

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

The rationality of inferring causation from correlational language

Permalink

<https://escholarship.org/uc/item/9p29w77n>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 46(0)

Authors

Lassiter, Daniel

Franke, Michael

Publication Date

2024

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

The rationality of inferring causation from correlational language

Daniel Lassiter (dan.lassiter@ed.ac.uk)

School of Philosophy, Psychology, and Language Sciences
University of Edinburgh

Michael Franke (michael.franke@uni-tuebingen.de)

Department of Linguistics
University of Tübingen

Abstract

Recent work shows that participants make asymmetric causal inferences from apparently symmetric correlational statements (e.g., “*A* is associated with *B*”). Can we make sense of this behavior in terms of rational language use? Experiment 1 investigates these interpretive preferences—what we call “PACE effects”—in light of theoretical and experimental pragmatics and psycholinguistics. We uncover several linguistic factors that influence them, suggesting that a pragmatic explanation is possible. Yet, since PACE effects do not show that correlational language leads to causal implicatures strong enough to influence action choice in practical decision contexts, Experiment 2 offers new evidence from an experiment that explicitly compares the effects of causal vs. correlational claims on decision-making. Our results suggest that causal inferences from correlation language are an intricate, but possibly rationalizable, feature of natural language understanding.

Keywords: causal inference, correlation, natural language pragmatics, rationality

Introduction

There is a well-known mantra that *correlation does not imply causation*. Statisticians and scientists have long decried the human tendency to ignore this rule (Adams et al., 2017; Huff, 1954; Seifert et al., 2022). On one popular theory, deep-seated cognitive biases cause us to hallucinate causation where only statistical conclusions are appropriate (Ariely & Jones, 2008; Tversky & Kahneman, 1982).

Taking a different perspective, causal learning is a core aspect of human cognition, crucial to reasoning, decision-making, categorization, and many other tasks (e.g. Danks, 2014; Rehder, 2017; Sloman, 2005). Given the importance of causal learning for so many other key areas of our cognition, it would be surprising if humans could not maintain the basic distinction between correlation and causation: a person unable to do this would have difficulty tying their shoes. Indeed, many theorists have argued that humans are able to combine statistical evidence with prior beliefs to draw sensible inferences about the structure and strength of causal relations (Gopnik et al., 2004; Griffiths & Tenenbaum, 2009).

The mere fact that evidence of correlation influences causal inference is not surprising or normatively problematic. By the common cause principle (Reichenbach, 1956), every non-accidental correlation is due to some sort of causal relationship. The existence of a correlation between *A* and *B* at least allows us to rule out that these variables are causally unconnected. The matter becomes more complex when we turn to

the linguistic expression of correlation. If a speaker uses an expression that *literally* conveys that *A* and *B* are correlated, do we take this to implicate a causal relation with a particular directionality? If so, is this inference rational?

Recent experimental work (Gershman & Ullman, 2023) showed that participants who are presented with a simple correlational claim, with no background information to inform priors, do not treat the various possible causal models symmetrically. Concretely, the study conducted a forced-choice experiment in which participants were presented with a correlational claim such as “*A* is associated with *B*” and were asked to choose either “*A* causes *B*” or “*B* causes *A*”. Since participants knew nothing about *A* and *B*, the reasoning above would lead us to expect purely random choice—50% for each option. Instead, the study found that the majority of participants selected “*B* causes *A*”. Strikingly, this pattern was reversed in a second experiment which used slightly different prompts, varying in details of their semantic content.

The experiments of Gershman & Ullman (2023) revealed an interesting phenomenon that we call **PACE effects** (for *P*references in *A*ssignments of *C*ause and *E*ffect). Gershman & Ullman concluded that they had found evidence of “causal implicature”, hinting that these phenomena were attributable to rationally intelligible features of natural language pragmatics (Grice, 1989). However, they did not offer a theory of how or why causal implicatures are generated. Assuming this general picture, this paper attempts to identify the detailed linguistic and pragmatic factors that generate PACE effects. In a series of studies partially replicating and extending Gershman & Ullman’s results, we show that PACE effects are influenced by grammatical, semantic, and discourse-level features of language use. We also present a second set of experiments, embedded in a practical decision-making context, that show that the use of associational language can generate causal inferences strong enough to influence action choice in concrete decision contexts. Experiment 2 indicates that causal inferences from non-causal language are not merely an artifact of the forced-choice format of the experiments.

We propose that PACE effects are related to the broader phenomenon of indirect speech acts (Clark, 1979), where a listener who has reason to think that a speaker is trying to convey causal information can use the form and meaning of the chosen sentence to infer what kind of causal information the speaker considers relevant. This account predicts the sensitiv-

ity of PACE effects to linguistic factors, and explains how it can be rational to infer causation from correlational language.

To be clear, it is not possible to show that the causal inferences studied here are *not* attributable to some sort of cognitive bias, since the hypothesis space of theories in this vein is fairly unconstrained. Rather, we steer towards an independently motivated explanation in terms of natural language pragmatics which offers an account of *why* specific linguistic and contextual factors play the role that they do.

Inferring implicit questions

Current pragmatic theory makes considerable use of the classic hypothesis that every declarative sentence is interpreted as the answer to a “Question under Discussion”, which may be explicitly posed or pragmatically inferred (Carlson, 1983; van Kuppevelt, 1995; Ginzburg, 1995; Roberts, 2012). QUD theory has been successful in explaining aspects of linguistic form and interpretation (Beaver et al., 2017). QUD theory has proved useful in psycholinguistic research, for instance by underwriting the discourse relations that are crucial for maintaining coherence and motivating many discourse-level phenomena (Kehler & Rohde, 2017). Bayesian models of language understanding have also made crucial use of QUDs to account for ways that context, plausibility, and speaker goals interact to influence pragmatic interpretation (Kao et al., 2014; Hawkins et al., 2015; Sumers et al., 2023). In these models, an (implicit or explicit) QUD determines **relevance**, which then influences pragmatic inference. For instance, Kao et al. (2014) construct a Rational Speech Act (RSA) model (Frank & Goodman, 2012) which accounts for hyperbolic interpretations of number expressions in terms of listeners jointly inferring the intended message and the QUD.

As Clark (1979) discusses, we frequently use QUD inference to make conversation more efficient by responding not to a question that a speaker has explicitly posed, but to a related question that we think practically relevant. This is a special case of bridging (Clark, 1977). Suppose a customer asks “Can I get a bottle of Kraken rum for \$8?” and the clerk responds “You can get a bottle of Black Seal rum for \$8”. The success of this conversational exchange depends on the speaker and listener’s ability to coordinate on some fairly complex inferences. The clerk’s failure to respond directly to the question, far from being uncooperative, functions to improve the efficiency of information exchange. Inferring that the customer wants to buy rum and has \$8 to spend, the clerk skips over the intermediate answer and likely follow-up question, leaving the listener to work out what is missing:

- **Customer:** Can I get a bottle of Kraken for \$8?
- Implicit: (**Clerk:** No, you can’t.)
- Implicit: (**Customer:** What kind of rum can I get for \$8?)
- **Clerk:** You can get a bottle of Black Seal for \$8.

Inferring causal questions

We hypothesize that PACE effects may be related to QUD inferences. Consider the piece of public messaging in Fig. 1,




Figure 1: Correlational language used to convey existence and direction of a causal relationship.

which is at the time of writing posted on elevator doors throughout the University of Edinburgh. The sign aims to promote a certain behavioral modification (“Take the stairs”) by providing summary evidence of a correlation. Of course, the fact that “active adults have greater cognitive function” could be explained in many ways that would not support the causal efficacy of the recommendation to take the stairs. Nevertheless, the intended message—that physical activity *causes* improved cognitive function—comes across clearly.

How do interpreters recover the implied causal relationship? Our guiding hypothesis extends Clark’s (1979) analysis of the liquor store conversation. The sign recommends taking the stairs. An astute listener would then respond “Why should I do that?”. This QUD casts the next sentence as an explanation of why one should comply with the previous sentence’s recommendation (cf. Hobbs, 1979; Asher & Lasnik, 2003; Kehler & Rohde, 2017). The correlational information in this sentence is compatible with models where **physical activity** \Rightarrow **cognitive function**, or where **cognitive function** \Rightarrow **physical activity**. However, only the former model would justify the author’s recommendation to take the stairs. Assuming the writer is attempting to efficiently convey information that would allow them to accomplish the stated persuasive goal, a listener can conclude that the evidence adduced suggests a causal relationship that does in fact support this goal. The writer’s choice to convey this message indirectly can be seen as a device to improve efficiency. Because the intended causal claim can be recovered pragmatically, citing a correlation allows the author to compress both the causal claim and its evidential basis into a single statement.

Experiment 1: PACE effects

Guided by the theoretical considerations above, we ran a series of experiments adapting Gershman & Ullman’s (2023) design by manipulating the form and content of the correlational statements provided to participants.

Experiment 1a: Grammatical manipulation

Linguistic research has shown that every sentence is associated with a **topic** determining what the sentence is intuitively “about”, with various syntactic and pragmatic effects (van Kuppevelt, 1995; Erteschik-Shir, 2007; Strawson, 1964; Reinhart, 1981). Crucially, the choice of topic places prag-

matic constraints on the QUD, because the topic must be Given in the discourse and will typically be mentioned in the QUD (Clark & Haviland, 1977; Vallduví, 2016). These discourse-level features of interpretation are reflected in the form of sentences as well. We tend to place Given before New information. In particular, grammatical subjects are default topics in English (Clark & Haviland, 1977): (1) is more naturally associated with the QUD in (2a) than the one in (2b). This is because, absent further context, we tend to treat the subject phrase “The old judge” as the sentence topic in (1).

- (1) The old judge took the bribe.
- (2) a. What did the old judge do?
b. Who took the bribe?

We hypothesized that the asymmetrical causal inferences noted by Gershman & Ullman (2023) were due to the asymmetrical grammatical form of their stimuli. Absent further context, participants tend to read “*A* is associated with *B*” as being *about* the subject *A*, and so to infer a QUD that involves this item. To explain previous results, it needs to be assumed that participants preferentially infer a QUD along the lines of “What is the cause of *A*?”. If so, a simple bridging inference takes us from “*A* is associated with *B*” to “*A* is caused by *B*”. If this reasoning is correct, we may predict that sentences with a coordinated noun-phrase as in “*A* and *B* are associated” should *not* give rise to same preference of assigning *A* as effect and *B* as cause.

Participants. We recruited 203 participants via Prolific (English native speaker, at least 10 previous studies, at least 90% approval rate). Participants were paid 0.30 British pounds for an estimated 2 minutes of work (median of actual time taken was close to one minute for all three experiments). All experiments reported in this paper were built using magpie (Franke et al., 2021).

Design, materials & procedure. In a one-shot experiment, each participant was presented with either the asymmetric “*A* is associated with *B*”, or the symmetric “*A* and *B* are associated”. Variables *A* and *B* were randomly filled with non-words used by Gershman & Ullman (2023): “Pneuben”, “Themaglin”, “Rebosen”, “Denoden”, “Agoriv”, “Flembers”, and “Ceflar”. Participants chose between “*A* causes *B*” and “*B* causes *A*”, with the order of choices randomized. Once they had made a choice, they filled in a brief demographic questionnaire and the experiment ended.

Predictions. We expected that the asymmetric form would replicate the main result of the previous study. Since the symmetric form does not provide clues suggesting that one item is Given, or that the sentence is “about” one rather than the other, we predicted that the tendency to associate *B* as the cause should be absent, or at least reduced.

Results. Data from our asymmetric condition indeed replicated previous findings: a majority (68%) chose *B* as the cause. This was significantly different from the chance level of .5 ($p < .001$ in a two-sided binomial test with $N = 94, k = 64$). By contrast, in the symmetric condition only 30%

chose this response. This was also significantly different from chance ($p < 10^{-4}$ with $N = 109, k = 33$). The overall pattern of results in Exp.1 was also highly significant in a χ^2 test ($\chi^2(1) = 27.4, p < 10^{-6}$).

Discussion. The results are consistent with our hypothesis that PACE effects are conditioned by discourse-level factors that relate to grammatical form. However, while consistent with our *ex ante* prediction, we did not anticipate the pronounced preference for “first-causes-second” in the symmetric condition. This result suggests that further factors involving sentence form also play a role. For the purpose of interpreting it, the key question is whether the factors relate to the predicate “are associated”, the conjoined noun phrase subject “*A* and *B*”, or an interaction between the two.

Results of a corpus study by Benor & Levy (2006) suggest that the preference in our “*A* and *B*” condition may be due to features of the conjoined subject, independent of the predicate “are associated”. Benor & Levy (2006) looked at ordering preferences in binomial constructions: e.g., “salt and pepper” vs. the rare and unnatural “pepper and salt”. They found that an “Iconic sequencing” constraint, with “cause before effect” as a subtype, was the “strongest and most frequently active” predictor of ordering preferences in binomial constructions. This accounts for many other binomials, e.g., “principal and interest” vs. the odd “interest and principal” and —tellingly—“cause and effect” vs. “effect and cause”. One possible explanation of the observed preference, then, is that knowledge of this production preference drove participants’ interpretations of the conjoined subjects. Lacking further clues from content or sentence form, our participants may have inferred that a speaker choosing “*A* and *B*” as subject would probably choose *A* as cause and *B* as effect.

Experiment 1b: Topic manipulation

To investigate our hypothesis that effects of linguistic form are mediated by discourse factors (e.g., QUD and topicality), we constructed a variant of the asymmetric condition in Exp.1 with an overt information-structural manipulation. Specifically, we tested the effects of a leading ‘Speaking of ...’ construction, which explicitly marks out its complement as a topic.

Participants. We recruited 221 participants via Prolific, with the same conditions and payment as in Exp.1a.

Design, materials & procedure. The design was identical to Exp.1a, except that we only tested the asymmetric frame, with the “Speaking of” manipulation independent of grammatical role. Participants saw one of the following frames:

- (3) a. “Speaking of *A*, it is associated with *B*.”
b. “Speaking of *B*, *A* is associated with it.”

Predictions. If topicality is crucial and the preferred QUD for a topic is “What is its cause?”, the prediction is that (3a) should suggest that *B* is the cause, much like the original asymmetric statement. Reversely, we predict that (3b) should show the opposite pattern.

Results and discussion. Participants who saw (3a) chose *B* as the cause 54% of the time, somewhat less than in the asymmetric condition of Exp.1a, while those who saw (3b) chose it at exactly the same rate as in the asymmetric condition of Exp.1a (68%). These rates were marginally different ($\chi^2(1) = 3.5889, p = .058$). These results are not as predicted by the topicality+QUD account. It is possible, however, that the “Speaking of ...” manipulation may not have marked topicality in the intended manner.

Experiment 1c: Adding category information

Inferring QUDs depends crucially on semantic and pragmatic plausibility. In Kao et al.’s (2014) RSA model of QUD inference, listeners assume that speakers are trying to be informative. This means that listeners should expect speakers to select and respond to QUDs whose answers are not obvious or already in common ground. We expected that providing category information should influence their responses by influencing the expected informativity of different QUDs.

Participants. We recruited 413 participants via Prolific, with the same conditions and payment as in Exps.1a and 1b.

Design, materials & procedure. The design was identical to Exps. 1a and 1b except for the stimuli. Participants saw one of the following at random:

- (4) a. “The drug *A* is associated with *B*.”
- b. “*A* is associated with the drug *B*.”
- c. “The disease *A* is associated with *B*.”
- d. “*A* is associated with the disease *B*.”

Predictions. On the QUD informativity hypothesis, when *B* is topical, a question about the causes of *B* should be dis-preferred if the answer is obvious. In this case, a QUD about *B*’s effects should be preferred. We expected participants to reason as follows. The proximate cause of a drug is usually an act of volition, namely taking the drug, and so not very conversationally relevant, but the effects of a drug are often uncertain and so of greater interest. By contrast, both causes and effects of a disease are often uncertain and of interest. So, we expected participants to be more evenly split in PACE effects when they knew that one of the items was a disease. We expected that these effects would modulate the grammatical/information-structural effects discussed above.

Results. When participants were told that an item was a drug, they overwhelmingly chose that item as the cause: 88% for (4a), and 89% for (4b). When told that one item was a disease, the responses were more mixed and were strongly influenced by grammatical form. Participants who saw (4c) chose “*A* causes *B*” 17% of the time, suggesting a slightly weaker but still active preference for the basic response pattern from Exp.1a. In contrast, those who saw (4d) were evenly split in whether to treat the disease as cause or effect (48% “*A* causes *B*”). The pattern of responses differed significantly ($\chi^2(1) = 22.3, p < 10^{-5}$).

Discussion. The effects of linguistic form were strongly modulated when participants were able to map the content

onto even skeletal causal knowledge. When participants were told that one item was a drug, the preference for informative QUDs overwhelmed the effects of grammatical position. The information that one item was a disease modulated, but did not eliminate, the effect of grammatical position. As compared to the asymmetric condition of Exp.1a, the information that the subject was a disease made participants even more likely to treat it as the effect. The information that the non-subject was a disease made participants more likely to treat this item as an effect, as compared to the item in the same grammatical position in Exp.1a.

This result is consistent with the QUD inference hypothesis, but does not support it uniquely. One suggestion is that a listener might believe, as a matter of world knowledge, that drugs are effects and diseases can be causes or effects, drawing inferences in a way is unrelated to conversational pragmatics. However, this non-pragmatic explanation strikes us as unlikely. In the real world, (taking) drugs and (having) diseases consistently have both causes and effects. The difference seems to lie in the *pragmatic issue* of how interesting the causes and effects are to talk and think about. A more plausible variant of this hypothesis is that participants might have a non-QUD-related reason for thinking that speakers are more likely to *talk about* the effects of drugs than their causes—derived from direct linguistic experience, for instance. This alternative account seems plausible, and further work would be required to test it against the QUD hypothesis. In any case, the alternative would continue to support the overall picture of PACE effects as a product of pragmatic reasoning.

Exp.1 general discussion

The studies reported in Exp.1 explored several manipulations of linguistic form and content, in order to get a clearer picture of the source and nature of PACE effects. The hypotheses investigated were inspired by work in theoretical and experimental/computational pragmatics, particularly the idea that listeners enrich the interpretation of an out-of-context sentence by attempting to work out what question the speaker is trying to address—in other words, what the sentence is *for* in the conversation. We found that grammatical position and semantic content exert a large influence on PACE effects. We did not find evidence for a mediating role of a topicality manipulation using “Speaking of ...”, but we did find a preference for cause before effect in binomial (“*A* and *B*”) constructions. All of these factors deserve further attention. More broadly, the studies in Exp. 1 suggest that PACE effects are highly sensitive to linguistic factors, but largely explicable based on assumptions about goal-driven communication.

Experiment 2: Causal implicature

The previous experiments demonstrated systematic PACE effects for language expressing correlation, but PACE effects, as such, do not yet provide strong evidence for *practically relevant* causal implicatures. PACE effects show that certain constructions make one causal direction between two variables more likely than the other, but that does not necessar-

ily imply that the degree of belief in causality as such is increased. In other words, PACE effects are in principle compatible with belief changes that are not *bona fide* examples of what we would want to call causal implicatures, i.e., practically relevant increases in belief in a causal relationship. We ran two experiments to test whether a statement like “A is associated with B”, can convey or suggest a causal connection to an extent that would affect practical decision making.

Participants. We recruited 200 participants via Prolific (English Native, at least 10 previous studies, at least 90% approval rate). For technical reasons, N=198 data sets were received for Exp. 2a, and N=203 for Exp. 2b. Participants took an average of 2.69 minutes for Exp. 2a, and 2.69 minutes for Exp. 2b. They were paid 0.40 British Pounds.

Design, materials & procedure. The experiments are one-trial forced-choice designs with four between-subject conditions. Each subject supplied one data point for the critical trials. In all conditions, participants read the same cover story about colonist in space (see Fig. 2). The task is to make a decision, whether to cultivate only one plant, which is relevant for survival, or two, the second of which may or may increase the chance of a higher yield of the first. The decision-making is based on information from a science team. The four experimental conditions differ in the report of the science team, i.e., the sentence in bold in Fig. 2. The conditions and reports were:

- **association:** “a high yield of xeliherb is associated with the presence of another plant called ralocrop.”
- **intervention:** a high yield of xeliherb was obtained whenever another plant called ralocrop was cultivated as well.
- **commonCause:** a high yield of xeliherb is associated with the presence of another plant called ralocrop. But they also found that ralocrop only grows on particularly fertile grounds.
- **deniedCausation:** a high yield of xeliherb is associated with the presence of another plant called ralocrop. But there is no evidence that the cultivation of ralocrop causes a better yield of xeliherb.

Experiments differed only in the exclusion (2a) or inclusion (2b) of additional information in the context description about a potential cost of cultivating the second plant.

Predictions. The scenario leaves underdetermined how large a potential gain from cultivating both plants is. Likewise, while Exp. 2b mentions a potential cost of cultivating both plants, the precise cost is also implicit. Nevertheless, it is reasonable to assume that, all else equal, the more likely a participant considers the relevant causal relation, the more likely they would select the “both” option.

The condition of main interest is *association*. The other conditions function as reference categories. We expect different degrees of belief in the relevant causal relation (ralocrop causes a high yield of xeliherb) in different conditions. Concretely, we expect higher choice rates of the “both” option in

comparison	Exp. 2a	Exp. 2b
denC < comC	0.87, [-0.36 ; 1.46]	0.98, [0.13 ; 1.68]
comC < ass	0.78, [-0.67 ; 1.58]	0.68, [-0.64 ; 1.08]
ass < inter	0.77, [-0.89 ; 1.85]	0.52, [-0.90 ; 0.93]
denC < ass	0.96, [-0.11 ; 2.08]	0.99, [0.22 ; 1.94]
denC < inter	1.00, [0.33 ; 2.87]	0.99, [0.29 ; 1.93]

Table 1: Full results of regression analysis.

the *association* condition than in the *deniedCausation* condition. The latter condition is intended to provide a lower bound on the strength of causal implicature. An upper bound on causal implicature strength is provided by the *intervention* condition, which is predicted to have the highest rates of the “both” choice, as it explicitly mentions effects observed after intervention. The *commonCause* condition makes only a very mild, suggestive gesture to a potential common cause of a high yield (fertile ground), and may thus be expected to lie in between the *deniedCausation* and the *association* condition. Since Exp. 2b additionally makes more salient a potential cost, we expect that choice rates for the “both” option are generally lower.

Results. Fig. 3 shows the proportions of “both” choices. We find that participants respond with a choice option indicating an increased degree of belief in a causal relation more often for the *association* condition than for the *deniedCausation* condition. However, in a Bayesian logistic regression model this contrast is only credible for Exp. 2b, i.e., when the cost of additionally cultivating another plant are salient. The posterior probability for contrast *deniedCausation* < *association* was 0.96 with 95% credible interval of difference [-0.11; 2.08] for Exp. 2a, and 0.99 ([0.26; 1.98]) for Exp.2b. Moreover, the contrast *association* < *intervention* is not credible in either experiment (0.76, [-0.82; 1.95] for Exp. 2a, and 0.52, [-0.92; 0.92] for Exp. 2b). The full results of the regression analyses are in Table 1.

Discussion. Exp.2 suggests that statements of association can influence practical decision making to a similar extent as descriptions of intervention effects. It could be objected that the data are in principle compatible with the idea that mere mentioning of “ralocrop” alone increases belief in a potential causal connection, and that a statement of association, while not actually *increasing* beliefs in a causal connection, merely only *decreases* beliefs to a certain extent. Nevertheless, the studies show a behavioral effect that hints at different degrees of causal beliefs triggered by different linguistic expressions, some with a clearly non-causal literal meaning.

Conclusion

People do draw causal inferences from correlational language. Is this due to a cognitive error, or a rational feature of language understanding? In two sets of experiments we provided support for a picture in which it is appropriate for lis-

Context: You are leading a group of colonists to a far away planet called Xelifan-3. To survive on Xelifan-3, the colonists need constant supply of the seeds of a plant called xeliherb, which grows sparsely and only on Xelifan-3. You will need to cultivate xeliherb to ensure survival of the colonists. Your science team have explored the planet and found that ...

... a high yield of xeliherb is associated with the presence of another plant called ralocrop.

However, the cultivation of ralocrop is costly (water, energy resources).

Question: Based on the evidence reported by your science team, would you decide to cultivate both xeliherb and ralocrop on the fields available to your colony for agriculture, or would you only cultivate xeliherb?

BOTH XELIHERB AND RALOCROP ONLY XELIHERB

Figure 2: Screenshot from main trial of Exp. 2b. For Exp. 2a the sentence “However, the cultivation ...” was omitted.




Figure 3: Average “both” choices for different conditions in Experiments 2a and 2b. Error bars are bootstrapped confidence intervals for the mean.

teners to draw rich, context-dependent causal inferences from correlational statements. We related the details of the context-dependence to details of the linguistic form and meaning of sentences, the discourse context in which they occur, and the practical decision context surrounding the discourse. These results put pressure on accounts based solely on cognitive biases. This is not because the latter could not account for the results, but because a rational pragmatic account can draw on independently motivated resources for explaining the detailed linguistic and pragmatic patterns observed.

Building on a preliminary hypothesis drawn from discourse pragmatics and psycholinguistics, the studies in Exp. 1 provided evidence that PACE effects (Preferences for Assigning Cause and Effect) are highly sensitive to the grammatical form and semantic content of sentences. We suggested that both factors can be related to discourse interpretation, in particular the effects of implicit QUDs and the contextual constraints on QUD inference. While a topic manipulation using “Speaking of ...” did not confirm strong predictions, we found more speculative evidence for order effects in binomial (“A and B”) constructions.

Exp. 2 provided preliminary evidence that associational language can convey causal information that is strong enough to influence decision-making. This result may be explained by the fact that the scenario we constructed involves a communicative interaction embedded in a practical decision context: participants can reasonably expect that the science team would only provide correlational information if it is practically relevant to the decision at hand. If so, the causal inferences that participants made in Exp. 2 are similar in nature to the intuitive causal interpretation of the public messaging in Fig.1. In both cases, only one direction of intended causal influence would be sufficient to rationalize a speaker’s choice to provide this particular piece of correlational evidence.

While correlation does not imply causation, correlational *statements* can. There is virtually no previous linguistic work on the correlation vs. causation issue, and the research reported here generates more questions than answers. However, it does point the way toward a more detailed understanding of the linguistic and contextual factors that influence the rational inference of causation from correlational language.

References

- Adams, R. C., Sumner, P., Vivian-Griffiths, S., Barrington, A., Williams, A., Boivin, J., ... Bott, L. (2017). How readers understand causal and correlational expressions used in news headlines. *Journal of experimental psychology: applied*, 23(1), 1.
- Ariely, D., & Jones, S. (2008). *Predictably irrational*. HarperCollins New York.
- Asher, N., & Lascarides, A. (2003). *Logics of conversation*. Cambridge University Press.
- Beaver, D. I., Roberts, C., Simons, M., & Tonhauser, J. (2017). Questions under discussion: Where information structure meets projective content. *Annual Review of Linguistics*, 3, 265–284.
- Benor, S. B., & Levy, R. (2006). The chicken or the egg? A probabilistic analysis of english binomials. *Language*, 233–278.
- Carlson, L. (1983). *Dialogue games: An approach to discourse analysis*. Springer.
- Clark, H. H. (1977). Bridging. In P. N. Johnson-Laird & P. C. Wason (Eds.), *Thinking*. Cambridge University Press.
- Clark, H. H. (1979). Responding to indirect speech acts. *Cognitive psychology*, 11(4), 430–477.
- Clark, H. H., & Haviland, S. E. (1977). Comprehension and the given-new contract. In R. O. Freedle (Ed.), *Discourse production and comprehension*. Ablex Publishing Company.
- Danks, D. (2014). *Unifying the mind: Cognitive representations as graphical models*. MIT Press.
- Erteschik-Shir, N. (2007). *Information structure: The syntax-discourse interface* (Vol. 3). OUP Oxford.
- Frank, M., & Goodman, N. (2012). Predicting pragmatic reasoning in language games. *Science*, 336(6084), 998.
- Franke, M., Ji, X., Ilieva, S., Rautenstrauch, J., & Klehr, M. (2021). *magpie: Minimal architecture for the generation of portable interactive experiments*. Retrieved from <https://magpie-experiments.org/>
- Gershman, S. J., & Ullman, T. D. (2023). Causal implicatures from correlational statements. *Plos One*, 18(5), e0286067.
- Ginzburg, J. (1995). Resolving questions, I. *Linguistics and Philosophy*, 18(5), 459–527.
- Gopnik, A., Glymour, C., Sobel, D. M., Schulz, L. E., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: Causal maps and bayes nets. *Psychological Review*, 111(1), 3.
- Grice, H. P. (1989). *Studies in the Way of Words*. Harvard University Press.
- Griffiths, T. L., & Tenenbaum, J. B. (2009). Theory-based causal induction. *Psychological review*, 116(4), 661.
- Hawkins, R. X., Stuhlmüller, A., Degen, J., & Goodman, N. D. (2015). Why do you ask? good questions provoke informative answers. In *Cogsci 2015*.
- Hobbs, J. R. (1979). Coherence and coreference. *Cognitive science*, 3(1), 67–90.
- Huff, D. (1954). *How to lie with statistics*. Norton.
- Kao, J. T., Wu, J. Y., Bergen, L., & Goodman, N. D. (2014). Nonliteral understanding of number words. *Proceedings of the National Academy of Sciences*, 111(33), 12002–12007.
- Kehler, A., & Rohde, H. (2017). Evaluating an expectation-driven question-under-discussion model of discourse interpretation. *Discourse Processes*, 54(3), 219–238.
- Rehder, B. (2017). Concepts as causal models: Categorization. In M. R. Waldmann (Ed.), *The oxford handbook of causal reasoning* (pp. 347–376). Oxford University Press New York, NY.
- Reichenbach, H. (1956). *The direction of time* (Vol. 65). Univ of California Press.
- Reinhart, T. (1981). Pragmatics and linguistics: An analysis of sentence topics. *Philosophica*, 27.
- Roberts, C. (2012). Information structure in discourse: Towards an integrated formal theory of pragmatics. *Semantics & Pragmatics*, 5, 1–69.
- Seifert, C. M., Harrington, M., Michal, A. L., & Shah, P. (2022). Causal theory error in college students' understanding of science studies. *Cognitive Research: Principles and Implications*, 7(1), 4.
- Sloman, S. A. (2005). *Causal models: How we think about the world and its alternatives*. OUP.
- Strawson, P. F. (1964). Identifying reference and truth-values. *Theoria*, 30(2), 96–118.
- Sumers, T. R., Ho, M. K., Griffiths, T. L., & Hawkins, R. D. (2023). Reconciling truthfulness and relevance as epistemic and decision-theoretic utility. *Psychological Review*.
- Tversky, A., & Kahneman, D. (1982). Causal schemas in judgments under uncertainty. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases* (pp. 117–128). Cambridge University Press.
- Vallduví, E. (2016). Information structure. In M. Aloni & P. Dekker (Eds.), *Cambridge handbook of formal semantics*. Cambridge University Press.
- van Kuppevelt, J. (1995). Discourse structure, topicality and questioning. *Journal of Linguistics*, 31(01), 109–147.