Lossy Compression and the Granularity of Causal Representation

David Kinney
Yale University
New Haven, CT 06510
david.kinney@yale.edu

Tania Lombrozo
Princeton University
Princeton, NJ 08544
lombrozo@princeton.edu

Abstract

A given causal system can be represented in a variety of ways. How do agents determine which variables to include in their causal representations, and at what level of granularity? Using techniques from information theory, we develop a formal theory according to which causal representations reflect a trade-off between compression and informativeness. We then show, across three studies (N=1,391), that participants' choices over causal models demonstrate a preference for more compressed causal models when all other factors are held fixed, with some further tolerance for lossy compressions.

1 Introduction and Formal Framework

Scientists often aim to produce causal models of the world that balance informativeness with compression. That is, they aim to model data-generating processes in a way that captures as much information about those processes as possible, while omitting cumbersome or unnecessary details. For example, epidemiologists might produce a model of cancer rates in a population that treats smoking as a binary variable representing whether or not a person smokes cigarettes, but without specifying the average number of cigarettes the person smokes per day, and omitting additional background variables such as the person's blood type. Ordinary agents face an analogous challenge: when representing the social and physical world around us, each of us must determine which variables to include in our causal models, and at what level of granularity. For example, a causal model of a toddler's tantrums could include whether they napped or not as a binary variable, or a finer-grained specification of the number of minutes they napped; it could include the time of day, or omit this variable entirely. Any such choice of variables instantiates a particular trade-off between informativeness and compression. How do people navigate this choice in building causal models of the world?

To formalize this trade-off, we begin by using Bayesian networks to represent causal structure. Let $\mathcal{V}_{\mathcal{P}}$ be a set of random variables that are all measurable with respect to the same probability space $\mathcal{P} = (\Omega, \Sigma, p)$. Let $\mathcal{E}$ be an acyclic set of ordered pairs, or edges, relating the variables in $\mathcal{E}$. The set of edges $\mathcal{E}$ allows us to define parent and descendant relations between variables in the obvious way. A **causal Bayes net** is a pair $\mathcal{G}_{\mathcal{P}} = (\mathcal{V}_{\mathcal{P}}, \mathcal{E})$ such that: i) according to the probability distribution p, all elements of $\mathcal{V}_{\mathcal{P}}$ are independent of their non-descendants, conditional on their parents (**Markov Condition**), ii) there is no set of edges $\mathcal{E}^* \subset \mathcal{E}$ such that $(\mathcal{V}_{\mathcal{P}}, \mathcal{E}^*)$ satisfies the Markov condition according to the probability distribution p (**Minimality Condition**), and iii) no variable in $\mathcal{V}_{\mathcal{P}}$ is a coarsening of or identical to any other variable in $\mathcal{V}_{\mathcal{P}}$ (**Co-possibility Condition**).

For any variable set $\mathcal{V}_{\mathcal{P}} = \{V_1, \ldots, V_n\}$, let $\mathcal{V}_{\mathcal{P}}(\omega)$ be the vector $[V_1(\omega), \ldots, V_n(\omega)]$. This allows us to define an equivalence relation $\sim_{\mathcal{V}_{\mathcal{P}}}$ such that $\omega \sim_{\mathcal{V}_{\mathcal{P}}} \omega'$ iff $\mathcal{V}_{\mathcal{P}}(\omega) = \mathcal{V}_{\mathcal{P}}(\omega')$. A **compression of** a causal Bayes net $\mathcal{G}_{\mathcal{P}} = (\mathcal{V}_{\mathcal{P}}, \mathcal{E})$ is any causal Bayes net $\hat{\mathcal{G}}_{\mathcal{P}} = (\hat{\mathcal{V}}_{\mathcal{P}}, \hat{\mathcal{E}})$ that: i) satisfies the Markov, Minimality, and Co-possibility conditions according to $\mathcal{P}$; ii) is such that for any $\omega, \omega' \in \Omega$, if

Information-Theoretic Principles in Cognitive Systems Workshop at the 37th Conference on Neural Information Processing Systems (NeurIPS 2023).

$\omega \sim_{\mathcal{V}_\mathcal{P}} \omega'$, then $\omega \sim_{\hat{\mathcal{V}}_\mathcal{P}} \omega'$; and iii) there exists a pair $\omega, \omega' \in \Omega$, such that $\omega \not\sim_{\mathcal{V}_\mathcal{P}} \omega'$ but $\omega \sim_{\hat{\mathcal{V}}_\mathcal{P}} \omega'$. In other words, a compression of a causal Bayes net is a second causal Bayes net that satisfies the necessary criteria for representing the same target system as the first, while also defining a strictly more general partition over possibility space. The compression-of relation between causal Bayes nets defines a lexical ordering between causal representations of the same target system, such that while we cannot definitively say, for *any* two causal Bayes nets, which is more compressed, we can sometimes say of two Bayes nets that one is more compressed than the other (namely, when one is a compression of the other).

Next, we define the degree to which a set of variables $\mathbf{C}$ is **informative** about an effect variable E, where both variables are embedded in a causal Bayes net $\mathcal{G}_P$. We begin by noting that for given causal Bayes net $\mathcal{G}_\mathcal{P}$, we can calculate the probability distribution over any variable V in the set $\mathcal{V}_\mathcal{P}$, given an intervention setting some set of variables $\mathbf{X}$ to some set of values $\mathbf{x}$ (denoted $do(\mathbf{x})$), using the following formula (Pearl [2000]):

$$p_{\mathcal{G}_\mathcal{P}}(v|do(\mathbf{x})) = \begin{cases} p(v|\operatorname{par}_{\mathcal{G}_\mathcal{P}}(V)) & \text{if } V \notin \mathbf{X} \\ 1 & \text{if } V \in \mathbf{X} \text{ and } v \text{ is consistent with } \mathbf{x} \\ 0 & \text{otherwise} \end{cases} \tag{1}$$

where $\operatorname{par}_{\mathcal{G}_\mathcal{P}}(V)$ denotes the values taken by the parents of V in $\mathcal{G}_\mathcal{P}$. This allows us to derive the probability distribution that would be defined over any variable in the causal Bayes net if any other variable were set to some value via an exogenous, "surgical" intervention on the data-generating system. This, in turn, allows us to define the **causal mutual information** between a set of causal variables $\mathbf{C}$ and an effect variable E, where both variables are embedded in a causal Bayes net $\mathcal{G}_P$:

$$\text{CMI}(\mathbf{C}, E, \mathcal{G}_P) = \sum_{\mathbf{c}, e} q(\mathbf{c}) p(e|do(\mathbf{c})) \log_2 \frac{p(e|do(\mathbf{c}))}{p(e)}, \tag{2}$$

where q is a distribution over possible interventions on $\mathbf{C}$ (see Ay and Polani [2008] and Pearl [1994] for a similar deployment of a distribution over interventions).

Suppose that a causal Bayes net $\hat{\mathcal{G}}_\mathcal{P}$ is a compression of some other causal Bayes net $\mathcal{G}_\mathcal{P}$. Both Bayes nets contain an effect variable E, but may contain different causal variables $\mathbf{C}$ and $\hat{\mathbf{C}}$. The amount of information about E that is lost as a result of replacing the less-compressed causal Bayes net with a more compressed alternative is given by the equation:

$$\mathcal{L}(\mathcal{G}_\mathcal{P}, \hat{\mathcal{G}}_\mathcal{P}, \mathbf{C}, \hat{\mathbf{C}}, E) = \text{CMI}(\mathbf{C}, E, \mathcal{G}_P) - \text{CMI}(\hat{\mathbf{C}}, E, \hat{\mathcal{G}}_P). \tag{3}$$

This quantity is non-negative whenever the variable set $\hat{\mathbf{C}}$ defines a strictly more general equivalence class on Ω than $\mathbf{C}$, and only takes the value zero in the case where compression does not result in any change in the informativeness of the casual variables with respect to the effect variable of interest. We hypothesized that as more information is lost due to compression (as measured by the function defined in Eq. 3), participants will evaluate claims consistent with the more compressed representation less positively relative to claims consistent with the less compressed representation. In what follows, we present an experiment confirming this hypothesis. Two further experiments, which in addition rule out alternative explanations of our findings, are reported in the supplemental materials.

1.1 Previous Work

From a theoretical perspective, the work that is closest to our framework consists of previous attempts to quantify properties of causal relationships in Bayesian networks using tools from information theory. These include specific attempts to measure various aspects of causal relationships, such as: proportionality, stability, power, abstraction, strength, or specificity using formalism from information theory Pocheville et al. [2017], Korb et al. [2011], Griffiths et al. [2015], Hoel [2017], Beckers and Halpern [2019], Bourrat [2021]. However, none of these approaches aim, as we do, to provide a unified account of our preference for more or less compressed causal representations in terms of information loss.

2 Experiment 1

We presented participants with a description of the results of controlled experiments on a fictional variety of mushroom, fly, or rock, and asked them to rate how good it would be to include various claims in a summary of the described results. These claims included more and less compressed causal claims. We manipulated the vignette used, the amount of information loss realized by the more compressed causal claim, and whether the compression was achieved by coarse-graining a single variable or removing a background variable from the representation.

2.1 Participants

Participants were 450 adults recruited via Prolific. 150 additional participants were excluded for failing comprehension checks or for rating poor causal claims non-negatively. For both studies, participation was restricted to users with a US-based IP address and a 95% rating based on at least 100 previous studies. Both studies were pre-registered, and IRB approval was obtained from the authors' university. Data, stimuli, and pre-registrations are available at `https://osf.io/zm6kr/?view_only=124c22b8b2dd4d64b44046c8784911db` (Experiments 1-3).

2.2 Materials and Procedures

Participants read a vignette in which they learned about a novel causal system, including the results of experiments involving that system. For example, in the insect vignette, participants were presented with the following four facts resulting from experiments on the "Bricofly" insect: a) x% of all Bricofly larvae raised in a warm, humid tank developed blue wings; b) 70% of all Bricofly larvae raised in a warm, dry tank developed blue wings; c) 1% of all Bricofly larvae raised in a cold, humid tank developed blue wings; d) 1% of all Bricofly larvae raised in a cold, dry tank developed blue wings.

The value of x was varied between subjects and set at either 70, 85, or 98. Participants were then asked to rate, on a scale from -3 (very bad) to 3 (very good), how good it would be to include each of the following statements in a summary of the findings of the descriptions given above:

- *Compressed:* Raising Bricofly larvae in a warm tank causes them to develop blue wings.
- *High:* Raising Bricofly larvae in a warm, humid tank causes them to develop blue wings.
- *Low:* Raising Bricofly larvae in a warm, dry tank causes them to develop blue wings.

The values of x correspond to information loss amounts for *Compressed* of 0, $.01$, and $.06$ respectively, assuming a uniform distribution over possible interventions. Participants were randomly assigned to one of eighteen possible conditions, which differed with respect to which of the three vignettes they were shown, whether they were asked to evaluate a compressed claim achieved by coarsening a causal variable or eliding a background variable, and the amount of information loss inherent in compression. Finally, participants were also asked to evaluate three poor causal claims, constructed by substituting the value of the primary causal factor (e.g., changing the warm tank to a cold tank). These were included to help anchor the scale and verify participant understanding; we do not discuss them further here in the interest of space.

2.3 Results

To test whether evaluation of less compressed causal claims relative to more compressed causal claims increased as a function of information loss due to compression, we computed (as pre-registered) two difference scores:

- V-A. The difference between the participant's evaluation of *Compressed* and their evaluation of *High* (e.g., the difference between the evaluation of 'Raising Bricofly larvae in a warm tank causes them to develop blue wings' and the evaluation of 'Raising Bricofly larvae in a warm, humid tank causes them to develop blue wings').
- V-B. The difference between the participant's evaluation of *Compressed* and a uniform average of their evaluations of *High* and *Low* (e.g., the difference between the evaluation of 'Raising Bricofly larvae in a warm tank causes them to develop blue wings' and the average evaluation of 'Raising Bricofly larvae in a warm, humid tank causes them to develop

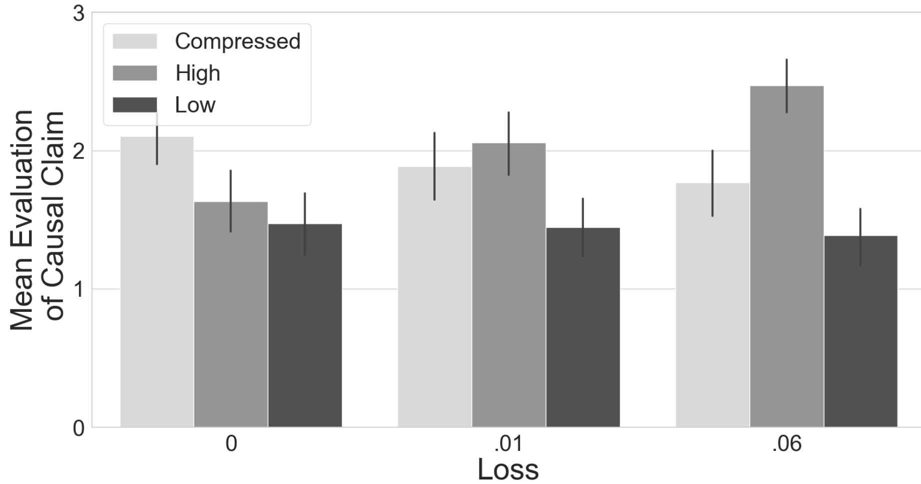

Figure 1: Mean evaluations of claims in Experiment 1, with bars showing 95% CIs. 'Loss' corresponds to information loss due to compression inherent in choosing *Compressed* over *High* and *Low*. Mixed ANOVA for each value of Loss found that at Loss=0, *Compressed* was rated more highly than both *High* ($\eta^2 = .025$, $p = .002$) and *Low* ($\eta^2 = .041$, $p < .001$). When Loss=.01, *High* was not rated significantly higher than *Compressed* ($\eta^2 = .005$, $p = .156$), but was rated higher than *Low* ($\eta^2 = .050$, $p < .001$). When Loss=.06, *High* was rated significantly higher than both *Compressed* ($\eta^2 = .058$, $p < .001$) and *Low* ($\eta^2 = .153$, $p < .001$).

blue wings' and 'Raising Bricofly larvae in a warm, dry tank causes them to develop blue wings').

We regressed these dependent variables against independent variables denoting the assigned vignette (Vignette), whether the more compressed claim was generated by coarsening a variable or eliding a background condition (Mode of Compression), and the amount of information loss (Loss), as well as all possible interactions. The regressions revealed that only Loss was a significant predictor of V-A ($\beta = -16.623$, $p < .001$) and V-B ($\beta = -9.852$, $p < .001$). As a sanity check, we also analyzed the difference between the participant's evaluation of *High* and their evaluation of *Low* (V-C). As expected, only Loss was a significant predictor of V-C, with the value of V-C increasing as the probability of the effect given the description of the cause in *High* increases with Loss ($\beta = 13.543$, $p < .001$). In an exploratory analysis, we measured the percentage of participants who strictly preferred *Compressed* to *High* across all three loss levels. This percentage was approximately 36% when Loss=0, 21% when Loss=.01, and 10% when Loss=.06.

3 Discussion and Conclusion

These results provide strong evidence in favor of the claim that participants' relative evaluations of more and less compressed causal claims are partially governed by the amount of information loss that is inherent in the more compressed causal claim. As can be seen in Fig. 1, which plots participants' absolute evaluations of each causal claim at each loss level, when there is no information loss, participants evaluate more compressed causal claims significantly more highly than less compressed causal claims, suggesting that people award simplicity and penalize unnecessary complexity in their evaluation of causal claims. When information loss is moderate, there is no significant difference between participants' evaluations of more and less compressed causal claims, suggesting that some participants prefer a compressed claim even when some information loss is inherent. That is, they are tolerant of some amount of *lossy compression* in the causal representation of their environment. In the supplemental materials, we report results from two further experiments. These results lend further confirmation to our hypothesis, while also ruling out alternative explanations due to Cheng [1997] and Lien and Cheng [2000]. In future work, we hope to explore connections between our results and the formalization of lossy compression found in rate distortion theory (e.g., Berger [1971]).

References

N. Ay and D. Polani. Information flows in causal networks. *Advances in complex systems*, 11(01): 17–41, 2008.

S. Beckers and J. Y. Halpern. Abstracting causal models. In *Proceedings of the aaai conference on artificial intelligence*, volume 33, pages 2678–2685, 2019.

T. Berger. *Rate distortion theory, a mathematical basis for data compression.* Prentice Hall, 1971.

P. Bourrat. Measuring causal invariance formally. *Entropy*, 23(6):690, 2021.

P. W. Cheng. From covariation to causation: A causal power theory. *Psychological review*, 104(2): 367, 1997.

P. E. Griffiths, A. Pocheville, B. Calcott, K. Stotz, H. Kim, and R. Knight. Measuring causal specificity. *Philosophy of science*, 82(4):529–555, 2015.

E. P. Hoel. When the map is better than the territory. *Entropy*, 19(5):188, 2017.

K. B. Korb, E. Nyberg, and L. R. Hope. A new causal power theory. In *Causality in the Sciences*, pages 628–652. Oxford University Press, 2011.

Y. Lien and P. W. Cheng. Distinguishing genuine from spurious causes: A coherence hypothesis. *Cognitive Psychology*, 40(2):87–137, 2000.

J. Pearl. A probabilistic calculus of actions. In *Uncertainty Proceedings 1994*, pages 454–462. Elsevier, 1994.

J. Pearl. *Causality*. Cambridge University Press, 2000.

A. Pocheville, P. E. Griffiths, and K. Stotz. Comparing causes—an information-theoretic approach to specificity, proportionality and stability. In *Proceedings of the 15th congress of logic, methodology and philosophy of science*, pages 93–102. College Publications London, 2017.

Supplemental Materials

4 Experiment 2

In Experiment 1, participants evaluated the three key causal claims (*Compressed*, *High*, and *Low*) on the same screen. This could have introduced unintended task demands. For instance, participants may have felt that endorsing *Compressed* was redundant with the endorsement of both *High* and *Low*, or that endorsing *Compressed* (when the option to select more fine-grained options was available) implied the causal irrelevance of the unspecified factor. To ensure that the results of Experiment 1 were robust to such considerations, we replicated the study with the amendment that participants were shown the same data twice, and asked first to evaluate *Compressed* and second to independently evaluate *High* and *Low*. Our results speak against an alternative interpretation of the results of Experiment 1 that uses the causal power theory of Cheng [1997].

4.1 Participants

483 adults were recruited via Prolific. 117 additional participants were excluded for failing comprehension checks or rating poor causal claims non-negatively.

4.2 Materials and Procedures

The procedure was identical to that used in Experiment 1 with three exceptions. First, as described above, participants were asked to evaluate *Compressed* as part of a separate task than their evaluation of *High* and *Low*. Second, sentence (b) in both descriptions used in the first experiment was amended to replace '70%' with '55%'. Analogous replacements were made for the other two vignettes. Third, the value of x in (a) and (b) was varied between subjects and set at either 55, 85, or 98, leading to information loss amounts of 0, .04, and .11 respectively. Thus, we replicated Experiment 1 for a different range of loss values.

4.3 Results

We performed the same regressions as in Experiment 1. Loss was a significant predictor of all three dependent variables (V-A: $\beta = -14.53$, $p < .001$, V-B: $\beta = -6.391$, $p < .001$, V-C: $\beta = 16.303$, $p < .001$). Fig. 2 shows the relationship between Loss and participants' absolute evaluations of *Compressed*, *High*, and *Low*. In an exploratory analysis, we measured the percentage of participants who strictly preferred *Compressed* to *High* across all three loss levels. This percentage was approximately 39% when Loss=0, 10% when Loss=.04, and 2% when Loss=.11.

The results of Experiment 2 also offer evidence against an alternative interpretation of our results. Specifically, the causal power theory (Cheng [1997]) holds that agents evaluate causal claims positively to the extent that they optimize the following quantity, which we express in terms of Pearl's do-calculus:[1]

$$\text{Power}(\mathbf{c}, e) = \frac{p(e|do(\mathbf{c})) - p(e|do(\neg\mathbf{c}))}{1 - p(e|do(\neg\mathbf{c}))}. \tag{4}$$

Table 1 shows crucial calculations of causal power. If evaluations of causal claims are primarily driven by differences in causal power, then we would expect that the difference between participants' evaluations of Compressed and their average evaluation of High and Low (i.e., the dependent variable V-B) should be positively correlated with value of Power(Comp) - AVG[Power(High),Power(Low)]. However, if we use the data from Experiment 2 to regress V-B against Power(Comp) - AVG[Power(High), Power(Low)], along with Mode of Compression, Vignette, and all interactions between these three variables, we observe a significant predictive relationship between V-B and Power(Comp) - AVG[Power(High), Power(Low)] going in the opposite direction ($\beta = -65.853$, $p = .001$) such that higher values of Power(Comp) - AVG[Power(High), Power(Low)] are associated with lower values of V-B. Thus, a causal power theory fails to predict a crucial dependent variable that our information loss theory is able to successfully predict.

[1] The original formulation of causal power given in Cheng [1997] is not stated in terms of Pearl's do-calculus. Instead, it is written $\text{Power}(\mathbf{c}, e) = \frac{p(e|d\mathbf{c}) - p(e|\neg\mathbf{c})}{1-p(e|\neg\mathbf{c})}$ (see p. 374, Eq. 8). We state causal contrast in these terms here to maintain formal consistency with our own measure of information loss.

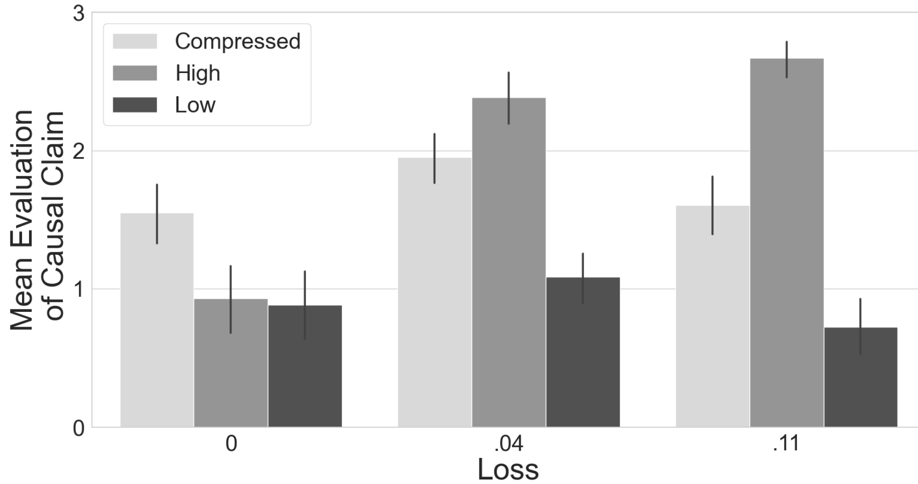

Figure 2: Mean evaluations of claims in Experiment 2, with bars showing 95% CIs. 'Loss' corresponds to information loss due to compression inherent in choosing *Compressed* over *High* and *Low*. Mixed ANOVA for each value of Loss found that at Loss=0, *Compressed* was rated more highly than both *High* ($\eta^2 = .044$, $p < .001$) and *Low* ($\eta^2 = .048$, $p < .001$). At Loss=.04, *High* was rated more highly than *Compressed* ($\eta^2 = .036$, $p < .001$) and *Low* ($\eta^2 = .220$, $p < .001$). At Loss=.11, *High* was rated more highly than *Compressed* ($\eta^2 = .179$, $p < .001$) and *Low* ($\eta^2 = .442$, $p < .001$).

p(Effect|High)	Power(Comp)	Power(High)	Power(Low)	Power(Comp) - AVG[Power(High),Power (Low)]
.55	.545	.444	.444	.101
.85	.697	.815	.366	.106
.98	.763	.975	.325	.112

Table 1: Causal power values for causal claims in Experiment 2.

4.4 Discussion

Nevertheless, our results in Experiments 1 and 2 remain subject to two salient concerns. First, Experiments 1-2 varied the granularity of a causal variable by adding or omitting qualifiers to a variable (e.g., warm, humid tank versus warm tank). More canonical manipulations of variable granularity involve a continuum that can be coarsened into discrete ranges (e.g., a scale with ten values that is coarsened into two ranges of values). Thus, the results of Experiments 1-2 leave open the possibility that these more canonical manipulations of granularity would show divergence between effects of information loss on across different modes of compression. Second, Experiments 1 and 2 were not designed to differentiate our account from another alternative hypothesis: that evaluations of more and less compressed causal claims do not reflect information loss, as our account suggests, but instead differences in causal contrast (consistent with Lien and Cheng [2000]). Experiment 3 was designed to addresses both of these concerns.

5 Experiment 3

Lien and Cheng [2000] develop an account of how people differentiate between genuine and spurious causes, and in so doing present results in keeping with the claim that when agents choose between candidate causal explanations of a given event, they choose the one that maximizes causal contrast, which is given by the following equation:[2]

$$\text{Cont}(\mathbf{c}, e) = p(e|do(\mathbf{c})) - p(e|do(\neg\mathbf{c})). \tag{5}$$

Applying this formula to the values from Experiments 1 and 2 reveals that as the probability of the effect for *High* increases, the difference in contrast between the compressed causal claim and the high causal claim decreases. So, it seems that our results in Experiment 1-2 might just as well be explained by the hypothesis that participants are basing their judgments on the difference in contrast between *Compressed* and *High* as they are by our hypothesis that participants are balancing compression against information loss. To distinguish between these two hypotheses, we ran an experiment using a similar paradigm to Experiments 1-2, but wherein participants were shown data sets for which information loss and causal contrast generated different qualitative predictions. This design allows us to test which of these two quantities is a more plausible candidate for the cue that participants are using to evaluate causal claims. Experiment 3 also differed from Experiments 1-2 in manipulating compression through coarsenings of a continuous quantity.

5.1 Participants

Participants were 458 adults recruited via Prolific. An additional 185 participants were excluded for failing comprehension checks or rating poor causal claims non-negatively.

5.2 Materials and Procedures

Participants read a vignette in which they learned about a novel causal system, including the results of experiments involving that system. As in Experiments 1 and 2, the fictional experiments involved either insects, mushrooms, or rocks. For example, in the insect vignette, participants assigned to the condition in which compression was achieved by coarsening a variable were presented with one of the data scenarios shown in Table 2, and asked to evaluate the following three causal claims on a scale from -3 to 3:

- *Compressed*: Raising Bricofly larvae in a moderate-temperature tank causes them to develop blue wings.
- *High*: Raising Bricofly larvae in a moderately warm tank causes them to develop blue wings.
- *Low*: Raising Bricofly larvae in a moderately cold tank causes them to develop blue wings.

Table 2 also shows the values of both information loss and the difference in contrast between *Compressed* and *High* for all three data sets. As can be seen from the table, Scenarios 2 and 3 both differ from Scenario 1 by the same amount with respect to the difference in contrast between *Compressed* and *High*, but only Scenario 2 differs from Scenario 1 with respect to information loss. Thus, if we believe that information loss and not causal contrast is affecting participants' evaluations of causal claims, then we would predict that participants will treat Scenarios 1 and 3 similarly, but treat Scenario 2 differently from both Scenarios 1 and 3.

5.3 Results

Figure 3 shows the results of Experiment 3 for all three scenarios across both modes of compression. As we were primarily concerned with differential evaluations of *Compressed* and *High* across different scenarios, we ran mixed ANOVA for the within-participants difference between *Compressed* and *High* in each scenario. We found that in Scenarios 1 and 3, Compressed was strictly preferred to High (Scenario 1: $\eta^2 = .019$, $p = .001$; Scenario 3: $\eta^2 = .032$, $p < .001$). This is consistent with participants favoring compression when information loss is zero. Moreover, the preference for the claim *Compressed* over the claim *High* did not differ across Scenarios 1 and 3 ($\beta = .041$, $p = .673$), which is consistent with the predictions of information loss, but not those of causal contrast (see Table 3).

Unlike Scenarios 1 and 3, in Scenario 2, the causal claim *High* was strictly preferred to the causal claim *Compressed* ($\eta^2 = .120$, $p < .001$). This is consistent with the hypothesis that compression

[2] As in the case of causal power, we have re-written Lien and Cheng's contrast measure in terms of Pearl's do-calculus. In the original formulation, causal contrast is written as $\text{Cont}(\mathbf{c}, e) = p(e|\mathbf{c}) - p(e|\mathbf{c})$

Scenario 1		Scenario 2		Scenario 3	
Tank Condition	% of Bricofly Developing Blue Wings	Tank Condition	% of Bricofly Developing Blue Wings	Tank Condition	% of Bricofly Developing Blue Wings
Extremely Cold Tank (0-24 degrees)	1%	Extremely Cold Tank (0-24 degrees)	1%	Extremely Cold Tank (0-24 degrees)	43%
Moderately Cold Tank (25-49 degrees)	70%	Moderately Cold Tank (25-49 degrees)	70%	Moderately Cold Tank (25-49 degrees)	70%
Moderately Warm Tank (50-74 degrees)	70%	Moderately Warm Tank (50-74 degrees)	98%	Moderately Warm Tank (50-74 degrees)	70%
Extremely Warm Tank (55-99 degrees)	1%	Extremely Warm Tank (55-99 degrees)	1%	Extremely Warm Tank (55-99 degrees)	43%
Loss	0	Loss	.06	Loss	0
Contr(Comp)-Contr(High)	.23	Contr(Comp)-Contr(High)	.09	Contr(Comp)-Contr(High)	.09

Table 2: Information loss and causal contrast for three different scenarios shown in Experiment 3.

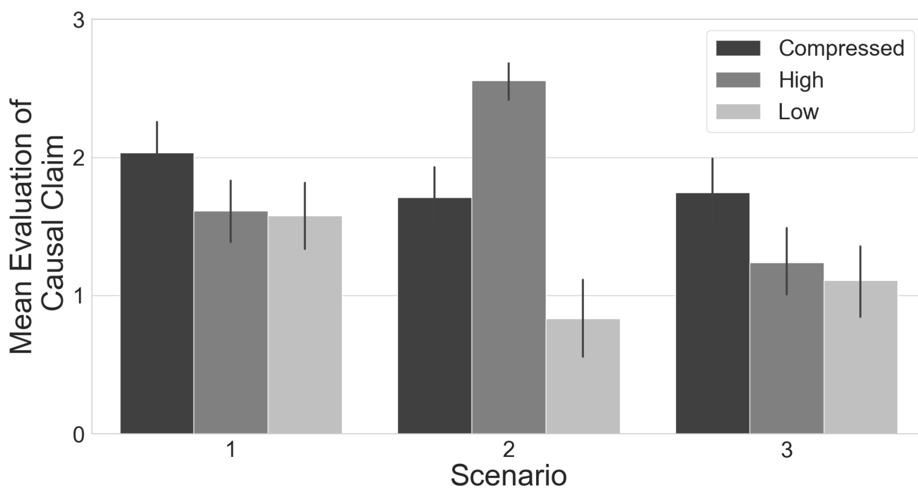

Figure 3: Mean evaluations of causal claims for three scenarios in Experiment 3.

trades off with information loss, such that less compressed causal claims may be favored when information loss is not negligible. Moreover, the difference in ratings between *Compressed* and *High* differed across Scenarios 2 and 3 ($\beta = -22.276$, $p < .001$); this is consistent with the predictions of information loss, but not with those of causal contrast (see Table 3).

5.4 Discussion

Although Lien and Cheng's causal contrast theory was developed as an account of how people differentiate between genuine and spurious causes, rather than how people determine a level of compression at which to represent the causal structure of their environment, causal contrast nevertheless offers a natural alternative to our own account of information loss in explaining why people might favor more or less compressed causal claims. Experiment 3 was designed to provide a direct test of the predictions of information loss versus those of causal contrast in explaining judgments like those elicited in Experiments 1-3. The results provide clear support for information loss: differences in information loss (holding differences in causal contrast fixed) predicted different patterns in ratings, while differences in causal contrast (holding information loss fixed) did not.