

A latent-mixture quantum probability model of causal reasoning within a Bayesian inference framework

Percy K. Mistry (pkmistry@uci.edu)

Jennifer S. Trueblood (jstruebl@uci.edu)

Joachim Vandekerckhove (joachim@uci.edu)

Department of Cognitive Sciences, University of California Irvine, Irvine, CA 92697-5100 USA

Emmanuel M. Pothos (e.m.pothos@gmail.com)

Department of Psychology, City University London, London EC1R 0JD UK

Abstract

We develop a quantum probability model that can account for situations where people’s causal judgments violate the properties of causal Bayes nets and demonstrate how the parameters of our model can be interpreted to provide information about underlying cognitive processes. We implement this model within a hierarchical Bayesian inference framework that allows us to systematically identify individual differences and also provide a latent classification of individuals into categories of causal and associative reasoners. Finally, we implement a basic normative causal Bayes net within the same inference framework that allows us to directly compare quantum and classical probability models using Bayes factors.

Keywords: Causal reasoning; quantum probability; Bayesian graphical models; causal Bayesian networks; individual differences; latent mixture models; violations of normative properties; Bayesian inference; associative reasoning

Introduction

We investigate a subset of causal reasoning paradigms where people have to make judgments about causal systems based on linguistic descriptions of the causal properties of the system, i.e. where precise statistical information is not presented. In these situations, people’s judgments often violate the properties of causal Bayes nets (Rottman & Hastie, 2014; Park & Sloman, 2013; Rehder, 2014; Fernbach & Sloman, 2009; Waldmann, Cheng, Hagmayer, & Blaisdell, 2008; Hagmayer & Waldmann, 2002). One important property of Bayes nets is the causal Markov condition, which states that any node in the network is conditionally independent of its non-effects, given its direct causes (Hausman & Woodward, 1999). To account for observed behavior, Bayes nets are often augmented with heuristic-like shortcuts (Fernbach & Rehder, 2013) or with the inclusion of variables that account for hidden aspects of the problem (that is, not part of the experimental paradigm, and assumed to be the result of participants’ mental process) such as alternative causes, shared disabling, mediating or enabling conditions between variables, and so on. (Rehder, 2014).

We propose that quantum probability models of causal reasoning can provide a more formal, principled alternative approach for explaining violations of the properties of classical probability models, such as causal Bayes nets. We

investigate the application and benefits of applying quantum probability models by using these models to explain the empirical results reported in Rehder (2014), which demonstrated violations of the causal Markov condition, as well as failure to exhibit robust discounting (this occurs when the presence of one cause makes the presence of another less likely in certain casual structures).

Description and Results of Experiments

In Rehder (2014), participants were taught one of the three causal network structures (common cause, chain or common effect) encompassing a set of relationships between three binary variables as shown in Figure 1, instantiated in either a domain-general (abstract) or domain-specific (economics, sociology, or meteorology) setting. The causal relations specified how variables were related, (e.g. in the economics domain, the variables were interest rates, trade deficits, and retirement savings, each of which could be large (high) or small (low); a relationship could take the form: low interest rates cause small trade deficits). Each individual relationship in the causal structure was described as being driven by independent causal processes.

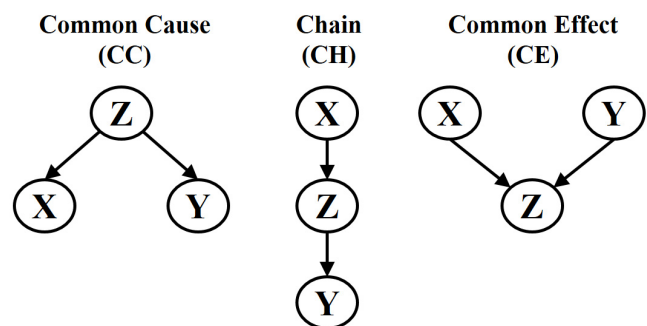


Figure 1: Causal structures considered in the experiments

The causal relationships were unidirectional, in the sense that if a particular (high or low) value of a cause variable facilitated the presence of an effect, the other binary value did not have the opposite effect (e.g. if low interest rates caused small trade deficits, high interest rates were causally unrelated to trade deficits). Once the causal structures had been taught, participants were asked to make a comparative inference on the probability of a target variable (Y) taking a

specific value (denoted y_1), between two different network states within a particular causal structure. The eight network states used in the comparison (situations A to H; see Table 1) represented different values of the remaining two variables X and Z , namely '0' (representing a state value that does not exert any causal influence), '1' (representing a state value that causally influences or is influenced), or '?' (unknown value). Participants were asked to compare states A vs B, B vs C, D vs E, F vs G and G vs H, and indicate which of the two situations provided stronger inferential support for the target variable (always referred to as Y for the rest of the paper), or whether support was equal for each variable.

Table 1: Normative Inferential Predictions

Situation	X	Z	CC	CH	CE
A	1	1			
B	?	1	A=B=C	A=B=C	C > B > A
C	0	1			
D	1	?	D >> E	D >> E	D = E
E	0	?			
F	1	0			
G	?	0	F=G=H	F=G=H	F=G=H
H	0	0			

For our analyses, we combined the data from four experiments in Rehder (2014) (i.e., experiments 2, 3, 4A and 4B). Across these four experiments, there were 315 participants. Each participant made twenty such comparative judgments, with the causal structure (common cause, chain, common effect) and domain of variables (economics, sociology, meteorology, and an abstract domain in one condition) as between-subject conditions.

Table 1 also shows the normative predictions for each possible pair of situations based on a causal Bayes net treatment of the inference problem (see Rehder, 2014 for a detailed analysis of the normative predictions). Two key properties included the causal Markov condition and discounting behavior in the common effect structure (the known presence of one cause makes the presence of an alternate cause less likely). Rehder (2014) found that a significant number of participants violated these two properties, and observed that about 23% of the 315 participants exhibited some form of associative reasoning, that is, a lack of sensitivity to causal direction, ignoring conditional independence stipulated by the causal Markov property or exhibiting anti-discounting behavior (i.e. judging the target cause as highly probable based on the presence of an alternative cause, which is opposite to normative expectation). These participants were classified as associative reasoners, whereas the rest were classified as causal reasoners.

Rehder (2014) accounted for these violations by suggesting a mixture of normative behavior with three additional inference strategies (conditional probabilities being assessed conjunctively, assumption of a hidden

disabling mechanism shared by the cause variables, assumption of an associative Markov random field network), which when appropriately weighted could reproduce behavior for both causal and associative reasoners. Each of these models had between three to five free parameters, with an additional three free parameters for the mixture weighting these models.

Specifying the Quantum Probability Model

We specify a 2-dimensional (2-d) quantum probability (QP) model to reflect the mental representations of the three binary variables X , Y and Z (Trueblood & Pothis, 2014). In this model, the three variables are deemed incompatible, that is these variables span separate subspaces within the 2-dimensional space. This implies that consideration of these variables cannot take place concurrently, but rather has to be sequential, so that order effects may arise when both variables need to be assessed (e.g., in a conjunction).

Accordingly, the two dimensions for each subspace ($\{x_1, x_0\}$, $\{y_1, y_0\}$, $\{z_1, z_0\}$) shown in Figure 2 represent the two values that each binary variable can take. Since the causal structures are unidirectional (that is, only one value affects the system causally), the values (for instance y_1 and y_0) are encoded such that y_1 always indicates the value that is causally linked (e.g. if low interest rates cause high deficits, low interest rates and high trade deficits are encoded as 1, high interest rates and low trade deficits, which do not influence or experience causal influence, are encoded as 0).

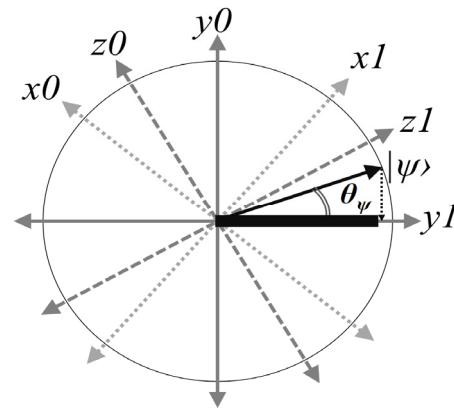


Figure 2: Representation of the 2-d QP model

The model specifies a unit length state vector $|\psi\rangle$ which represents the current state of belief held by an individual. The relative degree of rotation of the state vector from each basis vector (e.g. θ_ψ from the y_1 basis vector) defines the individual's belief in the probability of that variable. Thus, an individual's belief in the probability of a certain variable taking a certain value (e.g. $p(y_1)$ in Figure 2) can be obtained by projecting the state vector (black dotted line) on to the coordinate axis of interest and taking the squared value of the resulting amplitude (black bar).

Conditional probabilities (for example, $p(y_1|x_1)$) are measured by projecting the state vector onto the known (x_1) basis vector, normalizing it (making it unit length) to account for the fact that the state x_1 is known (squared amplitude =1), and then projecting this vector to the basis vector y_1 . The squared amplitude of the resulting projection along y_1 gives the conditional probability $p(y_1|x_1)$. Similarly, conjunctive probabilities (for example $p(x_1&y_1)$) are assessed by making successive projections to x_1 and y_1 , but without normalizing the intermediate projection.

We propose that this formulation can predict order effects, reciprocity (related to the inverse fallacy; Koehler, 1996), memoryless effects such as a lack of discounting, and violations of the Markov condition. In terms of assessed probabilities, order effects, for instance, allow the probability $p(y_1&x_1)$ to differ from $p(x_1&y_1)$. Reciprocity is a specific property of the 2-d model, where for two variables, e.g. Y and X, the conditional probabilities reciprocate, that is $p(y_1|x_1) = p(x_1|y_1)$. Memoryless effects refers to the fact that assessment of probabilities conditional or more than one variable reduce to the conditional probability based on the last updated information only, for instance, $p(y_1|x_1,z_1) = p(y_1|z_1)$ and $p(y_1|z_1,x_1) = p(y_1|x_1)$, which hence also leads to order effects. A lack of discounting can be explained by way of the memoryless property. Discounting refers to the fact that in a common effect scenario, considering classical probabilities, $p(y_1|x_1,z_1) < p(y_1|z_1)$, where z_1 is the common effect. The memoryless property however reduces $p(y_1|x_1,z_1)$ to either $p(y_1,z_1)$ or $p(y_1|x_1)$, depending on what projection order is used. In the former case, this leads to a lack of discounting. Unless the two bases are exactly at an angle of 45° to each other, the model also predicts violations of the Markov condition, that is, a violation of the fact that $p(y_1|x_1)$ should be equal to $p(y_1|x_0)$, if X and Y are conditionally independent.

To formulate the specific causal networks investigated we set the basis for Y (all inference by participants is made on the variable Y) as the standard basis, and two free parameters (θ_x and θ_z) denote rotations for the basis vectors of X and Z in the 2-dimensional space. The degree of rotation between the different bases determines the conditional and conjunctive probability relationships between them. The rotation is restricted to the first quadrant to reduce any identifiability issues (e.g. a rotation from y_1 of 30° and 330° would result in an identical projection onto y_1). The state vector is not a free parameter but is fixed at a neutral position of 45° to the standard basis, reflecting the assumption that people should not have any preconceived bias towards the presence or absence of the target variable to be inferred, and that the randomized configurations provide no information on the base rate of events.

The model separates two types of inference situations. In the first, inference on Y needs to be made with only one of the other two variables (either X or Z) being known and the other being unknown (situations B, D, E and G in the experiment, see Table 1). Here, the model specification for

$p(y_1)$ is given by $p(y_1 \& \text{Unknown} = 1 | \text{Known}) + p(y_1 \& \text{Unknown} = 0 | \text{Known})$. This is achieved by projecting the state vector $|\psi\rangle$ onto the basis vector representing the known variable value and normalizing it (unit length to reflect the conditional probability), then projecting the normalized vector on to the basis vector representing the unknown variable value '1'. The squared amplitude of this intermediate projection reflects $p(\text{Unknown} = 1 | \text{Known})$. This projection (without normalizing) is then projected again on the basis vector y_1 . This final projection is squared to get the first term on the left, that is, $p(y_1 \& \text{Unknown}=1 | \text{Known})$ in the above probability. The second term is obtained using a similar sequence of operations, with the intermediate projection being to the basis vector representing the unknown variable value '0' instead of '1'. The combination of these probabilities reflects the assumption that people consider the known information, then effectively integrate over the two conjunctive probabilities of the target variable (Y) being '1' and each possibility for the unknown binary variable.

The second type of inference relates to situations where both the remaining variables, X and Z, are known (situations A, C, F and H in the experiment). Here the model specification is similar to that described above, except that the intermediate projection is only made on the known value and is also normalized (projection rescaled to unit length to reflect calculation of conditional probability). Since both X and Z are known, the order of projections can vary and since the model exhibits memoryless properties, this can give rise to order effects between participants, such that participants can calculate either $p(y_1 | \text{Known X, Known Z}) = p(y_1 | \text{Known Z})$ or $p(y_1 | \text{Known Z, Known X}) = p(y_1 | \text{Known X})$. We include a free latent parameter that allows the model to infer the most likely order representation for each individual.

Table 2 lists the probability calculations and projection sequences for each of the situations. For A, C, F, and H, the individual differences in probability estimates can arise from the projection order or rotation parameters. For B, G, D, and E, individual differences arise from the rotation parameters.

Table 2: Probability specification in the 2-d model

Situation	X	Z	Probability Specification
A	1	1	$p(y_1 x_1, z_1)$ or $p(y_1 z_1, x_1)$
B	?	1	$p(y_1 \& x_1 z_1) + p(y_1 \& x_0 z_1)$
C	0	1	$p(y_1 x_0, z_1)$ or $p(y_1 z_1, x_0)$
D	1	?	$p(y_1 \& z_1 x_1) + p(y_1 \& z_0 x_1)$
E	0	?	$p(y_1 \& z_1 x_0) + p(y_1 \& z_0 x_0)$
F	1	0	$p(y_1 x_1, z_0)$ or $p(y_1 z_0, x_1)$
G	?	0	$p(y_1 \& x_1 z_0) + p(y_1 \& x_0 z_0)$
H	0	0	$p(y_1 x_0, z_0)$ or $p(y_1 z_0, x_0)$

The probability of the target variable under any situation A to H is calculated as $p(Y=1|\text{situation})$, based on the appropriate sequences of projections and normalization as

described above. Probabilities for the two situations (say, s_1 and s_2) being compared are calculated separately, assuming the same parameter value for the rotations X and Z under both, so that comparisons are based on a consistent set of beliefs about the entire causal system. The final choice proportions are predicted based on a softmax decision rule, so that choice ratio (\mathbf{h}) for s_1 versus s_2 is given by $\exp(\text{logit}(p(y_1|s_1))/\tau) / \sum_{s=s_1,s_2} \exp(\text{logit}(p(y_1|s))/\tau)$, where τ is the temperature parameter and is fixed to a constant of 1. We use this rule for consistency with the original study (Rehder, 2014) to ensure that any differences in prediction can be attributed to the underlying probability models.

Bayesian Latent Mixture Model

The quantum probability (QP) model requires inference on the rotation parameters for the X and Z bases, as well as the projection orders for each participant. We propose a hierarchical Bayesian latent mixture model to infer these parameters, which allows us to account for individual differences systematically, and build in a hyper-parameter for latent classification of participants between causal and associative reasoning. The latent classification is built by specifying a different set of priors for specific projection orders and rotation parameters.

Parameters θ_z and θ_x represent the rotation of the Z and X bases from the standard Y basis. Recollect that if the rotation $\theta_z < \pi/4$, then $p(y_1|z_1) > p(y_0|z_1)$, and that $p(y_1|z_1) > p(y_1|z_0)$. Thus, θ_z is modeled hierarchically with a probit transformation and scaled to span $[0, \pi/4]$. This range reflects the assumption that participants learn the causal structure of the networks (i.e. that z_1 has a causal influence on y_1), and it is unlikely that they will reverse the implied structure. The rotation θ_x is modeled with a prior range of $[0, \pi/2]$. Recollect that if the rotation $\theta_x = \pi/4$, then there is no causal influence of X . That is, $p(y_1|x_1) = p(y_1|x_0)$ and $p(y_1|x_1) = p(y_0|x_1)$. We allow θ_x to vary on both sides of $\pi/4$, thus allowing individuals to construct positive and negative causal influences.

All hyper-parameters are separately modeled for causal and associative categories, and a latent classification parameter (γ_i) is used to build a mixture model that classifies each individual into a causal or associative category. The binary projection orders for situations A, C, F and H are modeled as Bernoulli processes, with the Bernoulli prior (α) for each being dependent on the causal structure type and the latent classification. The latent classification parameter is itself modeled as a Bernoulli process with equal prior probability of classification to causal or associative categories.

The latent classification mechanism works because the model infers the projection orders that best explain the observed data, and since different orders are given different priors under each category (causal and associative), it selects the category that provides the highest posterior probability for the best projection order. We examine this mechanism in greater detail below. First, we describe how behavior similar to the normative prescriptions of a classical

probability model might be represented under the QP model. Note that when the values of X and Z are both known (situations A, C, F and H), the model exhibits a memoryless property, that is, the projection order of considering X first and then Z reduce to making an inference based on the last seen value of Z only, and vice versa. Ideal normative reasoning for the common cause and chain structures, for these situations (ACFH), then implies that the order of processing is X followed by Z (equivalent to processing only Z), where inference is made by a projection from z_1 to y_1 for situations A and C, and from z_0 to y_1 for situations F and H.

For the common effect structure, an ideal causal reasoner can be represented with a projection sequence from z_0 to y_1 for situations F and H, where no causal effect exists. However, for situations A and C where the value of $Z = 1$ (a common causal effect is known to exist), the projection sequence should include a final projection from the X vector to y_1 , accompanied by ensuring that the rotation for the X basis is significantly more than the rotation for the Z basis (ideally, also greater than $\pi/4$ to ensure that x_0 has a higher level of association with y_1 than x_1 does to reflect discounting). It should be noted that the causal reasoners identified in the cluster analysis in Rehder (2014) were not ideal causal reasoners, and although they demonstrated normative behavioral patterns for the most part, there were also some patterns of deviations. These deviations were more salient in the common effect structure.

To identify individual differences, we do not impose these strict restrictions on causal reasoners (indeed, very few participants, if any, would demonstrate alignment to the strictest criteria). Rather, we allow the model to flexibly account for all types of behavior ranging from highly normative to highly associative. For each structure, we then identify aspects of the model that potentially differentiate causal and associative reasoners along this continuum.

Under the common cause network, we propose that situation C reflects a unique situation where z_1 and x_0 suggest a potential conflict (although not normatively) if the causal structure is assumed to be deterministic. Note that in this experimental paradigm, situation F with z_0 and x_1 does not represent as strong a conflict since causal influence is unidirectional. Similarly, for the chain network structure, situation F represents a potentially conflicting situation if a deterministic causal influence from X (value x_1) to Z (value z_0) is expected, but does not occur. While multiple experiments ensured that deterministic and probabilistic versions of causal influence were tested, differences in behavior across these experimental conditions were not significant. Additionally, even probabilistic causal influence under situation C in the common cause and situation F in the chain network would reflect some form of conflict as measured against the expected direction of causal influence.

In our model, we use these potential sources of conflict to model the latent classification between causal and associative reasoners. Participants who are more likely to process X last under situation C in the common cause and

under situation F in the chain network are more likely to be classified as associative reasoners. Note that this mechanism is not necessary to provide a good fit to the data, but only for adding the latent classification mechanism to the model.

For the common effect structure, the projection orders are not envisaged to affect the classification, which is determined solely by the inferred rotation of the X and Z bases. This is achieved by considering a partition of the prior space for θ_x , with a prior space range of $[\theta_z, \pi/2]$ for causal reasoners (reflecting the fact that causal reasoners will not hold a higher association between the two causes than between the cause and the effect) and $[0, \pi/4]$ for associative ones (reflecting the fact that associative reasoners will hold some positive association between the two causes, leading to anti-discounting behavior). The prior space for θ_z itself is modeled through separate hierarchical distributions. Thus, the inference mechanism decides on a latent classification to the causal category if the resulting posterior for θ_x biased towards values higher than θ_z provides a better account of the data. Implementing these biases in the projection order for situations C and F in the common cause and chain structures respectively, and in the rotation parameters for the common effect structure as partitions of the prior space, enables the latent classification mechanism to categorize participants as causal or associative reasoners depending on which prior space provides the best posterior predictions.

Modeling the empirical data

We fit the model to the mean choice proportions for the 20 inference questions for each of the 315 participants by estimating parameters to the quantum probability model using the Bayesian inference model described above. A version of the normative causal graphical model (CGM) suggested in Rehder (2014) was also used to fit the data within a similar Bayesian inference framework. Figure 3 shows the mean choice proportions (i.e., the number of times a situation was judged to provide stronger inferential support for the presence of the unknown target variable) made by the participants (empirical), compared to the posterior predictive choice proportions generated by the best fitting QP and normative CGM models. Table 3 summarizes the performance of these models.¹

The posterior predictive mean choice proportions are used to calculate the correlation with the actual data and the mean square error (MSE). The Bayesian modeling framework allows us to capture the deviance information criteria (DIC) that assess both model fit and complexity, and use Bayes Factors to compare the QP and normative CGM models at an individual level. As shown in Table 3, the QP posterior predictions show an excellent correlation to empirical data across network structures and types of reasoners (ranging

from 90% to 95%) and provides a significantly better fit compared to the baseline normative CGM model even for the causal reasoners. The QP model yielded an MSE of 0.01 against 0.045 for the normative CGM model.

Table 3: Model Comparison and Performance

	DIC (lower is better)	
Causal Structure	QP	Normative CGM
Common Cause (CC)	1470	1936
Chain (CH)	1547	1998
Common Effect (CE)	1641	2366
	Correlation (model prediction and data)	
Structure / Reasoning	QP	Normative CGM
CC – Causal	92.6	77.2
CC – Associative	93.4	31.2
CH – Causal	93.3	75.9
CH – Associative	93.6	30.7
CE – Causal	90.3	52.5
CE – Associative	95.0	16.0

A Bayes factor (BF) analysis was carried out using the product space method (Lodewyckx, Kim, Lee, Tuerlinckx, Kuppens, & Wagenmakers, 2011). The analysis was inconclusive for 178 of the 315 participants since the BF in favor of the QP model for these participants ranged between 1/5 to 5. No participants had a $BF > 5$ in favor of the CGM model while 137 participants had a $BF > 5$ in favor of the QP model, of whom 103 had a $BF > 100$, showing a significant preference for the QP model.

Explaining Individual Differences

The latent classification mechanism provided a Bayes Factor comparing the evidence for each individual being a causal versus associative reasoner. Based on classifying a participant in a category if the $BF > 1$ in favor of that category, the model identified 265 of the 315 participants (84%) in accordance with the classification provided in the original study (Rehder, 2014). The model classified 32% of the participants as associative (as compared to 23% in the original study) without being provided any details about the base rate. Individual differences in the parameter space were thus successful in identifying cognitive differences. We discussed how projection orders can explain order effects and specifically, violations of the causal Markov condition. The projections orders for situations C and F in particular, for the common cause and chain conditions, were instrumental in the latent classification process. The posterior samples from the model shows that for the common cause structures, causal reasoners were inferred to make a final projection from Z onto y1 84% of the time and associative reasoners only 12% of the time for situation C. Similarly, for situation F under the chain structure, causal reasoners were inferred to make the normative projection from Z to y1 77% of the time, and associative reasoners only 33% of the time. The more frequent projection from X

¹ Overall the models and measured deviance show good convergence ($R < 1.1$), however examining the three individual level parameters across 315 participants shows that the MCMC chains show poor convergence ($R > 1.1$) for about 0.8% of the individual estimates.

to y_1 (67%) for the associative reasoners suggests that the potentially conflicting signal provided by X in these situations was not disregarded, as would be suggested by the causal Markov condition.

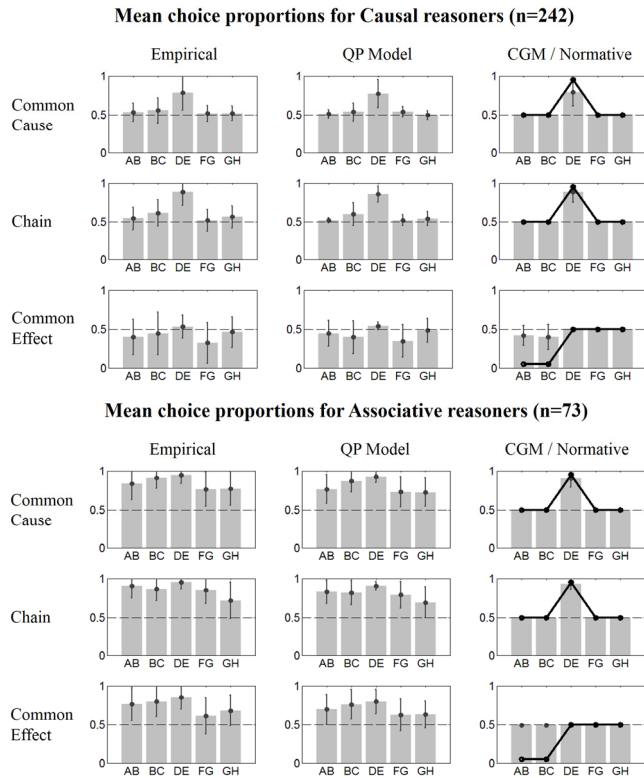


Figure 3: Mean and SD of the mean choice proportions (Normative: Bars indicate best fit normative CGM; Line plots indicate normative prescriptive predictions independent of data)

Individual differences in the common effect structures arise primarily due to differences in the inferred rotation parameters. The mean of the posterior samples for the rotation parameters was $\theta_x = 54^\circ$ and $\theta_z = 43^\circ$ for causal reasoners and $\theta_x = 23^\circ$ and $\theta_z = 27^\circ$ for associative reasoners. Note that the rotations do not reveal the direction of causality but the strength of the bidirectional association between the variables (lower rotation implies higher associative strength). Thus, the significantly lower inferred value for θ_x for associative reasoners signifies the increased influence of x_1 (as the projected value from x_1 to y_1 increases), which in situation A would lead to anti-discounting behavior as empirically observed. The higher inferred value of θ_x (specifically, greater than $\pi/4$) for causal reasoners implies discounting, since any projection from x_1 to y_1 would have a much smaller amplitude.

Conclusion

We showed how a QP model can account for violations of normative properties observed in a causal reasoning task, and how its parameters of the model could be interpreted from a cognitive perspective. Implementing a QP model

within a hierarchical Bayesian inference framework allowed us to develop such a latent classification model with a high level of accuracy as well as compute Bayes factors to compare the QP model to a baseline CGM model. Future work will involve implementing more sophisticated CGM and Bayes net models into a similar framework so that these can explicitly be tested against the QP model.

Acknowledgments

The authors thank Bob Rehder for sharing his data. PKM and JST were supported by NSF grant SES-1326275.

References

Fernbach, P. M., & Rehder, B. (2013). Cognitive shortcuts in causal inference. *Argument and Computation*, 4 (1), 64-88.

Fernbach, P. M., & Sloman, S. A. (2009). Causal learning with local computations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 35, 678-693.

Hagmayer, Y., & Waldmann, M. R. (2002). A constraint satisfaction model of causal learning and reasoning. *W. D. Gray & C. D. Schunn (Eds.), Proceedings of the twenty-fourth annual conference of the cognitive science society* (p. 405-410). Mahwah, NJ: Erlbaum.

Hausman, D. M., & Woodward, J. (1999). Independence, invariance, and the causal markov condition. *British Journal for the Philosophy of Science*, 50 (521-583).

Koehler, J. J. (1996). The base rate fallacy reconsidered: Descriptive, normative, and methodological challenges. *Behavioral and Brain Sciences*, 19 (1), 1-17.

Lodewyckx, T., Kim, W., Lee, M. D., Tuerlinckx, F., Kuppens, P., & Wagenmakers, E. J. (2011). A tutorial on Bayes factor estimation with the product space method. *Journal of Mathematical Psychology*, 55(5), 331-347.

Park, J., & Sloman, S. A. (2013). Mechanistic beliefs determine adherence to the markov property in causal reasoning. *Cognitive Psychology*, 67 (4), 186-216.

Rehder, B. (2014). Independence and dependence in human causal reasoning. *Cognitive Psychology*, 72, 54-107.

Rottman, B. M., & Hastie, R. (2014). Reasoning about causal relationships: Inferences on causal networks. *Psychological Bulletin*, 140 (1), 109-139.

Trueblood, J. S. & Poehos, E. M. (2014). A quantum probability approach to human causal reasoning. *P. Bello, M. Guarini, M. McShane, & B. Scassellati (Eds.), Proceedings of the 36th Annual Conference of the Cognitive Science Society*. (pp. 1616-1621). Austin, TX: Cognitive Science Society.

Waldmann, M. R., Cheng, P. W., Hagmayer, Y., & Blaisdell, A. P. (2008). Causal learning in rats and humans: A minimal rational model. *N. Chater & M. Oaksford (Eds.), The probabilistic mind. Prospects for bayesian cognitive science*. Oxford: University Press.