When Correlation means Causation: Pragmatic Factors modulate Causal Implicatures in Decision-Making Contexts

Hening Wang[1], **Daniel Lassiter**[2] **& Michael Franke**[1]
[1]Department of General Linguistics, University of Tübingen , {hening.wang, michael.franke}@uni-tuebingen.de
[2]School of Philosophy, Psychology, and Language Sciences, University of Edinburgh , dan.lassiter@ed.ac.uk

Abstract

Correlational statements (e.g., "*X* is associated with *Y*") can be interpreted as conveying a causal relationship. Recent work suggests that such interpretations may be legitimate pragmatic inferences, especially in decision-relevant contexts. Here, we investigate whether causal enrichment of correlational language affects memory recall and message passing, and how pragmatic features of the communicative context modulate it. Our findings support a pragmatic account of causal enrichment of correlational language.

Keywords: causal implicature; correlational language; pragmatic inference; decision making; belief updating

Introduction

Most "Stats 101" courses eagerly warn that *correlation does not mean causation*. Yet outside the classroom, correlational claims are rarely encountered as abstract, decontextualized statements. Instead, they appear in headlines, summaries of scientific findings, expert advice, or reports embedded in broader narratives. In such settings, interpretation is shaped not only by statistical knowledge, but also by background beliefs, assumptions about speakers' goals, and general world knowledge. In line with this, a growing body of evidence shows that people often interpret correlational statements in causal terms, even when no explicit causal language is used (Adams et al., 2017; Huff, 2023; Seifert et al., 2022). The theoretical question is what mechanism delivers this causal enrichment. There are at least two possible accounts, and they should be clearly separated.

On one account, the effect reflects stable *preferences in assignments of cause and effect* (Lassiter & Franke, 2024): systematic structural biases driven by linguistic features such as syntactic structure, topicality, and implicit-causality verb semantics (Bott & Solstad, 2014; Hartshorne, 2014). Gershman and Ullman (2023) document that participants reliably pick one causal direction over another given a bare association statement, and label such effects "causal implicatures." The mechanism in their framing, however, is closer to a structural bias than to a Gricean inference: the preference is described as automatic and linguistic rather than as a product of reasoning about a speaker's alternatives or communicative goals.

On a second account, causal enrichment of correlational language can be viewed as a genuine pragmatic inference: listeners reason about why a speaker chose an indirect (correlational) formulation in a given context, and enrich the statement with causal meaning when that enrichment makes the utterance maximally relevant to the listener's decision problem (Benz & van Rooij, 2007; Goodman & Frank, 2016; Sperber & Wilson, 1995). Following Reichenbach (1956), a correlational statement presupposes some causal structure, however indirect. When speaker and listener share the assumption that direct causation is the relevant question, the listener can read the speaker's choice to mention a correlation as a signal that causation is in view.

The two accounts make sharply different predictions for how causal enrichment should behave in communicative context. A structural bias account predicts roughly stable enrichment across speakers, contexts, and communicative goals, because the mechanism that drives it is local to the utterance. A pragmatic account predicts that enrichment varies systematically with features that matter for speaker reasoning: whether the message was directly addressed to the listener or merely overheard; whether the speaker's goal makes correlational information decision-relevant or merely incidental; and whether the listener, once they form a causal interpretation, carries it into downstream communication.

Consistent with the pragmatic account, Lassiter and Franke (2024) provide positive evidence for this view. Their experiment introduced a fictional space-colonization scenario in which a Science Team provided information to participants facing a binary decision (see Figure 1A–B). The scenario involved a plant called *xeliherb*, important for the space colonists' survival. In the critical condition, the Science Team presents an association statement of the form (1), where *ralocrop* is another plant.

(1) A high yield of *xeliherb* is associated with the presence of *ralocrop*.

Based on this information, participants chose to cultivate *both* plants or *only xeliherb*. The experiment manipulated (between-subjects) the information in the Science Team's report, ranging from explicitly denied causation, through correlational association (as in (1)), to explicit intervention-based causal descriptions. Since cultivation of *ralocrop* may be costly, choosing *both* is rational only to the extent that the decision-maker believes in a causal effect on the yield of *xeliherb*. Evidence for *causal belief* after the association statement comes from the observation that the rate of *both*-choices following (1) was higher than after explicit denial of a causal relationship, and indistinguishable from the interventionist

In G. Boleda, I. Dautriche, M. Maldonado, C. Saldana and T. Torrent (Eds.), *Proceedings of the 48th Annual Conference of the Cognitive Science Society*

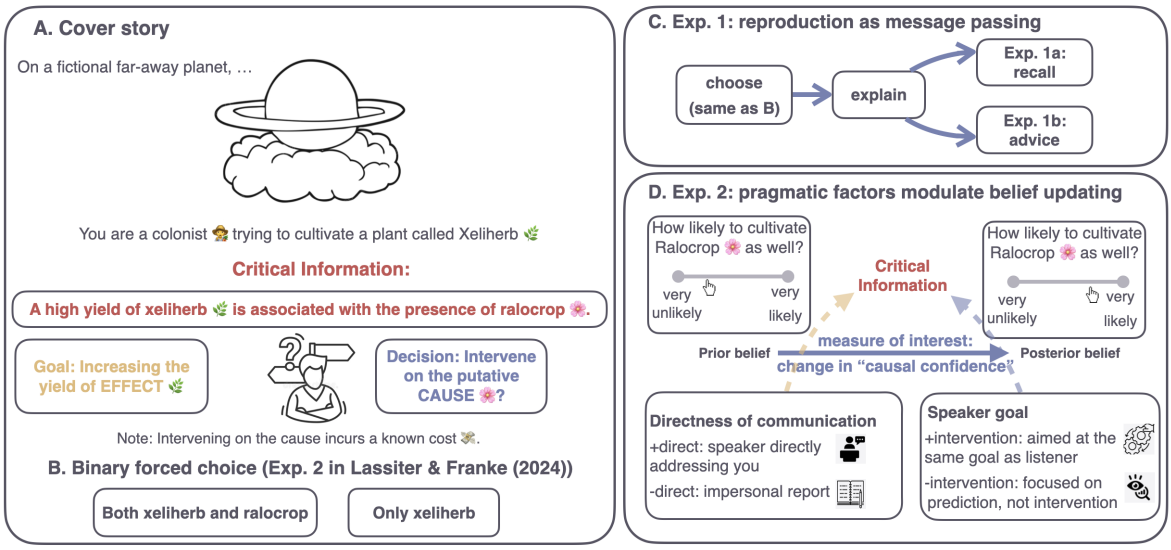

Figure 1: Overview of the experimental paradigm across studies. **(A) Cover story and critical information:** Participants were placed in a fictional space-colonization scenario in which they aimed to increase the yield of a target plant (*xeliherb*). They received correlational information stating that a high yield of *xeliherb* is associated with the presence of another plant (*ralocrop*). **(B) Binary forced-choice task:** Participants chose whether to cultivate both plants (interpreted as intervening on the putative cause) or only *xeliherb*. **(C) Experiment 1: message passing via free production:** After making the same choice as in (B), participants either recalled the information (Exp. 1a) or reproduced it as advice to another agent (Exp. 1b), allowing assessment of how correlational statements are reported from memory or communicated through message passing. **(D) Experiment 2: belief updating under manipulation of pragmatic factors:** Participants provided slider ratings for their willingness to cultivate *ralocrop* before and after the critical statement, interpretable as *causal belief change*. Pragmatic factors—directness of communication and speaker goal orientation—were manipulated to test how context modulates causal belief updating.

condition.

The present paper builds on the *xeliherb*-paradigm by addressing two shortcomings of that study and testing more nuanced predictions of the pragmatic account. First, previous work has not shown that the imputed causal belief from correlational language affects anything beyond the immediate forced-choice decision. Experiment 1 therefore examines how causal interpretations of correlational statements influence subsequent decisions, such as recalling the Science Team's report or advising future decision-makers. Second, the between-subjects forced-choice design provides only a coarse signal of the belief change triggered by correlational statements. Experiment 2 therefore directly measures belief updating within individuals and tests how the pragmatic factors identified above modulate causal enrichment. Across these experiments, we find that correlational statements can strengthen causal beliefs in decision-relevant contexts, that these beliefs are reflected in participants' memory reports and action-guiding communication, and that their impact depends systematically on pragmatic features of the communicative context. Crucially, we identify not only the conditions under which causal interpretations are strengthened, but also those under which they are attenuated—patterns that provide converging support for a pragmatic account of causal implicature grounded in reasoning about speakers' communicative intentions and goals.

Experiment 1

In the *xeliherb*-paradigm, choosing to cultivate *both* plants is consistent with a causal interpretation of the correlational information (see Figure 1B). Experiment 1 elicits participants' interpretations through free production after they underwent the same binary decision task as in prior work by Lassiter and Franke (2024). Experiment 1a focused on recall of the previously encountered information, targeting how participants reconstruct the message they understood themselves to have received. Experiment 1b required participants to reproduce the information as advice, thereby engaging communicative use under action-oriented goals. By comparing production patterns across these tasks, Experiment 1 examines whether participants who choose to intervene are also more likely to recall correlational statements using explicit causal language.

Participants We recruited a total of 99 participants via Prolific. All participants were native English speakers, had completed at least 10 previous studies, and had an approval rate of at least 90%; participants who had taken part in earlier pilot studies were excluded. Of these, $N = 49$ participated in Experiment 1a and $N = 50$ participated in Experiment 1b.

Participants completed a single-trial experiment lasting approximately 5 minutes and were compensated £0.90. This and all subsequent experiments offered above-average compensation according to the platform.

Procedure The procedure followed the general paradigm illustrated in Figure 1A–C. Participants first read the background story, followed by the critical correlational information. They then completed a binary forced-choice decision task in which they chose whether to cultivate *only xeliherb* (the effect) or *both xeliherb* and *ralocrop* (the effect and the putative cause). Participants were additionally asked to provide a brief justification for their choice.

Following the decision task, participants completed the free-production task. In Experiment 1a, participants were asked to recall and pass on the information they had just read. In Experiment 1b, participants were asked to reproduce the information in the form of advice intended for another agent facing a similar decision. In both sub-experiments, the produced message was framed as being passed on to a subsequent recipient within the same fictional scenario.

Materials After the decision task, participants produced a free-text message based on the information they had received. The wording of the prompt differed only with respect to the communicative goal. In Experiment 1a (recall), participants were instructed:

> *Due to atmospheric conditions and technical problems you are about to lose contact to an outpost on the planet. You have only very little time to send a short message to the colonists in this outpost before all contact breaks off for potentially a long time. The colonists in the outpost know about xeliherb and its importance. They know the costs associated with cultivation of ralocrop. What they do not know is what your Science Team reported. You cannot forward the original report. Please type what you recall from the Science Team's report into the text box, so that the other colonists can make a decision like you did on their own!*

In Experiment 1b (advice), participants were instructed:

> *There is also an outpost of your colonists on Xelifan-3. Unfortunately, due to atmospheric conditions, you can only send one message to the people in this outpost before all contact breaks off for potentially a long time. Tell the people in the outpost all they need to know about xeliherb and ralocrop, what they are good for, and how they might be related, so that they can decide what to cultivate for themselves!*

We expect that participants who interpret the correlational statement causally—as indexed by their decision to intervene—will more often reproduce the information in causal terms in both tasks.

Results Figure 2 summarizes the results from the free-production recall task in Experiment 1a and the advice task in Experiment 1b. All free-production responses were manually annotated into three categories: *causation*, *neutral*, and *none*, illustrated here with examples:

1. **None**: "Ralocrop needs a lot of water, so I would just focus on xeliherb for now."

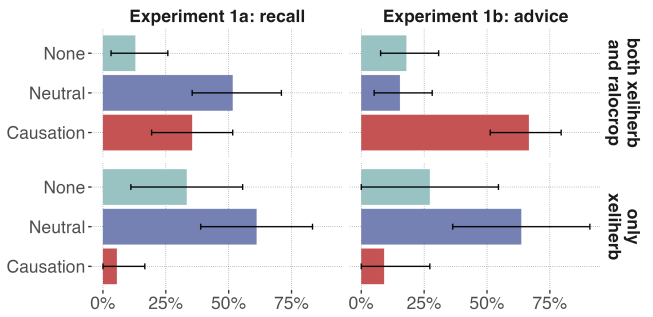

Figure 2: Proportions of annotated response categories in Experiments 1a,b, split by participants' initial choice in the decision-making trial (*both xeliherb and ralocrop* vs. *only xeliherb*). Bars are bootstrapped 95% confidence intervals.

2. **Neutral**[1]: "If ralocrop grows, then xeliherb grows too."
3. **Causation**: "Xeliherb needs ralocrop in order to grow."

For hypothesis testing, responses coded as *causation* were assigned a value of 1, and all other response categories were assigned a value of 0. We fit a generalized Bayesian regression model using `brms` (Bürkner, 2021) in R (R Core Team, 2024), including participants' initial choice (*both xeliherb and ralocrop* vs. *only xeliherb*) as a predictor. Inference was based on posterior expected predictions.

In Experiment 1a, the majority of recall responses were classified as *neutral* irrespective of participants' initial choice. However, participants who initially chose *both xeliherb and ralocrop* produced a higher proportion of *causation* responses than participants who chose *only xeliherb* (μ = 1.68, CrI [0.36, 3.08], posterior probability = .98).

In Experiment 1b, responses showed a clearer differentiation between choice groups. Participants who selected *only xeliherb* in the first decision-making task predominantly produced *neutral* responses. By contrast, participants who initially chose *both* most frequently produced *causation* responses, with a credibly higher rate than in the *only* group (μ = 2.37, CrI [1.07, 3.79], posterior probability = .99).

Discussion In both experiments, free-production responses varied systematically with participants' initial choices, indicating how correlational information was recalled and communicated under different task demands.

In Experiment 1a, recall responses largely reproduced the correlational formulation, which is expected given that the original message was itself neutral. Importantly, however, participants who chose *both* were more likely to produce causal paraphrases than those who chose *only*. One interpretation is that the recall task was not treated as a verbatim memory test, but rather as a request to convey the intended message or meaning as understood in context. On this view, the use

[1] We treat conditionals as semantically non-causal; causal readings, where they arise, are pragmatic enrichments (cf. Grusdt et al., 2022).

of causal language in recall responses reflects participants' interpretations of the correlational statement, rather than a failure to remember its surface form. The observed difference between choice groups is therefore consistent with the possibility that participants who opted to intervene interpreted the original statement as conveying causal information.

This contrast was more pronounced in Experiment 1b, where participants were asked to reproduce the information as advice for another agent. In this communicative context, participants in the *both* group most often produced explicitly causal messages, whereas participants in the *only* group largely refrained from doing so. Because the advice task requires tailoring information to another agent's presumed needs, these results suggest that causal interpretations were not only internally available but also selectively externalized when action relevance was salient.

Experiment 2

Previous experiments established that correlational statements can support causal interpretations, as reflected in binary intervention choices and in how participants reproduce and communicate information. However, these measures provide only indirect access to causal belief and do not track how strongly such beliefs are updated within individuals. In Experiment 2, we therefore move beyond categorical choices and introduce a belief-update paradigm that directly measures changes in participants' willingness to intervene relative to an individually measured baseline. This approach allows us to assess whether exposure to correlational information produces a graded shift in degrees of belief in causal power rather than merely influencing discrete responses. In addition, Experiment 2 systematically manipulates pragmatic features of the communicative context—specifically, directness of communication and speaker goal orientation—to test how reasoning about communicative intent modulates causal belief updating, as suggested by the patterns observed in Experiment 1.

Experiment 2 employed a between-participants factorial design with two manipulated factors: DIRECTNESS OF INFORMATION and SPEAKER GOAL ORIENTATION (see Figure 1D). DIRECTNESS OF INFORMATION determined whether the correlational information was presented as communicative input explicitly addressed to the participant (the *direct* condition) or encountered as impersonal evidence not tailored to a specific listener (the *indirect* condition). This manipulation varied the extent to which participants could treat the information as the result of audience design and deliberate communicative choice, as opposed to merely informational content. SPEAKER GOAL ORIENTATION manipulated the functional role that the correlational information played in the speaker's task. In the *intervention-oriented* condition, the speaker belonged to a cultivation team whose goal was to improve the yield of *xeliherb*, a context in which acting on potential causes is directly relevant. In the *diagnostic-oriented* condition, the speaker belonged to a localization team whose goal was to identify where *xeliherb* naturally grows, a context in which correlational information is relevant in a purely diagnostic role.

Participants We recruited a total of 404 participants via Prolific using the same eligibility criteria as in the previous experiment. Each participant completed a one-shot trial for a single instance of randomly assigned between-subject conditions.

Materials The materials were framed in the same space-colonization scenario as the previous experiment (Figure 1A). All participants were presented with the same critical correlational statement in (1).

Two aspects of the materials varied by condition. First, the correlational information was framed either as *direct communication* from a speaker explicitly addressing the participant or as *indirect*, impersonal information encountered without audience design. Second, the materials varied in the description of the science team's goal, corresponding to the assigned level of SPEAKER GOAL ORIENTATION. Aside from these manipulations, all other aspects of the materials were held constant across conditions. Non-linguistic contextual elements, including AI-generated artwork, were used to support immersion but were identical across conditions.

The SPEAKER GOAL ORIENTATION was introduced through the following descriptions:

> **Diagnostic-oriented (Localization):** In the early days[2] of deep space colonization, humans have relied on finding naturally growing *xeliherb*. Dedicated Science Teams for Localization accumulated evidence for where to locate the precious, but fragile and hard-to-discover plant.
>
> **Intervention-oriented (Cultivation):** While the discovery of naturally growing *xeliherb* is still important, farmers have started successfully cultivating *xeliherb* in controlled environments as well. Still, the herb proves difficult to cultivate, so that dedicated Science Teams for Cultivation investigate how best to optimize the yield of *xeliherb* farming.

The DIRECTNESS OF INFORMATION manipulation framed the correlational information either as directly addressed communication or as impersonal background information (a passive observation of a written record that happened a long time ago).

> **Direct communication:** You received a recent report from a Science Team for (Localization | Cultivation) who made a very interesting discovery. The report states:
>
> **Indirect communication:** Coincidentally, you found an old journal from a Science Team for (Localization | Cultivation) from the early days. Some scientists interested in (localizing | cultivating) *xeliherb* took the following note for themselves:

Procedure Participants first read the background story introducing the space-colonization scenario. They were then informed about the goal of the science team, corresponding

[2]The phrase "early days" refers to an earlier phase of the colonization effort, during which the emphasis was on discovering naturally occurring *xeliherb*. In contrast, the present-day context involves intentional cultivation, motivating the role of intervention-oriented science teams. This temporal contrast provides narrative grounding for the speaker goal manipulation without presenting the two teams as opposing or incompatible.

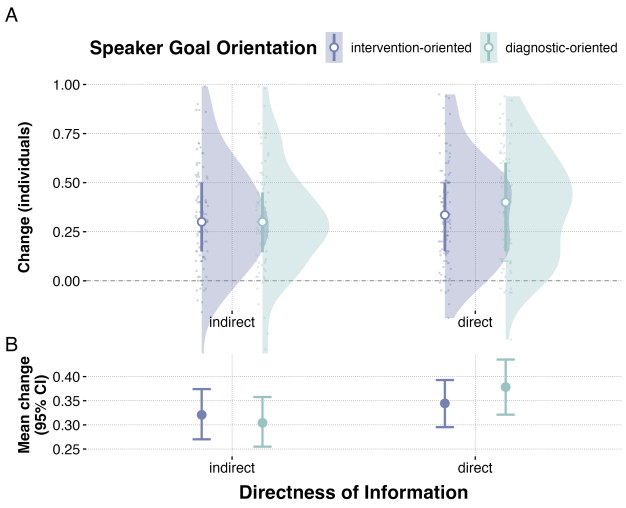

Figure 3: Experiment 2 results. **(A)** Per-cell distribution of individual change scores (posterior minus prior causal-confidence rating): half-violin densities with jittered individual points; white-filled pointranges indicate medians and interquartile ranges. The dashed grey line marks zero change. The y-axis is clipped to the central 99% of the data. **(B)** Zoomed view of condition means with bootstrapped 95% confidence intervals.

to their assigned level of SPEAKER GOAL ORIENTATION. To ensure comprehension, participants answered a control question about the team's goal and were required to reread the introduction if they responded incorrectly.

Participants next provided a *prior* judgment of their willingness to intervene by cultivating *ralocrop*, using a continuous slider, initially positioned at the neutral midpoint, ranging with labels for endpoints "very unlikely to intervene" and "very likely to intervene." Participants then received the critical correlational information, framed according to the assigned level of DIRECTNESS OF INFORMATION. Finally, they provided a *posterior* judgment using the same question and slider scale.

The primary dependent measure was the change between posterior and prior judgments (see Figure 1D). Because all information available at the time of the prior judgment was identical across conditions, and the pragmatic manipulations were introduced only with the critical information, this change score can be interpreted as a within-participant update in willingness to intervene associated with exposure to the correlational statement under different pragmatic framings.

Results Figure 3 summarizes the results. Data were analyzed using Bayesian hierarchical regression within the same general modeling framework as Experiment 1; difference scores were positive on average across all conditions, indicating that the correlational information made participants more likely to endorse an action that is rational under a belief in causation.

Slider responses were mapped on the unit interval and modeled with a zero–one–inflated Beta likelihood, capturing graded belief in the Beta component while treating exact boundary responses (0 and 1) as categorical commitments. We modeled responses as a function of DECISION ORDER (prior vs. posterior), DIRECTNESS OF INFORMATION, SPEAKER GOAL ORIENTATION, and their interaction, with by-participant random intercepts and random slopes for DECISION ORDER; the prior judgment thus served as an individually measured baseline.

Belief increased substantially from the prior to the posterior judgment ($\beta = 0.31$, 95% CrI $[0.28, 0.34]$), confirming systematic updating in response to the correlational information. Effects below are reported as posterior expected predictions on the 0–100 slider scale.

Crucially, the analysis revealed a credible interaction between DIRECTNESS OF INFORMATION and SPEAKER GOAL ORIENTATION ($\Delta = 0.07$, 95% CrI $[0.00, 0.13]$, posterior probability $P(\Delta > 0) = .97$, evidence ratio = 36.38). Decomposing this interaction revealed a systematic difference in directional tendencies across levels of DIRECTNESS OF INFORMATION. Under direct communication, posterior beliefs tended to be higher for diagnostic-oriented than for intervention-oriented speakers ($\mu = 0.03$, 95% CrI $[0.01, 0.08]$, $P(\mu > 0) = .94$), indicating moderate evidence for this contrast. Under indirect communication, the pattern was reversed, with posterior beliefs tending to be higher for intervention-oriented than for diagnostic-oriented speakers ($\mu = -0.02$, 95% CrI $[-0.07, 0.02]$), with marginal evidence in this direction ($P(\mu < 0) = .88$).

In addition, there was strong evidence for a main effect of DIRECTNESS OF INFORMATION, with higher posterior beliefs under direct than indirect communication ($\mu = 0.06$, 95% CrI $[0.02, 0.09]$, $P(\mu > 0) > .99$). No credible main effect of SPEAKER GOAL ORIENTATION was observed in isolation.

The model revealed substantial individual differences in both baseline beliefs and belief updating, with participants varying widely despite identical initial conditions. The per-cell distributions in Figure 3A nevertheless show that these effects are not driven by atypical participants: 89–91% of participants in every cell updated in the positive direction (7–9% negative, 2–3% zero). The crossover is also visible at the raw distributional level: under *direct* communication the diagnostic-oriented cell's median sits 6.5 points above the intervention-oriented cell's, with the 75th percentile 10 points above—a shift of the whole upper half of the distribution. Under *indirect* framing the medians are identical (30 points in both cells) and only the upper tail differs slightly.

Discussion Experiment 2 provides direct evidence that correlational statements can increase participants' willingness to intervene, indicating a genuine shift in belief in causality. Crucially, while a generic increase in belief in *some* causal relation is explained by assuming that participants (implicitly) applied Reichenbach's principle of *no correlation without causation*, the data suggest that participants' causal beliefs changed towards a more specific causal interpretation (that *ralocrop* cultivation influences the yield of *xeliherb*).

Moreover, the experiment suggests that changes in causal beliefs were systematically modulated by pragmatic features of the communicative context. In particular, belief updates were larger under direct communication—when information was explicitly addressed to the participant—than under indirect exposure. This main effect suggests that listeners are sensitive to audience design and to the inferential pressures associated with communicative intent.

When correlational information is presented as having been formulated to support a listener's decision, it is reasonable for listeners to assume that the speaker has selected information that is generally relevant (Sperber & Wilson, 1995), and in particular relevant for practical decision making (Benz & van Rooij, 2007). As a result, listeners may enrich correlational statements with causal meaning in order to render them useful for action. By contrast, when the same information is encountered as impersonal or self-directed—such as a report or note not tailored to a specific decision problem—the pragmatic pressure to infer causal structure is weaker. In these cases, correlational statements can be interpreted as merely descriptive, and belief updating is correspondingly reduced.

Beyond this main effect, the interaction between communicative directness and speaker goal orientation further supports the pragmatic enrichment view by revealing not only when causal belief updating is strengthened, but also when it is attenuated. Under direct communication, correlational statements attributed to diagnostic-oriented speakers elicited larger belief updates than those attributed to intervention-oriented speakers. This asymmetry is consistent with pragmatic reasoning over alternative utterances. When a speaker whose role already presupposes intervention nevertheless opts for a weak correlational formulation, listeners may reason more cautiously: had the speaker been confident in a causal relationship, a stronger or more directive formulation would have been communicatively preferable. This inference may block causal enrichment and attenuate belief updating.

By contrast, when a diagnostic-oriented speaker—whose own task does not presuppose intervention—directly communicates correlational information to a listener who is known to care about acting, the choice to mention a manipulable factor can be interpreted as intentional and action-relevant. Despite the lack of goal alignment, listeners may reason that the speaker selected this information precisely because it bears on the listener's decision problem, and enrich the correlational formulation accordingly.

In indirect contexts, where the same information is encountered as impersonal evidence rather than as communicative input, opportunities for pragmatic reasoning are substantially reduced. In the absence of clear audience design, listeners have less reason to consider alternative utterances or to infer speaker intentions, and belief updating correspondingly becomes less sensitive to speaker goal orientation. Nevertheless, we observe marginal evidence for a small tendency in the opposite direction: when the information originates from an intervention-oriented context, belief updating is slightly higher, possibly reflecting alignment between the inferred function of the information and the listener's own decision goals.

Conclusion

The present study revisits a tension in causal reasoning: although correlational statements are semantically compatible with multiple causal structures, listeners often enrich them to carry more specific causal information. Across two experiments, we show that the same correlational statement produces systematically different degrees of causal commitment depending on its communicative context—belief updates were stronger under direct than under indirect communication (Experiment 2), participants reformulated such statements as causal more often when advising others than when recalling them for themselves (Experiment 1), and the speaker-goal contrast reversed sign as directness was reduced (Experiment 2). This constellation is difficult to capture with a context-insensitive bias or heuristic over the surface form of correlational statements, which would have to postulate distinct biases for distinct framings, but follows naturally from a pragmatic account in which causal implicatures arise through reasoning about why a particular speaker, in a particular communicative setting, chose a correlational formulation rather than a more directly causal alternative.

The finding that speakers are more likely to use causal language when advising others (Experiment 1), and that listeners are more likely to draw causal inferences when the information is communicated directly (Experiment 2), mirrors everyday patterns of word-of-mouth and social-media communication. In such contexts, correlational findings are often encountered directly and then passed on as recommendations or advice. When individuals choose to act on this information, they may reframe it in causal terms for others, thereby amplifying its apparent causal force through successive message passing. Crucially, this transformation does not require intentional distortion; instead, it can arise naturally from pragmatic reasoning about what information is most relevant or useful for another person's decision-making context.

Several limitations remain. First, while the current paradigm captures decision-making behavior in response to correlational information, it does not fully disentangle causal belief from utility-based reasoning. Future work could address this by separately modeling inferred causal structure and subjective cost–benefit considerations. Second, our experiments focused primarily on pragmatic cues related to speaker goals and communicative framing; exploring how inferences are shaped by the listener's own goals and relevance constraints would offer a valuable extension. Finally, formalizing the observed effects within probabilistic pragmatic frameworks (Degen, 2023; Goodman & Frank, 2016; Scontras et al., 2021), in particular extending previous like-minded models of causal enrichment of conditionals (Grusdt et al., 2022), could yield more precise, model-based predictions about when and how causal enrichment arises from correlational language.

Acknowledgments

This research is part of the "Communicating Causality" project, supported by the UK Arts and Humanities Research Council (grant number AH/Z507490/1, PI: Daniel Lassiter); and the Deutsche Forschungsgemeinschaft (PI: Michael Franke) under project ID 547376651. The work of Hening Wang and Michael Franke for this project was supported by the project LMBayes (grant: Leibniz Collaborative Excellence 22-20, PI Anton Benz).

References

Adams, R. C., Sumner, P., Vivian-Griffiths, S., Barrington, A., Williams, A., Boivin, J., Chambers, C. D., & Bott, L. (2017). How readers understand causal and correlational expressions used in news headlines. *Journal of Experimental Psychology: Applied*, *23*(1), 1.

Benz, A., & van Rooij, R. (2007). Optimal assertions and what they implicate. *Topoi*, *26*, 63–78. https://doi.org/doi:10.1007/s11245-006-9007-3

Bott, O., & Solstad, T. (2014). From Verbs to Discourse: A Novel Account of Implicit Causality. In B. Hemforth, B. Mertins, & C. Fabricius-Hansen (Eds.), *Psycholinguistic Approaches to Meaning and Understanding across Languages* (pp. 213–251). Springer International Publishing. https://doi.org/10.1007/978-3-319-05675-3_9

Bürkner, P.-C. (2021). Bayesian item response modeling in R with brms and Stan. *Journal of Statistical Software*, *100*(5), 1–54. https://doi.org/10.18637/jss.v100.i05

Degen, J. (2023). The rational speech act framework. *Annual Review of Linguistics*, *9*(1), 519–540. https://doi.org/10.1146/annurev-linguistics-031220-010811

Gershman, S. J., & Ullman, T. D. (2023). Causal implicatures from correlational statements. *PLOS ONE*, *18*(5). https://doi.org/10.1371/journal.pone.0286067

Goodman, N. D., & Frank, M. C. (2016). Pragmatic language interpretation as probabilistic inference. *Trends in cognitive sciences*, *20*(11), 818–829.

Grusdt, B., Lassiter, D., & Franke, M. (2022). Probabilistic modeling of rational communication with conditionals. *Semantics and Pragmatics*, *15*, 13:1–59. https://doi.org/10.3765/sp.15.13

Hartshorne, J. K. (2014). What is implicit causality? *Language, Cognition and Neuroscience*, *29*(7), 804–824. https://doi.org/10.1080/01690965.2013.796396

Huff, D. (2023). *How to lie with statistics*. Penguin UK.

Lassiter, D., & Franke, M. (2024). The rationality of inferring causation from correlational language. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *46*.

R Core Team. (2024). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing. Vienna, Austria.

Reichenbach, H. (1956). *The Direction of Time*. Dover.

Scontras, G., Tessler, M. H., & Franke, M. (2021). A practical introduction to the rational speech act modeling framework. https://doi.org/10.48550/ARXIV.2105.09867

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.

Sperber, D., & Wilson, D. (1995). *Relevance: Communication and cognition* (2nd ed.). Blackwell.