Contents lists available at ScienceDirect

Cognition

journal homepage: www.elsevier.com/locate/cognit

Short Communication

The latent scope bias: Robust and replicable

Sangeet Khemlani [a,*], Samuel G.B. Johnson [b,e,f], Daniel M. Oppenheimer [c], Abigail B. Sussman [d]

[a] Navy Center of Applied Research in Artificial Intelligence, US Naval Research Laboratory, United States
[b] Department of Psychology, University of Waterloo, United States
[c] Department of Social and Decision Sciences, Carnegie Mellon University, United States
[d] Booth School of Business, University of Chicago, United States
[e] School of Management, University of Bath, United Kingdom
[f] Department of Psychology, University of Warwick, United Kingdom

ARTICLE INFO

Keywords:
Latent scope
Explanatory reasoning
Bias
Causation
Inferred evidence

ABSTRACT

People appear to prefer explanations that minimize unobserved effects, a pattern known as the *latent scope bias* in explanatory reasoning. A recent set of studies published in *Cognition* argues that the bias can be elicited only in certain narrow conditions and with certain tasks, such as a forced-choice task (Stephan, 2023). This commentary assesses the robustness of the bias in two ways: it weighs the most recent discoveries against previous research, and it presents two new studies using the most general possible elicitation task, i.e., spontaneous written responses to problems designed to test for a latent scope bias. Across 35 previous studies, 7 studies published in Stephan (2023), and 2 new studies described herein, the overwhelming majority of studies showed that people preferred narrow latent scope explanations over broad ones. This analysis led us to conclude that the bias is both robust and replicable. Taken together, Stephan's (2023) contribution and our new analyses advance our understanding of explanatory reasoning behavior.

1. Introduction

Consider the following problem:

Two suspects are implicated in a murder, Plum and Mustard.
- *Plum's actions would cause a **dent** in the candlestick, **mud** on the drawing room carpet, and a **file** stolen from the library.*
- *Mustard's actions would cause a **dent** in the candlestick and **mud** on the drawing room carpet.*

*Evidence shows a **dent** in the candlestick, but we don't know whether or not there's **mud** or a stolen **file**.*
Which suspect(s) committed the murder?

Latent scope refers to the number of unobserved and unverified effects of a given cause: in this problem, Plum's actions can explain the mud or the stolen file, should they be discovered, whereas Mustard's actions can explain only the mud – hence, the explanation that Plum committed the crime has broader latent scope. In a series of studies, we found that people prefer explanations with *narrower* latent scope, e.g., they are more likely to think Mustard committed the crime (Johnson, Rajeev-Kumar, & Keil, 2016; Khemlani, Sussman, & Oppenheimer, 2011; Sussman, Khemlani, & Oppenheimer, 2014). While people ordinarily seem to be fluent and adept explainers, a bias toward narrow latent scope explanations defies the laws of probability.

A recent paper by Stephan (2023) described new tests of the bias. It concluded that the bias isn't as robust as had been previously claimed. Below, we address the five main concerns Stephan raised in his analysis, and explain why his results had only a modest impact on our own understanding of the bias.

1.1. Response format

Some previous demonstrations of the latent scope bias used a forced-choice format in which people had to choose either the narrow or broad explanation. Since the correct answer to the problem is that both are equally likely, this methodology could inflate or produce erroneous error rates by forcing people to commit to a response when they would otherwise choose both. Stephan (2023) presented participants with problems like the one that opens this manuscript, but varied whether they responded through a forced-choice format or through a scale that included the option to state that both explanations are equally likely. The forced-choice format produced a large bias, but a scale format

* Corresponding author at: Navy Center for Applied Research in Artificial Intelligence, US Naval Research Laboratory, Washington, DC 20375, United States.
E-mail address: sangeet.s.khemlani.civ@us.navy.mil (S. Khemlani).

https://doi.org/10.1016/j.cognition.2024.105872
Received 22 May 2024; Received in revised form 27 June 2024; Accepted 1 July 2024

eliminated (Experiment 1) or reduced the bias (Experiments 2 and 3). He concluded that latent narrow scope bias was largely an artifact of experimental design, and suggested that it is only robust under narrow prescribed circumstances, particularly in forced-choice tasks. Yet, in previous work, a variety of measures produced the narrow latent scope bias (Table A1 in the Appendix), including many using scales. These studies do not uniformly show that the forced-choice task amplifies the bias, and indeed, they demonstrate that the bias persists using a wide variety of tasks.

One reason to use many different sorts of task to test the robustness of the bias is because each task corresponds to a different kind of real-world reasoning context, including forced-choice tasks. For example, decision makers are often forced to commit to one choice or another. Even when a person *could* make a graded judgment of likelihood, they often fail to do so in practice: People often treat a likely explanation as certain when using that explanation to make further predictions (Murphy & Ross, 1994), including in the case of latent scope (Johnson, Merchant, & Keil, 2020). Hence, the choice of the task alone cannot explain away scope biases. We return to this point below.

1.2. Question wording

Some previous demonstrations asked participants which explanation was most "satisfying" rather than "likely." The "satisfaction" wording is more ambiguous; one interpretation in particular ("which explanation would have better consequences?") is problematic for stimuli such as diseases where all of the unobserved effects are undesirable. Stephan (2023) finds that the bias is larger for "satisfying" than for "likely" (Experiment 2), particularly for effects with negative valence. However, the effect of the prompt does not seem to generalize. Many previous experiments (Table A1) show narrow latent scope preferences using both phrasings; prior experiments also found biases for effects without a clear valence (Table A1).

Perhaps the most ecologically valid way of assessing the bias is to ask participants to generate responses instead of evaluating one explanation against another. A previous study did so (Khemlani et al., 2011, Experiment 1a), but the study had weaknesses which could muddy the interpretation. To test whether the response format or the question wording could diminish the bias, we ran two experiments that conceptually replicated Khemlani et al. (2011, Experiment 1a) using cleaner materials and instructions. They provided participants with materials related to diseases that could cause various symptoms, and described a patient who was exhibiting some symptoms with certainty, while other symptoms remained unconfirmed (for complete stimuli, see https://osf.io/nwkrz/). The studies asked participants to describe "What is the most satisfying" (Experiment 1) and "What is the most likely" (Experiment 2) explanation of the patient's symptoms by typing whatever they wished into a text-box interface. Responses were coded as supporting a broad scope explanation, a narrow scope explanation, or as noncommittal by independent raters blind to the manipulations in the study; their codes revealed high interrater reliability.[1] Table 1 shows the results of both studies. Participants indicated that the narrow latent scope disease best explained a patient's symptoms on the majority of responses to relevant problems, and in fact the effect was no stronger when using the "satisfying" rather than "likely" dependent measure.

1.3. Individual differences

Stephan also observes that, even if participants are biased toward latent scope explanations on average, this conclusion is qualified by large individual differences in the magnitude of the bias (see also Tsukamura, Wakai, Shimojo, & Ueda, 2022). In most of Stephan's studies, the bias was driven by a minority of participants, except when participants responded as a forced choice. His results accord with previous studies, which also found small biases with high variance (see available data: https://osf.io/x8c5u/). However, the open-ended studies reported in the present manuscript reveal that most responses aligned with the narrow scope bias (55% in Experiment 1; 62% in Experiment 2) – and 127 out of the 187 participants across both studies exhibited the bias (binomial test, $p < .001$, assuming a prior probability of 0.5). Regardless, individual differences explicate, but do not undermine, aggregate analyses. Biases can be important when they are rare but large (as a minority may commit a very large error) or small but common (as they can be amplified through channels such as binary choices).

Table 1
The results of two replication studies of Khemlani et al. (2011; Experiment 1a): Experiment 1 ($N = 92$) prompted for satisfying explanations, and Experiment 2 ($N = 95$) for likely explanations.

	Response type		
Problem type	Broad	Noncommittal	Narrow
*Experiment 1: "What is the most **satisfying** explanation …?"*			
Experimental	8	37	**55**
Control	92	6	2
*Experiment 2: "What is the most **likely** explanation …?"*			
Experimental	6	32	**62**
Control	94	4	2

The table shows percentages of different types of responses as a function of whether the problem was an experimental problem designed to test the presence of a narrow latent scope bias, or else a control problem whose only sensible response was a broad scope explanation. Natural responses could be coded as: noncommittal (e.g., responses such as "both" or "we can't be sure" or "more information is needed"); or else in support of a broad or narrow latent scope explanation. Coders were blind to the problem type and structure.

1.4. Pragmatics

Stephan (2023) argues that the observed bias may result from the inference that the absence of evidence is evidence of absence. This inference can be sensible in some contexts, but in many situations it is unfounded, such as in cases wherein evidence can be obscured, misplaced, forgotten, or not immediately available. In some of Stephan's (2023) studies, he shows that the preference for narrow latent scope is larger when the missing evidence is more "diagnosable," that is, more likely to be truly absent when it isn't observed. For example, a dented candlestick would be hard to miss, but a stolen file would not—so a rational observer would avoid an explanation that predicts an unobserved dented candlestick, but would be more agnostic about an explanation that predicts a missing file.

The very first paper on this topic considered this issue (Khemlani et al., 2011), and previous work used converging approaches to address it. For example, while 28 out of 33 of the materials used in Khemlani et al. (2011, Experiments 1a-d) had *high* diagnosability (e.g., "red bumps"), Khemlani et al. (2011, Experiment 2) used only materials with *low* diagnosability (e.g., chemicals identifiable only through specific tests). Moreover, like Stephan's studies, many previous studies have used different methods to block pragmatic inferences, including giving plausible explanations for the absence of evidence and using visual or physical stimuli (rather than verbal descriptions) where the missing evidence is occluded (see Table A1). These studies find narrow latent scope preferences, just like Stephan's Experiment 3.

1.5. Inferred evidence

One explanation for the bias is *inferred evidence*—that people rely on the base rates of the unobserved evidence to infer whether it is present or absent even when this base rate is not relevant. For example, most

[1] We thank Jeremy Fede, Kalyan Gupta, and Knexus Research Corporation for helping us conduct these studies; and we thank Berke Aydas, Jieyi Chen, Aleksaundra Handrinos, and Brett Nyman for help in coding their data.

candlesticks are not dented, so a reasoner might infer that *this* candlestick is not dented either. This typically adaptive reasoning pattern is a mistake when you know that either the broad or narrow explanation is true, and their prior probabilities are known. Stephan (2023) appears to be skeptical of the inferred evidence hypothesis as a general account of the bias, because participants did not tend to appeal to how rare an observed feature is in their justifications. Nevertheless, the hypothesis was not intended to account for people's explicit, conscious reasoning, but rather their underlying assumption patterns. These assumptions may be difficult to verbalize (see, e.g., Bear, Bensinger, Jara-Ettinger, Knobe, & Cushman, 2020; Cushman, Young, & Hauser, 2006; Nisbett & Wilson, 1977) and thus verbal reports are at best provisional evidence against potential inference patterns, particularly given experimental demonstrations of those patterns (including in Stephan's own Experiment 3).

2. Conclusion

Stephan's (2023) analysis of latent scope preferences prompted a close look at the strength of the bias. We learned from the prior literature: Table A1 taught us that there were previously 35 studies of the narrow latent scope preference; they used a range of tasks, question wordings, contents, and pragmatic pressures (or lack thereof). They demonstrated several moderators and substantial individual differences, much like any other psychological phenomenon – but they showed a consistent preference for narrow latent scope explanations. To this evidence base, we can add 7 new studies that Stephan (2023) conducted, which mostly found a narrow latent scope preference, even in a study (his Experiment 3) that controlled for all of Stephan's concerns simultaneously. We also learned from a new pool of participants: in two new studies that asked them to generate open-ended responses, participants preferred narrow latent scope explanations when we prompted for satisfying (Experiment 1) or likely explanations (Experiment 2). After reviewing 44 studies on the narrow scope bias, we learned that it survived even stringent testing — it is both replicable and robust.

CRediT authorship contribution statement

Sangeet Khemlani: Writing – review & editing, Writing – original draft, Validation, Software, Resources, Project administration, Investigation, Funding acquisition, Conceptualization, Methodology. **Samuel G.B. Johnson:** Writing – review & editing, Writing – original draft, Validation, Resources, Methodology, Investigation, Conceptualization. **Daniel M. Oppenheimer:** Writing – review & editing, Writing – original draft, Investigation, Conceptualization, Validation, Resources. **Abigail B. Sussman:** Writing – review & editing, Writing – original draft, Validation, Resources, Investigation, Conceptualization, Validation, Resources.

Data availability

The experimental materials, data, analysis scripts, and coding rubric are available online: https://osf.io/nwkrz/

Appendix A. Appendix

Below we list every paper of which we are aware that has contrasted preferences for broad versus narrow latent scope explanations. As the latent scope bias was an ancillary point to many of these papers, we list only those studies that included this contrast. In cases where a portion of a conference proceeding or dissertation was superseded by a journal article, we include each study only once.

Table A1
Experiments examining latent scope bias. Each row describes an experiment; the theme/context of the premises used; the task format and salient wording in the prompt; the valence of the effects and the framing used to describe unconfirmed evidence; and the experiment's key finding.

Experiment		Context	Measure		Latent Evidence		Key Finding
			Task format	Wording	Valence	Framing	
KSO11	1a	Fantasy	Free response	"satisfying"	Negative (magical maladies)	Unexplained	Narrow preference
	1b	Fantasy	Separate 7-point scales	"likely"	Negative (magical maladies)	Spell that reveals some, but not all, symptoms	Narrow preference
	1bR[1]	Fantasy	Separate 7-point scales	"likely"	Negative (magical maladies)	Spell that reveals some, but not all, symptoms	Narrow preference (ruled out scale distortion due to control condition)
	1c	Fantasy	Separate 7-point scales	"likely" (about latent evidence)	Negative (magical maladies)	Spell that reveals some, but not all, symptoms	No preference (unobserved effects not inferred as absent)
	1d	Fantasy	Forced choice	"likely"	Negative (magical maladies)	Spell that reveals some, but not all, symptoms	Narrow preference (base rates of causes equated)
	2	Medical	Separate 7-point scales	"likely"	Negative (abnormal blood results)	Lab has gotten backed up, so not all results have arrived	Narrow preference (no effect of cause base rates)
	3	Participant-generated	Forced choice	"satisfying"	Various	None (measured rather than manipulated)	Narrow preference (in naturalistic setting)
SKO14	1	Social categories	Forced choice	"likely"	Various (most neutral)	Unexplained	Narrow preference
	2	Monster categories	Forced choice	which food to feed monster	Neutral (bodily features)	Visually occluded	Narrow preference
JJTK14	1b	Various	Separate 9-point scales	causal strength	Various (most negative)	Unexplained	Narrow preference (not necessarily non-normative compared to control)
	2	Various	Separate 9-point scales	"satisfying"	Various (most negative)	Unexplained	Narrow preference (not necessarily non-normative compared to control)
JMK15	1	Various	Separate 11-point scales	"likely" that generalization is true	Various (most neutral)	Results not yet back from lab	Narrow preference

(*continued on next page*)

Table A1 (*continued*)

Experiment		Context	Measure		Latent Evidence		Key Finding
			Task format	Wording	Valence	Framing	
	2	Various	Separate 11-point scales	argument strength	Various (most neutral)	Unexplained	Narrow preference
J16	–	Current Events	Bipolar 11-point scale	"most probable"	Negative	"it is too early to tell"	Narrow preference only when latent effect is rare; low estimates of base rate if not explicitly provided
JRK16	1	Various	Bipolar 11-point scale	"satisfying"	Negative	Unexplained	Narrow preference (only if latent evidence is rare)
	2	Medical	Bipolar 11-point scale	"likely"	Negative	Blood test had not yet come back from the lab	Narrow preference (only if latent evidence is rare; adjusted statistically for inferences about priors, evidence independence)
	3	Medical	Bipolar 11-point scale	"likely"	Negative	Manipulated (blood test had not yet come back from lab or no explanation)	Narrow preference (no effect of evidence framing)
	4	Various	Ranking	relevance judgments of base rates	Negative	Unexplained	Base rate of latent evidence seen as most relevant
	5 A	Animal categories	Bipolar 11-point scale	which species	Neutral	Unexplained	Narrow preference (only when latent feature is rare)
	5B	Animal categories	0–100% probability scale	probability of latent feature	Neutral	Unexplained	Latent feature deemed less likely when base rate was low in other categories
	6	Various	Bipolar 11-point scale	which type	Neutral	Unexplained (variability in "diagnosability")	Narrow preference (larger for evidence with lower implicit base rates)
	7	Magical, Medical	Bipolar 11-point scale	"satisfying"	Negative (magical traces or abnormal blood results)	Manipulated various reasons (e.g., equipment failure, test does not exist)	Narrow preference (stronger when reasons preclude ever finding the evidence)
JKK16a	1	Social categories	Bipolar 11-point scale	"likely"	Positive, negative, neutral	"no one has told you whether"	Narrow preference
JZK16b	1	Products	Bipolar 11-point scale	which caused or which to buy	Various (most neutral)	Various explanations	Narrow preference for cause but not for choice or "choice-implying" causal diagnoses
	2	Products	Bipolar 11-point scale	which caused then which to buy	Various (most neutral)	Various explanations	Narrow preference for cause "locks in" narrow preference for choice
J17	33a	Mental-state inference	Bipolar 11-point scale	"best explanation"	Neutral	Various explanations	Narrow preference
	33b	Mental-state inference	Bipolar 11-point scale	"best explanation"	Neutral	Various explanations	No significant difference
	36	Products	Bipolar 11-point scale	which caused	Various (most neutral)	Various explanations	Narrow preference for cause but not "choice-implying" causal diagnoses
	37	Products	Unipolar 11-point scale for one explanation	which caused or which to buy	Various (most neutral)	Various explanations	Narrow preference for both causes and choices in separate evaluation
JJKK17	1	Magical Animal Transformations	Forced choice	which caused and which were possible	Neutral	Visually occluded	Narrow preference in 5–8-year-olds
	2	Magical Animal Transformations	Forced choice	state of latent evidence and which caused	Neutral	Visually occluded	Both causes were broad; preferred explanation that matched guess about latent evidence
	3	Physical machine toy	Forced choice	which caused	Neutral	Physically occluded	Narrow preference in 4- and 5-year-olds
	4	Physical machine toy	Forced choice	which caused	Neutral	Physically occluded	Wide preference when cause base rates favor wide
JMK20	2	Ecological systems	Forced choice	"satisfying"	Neutral	Unexplained	Narrow preference; those with this preference acted as if the narrow explanation was certain when making predictions
TWSU22	–	Medical	Separate 7-point scales	"likely"	Negative	Unexplained	Narrow preference; Bayesian model supports inferred evidence
S23	1a	Social categories	Forced choice	"likely"	Neutral	Unexplained (varied "diagnosability")	Narrow preference
	1b	Social categories	Bipolar 11-point scale	"likely"	Neutral	Unexplained (varied "diagnosability")	Narrow preference for high diagnosability but not low diagnosability
	1c	Social categories	Bipolar 11-point scale	"likely"	Neutral	Unexplained (varied "diagnosability")	Narrow preference for high diagnosability but not low diagnosability
	2a	Medical	Bipolar 11-point scale	"probable" or "satisfying"	Manipulated (negative or positive)	"don't know yet"	Marginal or null results for "probable"; narrow bias with

(*continued on next page*)

Table A1 (*continued*)

Experiment		Context	Measure		Latent Evidence		Key Finding
			Task format	Wording	Valence	Framing	
							negative (but not positive) evidence for "satisfying"
	2b	Medical	Bipolar 11-point scale	"probable" or "satisfying"	Negative	"don't know yet"	Narrow preference (larger in "satisfying" conditions)
	3	Animal categories	Bipolar 11-point scale	"probable"	Neutral	Visually occluded	Narrow preference (only when latent feature is rare)
	S[2]	Medical	Bipolar 11-point scale	"likely"	Negative	Manipulated (coffee spill on lab results or no explanation)	Narrow preference significant only when no explanation given
KJOS24	1	Fantasy	Free response	"satisfying"	Negative	Unexplained	Narrow preference
KJOS24	2	Fantasy	Free response	"likely"	Negative	Unexplained	Narrow preference

Note. Feature valence is a judgment call in some cases, since this feature was not manipulated explicitly in articles prior to S23.
KSO11 = Khemlani et al., 2011.
SKO14 = Sussman et al., 2014.
JJTK14 = Johnson, Johnston, Toig, & Keil, 2014.
JMK15 = Johnson, Merchant, & Keil, 2015.
JRK16 = Johnson, Rajeev-Kumar, & Keil, 2016.
JKK16a = Johnson, Kim, & Keil, 2016.
JZK16b = Johnson, Zhang, & Keil, 2016.
J16 = Johnson, 2016.
J17 = Johnson, 2017.
JJKK17 = Johnston, Johnson, Koven, & Keil, 2017.
JMK20 = Johnson et al., 2020.
TWSU22 = Tsukamura et al., 2022.
S23 = Stephan, 2023.
KJOS24 = Khemlani, Johnson, Oppenheimer, & Sussman, 2024 (i.e., the current article).
[1] Additional replication study documented in a footnote.
[2] Additional study reported in General Discussion.

References

Bear, A., Bensinger, S., Jara-Ettinger, J., Knobe, J., & Cushman, F. (2020). What comes to mind? *Cognition, 194*, Article 104057.

Cushman, F., Young, L., & Hauser, M. (2006). The role of conscious reasoning and intuition in moral judgment: Testing three principles of harm. *Psychological Science, 17*, 1082–1089.

Johnson, S. G. B. (2016). Explaining December 4, 2015: Cognitive science ripped from the headlines. In A. Papafragou, D. Grodner, D. Mirman, & J. C. Trueswell (Eds.), *Proceedings of the 38th Annual Conference of the Cognitive Science Society* (pp. 63–68). Austin, TX: Cognitive Science Society.

Johnson, S. G. B. (2017). *Cognition as Sense-Making*. Unpublished doctoral dissertation. Yale University.

Johnson, S. G. B., Johnston, A. M., Toig, A. E., & Keil, F. C. (2014). Explanatory scope informs causal strength inferences. In P. Bello, M. Guarini, M. McShane, & B. Scassellati (Eds.), *Proceedings of the 36th Annual Conference of the Cognitive Science Society* (pp. 2453–2458). Austin, TX: Cognitive Science Society.

Johnson, S. G. B., Kim, H. S., & Keil, F. C. (2016). Explanatory biases in social categorization. In A. Papafragou, D. Grodner, D. Mirman, & J. C. Trueswell (Eds.), *Proceedings of the 38th Annual Conference of the Cognitive Science Society* (pp. 776–781). Austin, TX: Cognitive Science Society.

Johnson, S. G. B., Merchant, T., & Keil, F. C. (2015). Argument scope in inductive reasoning: Evidence for an abductive account of induction. In D. C. Noelle, R. Dale, A. S. Warlaumont, J. Yoshimi, T. Matlock, C. D. Jennings, & P. P. Maglio (Eds.), *Proceedings of the 37th Annual Conference of the Cognitive Science Society* (pp. 1015–1020). Austin, TX: Cognitive Science Society.

Johnson, S. G. B., Merchant, T., & Keil, F. C. (2020). Belief digitization: Do we treat uncertainty as probabilities or as bits? *Journal of Experimental Psychology: General, 149*, 1417–1434.

Johnson, S. G. B., Rajeev-Kumar, G., & Keil, F. C. (2016). Sense-making under ignorance. *Cognitive Psychology, 89*, 39–70.

Johnson, S. G. B., Zhang, M., & Keil, F. C. (2016). Decision-making and biases in causal-explanatory reasoning. In A. Papafragou, D. Grodner, D. Mirman, & J. C. Trueswell (Eds.), *Proceedings of the 38th Annual Conference of the Cognitive Science Society* (pp. 1967–1972). Austin, TX: Cognitive Science Society.

Johnston, A. M., Johnson, S. G. B., Koven, M. L., & Keil, F. C. (2017). Little Bayesians or little Einsteins? Probability and explanatory virtue in children's inferences. *Developmental Science, 20*, Article e12483.

Khemlani, S., Sussman, A. B., & Oppenheimer, D. M. (2011). Harry Potter and the sorcerer's scope: Latent scope biases in explanatory reasoning. *Memory & Cognition, 39*, 527–535.

Murphy, G. L., & Ross, B. H. (1994). Predictions from uncertain categorizations. *Cognitive Psychology, 27*, 148–193.

Nisbett, R. E., & Wilson, T. D. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review, 84*, 231–259.

Stephan, S. (2023). Revisiting the narrow latent scope bias in explanatory reasoning. *Cognition, 241*, Article 105630.

Sussman, A. B., Khemlani, S., & Oppenheimer, D. M. (2014). Latent scope bias in categorization. *Journal of Experimental Social Psychology, 52*, 1–8.

Tsukamura, Y., Wakai, T., Shimojo, A., & Ueda, K. (2022). Cognitive modeling for the probabilistic reasoning of causal explanations under uncertainty. In J. Culbertson, A. Perfors, H. Rabagliati, & V. Ramenzoni (Eds.), *Proceedings of the 44th Annual Conference of the Cognitive Science Society* (pp. 818–823). Austin, TX: Cognitive Science Society.