal reasoning. *Cognitive Psychology*, 67(4), 186-216.
- Pearl, J. (1988). *Probabilistic reasoning in intelligent systems*. San Francisco, CA: Morgan Kaufmann.
- Read, S. J., & Marcus-Newhall, A. (1993). Explanatory coherence in social explanations: A parallel distributed processing account. *Journal of Personality and Social Psychology*, 65(3), 429-447.
- Rogers Commission (1986). *Report of the Presidential Commission on the Space Shuttle Challenger Accident*. Washington, D.C.
- Rozenblit, L., & Keil, F. C. (2002). The misunderstood limits of folk science: An illusion of explanatory depth. *Cognitive Science*, 26, 521-562.
- Strevens, M. (2004). The causal and unification approaches to explanation unified—causally. *Noûs*, 38(1), 154-176.
- Thagard, P. R. (1978). The best explanation: Criteria for theory choice. *The Journal of Philosophy*, 75(2), 76-92.
- Thagard, P. (1989). Explanatory coherence. *Behavioral and Brain Sciences*, 12, 435-502.
- Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review*, 90, 293-315.
- Van Overwalle, F. (1998). Causal explanation as constraint satisfaction: A critique and a feedforward connectionist alternative. *Journal of Personality and Social Psychology*, 74(2), 312-328.
- Vasilyeva, N., & Lombrozo, T. (2015) Explanations and causal judgments are differentially sensitive to covariation and mechanism information. In Noelle, D. C., Dale, R.,

Warlaumont, A. S., Yoshimi, J., Matlock, T., Jennings, C. D., & Maglio, P. P. (Eds.), *Proceedings of the 37th Annual Meeting of the Cognitive Science Society*. (p. 2475-2480). Austin, TX: Cognitive Science Society.

Zemla, J. C., Sloman, S., Bechlivanidis, C., & Lagnado, D. A. (2017, in press). Evaluating everyday explanations. *Psychonomic Bulletin & Review*, 1-13.