

Running head: QUANTUM MODELS OF HUMAN CAUSAL REASONING

Quantum Models of Human Causal Reasoning

Jennifer S. Trueblood

University of California, Irvine

Percy K. Mistry

University of California, Irvine

Abstract

Throughout our lives, we are constantly faced with a variety of causal reasoning problems. A challenge for cognitive modelers is developing a comprehensive framework for modeling causal reasoning across different types of tasks and levels of causal complexity. Causal graphical models (CGMs), based on Bayes' calculus, have perhaps been the most successful at explaining and predicting judgments of causal attribution. However, some recent empirical studies have reported violations of the predictions of these models, such as the local Markov condition. In this chapter, we suggest an alternative approach to modeling human causal reasoning using quantum Bayes nets. We show that our approach can account for a variety of behavioral phenomena including order effects, violations of the local Markov condition, anti-discounting behavior, and reciprocity.

Keywords: Causal reasoning, causal graphical models, quantum Bayes nets, Markov condition, order effects

Quantum Models of Human Causal Reasoning

How do people reason about causes and effects? If you wake up in the morning with a stomach ache, how do you infer that it was the seafood you had for dinner rather than stress that caused your stomach to hurt? Human causal reasoning has intrigued scholars as far back as Hume and Kant and currently involves researchers from a variety of fields including cognitive science, developmental psychology, and philosophy. Many researchers approach this topic by developing models that can explain the processes by which people reason about causes and effects. In this chapter, we review these modeling approaches and comment on their strengths and weaknesses. We then introduce a new approach based on quantum probability theory.

Classical probability models of causal reasoning

Some of the first models of causal reasoning were centered around the idea that people use the covariation between causes and effects as a basis for causal judgments (Jenkins & Ward, 1965; Kelley, 1973). These approaches trace their roots back to Hume (1987) and are based on the idea that causation is inferred from the constant conjunction of events as perceived by our sensory system. While models based on covariation can account for many situations, they all face the same ultimate problem - covariation does not necessarily imply causation. As such, these models cannot account for situations where covariational relations are not perceived as causal. For example, we would never think that ice cream consumption causes shark attacks even though shark attacks increase at the same time as ice cream sales (because both increase during summer). To overcome this issue, Cheng (1997) and Novick and Cheng (2004) combined covariational information with domain specific prior knowledge to create the *power PC theory*. According to this theory, reasoners infer causal relations in order to understand observable regularities

between events. The model can explain why covariation sometimes reveals causation but other times does not. While power PC theory has been able to account for many behavioral findings, some studies have shown that people's causal judgments deviate significantly from the predictions made by the model (Lober & Shanks, 2000; White, 2005).

Another approach to modeling causal reasoning uses causal graphical models (CGMs), which represent causal relationships using Bayes' calculus (Kim & Pearl, 1983; Pearl, 1988). CGMs are quite successful at explaining and predicting causal judgments, and their predictions are generally accepted as normative. CGMs can account for causal inferences driven by intervention-based, observational, and counterfactual processes (Hagmayer, Sloman, Lagnado, & Waldmann, 2007), which are often difficult to discriminate in traditional probabilistic models. They have also been used to explain causal learning, where people learn the relationship between variables through observation or intervention (Griffiths & Tenenbaum, 2005, 2009). Some researchers have even combined different probabilistic approaches by integrating power PC theory with CGMs (Griffiths & Tenenbaum, 2005; Lu, Yuille, Liljeholm, Cheng, & Holyoak, 2008). Beyond causal reasoning and learning, CGMs have been applied to decision making (Hagmayer & Sloman, 2009), classification (Rehder & Kim, 2009, 2010) and structured knowledge (Kemp & Tenenbaum, 2009).

While CGMs have been quite successful in accounting for human causal reasoning, several recent empirical studies have reported violations of the predictions of these models. All CGMs must obey a condition called the local Markov property, which states that if we know about all the possible causes of some event Z , then the descendants (i.e., effects) of Z may give us information about Z , but the non-descendants (i.e., noneffects) cannot give us any more information about Z . Recently, several studies have provided evidence that people's causal inferences often violate the local Markov condition (Rottman & Hastie,

2014; Park & Sloman, 2013; Rehder, 2014; Fernbach & Sloman, 2009; Waldmann, Cheng, Hagmayer, & Blaisdell, 2008; Hagmayer & Waldmann, 2002). Relatedly, other studies have shown people often ignore relevant variables. For example, Fernbach, Darlow, and Sloman (2010) found that people ignore alternative causes in predictive causal reasoning (i.e., reasoning about an effect given information about causes), but not in diagnostic causal reasoning (i.e., reasoning about causes given information about the effect).

To overcome the issues mentioned above, CGMs are often elaborated through the inclusion of hidden variables (i.e., latent variables that are not explicitly part of the causal system being studied, but are added to the mental reconstruction of the causal system by individuals as part of their reasoning process). While these elaborated CGMs often provide good accounts of data (Rehder, 2014), they are difficult to conclusively test. Further the inclusion of hidden variables is typically post hoc, added when a basic CGM fails to capture data. As an alternative approach, we suggest expanding the set of probabilistic rules of basic CGMs by using quantum probability theory (Trueblood & Pothos, 2014). Our approach can be considered a generalization of Bayesian causal networks. The essential idea is that any CGM can be generalized to a quantum Bayes net by replacing the probabilities in the classic model with probability amplitudes in the quantum model (Tucci, 1995; Busemeyer & Bruza, 2012). In the next sections, we review CGMs in more detail and introduce quantum Bayes nets as generalizations of these models.

Causal graphical models

CGMs describe causal relationships as directed acyclic graphs (DAGs) representing a set of random variables and their conditional dependencies. For example, suppose that either an unusual dinner or the presence of stress can cause your stomach to hurt. In this example, the three variables - stomach ache, dinner, and stress - are represented as nodes in the DAG (Figure 1). Edges between the nodes represent conditional dependencies. In

Figure 1, edges connect dinner and stomach ache as well as stress and stomach ache. Nodes that are not connected by an edge are conditionally independent. In our example, dinner and stress are conditionally independent and thus there is no edge connecting them.

The probability of a node taking a particular value is determined by a probability function that takes as input the values of any parent nodes. These probabilities are specified in conditional probability tables. Consider the stomach scenario where all three variables have two possible values: A = stomach ache is present (true/false), D = dinner is unusual (true/false), S = stress is present (true/false). The CGM can answer questions such as “What is the probability that dinner was unusual, given that your stomach aches?” by using the formula for conditional probability:

$$p(D = t|A = t) = \frac{p(A = t, D = t)}{p(A = t)} = \frac{\sum_{j \in \{t, f\}} p(A = t, D = t, S = j)}{\sum_{i, j \in \{t, f\}} p(A = t, D = i, S = j)} \quad (1)$$

where the joint probability function

$p(A = t, D = i, S = j) = p(A = t|D = i, S = j)p(D = i)p(S = j)$ because D and S are conditionally independent. We can now calculate the desired probability $p(D = t|A = t)$ using the conditional probability tables in Figure 1:

$$p(D = t|A = t) = \frac{(.8 \times .1 \times .6) + (.7 \times .1 \times .4)}{(.8 \times .1 \times .6) + (.7 \times .1 \times .4) + (.2 \times .9 \times .6) + (.1 \times .9 \times .4)} \approx .346 \quad (2)$$

The probabilities of other combinations of variables (e.g., $p(D = f|A = f)$) follow similar calculations.

All CGMs obey the local Markov property, which states that any node in a Bayesian network is conditionally independent of its non-descendants (i.e., noneffects) given its parents (i.e., direct causes). Consider a situation where a variable X causes Y and Z (represented by the DAG: $Y \leftarrow X \rightarrow Z$). The local Markov property implies that if you

know X , then Y provides no additional information about the value of Z .

Mathematically, we have $p(Z|X) = p(Z|X, Y)$. There is empirical evidence that people's causal judgments do not always obey the local Markov property. For example, Rehder (2014) presented participants with causal scenarios involving three variables (e.g., an economic scenario with variables describing interest rates, trade deficits, and retirement savings) and asked them to infer the value of an unknown target variable given information about one or two of the remaining variables. Rehder found that information about non-descendants influenced judgments even when the values of the parent nodes (i.e., direct causes) was known, showing a direct violation of the local Markov property. (We discuss this experiment in more detail in a later section.)

In order to account for the observed violations of the local Markov property, Rehder (2014) augmented CGMs by including an additional variable that served as either a shared disabler, shared mediator, or shared cause. For example, in a common cause structure where X causes Y and Z (i.e., $Y \leftarrow X \rightarrow Z$), the structure can be elaborated in several different ways by the inclusion of a fourth variable W as shown in Figure 2. While such an approach can provide a good account of data, it is difficult to conclusively test because participants are never questioned about the hidden variable W . As an alternative approach, we propose generalizing CGMs to quantum Bayes nets.

Quantum Bayes nets

In our quantum Bayes nets, we replace the classical probabilities in the conditional probability tables of a CGM with quantum probability amplitudes as proposed by Tucci (2012, 1995) and Moreira and Wichert (2014). Consider the situation where there are two causally related variables X and Y such that $X \rightarrow Y$. Further assume that these two variables are binary (true/false). In quantum probability theory, the observables X and Y

are represented by Hermitian operators:

$$X = x_t P_{x_t} + x_f P_{x_f} \quad (3a)$$

$$Y = y_t Q_{y_t} + x_f Q_{y_f} \quad (3b)$$

where x_i and y_i are eigenvalues and P_{x_i} and Q_{y_i} are projectors onto corresponding eigen-subspaces. The probability of a concrete value, such as x_t (we use the notation x_t as shorthand for $X = t$), is given by Born's rule:

$$p_\rho(x_t) = \langle P_{x_t} \psi | \psi \rangle = \|P_{x_t} \psi\|^2 \quad (4)$$

where ψ is a pure state and $\rho = |\psi\rangle\langle\psi|$ is the corresponding density operator. Suppose we want to answer the question “What is the probability Y is false, given that X is true?”. In this situation, we first calculate the output state ρ_{x_t} as defined by the projection postulate (see the chapter *A Brief Introduction to Quantum Formalism*) and then apply Born's rule:

$$p_\rho(Y = f | X = t) = p_{\rho_{x_t}}(y_f) = \langle Q_{y_f} \psi_{x_t} | \psi_{x_t} \rangle = \|Q_{y_f} \psi_{x_t}\|^2. \quad (5)$$

We then use these conditional probabilities for our network rather than the classical ones used in a CGM.

Consider the stomach ache scenario again. In the classical model, in order to answer the question “What is the probability that dinner was unusual, given that your stomach aches?”, we needed to calculate joint probabilities such as $p(A = t, D = i, S = j)$. We can determine these probabilities from the conditional probability tables of the CGM by writing $p(A = t, D = i, S = j) = p(A = t | D = i, S = j)p(D = i)p(S = j)$. We take a similar approach in our quantum Bayes net. First, let the three observables, stomach ache, dinner, and stress, be represented by Hermitian operators A , D , and S with the

respective projectors P , Q , and R . Now, we define joint probabilities by Born's rule:

$$p_\rho(A = t, D = i, S = j) = p_{\rho_{d_i, s_j}}(a_t) p_\rho(d_i) p_\rho(s_j) = \|P_{a_t} \psi_{d_i, s_j}\|^2 \|Q_{d_i} \psi\|^2 \|R_{s_j} \psi\|^2 \quad (6)$$

where the output state is given by

$$\psi_{d_i, s_j} = \frac{R_{s_j} Q_{d_i} \psi}{\|R_{s_j} Q_{d_i} \psi\|}. \quad (7)$$

If the observables D and S do not commute, then the output state will depend on the order in which these two variables are considered so that $\psi_{d_i, s_j} \neq \psi_{s_j, d_i}$. As a consequence, $p(A = t | D = i, S = j) \neq p(A = t | S = j, D = i)$.

Figure 3 shows a quantum Bayes net generalization of the stomach ache CGM shown in Figure 1. For this example, the probabilities in the CGM have been replaced by probability amplitudes in the quantum Bayes net. These amplitudes are related to classical probabilities by taking the squared magnitude of the amplitudes. For example, the probability that dinner was unusual is given by

$$\begin{aligned} p(D = t) &= \|.3162e^{i\theta dt}\|^2 = (.3162e^{i\theta dt})(\overline{.3162e^{i\theta dt}}) = (.3162e^{i\theta dt})(.3162e^{-i\theta dt}) \\ &= (.3162)^2 e^{i(\theta dt - \theta dt)} = .1 \end{aligned} \quad (8)$$

which is the same as the classical probability in the CGM. Note that the term $e^{i\theta dt}$ is simply the phase of the amplitude.

When determining the conditional probabilities of the stomach ache given information about dinner and stress, the order in which dinner and stress are considered matters. In the quantum Bayes net in Figure 3, we assume that information about dinner is always processed before information about stress. Psychologically, we would say that an individual thinks about dinner and stress separately, always starting with dinner. In the

figure, we used a thick border on the dinner node to indicate that this variable is processed first. If we wish to switch the order and have stress processed before dinner, then we would need to define a different set of conditional probabilities. In other words, we have two different conditional probability tables describing the probability of the stomach ache given information about dinner and stress - one table describing the situation where dinner is considered before stress (as shown in Figure 3) and another table describing the situation where stress is considered before dinner (not shown in the figure). Even though there are two different conditional probability tables for the quantum version of the stomach ache scenario, these tables are related to one another. In quantum probability theory, noncommutative observables (such as dinner and stress) are related by a unitary transformation, which preserves lengths (the state vector must have length equal to one) and inner products.

For the stomach ache scenario, we started with a CGM and generalized this to a quantum Bayes net by designating a processing order (dinner before stress) and changing the classical probabilities into probability amplitudes. Note that our decision that dinner should be processed before stress was arbitrary. We could have easily specified the reverse order (stress processed before dinner). Thus, there are at least two different ways to generalize the CGM in this example. In general, there will often be multiple ways to generalize a CGM to a quantum Bayes net. As a consequence, if we start with a quantum Bayes net, it is not necessarily the case that we can derive a well-defined CGM. The conditional probability tables of a quantum Bayes net will always have classical probability analogs, which are derived by squaring the probability amplitudes in the quantum tables. However, when a quantum Bayes net involves noncommutative observables, the corresponding CGM is ill-defined. This is because noncommutative observables result in different conditional probabilities tables for the same causal situation. This is not allowed in a traditional CGM. Thus, the behavior of a quantum

Bayes net will often be fundamentally different than the behavior of a CGM.

Implications

Noncommutative quantum Bayes nets make several interesting predictions about human behavior. In the next sections, we discuss these predictions and supporting empirical evidence.

Order effects

Quantum Bayes nets with noncommutating observables naturally predict order effects. Consider a causal scenario where X and Y cause Z (represented by the DAG: $X \rightarrow Z \leftarrow Y$). In an experiment, participants might be asked to judge $p(Z|X, Y)$ where information about X precedes information about Y . An order effect occurs when final judgments depend on the sequence of information so that $p(Z|X, Y) \neq p(Z|Y, X)$. Classical probability models such as CGMs have difficulty accounting for order effects due to the commutative property because $p(X, Y|Z) = p(Y, X|Z)$ implies $p(Z|X, Y) = p(Z|Y, X)$ by Bayes' rule. To account for order effects, classical probability models need to introduce extra events such as O_1 that X is presented before Y and O_2 that Y is presented before Z . Then, it is possible that $p(Z|X, Y, O_1) \neq p(Z|X, Y, O_2)$. However, without a theory about O_1 and O_2 , this approach simply redescribes the empirical result. Further, in many empirical studies of order effects, the order of presentation is randomly determined so that order information such as O_1 and O_2 is irrelevant.

A large number of empirical studies have shown that order of information plays a crucial role in human judgments (Hogarth & Einhorn, 1992). Order effects arise in a number of different situations ranging from judging the guilt of a defendant in a mock trial (Furnham, 1986; Walker, Thibaut, & Andreoli, 1972) to judging the likelihood of selecting balls from urns (Shanteau, 1970). Recently, Trueblood and Busemeyer (2011)

found evidence for order effects in causal reasoning. In this experiment, participants made causal judgments about ten different scenarios where there was a single effect and two binary (present/absent) causes. For example, in one scenario, participants were asked about the likelihood of a fictitious person, Mary, losing weight over the next month (the effect) given that she did not make any changes to her diet (absent cause) and began an exercise program (present cause).

The participants ($N = 113$) provided likelihood judgments of the effect (e.g., Mary losing weight) on a 0 to 100 scale at three different times: (1) before reading either cause, (2) after reading one of the causes, and (3) after reading the remaining cause. Participants judged the present cause before the absent cause for a random half of the scenarios. The order of the causes was reversed (i.e., absent followed by present) for the other half of the scenarios. The results of the experiment showed a large, significant order effect ($p < 0.001$) across the ten scenarios. The presence of order effects in causal judgments provides support for quantum Bayes nets with noncommutating observables.

Violations of the local Markov condition

The local Markov condition of CGMs stipulates that any node in a DAG is conditionally independent of its nondescendants when its direct causes are known. For example, in the common effect structure $X \rightarrow Z \leftarrow Y$, this property implies that the two causes X and Y are conditionally independent. In other words, if X and Y are binary, then $p(Y = i|X = t) = p(Y = i|X = f)$ for $i \in \{t, f\}$ and similarly when X and Y are swapped. In a quantum Bayes net where X and Y do not commute, there is a natural dependency between these two variables. That is, knowing the value of X influences our beliefs about Y . This dependency can result in violations of the local Markov condition so that $p(Y = i|X = t) \neq p(Y = i|X = f)$. By the definition of conditional probability, $p(Y = i|X = j) = p(Y = i, X = j)/p(X = j)$. In a CGM, X and Y are independent so

that the joint probability $p(Y = i, X = j) = p(Y = i)p(X = j)$. Thus,

$p(Y = i|X = j) = p(Y = i)$ for all i, j . In a quantum Bayes net,

$p(Y = i, X = j) = \|Q_{y_i}P_{x_j}\psi\|^2$. If X and Y do not commute, then it is clearly the case that $\|Q_{y_i}P_{x_j}\psi\|^2 \neq \|Q_{y_i}\psi\|^2\|P_{x_j}\psi\|^2$, leading to violations of the local Markov condition.

Rehder (2014) empirically demonstrated that people often violate the local Markov condition in their causal judgments. In his task, participants were given two causal situations with an unknown target variable and were asked to select the situation where the target variable was more probable. For example, in the common effect structure $X \rightarrow Z \leftarrow Y$, participants had to decide whether the target variable Y was more likely to be true in a situation where $X = t$ or in a situation where $X = f$. According to CGMs, we expect participants' choice proportions to be equal on average because $p(Y = t|X = t) = p(Y = t|X = f)$. However, Rehder (2014) found that on average people selected the causal situation where $X = t$ more often than the one where $X = f$, suggesting people judged $p(Y = t|X = t) > p(Y = t|X = f)$. He also showed similar violations with other causal structures such as chain structures ($X \rightarrow Y \rightarrow Z$) and common cause structures ($Y \leftarrow X \rightarrow Z$).

Anti-discounting behavior

Noncommutating observables can also account for anti-discounting behavior in causal reasoning. The term discounting refers to the situation where one cause casts doubts on another cause. In the common effect structure $X \rightarrow Z \leftarrow Y$, knowing the value of X could cast doubt on the value of Y such that $p(Y|Z, X) < p(Y|Z)$. In many causal scenarios, discounting is considered normatively correct (Morris & Larrick, 1995). For example, it is normatively correct to judge $p(Y = t|Z = t) > p(Y = t|Z = t, X = t)$ because knowing $X = t$ sufficiently explains the value of the effect $Z = t$ and consequently renders the other cause Y redundant. When X is unknown, as in $p(Y = t|Z = t)$, there is

a greater chance the effect was brought about by Y .

Rehder (2014) found that many people display anti-discounting behavior. That is, people judge an unknown target cause Y as highly likely based on the presence of the alternative cause $X = t$, resulting in judgments where

$p(Y = t|Z = t) < p(Y = t|Z = t, X = t)$. Similar to violations of the local Markov

property, quantum Bayes nets can explain anti-discounting behavior by the

noncommutativity of X and Y , which produces a causal dependency between these variables.

Reciprocity

The term reciprocity describes the situation where a person judges the probability of one variable given another to be the same as the probability when the variables are swapped, $p(X|Y) = p(Y|X)$. This phenomenon is similar to the inverse fallacy (Koehler, 1996; Villejoubert & Mandel, 2002) where people equate posterior and likelihood probabilities. If H represents a hypothesis and D represents data, the inverse fallacy occurs when $p(H|D) = p(D|H)$, where the first term is the posterior and the second term is the likelihood. The inverse fallacy has been observed in a number of different medical judgment problems, where clinicians are asked to judge the likelihood of a disease based on a set of symptoms (Meehl & Rosen, 1955; Hammerton, 1973; Liu, 1975; Eddy, 1982). The fallacy has also been demonstrated in the famous *taxicab problem* (Kahneman & Tversky, 1972), where individuals are asked to judge the likelihood that a cab was in a wreck given its color (blue or green). Results of this experiment showed that most people judged $p(H|D)$ as $p(D|H)$.

The *law of reciprocity* (Peres, 1998) in quantum probability theory stimulates that if two events X and Y are represented by single dimensional subspaces, then $p(X|Y)$ is equivalent to $p(Y|X)$. We illustrate this result in Figure 4 for the probabilities $p(x_t|y_t)$

and $p(y_t|x_t)$. In the left panel, we calculate $p(x_t|y_t)$ by first projecting the state ψ onto the y_t subspace and then normalizing to produce the output state ψ_{y_t} . This new state is then projected onto the x_t subspace and the conditional probability is the length of this projection squared as defined by Born's rule (represented by the thick black bar in the figure). In the right panel, we calculate $p(y_t|x_t)$ following a similar procedure. First, we project the state onto the x_t subspace and then normalize to produce the output state ψ_{x_t} . This revised state is then projected onto the y_t subspace and the conditional probability is the length of the projection squared. As shown in the figure, the two conditional probabilities are the same (i.e., the thick black bars are the same length). Note that not all quantum models can account for reciprocity and the inverse fallacy. Only quantum models that make the specific assumption that different outcomes are represented by single dimensional subspaces can explain these findings.

Conclusions

One could argue that CGMs have been one of the most successful approaches in modeling human causal reasoning. These models can account for casual deductive and inductive reasoning in a large number of situations. Besides causal reasoning, CGMs have been applied to a variety of other domains including classification (Rehder, 2003; Rehder & Kim, 2009, 2010) and decision-making (Hagmayer & Sloman, 2009).

However, there has been recent evidence that people's judgments often deviate from the rules of CGM's. There are at least two possible ways to modify CGMs in order to account for these findings. One method involves elaborating CGMs through the inclusion of additional nodes and edges in the network. These hidden variables provide flexibility to the models and help them accommodate a wider range of human behavior. However, the addition of hidden variables to a CGM is often ad hoc and these additional variables are difficult to conclusively test.

In this chapter, we suggest an alternative approach using quantum probability theory. Instead of elaborating a CGM with extra nodes and edges, we suggest changing the probabilistic rules used to perform inference. In our approach, we replace the classical probabilities of a CGM with quantum ones to yield quantum Bayes nets. By using quantum probabilities, we allow for variables to be noncommutative. We show that quantum Bayes nets with noncommutative observables can account for a variety of different behavioral phenomena including order effects, violations of the local Markov condition, anti-discounting behavior, and reciprocity.

Quantum probability theory has successfully explained numerous findings in cognition and decision-making including violations of the sure thing principle (Pothos & Busemeyer, 2009), interference effects in perception (Conte, Khrennikov, Todarello, Federici, & Zbilut, 2009), conjunction and disjunction fallacies (Busemeyer, Pothos, Franco, & Trueblood, 2011), violations of dynamic consistency (Busemeyer, Wang, & Trueblood, 2012), interference of categorization on decision-making (Busemeyer, Wang, & Lambert-Mogiliansky, 2009), and order effects in survey questions (Wang & Busemeyer, 2013). We feel that quantum probability theory also has great potential to explain human causal reasoning. The results we considered here make us optimistic about this approach in the future.

References

- Busemeyer, J. R., & Bruza, P. D. (2012). *Quantum models of cognition and decision*. Cambridge, UK: Cambridge University Press.
- Busemeyer, J. R., Pothos, E., Franco, R., & Trueblood, J. S. (2011). A quantum theoretical explanation for probability judgment errors. *Psychological Review*, *118*, 193-218.
- Busemeyer, J. R., Wang, Z., & Lambert-Mogiliansky, A. (2009). Comparison of markov and quantum models of decision making. *Journal of Mathematical Psychology*, *53*, 423-433.
- Busemeyer, J. R., Wang, Z., & Trueblood, J. S. (2012). Hierarchical bayesian estimation of quantum decision model parameters. In J. R. Busemeyer (Ed.), *Qi 2012, Incs 7620*. Berlin, Germany: Springer-Verlag.
- Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, *104*, 367-405.
- Conte, E., Khrennikov, Y. A., Todarello, O., Federici, A., & Zbilut, J. P. (2009). Mental states follow quantum mechanics during perception and cognition of ambiguous figures. *Open Systems & Information Dynamics*, *16*, 1-17.
- Eddy, D. (1982). Probabilistic reasoning in clinical medicine: Problems and opportunities. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases* (p. 249-267). Cambridge: Cambridge University Press.
- Fernbach, P. M., Darlow, A., & Sloman, S. A. (2010). Neglect of alternative causes in predictive but not diagnostic reasoning. *Psychological Science*, *21* (3), 329-336.
- Fernbach, P. M., & Sloman, S. A. (2009). Causal learning with local computations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *35*, 678-693.

- Furnham, A. (1986). The robustness of the recency effect: Studies using legal evidence. *Journal of General Psychology, 113*, 351–357.
- Griffiths, T. L., & Tenenbaum, J. B. (2005). Structure and strength in causal induction. *Cognitive Psychology, 51*, 334–384.
- Griffiths, T. L., & Tenenbaum, J. B. (2009). Theory-based causal induction. *Psychological Review, 116*(4), 661–716.
- Hagmayer, Y., & Sloman, S. A. (2009). Decision makers conceive of themselves as interveners. *Journal of Experimental Psychology: General, 128*, 22–38.
- Hagmayer, Y., Sloman, S. A., Lagnado, D. A., & Waldmann, M. R. (2007). Causal reasoning through intervention. In A. Gopnik & L. Schulz (Eds.), *Causal learning: Psychology, philosophy, and computation* (p. 86–100). Oxford: Oxford University Press.
- Hagmayer, Y., & Waldmann, M. R. (2002). A constraint satisfaction model of causal learning and reasoning. In 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.
- Hammerton, M. (1973). A case of radical probability estimation. *Journal of Experimental Psychology, 101*, 252–254.
- Hogarth, R. M., & Einhorn, H. J. (1992). Order effects in belief updating: The belief-adjustment model. *Cognitive Psychology, 24*, 1–55.
- Hume, D. (1987). *A treatise of human nature (2nd ed.)*. Oxford, England: Clarendon Press. (Original work published 1739).
- Jenkins, H. M., & Ward, W. C. (1965). Judgment of contingency between responses and outcomes. *Psychological Monographs: General and Applied, 79*, 1–17.
- Kahneman, D., & Tversky, A. (1972). On prediction and judgment. *ORI Research Monographs, 12*.

- Kelley, H. H. (1973). The processes of causal attribution. *American Psychologist*, *28*, 107-128.
- Kemp, C., & Tenenbaum, J. B. (2009). Structured statistical models of inductive reasoning. *Psychological Review*, *116*, 20-58.
- Kim, J. H., & Pearl, J. (1983). A computational model for causal and diagnostic reasoning in inference systems. In *Proceedings of the 8th international joint conference on artificial intelligence (ijcai)* (pp. 190–193).
- Koehler, J. J. (1996). The base rate fallacy reconsidered: Descriptive, normative, and methodological challenges. *Behavioral and Brain Sciences*, *19*(1), 1-17.
- Liu, A. Y. (1975). Specific information effect in probability estimation. *Perceptual and Motor Skills*, *41*, 475-478.
- Lober, K., & Shanks, D. R. (2000). Is causal induction based on causal power? critique of cheng (1997). *Psychological Review*, *107*(1), 195-212.
- Lu, H., Yuille, A. L., Liljeholm, M., Cheng, P. W., & Holyoak, K. J. (2008). Bayesian generic priors for causal learning. *Psychological review*, *115*(4), 955.
- Meehl, P., & Rosen, A. (1955). Antecedent probability and the efficiency of psychometric signs of patterns, or cutting scores. *Psychological Psychological Bulletin*, *52*, 194-215.
- Moreira, C., & Wichert, A. (2014). Interference effects in quantum belief networks. *Applied Soft Computing*, *25*, 64–85.
- Morris, M. W., & Larrick, R. P. (1995). When one cause casts doubt on another: A normative analysis of discounting in causal attribution. *Psychological Review*, *102*(2), 331-355.
- Novick, L. R., & Cheng, P. W. (2004). Assessing interactive causal influence. *Psychological Review*, *111*(2), 455–485.
- Park, J., & Sloman, S. A. (2013). Mechanistic beliefs determine adherence to the markov

- property in causal reasoning. *Cognitive Psychology*, 67(4), 186–216.
- Pearl, J. (1988). *Probabilistic reasoning in intelligent systems: Networks of plausible inference*. Morgan Kaufmann.
- Peres, A. (1998). *Quantum theory: Concepts and methods*. New York: Kluwer Academic.
- Pothos, E. M., & Busemeyer, J. R. (2009). A quantum probability explanation for violations of 'rational' decision theory. *Proceedings of the Royal Society B*, 276(1165), 2171-2178.
- Rehder, B. (2003). Categorization as causal reasoning. *Cognitive Science*, 27, 709-748.
- Rehder, B. (2014). Independence and dependence in human causal reasoning. *Cognitive Psychology*, 72, 54–107.
- Rehder, B., & Kim, S. (2009). Classification as diagnostic reasoning. *Memory and Cognition*, 37, 715-729.
- Rehder, B., & Kim, S. (2010). Causal status and coherence in causal-based categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 36, 1171-1206.
- Rottman, B. M., & Hastie, R. (2014). Reasoning about causal relationships: Inferences on causal networks. *Psychological Bulletin*, 140(1), 109–139.
- Shanteau, J. C. (1970). An additive model for sequential decision making. *Journal of Experimental Psychology*, 85, 181–191.
- Trueblood, J. S., & Busemeyer, J. R. (2011). A quantum probability account of order effects in inference. *Cognitive Science*, 35, 1518-1552.
- Trueblood, J. S., & Pothos, E. M. (2014). A quantum probability approach to human causal reasoning. In P. Bello, M. Guarini, M. McShane, & B. Scassellati (Eds.), *Proceedings of the 36th annual conference of the cognitive science society* (p. 1616-1621). Austin, TX: Cognitive Science Society.
- Tucci, R. R. (1995). Quantum bayesian nets. *International Journal of Modern Physics*,

B, 9, 295-337.

- Tucci, R. R. (2012). An introduction to quantum bayesian networks for mixed states. *arXiv preprint arXiv:1204.1550*.
- Villejoubert, G., & Mandel, D. R. (2002). The inverse fallacy: An account of deviations from bayes's theorem and the additivity principle. *Memory & Cognition*, 30(2), 171–178.
- Waldmann, M. R., Cheng, P. W., Hagmayer, Y., & Blaisdell, A. P. (2008). Causal learning in rats and humans: A minimal rational model. In N. Chater & M. Oaksford (Eds.), *The probabilistic mind. prospects for bayesian cognitive science* (pp. 453–484). Oxford: Oxford: University Press.
- Walker, L., Thibaut, J., & Andreoli, V. (1972). Order of presentation at trial. *Yale Law Journal*, 82, 216–226.
- Wang, Z., & Busemeyer, J. R. (2013). A quantum question order model supported by empirical tests of an a priori and precise prediction. *Topics in Cognitive Science*, 5, 689-710.
- White, P. A. (2005). The power pc theory and causal powers: Comment on cheng (1997) and novick and cheng (2004). *Psychological Review*, 112(3), 675–682.

Author Note

JST and PKM was supported by NSF grant SES-1326275.

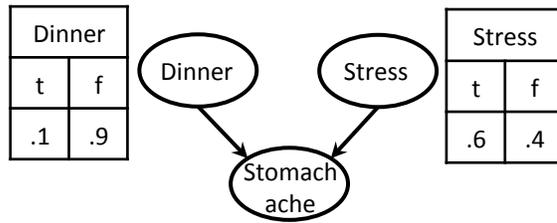
Figure Captions

Figure 1. A CGM of the stomach ache scenario. There are two possible causes of a stomach ache - an unusual dinner or stress. The three variables are represented as a DAG with conditional probability tables.

Figure 2. Three different ways to elaborate a common cause structure with an additional variable W .

Figure 3. A quantum Bayes net generalization of the stomach ache scenario. The dinner node in the DAG has a thick border to indicate that it is considered before stress. The tables contain probability amplitudes rather than probabilities. These amplitudes were determined from the CGM shown in Figure 1.

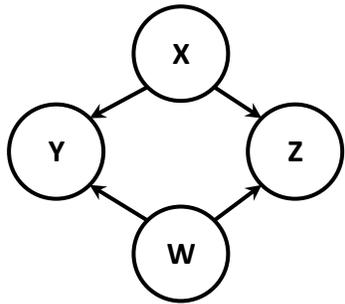
Figure 4. The *law of reciprocity* in quantum probability theory. In the left panel, $p(x_t|y_t)$ is calculated by a series of two projections. First, ψ is projected onto the y_t subspace (labelled projection 1) and then normalized to yield the output state ψ_{y_t} . Next, the output state is projected onto the x_t subspace (labelled projection 2), and the probability is calculated by squaring the length of the projection (represented by the thick black line). In the right panel, the series of projections is reversed to calculate $p(y_t|x_t)$.



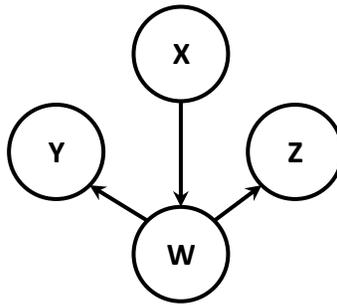
Dinner	
t	f
.1	.9

Stress	
t	f
.6	.4

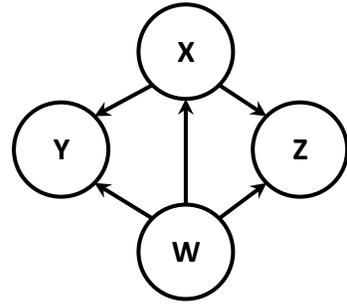
		Stomach ache	
Dinner	Stress	t	f
t	t	.8	.2
t	f	.7	.3
f	t	.2	.8
f	f	.1	.9



Shared Disabler

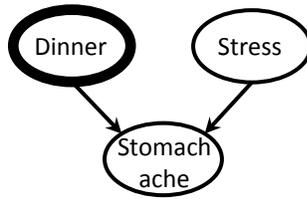


Shared Mediator



Shared Cause

Dinner	
t	f
$.3162e^{i\theta dt}$	$.9487e^{i\theta df}$



Stress	
t	f
$.7746e^{i\theta st}$	$.6325e^{i\theta sf}$

		Stomach ache	
Dinner	Stress	t	f
t	t	$.8944e^{i\theta at dt,st}$	$.4472e^{i\theta af dt,st}$
t	f	$.8367e^{i\theta at dt,sf}$	$.5477e^{i\theta af dt,sf}$
f	t	$.4472e^{i\theta at df,st}$	$.8944e^{i\theta af df,st}$
f	f	$.3162e^{i\theta at df,sf}$	$.9487e^{i\theta af df,sf}$

