

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Preferences for descriptiveness and co-explanation in complex explanations

Permalink

<https://escholarship.org/uc/item/08m8d3cg>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 45(45)

Authors

Hattersley, Michael

Orchinik, Reed

Ludvig, Elliot Andrew

et al.

Publication Date

2023

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Individual Preferences for descriptiveness and co-explanation in evaluating explanations

Michael Hattersley (michael.hattersley@warwick.ac.uk)

Department of Psychology, University of Warwick

Reed Orchinik (orchinik@mit.edu)

Sloan School of Management, MIT

Elliot A. Ludvig (e.ludvig@warwick.ac.uk)

Department of Psychology, University of Warwick

Rahul Bhui (rbhui@mit.edu)

Sloan School of Management, MIT

Abstract

Good explanations can be distinguished from bad ones in different ways, for instance by how much of the available information they can explain (i.e., maximise the likelihood of) the available data. Here, we consider two different components of likelihood: descriptiveness (the likelihood of the individual data points) and co-explanation (the likelihood of the specific subset of data under consideration). We consider whether people prefer explanations that are high in descriptiveness vs. co-explanation. Moreover, we consider whether people who endorse conspiracy theories prefer explanations for either quality. In a medical diagnosis task, participants make binary choices between two fictional disease variants: one higher in descriptiveness versus another higher in co-explanation. Overall, participants displayed a weak preference for descriptiveness. This preference, however, did not vary across increasing levels of descriptiveness. Moreover, such preferences were unrelated to conspiracy mentality. Thus, both explanatory virtues may play a role in the appeal of likely explanations.

Keywords: Complex Explanations; Explanatory Virtues; Descriptiveness; Co-explanation; Conspiracy Belief;

Introduction

Good explanations typically possess a number of qualities that distinguish them from bad explanations. For instance, a good hypothesis will (amongst other things) maximise the likelihood of the available data. Imagine that a colleague has been absent from the office for the past five days. The hypothesis that she has the flu makes her continued absences highly likely, whereas the hypothesis that she was stuck in traffic makes those absences unlikely. A hypothesis may explain data in different ways, however. They could maximise the likelihood of as many individual datapoints as possible, or they could maximise the likelihood of the whole dataset, over and above the individual datapoints. For example, the hypothesis that the colleague went to the cinema could fit the individual datapoints in isolation (i.e., given a cinema trip, a work absence is likely), but not the overall set of absences (i.e., given the cinema hypothesis, five consecutive absences are unlikely). By contrast, the hypothesis that the colleague was fired does not make individual absences likely, but would make prolonged absence likely. Here, we consider whether people prefer to maximise one of these two components of likelihood, referred to as descriptiveness and co-explanation

respectively (Wojtowicz & DeDeo, 2020), over the other. Specifically, we test whether people prefer explanations that are high in either descriptiveness or co-explanation.

Generally, explanations that provide high fit to the available information are more appealing to people than unlikely ones. Because people are highly motivated to make sense of their environment and maintain a sense of predictability (Hohwy, 2013; Chater & Loewenstein, 2016), hypotheses that can explain away information are potentially adaptive for minimising epistemic uncertainty. Sometimes, however, an explanation provides high model fit artificially, rather than with genuine explanatory power. In such cases, high model fit comes at the expense of other desirable explanatory virtues. For example, good explanations tend to be commensurate with background knowledge (Wojtowicz & DeDeo, 2020). For instance, the hypothesis that the absentee colleague had the flu is intuitively preferable to the hypothesis that she was abducted by aliens. Similarly, good explanations tend to be simpler than bad ones, consistent with Occam's Razor. For instance, the hypothesis that the colleague had flu is preferable to the compound hypothesis that she had both the flu and food poisoning. Indeed, people usually prefer simple explanations to complex ones (Lombrozo, 2007). Finally, good explanations also generalise well to new information, rather than being overly specific to immediate information. If the absent colleague returns to work the next day, the hypothesis that she had the flu would generalise better than the hypothesis that she was fired.

Sometimes, explanations with high model fit are appealing even when they lack in other explanatory virtues. A prominent example of this is belief in conspiracy theories. A conspiracy theory can be defined as an explanation for an event or state of the world that involves secretive groups of people engaging in collective action towards a desired and usually malevolent goal (Bale, 2007). Most adults are willing to endorse at least one conspiracy theory (Goertzel, 1994; Oliver & Wood, 2014; Walter & Drochon, 2022), despite their relatively low prior probability and their extreme complexity in comparison with simpler official accounts (Marsh, Coachys, & Kleinberg, 2022). For instance, belief that the

Apollo Moon Landings were staged is fairly common (Swami et al., 2013), even though the level of secrecy required for this to be true is untenably high, given that NASA employed over 400,000 people in 1965 (see (Grimes, 2016)). Conspiracy theories may be appealing to people because they provide high model fit to immediate data, albeit at the expense of wider generalisability (Hattersley, Brown, Michael, & Ludwig, 2022). They do so not by possessing genuine explanatory power, however, but by virtue of their extreme complexity and specificity to the present data (Grimes et al., 2016). Indeed, people rate conspiracy theories—but not official accounts—as more believable when they also rate them as more complex (Marsh et al., 2022). Thus, conspiracy theories offer an ability to maximise the likelihood of information, but at the expense of other desirable explanatory virtues such as simplicity and compatibility with background priors.

The likelihood component can be broken down into two components: descriptiveness and co-explanation (Wojtowicz & DeDeo, 2020). The descriptiveness of an explanation refers to how likely that explanation makes each of the data points at an individual level. Descriptiveness is calculated by multiplying the marginal likelihoods of each individual data point, given the explanation. Co-explanation, by contrast, refers to how well an explanation predicts the overall set of data, over and above its ability to predict individual data points (i.e., over and above its descriptiveness) (Wojtowicz & DeDeo, 2020). Co-explanation is calculated by dividing the probability of that overall set of data by the descriptiveness of that explanation (see Table 1 for formula). Thus, a highly descriptive explanation makes individual data points highly likely, whereas an explanation with high co-explanatory power makes the specific pattern of data points highly likely.

Table 1: *Formulae for calculating the Descriptiveness and Co-explanation of a Hypothesis (Wojtowicz & DeDeo, 2020).* H = hypothesis; D = data.

Descriptiveness	Coexplanation
$Desc(H) = \prod_{i=1}^n p(D_i H)$	$Coex(H) = \frac{p(D H)}{Desc(H)}$

To illustrate how descriptiveness and co-explanation work, consider a physician encountering a patient with symptoms A, B, and C, who may have either Disease X or Disease Y. Symptoms A, B, and C are each common amongst Disease X patients, but no patient has ever had all three at once. Thus, for the patient under consideration, the Disease X hypothesis is highly descriptive but lacking in co-explanation. Alternatively, symptoms A, B, or C rarely occur in isolation amongst Disease Y patients, but often occur together. In this case, the Disease Y hypothesis is low in descriptiveness, but high in co-explanation. Thus, a good explanation can maximise the likelihood of the data either by explaining as many individual data points as possible (descriptiveness), or by explaining the specific co-occurrence of all the data points (co-explanation), or both. Either of these two explanatory virtues may underlie

the appeal of good explanations. On one hand, people may prefer descriptive explanations that maximise the likelihood of as many data points as possible. People will rate a conclusion as more believable when it coheres with multiple pieces of evidence (Harris & Hahn, 2009). On the other hand, people may instead prefer explanations high in co-explanation, that connect data points together.

There could also be individual differences in preferences for descriptiveness vs. co-explanation, such as amongst people who believe in conspiracy theories. Endorsement of conspiracy beliefs is linked with a wide variety of individual difference measures, including anxiety (Grzesiak-Feldman, 2013), paranoid ideation (Darwin, Neave, & Holmes, 2011), aversion to uncertainty (Liekfett, Christ, & Becker, 2021), and less reflective thinking styles (Swami, Voracek, Stieger, Tran, & Furnham, 2014; Hattersley et al., 2022). One possibility is that people with conspiracy beliefs prefer explanations that are more descriptive. Conspiracy theories are remarkable for their focus on information not otherwise explained by official accounts (Keeley, 1999; Brotherton, 2015; Wojtowicz & DeDeo, 2020). For example, 9/11 conspiracy theorists frequently highlight the BBC’s premature report of Building 7’s collapse, even as it stood visibly intact behind the reporter (Brotherton, 2015). By contrast, people generally tend to be good at discounting outliers (Dannals & Oppenheimer, 2022). On the other hand, people with conspiracy beliefs may prefer explanations high in co-explanation. Conspiracy theories (e.g., about the Moon Landing) often suggest vast patterns that connect seemingly unrelated data points together (e.g., lighting in the video footage, odd looking shadows in photographs). Consistent with this, people with conspiracy beliefs tend to perceive illusory patterns in random stimuli, such as coin toss sequences or abstract art (van Prooijen, Douglas, & de Inocencio, 2018; Hartmann & Müller, 2023). Thus, a preference for either descriptiveness or co-explanation could be plausibly linked to conspiracy mentality.

The present work considers two questions. Firstly, do individuals generally prefer explanations that are high in descriptiveness to explanations high in co-explanation (or vice versa)? Secondly, are these preferences linked to conspiracy belief? We consider these questions using a medical-diagnosis task inspired by previous literature (e.g., (Medin, Altom, Edelson, & Freko, 1982; Lombrozo, 2007)). In such tasks, participants are shown patients with specific symptom profiles (e.g., A, B, C), and must diagnose the patient with one of two diseases (e.g., X and Y), each associated with particular symptom profiles or with a probability distribution of symptom profiles. In our task, we keep constant the likelihood of the data under consideration given either explanation under consideration in each trial. As a consequence, descriptiveness and co-explanation are inversely related. Thus, it becomes possible to directly trade off descriptiveness and co-explanation against each other. If participants have a general preference for either of these explanatory virtues, and

if such preferences vary due to conspiracy belief, then they should choose explanations high in their preferred explanatory virtue.

Method¹

Participants

172 psychology undergraduates at a large UK university completed the study in return for course credit. 6 participants were excluded from analysis for providing incomplete surveys. Another 57 were excluded for failing attention (4) or comprehension (53) checks, leaving 109 participants (11 Male, 95 Female, 1 undisclosed, 1 Non-Binary, 1 demigirl. Mean age = 18.9, SD = 0.71, 3 undisclosed). For robustness, we reran all reported analyses using all 162 participants who completed the survey and passed the attention check, the results of which are reported on OSF. The experiment was given approval by the local University Research Ethics Committee.

Materials

The experiment was conducted in Qualtrics. The medical diagnosis task consisted of 10 trials with an identical setup (see Figure 1 for an example trial). Each trial included a description of a fictional alien patient, including their symptoms, the virus for which the alien tested positive, and the two variants of that virus. Also included were two symptom profiles, one for each variant. Each profile consisted of a historical frequency distribution over all combinations of three possible symptoms. These distributions were represented graphically with one historical patient being presented as one dot. On each distribution, there were 19 dots.

Notably, the alien patient always presented with all three symptoms. The frequency distribution of symptom profiles were always constructed such that no previous patient had more than 2 symptoms. Thus, the new patient was presented to participants as the first to be encountered with all three symptoms present. This design feature brought the tradeoff between descriptiveness and co-explanation to the forefront, by holding fixed the likelihood of all three symptoms being present. This likelihood could also have been held fixed by having an equal non-zero number of three-symptom patients in each distribution. We opted against this, however, out of concern that participants would consider a three-symptom case as equally likely given either hypothesis and regress to the middle of the scale. Instead, by having the patient under consideration as the first three-symptom case, the likelihood of such a case would appear unknown.

For each distribution, descriptiveness and co-explanation were calculated post-hoc, conditional on the three-symptom patient under consideration being diagnosed with either variant of the virus. Thus, they were calculated as if the three-symptom patient was the 20th dot in the distribution, alongside the 19 that were visible to participants. In Figure 1, for example, the distribution for variant G-a shows 13 patients




Figure 1: Example distribution from our custom Medical Diagnosis Task, as shown in Qualtrics. In this trial, a "G-virus" patient with three symptoms (Wrinkled Nose, Nausea, and Tight Chest) is being diagnosed with either variant "G-a" (green) or "G-b" (pink). Each row represents a specific symptom profile, and each dot represents a patient with that symptom profile. In this example, the third row contains two green dots, which signify two "G-a" patients with Nausea and no other symptoms.

with Nausea (2 with only Nausea, 4 with Wrinkled Nose & Nausea, 6 with Nausea & Tight Chest). When including the three-symptom patient under diagnosis, this makes 14/20 patients with Nausea. The marginal probability of Nausea, therefore, is 0.7. The post-hoc descriptiveness of variant G-a is 0.175 ($\text{Desc}(G-a) = P(\text{Nausea}) * P(\text{Wrinkled Nose}) * P(\text{Tight Chest}) = 0.7 * 0.5 * 0.5 = 0.175$), 0.05 higher than that of variant G-b ($\text{Desc}(G-b) = 0.5 * 0.5 * 0.5 = 0.125$). By contrast, the co-explanation for variant G-a rounds up to 0.29 ($\text{Coex}(G-a) = P(\text{Nausea}, \text{Wrinkled Nose}, \text{Tight Chest} \mid G-a) / \text{Desc}(G-a) = 0.05 / 0.175 = 0.28571$), less than that of variant G-b ($\text{Coex}(G-b) = 0.05 / 0.125 = 0.4$).

The distributions systematically varied in descriptiveness (and, thus, co-explanation). Two sets of distributions had equal descriptiveness while there were an additional two sets for each of the following levels of difference in descriptiveness between the distributions: 0.05, 0.10, 0.15, 0.20. For each difference level, one of the two sets of distributions contained the same number of asymptomatic patients across variants, and the other varied the number. To construct the distributions, we randomly generated sets and selected the first ten that met the inclusion requirements.

To measure conspiracy mentality, we used the Generic Conspiracist Beliefs scale (GCB; Brotherton, French, & Pickering, 2013), which contains 15 conspiratorial statements (e.g., "the government uses people as patsies to hide its involvement in criminal activities").

Procedure

At the beginning of the main task, participants were asked to imagine that they were an intergalactic physician specializing in alien diseases. Participants were told that ten alien patients had been found with a novel combination of symptoms, and their help was needed to identify the correct one

¹All materials are provided at: <https://tinyurl.com/4dpzgviewonly>

of two possible variants of an infection. As a measure of task comprehension and to familiarise participants with the experiment layout, participants were given a practice trial. Specifically, participants were asked to indicate—on a 1-to-10 scale—how many patients with one of the two variants had a particular pairwise symptom profile (i.e., ‘How many cases (i.e., circles) are there of patients with the c variant, with the symptom profile: “Stuffy Nose & Itchy Feet”?’). Participants who answered incorrectly were excluded from analysis (See “Participants” section).

Next, participants completed the main tasks, with the 10 trials presented in a randomized order. Participants had to guess which of the two possible variants the patient had caught, reporting their diagnosis on a slider from 0 to 10. A response of 0 represented absolute certainty that the left variant was the correct diagnosis, and 10 represented absolute certainty that the right variant was correct, with 5 representing indifference. Responses were reverse-coded where appropriate, such that a higher numerical score always indicated a preference for the more descriptive option. The evaluation ordering was random for the two sets where there was no difference in descriptiveness. At the end of the task, participants were provided with a free text box and asked to describe any strategies that they used to make their decisions.

Participants then completed the 15 items of the GCB (Generic Conspiracist Belief) questionnaire. Responses were given on Likert scales from 1 (Strongly Disagree) to 7 (Strongly Agree), with 4 representing indifference. For extra granularity, responses were given to two decimal places. We included an extra item as an attention check (‘This question is an attention check. Please select/indicate the number “5”.’). Participants who responded incorrectly on the attention check were excluded from analysis. Item order was randomised, except the attention check which was always presented halfway through (i.e., item 8 of 15). Finally, we requested demographic information (age, gender, religious orientation, and political beliefs) and provided a full debrief.

Results

First, we evaluate whether there is a general preference for descriptiveness at the individual level. We constructed a variable for descriptiveness ratio which is simply the proportion of times the more descriptive distribution was chosen out of the 8 trials with differences in descriptiveness ($M = 0.64$, $SD = 0.30$). Figure 2 shows that, overall, participants showed a slight preference for the more descriptive distributions, as confirmed by a one-sample t-test against 0.5 (which would indicate no preference) [$t(108) = 4.65$, $p < .001$, $d = 0.45$].

Next, we examined responses to the free text question, and considered whether descriptiveness preference varied according to the participant’s strategy. We categorised participants into three groups based on their responses to the free text question (see Table 2 for a summary of categories and their frequency). The Co-occurrence category—which was the most popular—refers to participants who described a tendency to




Figure 2: *Distribution of Descriptiveness Ratio (Left) and Descriptiveness and Descriptiveness Preference Scores (Right). In both cases, high scores indicate preference for the descriptive explanation, and low scores for the co-explanatory explanation, with 0.50 (Left) and 0 (Right) representing indifference. Overall, participants showed a slight preference for descriptiveness.*

choose the variant with more instances where symptoms co-occurred (i.e., more two-symptom cases). The Symptomatic Cases category was for participants who described a tendency to choose the variant with more symptomatic (as opposed to asymptomatic) cases. Finally, the Other category included all other participants, who either included a different strategy or none at all (a full list of responses provided in the free text question is provided as supplementary material on OSF).

Table 2: *Task strategies reported by participants in the free text question, and the number of participants grouped in each.*

Strategy	Count / Percentage
Co-occurrence	45 (41.3%)
Symptomatic Cases	30 (27.5%)
Other	34 (31.2%)

To see whether these groups of participants differed in their preferences, we ran a one-way Analysis of Variance with Task Strategy as a three-level categorical variable and descriptiveness ratio as the dependent variable. There was a significant main effect of task strategy on the descriptiveness ratio [$F(2,106) = 17.34$, $p < .001$]. Figure 3 shows a strong preference for descriptiveness amongst participants who indicated the Co-occurrence strategy ($M = 0.80$, $SD = 0.22$), a mild preference amongst those who chose Other strategies ($M = 0.59$, $SD = 0.28$), and a slight preference for Co-explanation amongst those who used a Symptomatic Cases strategy ($M = 0.44$, $SD = 0.32$). Post-hoc Tukey tests confirmed that participants who used a Co-occurrence strategy had a stronger preference for descriptiveness than those in




Figure 3: *Descriptiveness Ratio (Left) and Descriptiveness Preference Score (Right) by strategy. In both cases, high scores indicate preference for the descriptive explanation, and low scores for the co-explanatory explanation, with 0.50 (Left) and 0 (Right) representing indifference. Participants using a strategy based on Co-occurrence showed the strongest preferences for descriptiveness, whereas participants using the Symptomatic Cases showed the strongest (albeit slight) preference for co-explanation.*

the Symptomatic Cases [$p < .001$] and Other categories [$p = .002$], but these latter two groups did not differ from one another [$p = .06$]. Thus, the general preference for descriptiveness seemed driven by participants who indicated they chose a variant based on the number of general co-occurrences.

Next, we explore whether this preference for descriptiveness varies between items, and whether it grows stronger as the difference in descriptiveness increases. Figure 4 plots an individual's item evaluation against the difference in descriptiveness between the two items, and shows a preference for the more descriptive option, which is highest at low levels of difference in descriptiveness. A linear regression with participant-clustered standard errors showed, however, that this negative relationship was not statistically significant [$b = -2.12, p = .17$].

Finally, we considered whether preferences for descriptiveness vs. co-explanation varied due to conspiracy belief. Responses to the 15 GCB items were averaged to produce a score for conspiracy mentality ($M = 3.75, SD = 1.07$), which had high internal reliability (Cronbach's $\alpha = 0.91$). As shown in Figure 5, conspiracy mentality was unrelated to descriptiveness ratio [$r(107) = -0.07, p = .48$]. Thus, there does not appear to be a relationship between preferences for descriptiveness and conspiracy mentality.

Discussion

We found evidence that people have a weak preference for explanations that are descriptive over those high in co-explanatory power. Moreover, this preference was consistent across all trials, regardless of the difference in descriptive-




Figure 4: *Correlation between Descriptiveness Difference and individual evaluations of each item. A higher score indicates a preference for descriptive explanations, a lower score a preference for co-explanation, and 5 representing indifference. Item evaluation did not change as a function of Descriptiveness Difference. Error band = 95% CI.*

ness, and was strongest amongst participants who made selections based on which variant had more pairwise occurrences. Finally, conspiracy mentality had no impact on these results.

Taken at face value, these results might suggest that people prefer to explain as many individual data points as possible, rather than to explain their overall co-occurrence. This interpretation is tentative, however, for two reasons. Firstly, the preference for descriptiveness was both mild and relatively consistent across all trials and was mildly higher when differences in descriptiveness were small. Given a strong preference for descriptiveness, we would expect the preference for descriptive explanations to increase across trials, as a function of the increased difference in descriptiveness between the two variants. Although we only find a weak preference, this is still compatible with existing literature. For instance, people are generally good at discounting outliers (Dannals & Oppenheimer, 2022). Thus, any preference for descriptiveness is not strong enough to compel people to explain outliers.

Secondly, we find that a large number of participants (41.3%) indicated in the free-text question that they picked the variant with the most pairwise co-occurrences. These participants showed the strongest preference for descriptiveness. As co-explanation is about the specific co-occurrence of all the data points (in this case, the co-occurrence of ABC), one might prefer explanations with high co-explanatory power because of a general preference for co-occurrence. In all our trials, however, the likelihood of all three symptoms given either variant was kept constant. Consequently, the variant with the most pairwise co-occurrences was always the more descriptive option. Recall that descriptiveness is calculated by multiplying the marginal likelihoods of all individual data points. Each pairwise case is actually double-counted, how-




Figure 5: *Correlations between Conspiracy Mentality—as measured by the Generic Conspiracist Beliefs scale (GCB; Brotherton, French, & Pickering, 2013)—and measures of preference for description. Left: GCB and Descriptiveness Ratio score. Right: GCB and Descriptiveness Preference Score. Neither correlation reached statistical significance. Error band = 95% CI.*

ever, because the likelihood of a specific pairwise combination of symptoms (e.g., P(A&B — Variant) is included in the marginal likelihoods of both symptoms individually (i.e., P(A — Variant), P(B — Variant)). Thus, if the likelihood of all three symptoms is kept constant, then whichever variant has the most cases with pairwise symptoms will always have higher descriptiveness and lower co-explanation. What this means is that an individual could pick the option that is mathematically more descriptive, for reasons that are conceptually related to co-explanation.

One possible explanation as to why the preference for descriptiveness is consistently weak across trials is that people generally value both descriptiveness and co-explanation. Recall that the likelihood of a given dataset given an explanation breaks down into both descriptiveness and co-explanation (Wojtowicz & DeDeo, 2020). If likelihood is held constant, then descriptiveness and co-explanation are inversely related. This means that in trials with a large difference in descriptiveness, the explanation with very high descriptiveness also has very low co-explanation. By contrast, in trials with minimal difference, the slightly more descriptive explanation still has some co-explanation, albeit slightly less than the other variant. People may slightly prefer descriptive explanations, but not so much that they are insensitive to low co-explanation. Thus, in trials with a greater difference between explanations in descriptiveness (and co-explanation), the low degree of one virtue may be as salient to participants as the high degree of another. As such, very low co-explanation cancels out higher descriptiveness (and vice versa). Consistent with this explanation, the preference for descriptiveness did not strengthen as the descriptiveness difference decreased, nor did it ever become a preference for co-explanation. Thus, participants

may value both qualities highly, but neither is valued at the expense of the other. Rather, both are highly valued, and descriptiveness slightly more so.

One potential methodological issue concerns the decision to fix the likelihood given the two variants in each trial. On one hand, as discussed earlier, this has consequences for the nature of the distributions, such that the variant with more general co-occurrences will have higher descriptiveness. On the other hand, the alternative (i.e., different likelihoods) is also problematic. If the data are more likely given one variant than another, then participants would just choose that variant. In theory, such a choice could still be motivated by either descriptiveness or co-explanation, because valuing either would entail valuing likelihood by extension, but it would be extremely difficult—if not impossible—to determine which one, purely on the basis of participants choosing the more likely variant. Thus, there is a potentially difficult tradeoff at play. One possible solution is to have different likelihoods, but to keep the difference in likelihoods constant across trials. In one trial, the two explanations could have equal co-explanation, but one could be higher in descriptiveness than another. That is, the difference in likelihood is entirely due to descriptiveness. In another trial, the descriptiveness could be fixed but the co-explanations could differ, and the difference in likelihood is entirely because of co-explanation. From there, one might investigate whether preference for the more likely explanation is stronger in one of these instances versus another. This possibility may be of interest to future research.

Finally, preferences did not vary as a function of conspiracy mentality. Thus, people with strong conspiracy beliefs may find explanations appealing when they maximise likelihood (see Hattersley et al., 2022), regardless of whether they do so by having high descriptiveness or co-explanation. This suggests that people with stronger conspiracy beliefs do not hold much preference either way. Another possibility is that preferences for descriptiveness over co-explanation (or vice versa) may be context specific. This is particularly interesting in relation to those with conspiracy beliefs. People with conspiracy beliefs tend to be intolerant of uncertainty (Liekfett et al., 2021) and averse to ambiguity (Piccillo & van den Hurk, 2021). In uncertain contexts, people with conspiracy beliefs may be attracted to highly descriptive hypotheses that can “explain away” outliers, or they may be drawn to co-explanatory explanations that create connections between unstructured observations (Wojtowicz & DeDeo, 2020). Consistent with this idea, stronger belief in conspiracy theories is associated with illusory pattern perception in abstract stimuli, but not structured stimuli (van Prooijen et al., 2018). The idea that context could affect explanatory preferences is interesting, given that conspiracy theories often thrive during or following impactful societal events (van Prooijen & Douglas, 2017), such as pandemics (Douglas, 2021), social regime changes (Krekó, 2019), and elections (Enders et al., 2021). Future research, therefore, may wish to consider the potential role for context on explanatory preferences.

Acknowledgments

We thank Gordon D.A. Brown for his helpful feedback throughout this project. We gratefully acknowledge financial support from the MIT Sloan School of Management, the Office of Naval Research, and the Jephtha H. and Emily V. Wade Fund.

References

- Bale, J. M. (2007). Political paranoia v. political realism: on distinguishing between bogus conspiracy theories and genuine conspiratorial politics. *Patterns of Prejudice*, 41(1), 45–60. doi: 10.1080/00313220601118751
- Brotherton, R. (2015). *Suspicious minds: why we believe conspiracy theories*. New York, NY: Bloomsbury Sigma.
- Brotherton, R., French, C. C., & Pickering, A. (2013). Measuring Belief in Conspiracy Theories: The Generic Conspiracist Beliefs Scale. *Frontiers in Psychology*, 4. doi: 10.3389/fpsyg.2013.00279
- Chater, N., & Loewenstein, G. (2016). The under-appreciated drive for sense-making. *Journal of Economic Behavior & Organization*, 126(B), 137–154. doi: 10.1016/j.jebo.2015.10.016
- Dannals, J. E., & Oppenheimer, D. M. (2022). How people deal with outliers. *Journal of Behavioral Decision Making, Early View*, e2303. doi: 10.1002/bdm.2303
- Darwin, H., Neave, N., & Holmes, J. (2011). Belief in conspiracy theories, the role of paranormal belief, paranoid ideation and schizotypy. *Personality and Individual Differences*, 50(8), 1289–1293. doi: 10.1016/j.paid.2011.02.027
- Enders, A. M., Uscinski, J. E., Klofstad, C. A., Premaratne, K., Seeling, M. I., Wuchty, S., ... Funchion, J. R. (2021). The 2020 presidential election and beliefs about fraud: continuity or change? *Electoral Studies*, 72. doi: 10.1016/j.electstud.2021.102366.
- Goertzel, T. (1994). Belief in Conspiracy Theories. *Political Psychology*, 15(4), 731–742. doi: 10.2307/3791630
- Grimes, D. R. (2016). On the Viability of Conspiratorial Beliefs. *PLOS ONE*, 11(1), e0147905. doi: 10.1371/journal.pone.0147905
- Grzesiak-Feldman, M. (2013). The effect of high-anxiety situations on conspiracy thinking. *Personality and Individual Differences*, 32(32), 100–118. doi: 10.1007/s12144-013-9165-6
- Harris, A. J., & Hahn, U. (2009). Bayesian rationality in evaluating multiple testimonies: Incorporating the role of coherence. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 35(5), 1366–1373. doi: 10.1037/a0016567
- Hartmann, M., & Müller, P. (2023). Illusory perception of visual patterns in pure noise is associated with covid-19 conspiracy beliefs. *i-Perception*, 14(1), 20416695221144732.
- Hattersley, M., Brown, G. D. A., Michael, J., & Ludvig, E. A. (2022). Of tinfoil hats and thinking caps: Reasoning is more strongly related to implausible than plausible conspiracy beliefs. *Cognition*, 218, 104956. doi: 10.1016/j.cognition.2021.104956
- Hohwy, J. (2013). *The predictive mind*. Oxford, UK: Oxford University Press.
- Keeley, B. L. (1999). Of Conspiracy Theories. *Journal of Philosophy*, 96(3), 109–126. doi: 10.2139/ssrn.1084585
- Krekó, P. (2019). “the stolen transition” –conspiracy theories in post-communist and post-democratic hungary. *Social Psychological Bulletin*, 14(4), 1–13. doi: 10.32872/spb.v14i4.2323
- Liekefett, L., Christ, O., & Becker, J. (2021). Can conspiracy beliefs be beneficial? longitudinal linkages between conspiracy beliefs, anxiety, uncertainty aversion, and existential threat. *Personality and Social Psychology Bulletin*. doi: 10.1177/0146167221060965
- Lombrozo, T. (2007). Simplicity and probability in causal explanation. *Cognitive Psychology*, 55(3), 232–257. doi: 10.1016/j.cogpsych.2006.09.006
- Marsh, J. K., Coachys, C., & Kleinberg, S. (2022). The Compelling Complexity of Conspiracy Theories. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 44(44). Retrieved from <https://escholarship.org/uc/item/0sw059jb>
- Medin, D. L., Altom, M. W., Edelson, S. M., & Freko, D. (1982). Correlated symptoms and simulated medical classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 8(1). doi: 10.1037/0278-7393.8.1.37
- Oliver, J. E., & Wood, T. J. (2014). Conspiracy Theories and the Paranoid Style(s) of Mass Opinion. *American Journal of Political Science*, 58(4), 952–966. doi: 10.1111/ajps.12084
- Piccillo, G., & van den Hurk, J. (2021). The surprising effect of social distancing on our perception: Coping with uncertainty. *Covid Economics*, 20, 132–146. doi: 10.1177/0146167221060965
- Swami, V., Pietschnig, J., Tran, U. S., Nader, I. W., Stieger, S., & Voracek, M. (2013). Lunar Lies: The Impact of Informational Framing and Individual Differences in Shaping Conspiracist Beliefs About the Moon Landings. *Applied Cognitive Psychology*, 27(1), 71–80. doi: 10.1002/acp.2873
- Swami, V., Voracek, M., Stieger, S., Tran, U. S., & Furnham, A. (2014). Analytic thinking reduces belief in conspiracy theories. *Cognition*, 133(3), 572–585. doi: 10.1016/j.cognition.2014.08.006
- van Prooijen, J.-W., & Douglas, K. M. (2017). conspiracy theories as part of history: the role of societal crisis situations. *Memory Studies*, 10(3), 323–333. doi: 10.1177/1750698017701615
- van Prooijen, J.-W., Douglas, K. M., & de Inocencio, C. (2018). Connecting the dots: illusory pattern perception predicts beliefs in conspiracies and the supernatural. *European Journal of Social Psychology*, 48(3), 320–335. doi: 10.1002/ejsp.2331

- Walter, A. S., & Drochon, H. (2022). Conspiracy Thinking in Europe and America: A Comparative Study. *Political Studies*, 70(2), 483–501. doi: 10.1177/0032321720972616
- Wojtowicz, Z., & DeDeo, S. (2020). From Probability to Consilience: How Explanatory Values Implement Bayesian Reasoning. *Trends in Cognitive Sciences*, 24(12), 981–993. doi: 10.1016/j.tics.2020.09.013