

The role of causal models in multiple judgments under uncertainty

Brett K. Hayes^{a*}, Guy E. Hawkins^a, Ben R. Newell^a, Martina Pasqualino^a, Bob Rehder^b

^a School of Psychology, The University of New South Wales, NSW, 2052, Australia

^b Department of Psychology, New York University, 6 Washington Place, New York, NY
10003, USA

* **Corresponding author:** School of Psychology, The University of New South Wales.
NSW, 2052, Australia, Fax: +61 2 9385 3641, email address: b.hayes@unsw.edu.au

Abstract

Two studies examined a novel prediction of the causal Bayes net approach to judgments under uncertainty, namely that causal knowledge affects the interpretation of statistical evidence obtained over multiple observations. Participants estimated the conditional probability of an uncertain event (breast cancer) given information about the base rate, hit rate (probability of a positive mammogram given cancer) and false positive rate (probability of a positive mammogram in the absence of cancer). Conditional probability estimates were made after observing one or two positive mammograms. Participants exhibited a *causal stability effect*: there was a smaller increase in estimates of the probability of cancer over multiple positive mammograms when a causal explanation of false positives was provided. This was the case when the judgments were made by different participants (Experiment 1) or by the same participants (Experiment 2). These results show that beliefs about the generative causes of observed statistics affect the interpretation of dependencies between multiple observations in intuitive judgment tasks.

Keywords: Causal models, Bayes nets, Judgment under uncertainty, Intuitive statistics

1. Introduction

Causal knowledge plays a central role in cognition. Such knowledge has profound effects on the way people learn contingencies (Cheng, 1997; Gopnik Glymour, Sobel, Schultz, & Kushnir, 2004; Griffiths & Tenenbaum, 2005; Waldmann, Hagmayer & Blaisdell, 2006), categorize (Ahn & Kim, 2001; Hayes & Rehder, 2012; Rehder & Kim 2010; Sloman, Love, & Ahn, 1998), reason (Fernbach, Darlow, & Sloman, 2011; Holyoak, Lee, & Lu, 2010; Kemp & Tenenbaum, 2009; Kemp, Shafto, & Tenenbaum, 2012; Rehder, 2006; Sloman, 2005), make decisions (Hagmayer & Sloman, 2009), and remember (Shank & Abelson, 1995). This paper examines how causal knowledge affects the way people interpret statistical information in judgments under uncertainty.

Many judgments under uncertainty require the evaluation of statistical information to arrive at an estimate of the probability of an outcome (Newell, 2013). Performance on such tasks is often poor with participants generating estimates that deviate considerably from a normative Bayesian solution (see Barbey & Sloman, 2007; Koehler, 1996; Tversky & Kahneman, 1974, for reviews). Early work suggested that causal knowledge might have an important role in such judgments (Ajzen, 1977; Tversky & Kahneman, 1980). For example, when base rate information is seen as causally relevant to a judgment it is less likely to be neglected (Ajzen, 1977; Bar-Hillel, 1980). Ajzen (1977) attributed this to a “causality” heuristic which increases the salience of statistical information associated with causal mechanisms relative to non-causal statistics.

A more elaborate account of the role of causal knowledge on judgments under uncertainty is suggested by causal Bayes net models of cognition (Griffiths & Tenenbaum, 2005; Rehder, 2010; Sloman, 2005; Waldmann, 1996). The causal Bayes net approach assumes that people make inferences by constructing hypotheses about the causal relations between the components of a judgment problem (often referred to as “variables”) and

updating these hypotheses in the light of observed data. In this approach, hypothesized causal relations can be represented as Bayesian networks – directed acyclic graphs in which the nodes represent variables in the system and linking arrows represent the causal relations between these variables. This graph structure can be used to infer conditional and joint probabilities over the variables in the network. Such an approach has been successfully applied to the explanation of key phenomena in domains such as learning about causal systems (Kemp, Goodman & Tenenbaum, 2010; Sloman & Lagnado, 2005; Waldmann, et al., 2006), reasoning (Kemp & Tenenbaum, 2009) and conceptual development (Gopnik et al., 2004; Tenenbaum, Kemp, Griffiths & Goodman, 2011).

Krynski and Tenenbaum (2007) proposed that the causal Bayes net approach could also lead to a better understanding of how people make judgments under uncertainty (also see Bes, Sloman, Lucas & Raufaste, 2012). Krynski and Tenenbaum suggest that failures to arrive at normative probability estimates in such judgments reflect a difficulty in mapping the statistics given in the problem to the relevant components of intuitive causal models. They show that performance on such problems may be improved by framing the task so that this mapping is more transparent (e.g., by providing a causal explanation of key statistics).

1.1 The causal Bayes net approach and uncertain judgments with multiple observations

The current studies focus on a different implication of the approach to judgments under uncertainty; namely how causal beliefs affect the interpretation of evidence arising from multiple observations. Our key intuition is that statistically equivalent information provided over multiple observations will be interpreted differently depending on one's causal beliefs about the variables generating the observations. In particular, we predict that repeated observations with no obvious cause will be treated as independent, stochastic events. When a common causal explanation is available for these observations however, the repeated

observations will be seen as dependent, arising from the same generative process. These ideas were suggested by Krynski and Tenenbaum (2007) but not tested empirically.

As a simple example, consider a situation where you turn on your laptop computer and attempt to launch your internet browser but receive an error message saying that no connection could be established. There are at least two interpretations of this event. It could reflect the somewhat random fluctuation in strength of wireless server signals as the laptop is moved around. Alternately, there may be a more serious and stable underlying cause, such as a failure of the remote server or suspension of your server account. In the absence of an obvious cause you may favor the former hypothesis and make several attempts to restart the browser. In this case, you are treating browser failures as independent events such that each subsequent attempt is as likely to succeed as the first. On the other hand if, after the first try, you recall that you have not paid the most recent bill from your internet provider, the browser failure may be seen as evidence for an alternative cause (i.e. that your account has been suspended). In this case it seems futile to continue restarting the browser because the underlying cause of the failure remains unchanged and will lead to the same outcome on each occasion. Hence, browser failures are interpreted differently depending on your mental model of the causal dynamics of the situation.

1.2 Modeling multiple observations in the mammogram problem

We now examine how this approach can be applied to a classic problem in judgment under uncertainty. In the “mammogram problem,” the task is to estimate the conditional probability that a woman has cancer given that she has received a positive mammogram (cf. Eddy, 1982; Gigerenzer & Hoffrage, 1995; Krynski & Tenenbaum, 2007). To derive this estimate, participants are given information about the base rate (probability of breast cancer in the target population), the “hit” rate (probability of obtaining a positive mammogram given cancer), and the false positive rate (probability of obtaining a positive result in the absence of

cancer). In our variant of this problem (hereafter the “double mammogram problem”) participants make conditional judgments after observing one positive mammogram and/or after two positive mammograms, assuming independence between the two test results.

Crucially we manipulated causal beliefs about the false positive rate (cf. Krynski & Tenenbaum, 2007). Our non-causal condition was similar to most previous studies using the mammogram problem, in that no cause of false positives was offered. In the causal condition, an alternative probabilistic cause of false positives (a benign cyst) was suggested; a positive mammogram could thus be seen as a “common effect” of two different causes (cancer or cyst). Figure 1 shows the detailed scenarios for each condition.

Doctors often encourage women at age 50 to participate in a routine mammography screening for breast cancer. To ensure a thorough assessment the clinic has adopted a policy whereby all women are required to complete two mammogram screenings in a single appointment. These screenings are carried out one straight after the other, but on two separate scanning machines, ‘Machine A’ and ‘Machine B’. The results of the test on each machine are independent of one another.

From past statistics, the following is known about a single mammogram screening:

2% of the women had breast cancer at the time of the screening.

Of those with breast cancer, 80% received a positive result on a mammogram.

[Non-causal version] Of those without breast cancer, 15% received a positive result on the mammogram.

[Causal version] 30% of the women had a benign cyst at the time of screening. Of those with a benign cyst, 50% received a positive test on the mammogram.

All others received a negative result.

These statistics are true for each machine.

Test questions

Single positive question

Suppose a woman gets a positive result in the mammogram screening on machine A. Without knowing any other symptoms, what are the chances she has breast cancer?

Double positive question

Suppose a woman gets a positive result in the mammogram screening on machine A, and a positive result in the screening on machine B. Without knowing any other symptoms, what are the chances she has breast cancer?

Figure 1. Single positive and double mammogram problems: non-causal and causal versions.

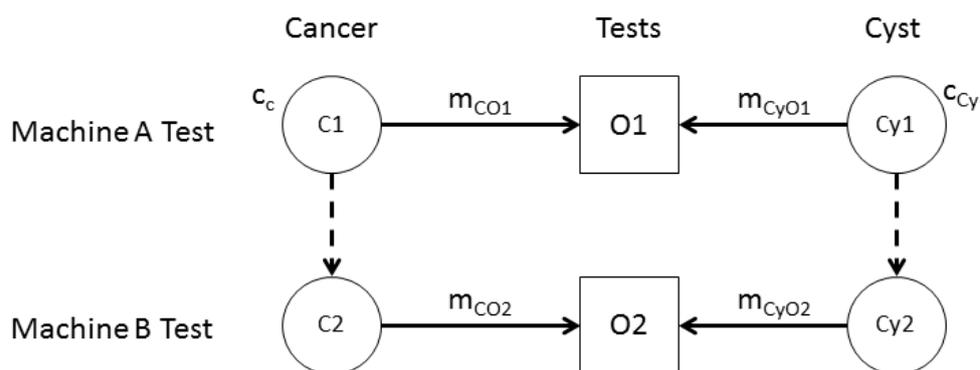
The different attributions for false positives in the respective conditions should lead to the construction of different intuitive causal models of the problem as shown in Figure 2. The Figure illustrates the relations between the causal variables of cancer (C) and cyst (Cy) or an unknown alternative cause of false positive tests (U), the respective base rates of these variables (c_C and c_{Cy}), and the probability that cancer (m_{C,O_j}), a cyst (m_{Cy,O_j}) or an unknown cause (b_j) will generate a positive test (O1 or O2). The dashed lines in the Figure represent instantiations of cancer and cyst on different mammogram tests.

The crucial question that we examined was how causal framing of false positives affects judgments about cancer probability in the light of multiple observations of positive mammogram tests. We hypothesize that people will see biological causes for a positive test (cancer or cyst) as relatively stable over multiple observations. Hence in the causal case in Figure 2, identifying that cancer (or cyst) is responsible for O1 *suggests that it is also the cause of O2*. In the non-causal condition however, two false positives are not assumed to have a common cause. As in the internet connection example, false positives at each observation might be interpreted as independent stochastic events.

Under these assumptions we expect different conditional probability estimates in the causal and non-causal conditions. At a qualitative level, the prediction for the causal condition is that the judged probability of a positive test being due to cancer should remain relatively stable across the successive observations O1 and O2. After observing O2 there may be some increase in the perceived probability of cancer but this increase should be

relatively small. This is because cyst remains a viable alternative explanation of O2 as well as O1 (see below for a quantitative demonstration). In contrast, in the non-causal condition, the probability of two statistically independent false positives is substantially lower than the probability of a single false positive. Hence, there should be a larger increase in estimates of the conditional probability of cancer from O1 to O2.¹

(a) Causal Condition



(b) Non-causal Condition

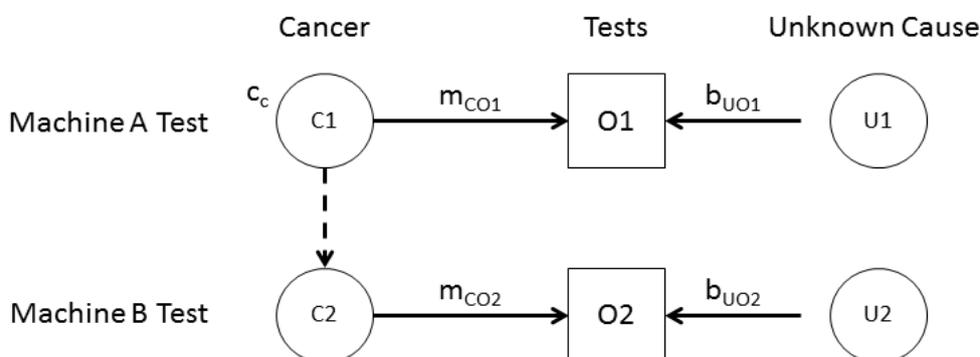


Figure 2. Graphical models of relations between causal nodes (circles) and observed effects (squares) in the causal (a) and non-causal (b) versions of the mammogram problem over two independent mammogram tests.

These same qualitative predictions emerged when we applied Bayes net formalisms to the graphical models in Figure 2. Our approach was similar to that used in a number of previous Bayes net applications (e.g., Fernbach et al., 2011; Griffiths & Tenenbaum, 2005;

Rehder, 2010) with the novel feature of allowing the relevant variables in the graphical model to be represented twice, once for each of the two mammogram tests administered (on machines A and B). This aspect of the model is similar to the assumptions made in “dynamic” Bayes nets (cf. Neapolitan, 2004; Rottman & Keil, 2012). The details of this model are given in the Appendix. The key prediction that emerges is that there should only be a modest increase in the judged conditional probability of cancer from O1 to O2 in the causal condition ($\approx 5\%$), as compared with a much larger increase in the non-causal case ($\approx 27\%$).² Note that we do not assume that participants’ probability estimates for each positive test will closely match those predicted by the model. Previous work with similar judgment problems shows that people typically over-estimate the conditional probability of cancer relative to normative values (e.g., Barbey & Sloman, 2007; Gigerenzer & Hoffrage, 1995) even when relevant causal information is present (Krynski & Tenenbaum, 2007). Nevertheless, even if probability estimates are generally higher than those predicted by the Bayes net model, we expected that the non-causal group would show a more substantial *change* in estimates across observations than the causal group.

2. Experiment 1

Experiment 1 tested these predictions of the causal Bayes net model. Different groups were presented with the non-causal or causal version of the mammogram problem and asked to estimate the probability of cancer following observation of a single positive mammogram, or two positive mammogram tests.

2.1 Method

2.1.1 Participants.

One hundred undergraduate Psychology students ($M_{AGE} = 19.6$ years) participated for course credit. Equal numbers were randomly allocated to the four experimental conditions.

2.1.2 Design and Procedure.

The study had a 2 (causal framing) x 2 (number of positive mammograms observed) design with both factors manipulated between participants. The causal framing instructions were adapted from those of Krynski and Tenenbaum (2007, Experiment 2).

The problem description (non-italicized text in Figure 1) was presented on a computer screen. After 15 s participants were presented with an open-ended question requesting an estimate of the chance of cancer in a woman with either one positive mammogram (on “Machine A”) or two positive mammograms (on “Machines A and B”, see Figure 1 for details). Note that in the single-positive mammogram condition participants were told that two mammogram tests could be carried out but were only given the results of the first test. Note also that our predictions are predicated on the assumption that having a second mammogram was not perceived as conditional on the outcome of the first mammogram. Hence our cover story emphasized that two tests were mandated for all women, regardless of test outcomes.³

The format of the requested probability estimate was a percentage chance of cancer (0-100%). After an estimate was entered, the cancer estimation question was repeated together with five alternative “answers that people commonly give to this question”; 2%, 10% (the correct option), 42%, 66%, 80%, listed in random order. Participants used a mouse to click on the option they thought was “closest to the correct answer”.

An on-screen calculator was available throughout the test procedure. After the mammogram task participants completed a computer-based version of the 4-item Berlin Numeracy Test (Cokely, Galesic, Schulz, Ghazal & Garcia-Retamero, 2012). There was no time limit on any part of the procedure.

2.2 Results and Discussion

The four experimental conditions did not differ in the likelihood of accessing a

calculator ($p > .30$). A preliminary analysis found a marginally significant interaction between causal framing and number of positive mammograms on numeracy scores, $F(1,96) = 3.88, p = .05$. The non-causal, one mammogram condition showed higher numeracy ($M = 2.80$ out of a maximum possible 4, $SE = 0.22$) than the corresponding causal condition ($M = 2.44, SE = 0.16$), but this trend reversed for participants who observed two mammograms (non-causal: $M = 2.12, SE = 0.21$; causal: $M = 2.56, SE = 0.22$).

Hence, the probability estimates shown in Figure 3 were entered into a 2 (causal framing) x 2 (number of observed positive mammograms) ANCOVA with numeracy as a covariate. Mean open-ended estimates differed as a function of task framing, with lower estimates in the causal ($M = 23.75, SE = 4.04$) than the non-causal condition ($M = 40.31, SE = 5.09$), $F(1,95) = 7.04$, partial $\eta^2 = .07, p = .009$. Those who observed two positive mammograms generally gave higher probability estimates ($M = 40.28, SE = 5.02$) than those who observed a single positive mammogram ($M = 23.79, SE = 4.20$), $F(1,95) = 6.65$, partial $\eta^2 = .07, p = .01$. Crucially, these effects were qualified by a significant interaction, $F(1,95) = 4.44$, partial $\eta^2 = .05, p = .04$.⁴ Figure 3 shows that as predicted by the causal Bayes net account, there was a larger increase in probability estimates from single to double mammogram cases in the non-causal than in the causal condition. Follow-up tests confirmed that there was a reliable increase in estimates from one to two positive mammograms in the non-causal (Tukey HSD, $q = 4.73, p < .01$) but not in the causal condition ($q = 0.28$).

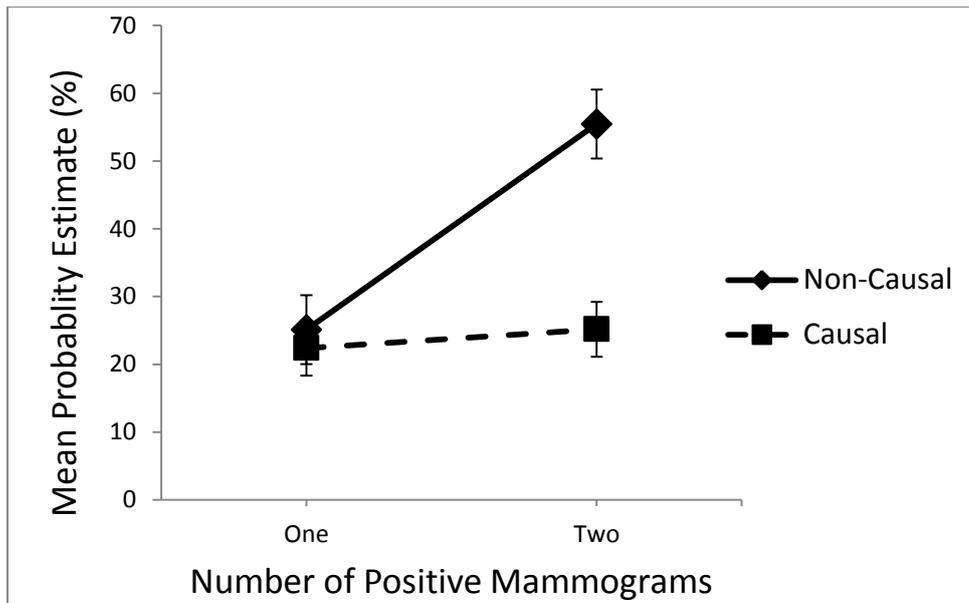


Figure 3. Experiment 1. Mean probability estimates (with standard error bars).

The probability estimate data was also analyzed by computing Bayes factors (BF) for independent groups t-tests comparing probability estimates at O1 and O2, separately for the causal and non-causal groups. An important advantage of this type of analysis is that it quantifies the relative odds that the data were generated from the null hypothesis (i.e., no difference in estimates at O1 and O2) compared to the alternative hypothesis of a reliable difference between means. Bayes factors were computed using the methods of Rouder, Speckman, Sun, Morey and Iverson (2009). Note that these BFs are scaled such that a BF with a value above 1.0 indicates evidence in favor of the null, and a BF less than 1.0 indicates evidence in favor of the alternative. For the non-causal group $BF = 0.08$, indicating that the alternative hypothesis of a difference between O1 and O2 estimates was approximately 12 times more likely than the null hypothesis of no difference. In contrast, for the causal group $BF = 8.56$, indicating that the null was approximately eight times more likely than the alternative.

Table 1 gives the frequency with which participants in the non-causal and causal conditions selected the various response alternatives on the forced choice question. The crucial issue here is whether the distribution of choices across alternatives differed for the

first and second positive mammograms. A significant change was found in the non-causal condition, with more choice of the highest probability estimates (66% or 80%) after the second positive test, $\chi^2(3) = 13.94, p = .003$. No such change was found in the causal condition, $\chi^2(3) = 0.80, p = .85$.

Table 1

Experiment 1. Frequency of Forced Choice Probability Choices Following One or Two Positive Mammograms in the Non-causal and Causal conditions

	Response Options			
	“2%”	“10%” (Correct)	“42%”	“66%/80%”
Non-causal				
One positive	10	7	2	6
Two positive	2	2	5	16
Causal				
One positive	11	6	4	4
Two positive	11	4	4	6

Both the open-ended probability and multiple choice results support the causal Bayes net prediction that supplying an alternative stable cause (cyst) for an event (positive mammogram) changes the way that people interpret the evidence from multiple observations.

In the causal condition positive mammogram tests could be seen as the product of one of two stable causal mechanisms (cancer or cyst). These same mechanisms could be invoked to explain both the first and second mammogram results (i.e. the relevant causal variables at O2 were correlated with those at O1). Hence, we found “causal stability” in estimates of the conditional probability of the target cause (cancer) across observations. In the non-causal condition the second positive result strengthens the belief in a known familiar cause (cancer) and weakens belief that both results were false positives generated by a stochastic process. The result was an increase in the estimated probability of cancer after observing two positive test results.

Although these qualitative data patterns are consistent with the Bayesian predictions, the estimates generated by participants were generally well above those predicted by the Bayes net model. For example, the average solution generated by non-causal participants at O2 (55.5%) was well above the predicted value ($\approx 37\%$). These data are not too surprising given the long history of work showing over-estimation of conditional probabilities in the mammogram problem (e.g., Barbey & Sloman, 2007; Eddy, 1982; Gigerenzer & Hoffrage, 1995). More crucially though the predicted level of *change* in estimates from O1 to O2 ($\approx 5\%$ increase in the causal condition, $\approx 27\%$ increase in the non-causal condition) was close to the observed pattern ($\approx 3\%$ in the causal, $\approx 30\%$ in the non-causal condition).

3. Experiment 2

This experiment aimed to further examine the robustness of the causal stability effect. In this case we again presented a scenario where two independent mammogram tests were performed. However, each participant now made two estimates; an initial estimate after observing one positive mammogram, and again after observing two positive mammograms. Notably this design allowed us to examine stability and change in *individual* probability estimates after observing each test result. Our causal Bayesian model predicts that

individuals in the non-causal group should increase their estimates across successive observations while individuals in the causal group should give similar estimates on each occasion.

3.1 Method

3.1.1 Participants

Fifty Psychology undergraduates ($M_{AGE} = 21.9$ years) participated for course credit. Equal numbers were randomly allocated to the non-causal and causal conditions.

3.1.2 Design and Procedure

The procedure was similar to Experiment 1 except that all participants made two estimates of the conditional probability of cancer; i) following a positive mammogram on Machine A, and ii) following a second positive mammogram, on Machine B. The two estimates were made sequentially, 10 s apart. As in the previous experiment, the instructions made it clear that each mammogram test was conducted using different but equally reliable machines and that the results of each test should be treated as independent (see Figure 1). The forced choice question was not presented because of concerns that responding to such a question after the first positive mammogram could bias the second probability estimate.

3.2 Results and Discussion

Preliminary analyses confirmed that the non-causal and causal groups did not differ in numeracy (non-causal: $M = 2.48$, $SE = 0.25$; causal: $M = 2.36$, $SE = 0.24$) or in their likelihood of accessing the on-screen calculator ($p > .5$).

Figure 4 shows the mean probability estimates generated in each condition. These were entered into a 2 (causal framing) x 2 (number of positive mammograms observed) ANOVA with repeated measures on the second factor. Lower estimates were given in the causal ($M = 26.6$, $SE = 5.83$) as compared to the non-causal condition ($M = 56.53$, $SE = 5.83$), $F(1,48) = 13.18$, partial $\eta^2 = .22$, $p = .001$.

A marginal main effect of the number of positive mammograms, $F(1,48) = 3.59$, partial $\eta^2 = .07$, $p = .07$, and a marginal interaction between mammogram number and causal framing were found, $F(1,48) = 3.05$, partial $\eta^2 = .06$, $p = .09$. Figure 4 shows the predicted trend towards a larger increase in probability estimates from one to two positive mammograms in the non-causal than in the causal condition.

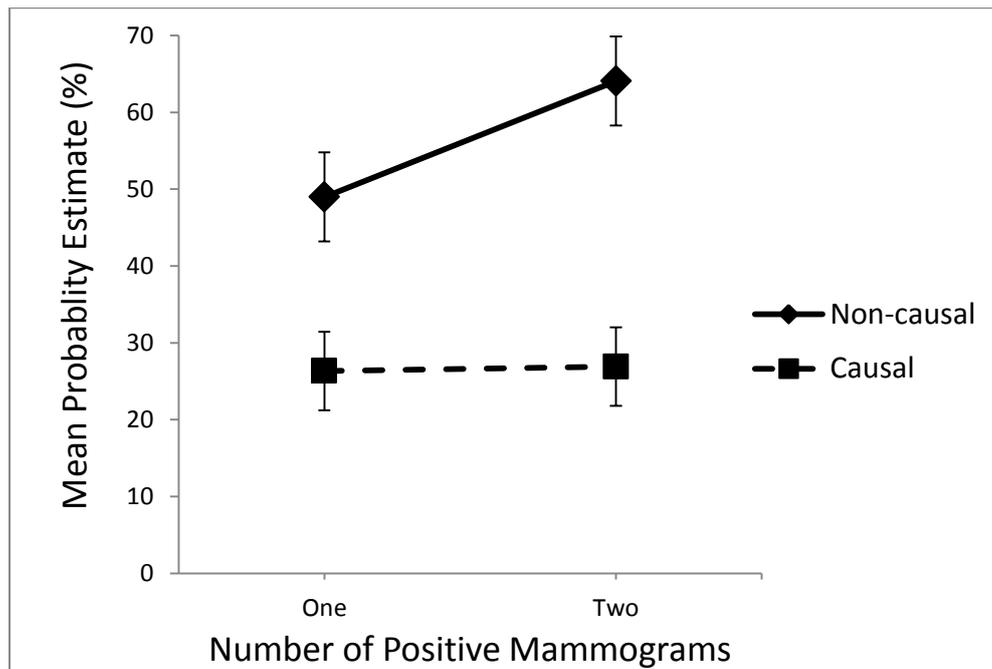


Figure 4. Experiment 2. Mean probability estimates (with standard error bars).

We again estimated Bayes Factors (Rouder, et al., 2009) to examine the extent to which the probability estimate data favored the null hypothesis of no change in estimates between O1 and O2, or the alternative. For the causal group $BF = 6.47$, representing strong evidence in favor of the null of no change in estimates. For the non-causal group $BF = 0.58$, indicating that the alternative hypothesis of a change in estimates was nearly twice as likely as the null.

Notably, in this experiment we were also able to examine patterns of change in individual probability estimates following each observation of a positive mammogram. Individual estimates for single and double mammograms were classified into one of three

groups; estimates that increased across multiple scans, estimates that decreased, or estimates that showed no change. “No change” was operationally defined as an estimate on the second scan that was within +/-2% of the first estimate.⁵ Table 2 shows support for our predictions. In the non-causal condition the majority of participants (60%) increased their probability estimates from O1 to O2, as compared with only 24% of participants in the causal condition, $\chi^2(2) = 6.68, p = .04$. In the causal condition the most common response was to give the same probability estimate after each positive test.

Table 2

Experiment 2. Frequency of Individual Patterns of Probability Estimates across First (O1) and Second Mammograms (O2) (Note: “same estimate = +/- 2% change from O1 to O2)

	Estimate decreased from O1 to O2	Same estimate on O1 and O2	Estimate increased from O1 to O2
Non-causal	4	6	15
Causal	7	12	6

Overall these results were consistent with the causal Bayes net predictions. Those in the causal condition showed greater stability in their conditional probability estimates from one to two positive mammograms than those in the non-causal condition. In terms of mean group estimates, the causal stability effect was somewhat smaller than that in Experiment 1, where participants made estimates after observing either one or two positive mammograms. This experiment however, also allowed us to examine patterns of change and stability in individual probability estimates. The analysis of this individual data showed unequivocal support for the experimental predictions with more than twice as many participants in the

non-causal condition increasing their estimates over successive observations as compared with the causal group.

In contrast to Experiment 1, we found mean probability estimates at O1 differed in the causal and non-causal groups, with the non-causal group in Experiment 2 further from the normative value. This is surprising given that at O1 participants in each experiment were making judgments based on identical information. A possible explanation was the relatively high numeracy scores found for the non-causal group in Experiment 1, which may have led to more accurate (i.e. lower) probability estimates in that group, relative to the non-causal group in this experiment (also see Footnote 3).

The finding that probability estimates in the causal condition were generally lower than those in the non-causal condition in Experiment 2 is similar in some respects to the causal facilitation effect reported by Krynski and Tenenbaum (2007). However, unlike Krynski and Tenenbaum, the percentage of participants generating the normative answer at O1 (within +/- 2%) did not differ significantly across conditions (causal = 12%; non-causal = 4%, $\chi^2(1) = 2.38, p = .12$). These data are consistent with other recent attempts to replicate the causal facilitation effects reported by Krynski and Tenenbaum (2007). McNair and Feeney (2013) and Hayes, Newell and Hawkins (2013) found that provision of relevant causal information can reduce the size of over-estimation errors in the mammogram problem, but does not necessarily increase the rate of normative responding. For our current purposes however, the most notable result is that the analysis of individual data patterns shows that the causal stability effect is maintained even when there are differences between initial probability estimates in the non-causal and causal conditions.

4. General Discussion

These experiments aimed to test predictions of a causal Bayes net model of judgments under uncertainty with multiple observations. According to this approach people formulate

causal models of judgment problems and attempt to incorporate given statistical information into these models (Krynski & Tenenbaum, 2007). We tested a novel prediction of this approach concerning the impact of beliefs about alternate sources of evidence across multiple observations. If an alternative explanation for a target event (e.g., a positive mammogram) appeals to a causal mechanism that is believed to persist over time (e.g., a benign cyst), it can serve as an explanation for multiple observations of the target. Alternately, if the mechanism that gives rise to the event is unknown, then such events are likely to be treated as stochastic, with an independent probability of occurrence at each observation. For the mammogram problem used in the current experiments this led to the prediction that estimates of the conditional probability of cancer would remain relatively stable across multiple observations when false positives were attributed to a stable alternative cause but would increase more sharply when no alternative causal explanation was available.

This prediction of causal stability was supported in both experiments where participants made probability estimates after observing one and/or two positive mammogram results. Probability estimates in the causal condition remained relatively stable over successive observations whereas those in the non-causal condition increased markedly. In Experiment 2, more than twice as many participants in the non-causal condition increased their probability judgments after observing a second positive test result than in the causal condition. The majority of participants in the causal condition gave similar estimates across successive observations.

These results show that the impact of causal beliefs on judgments under uncertainty goes well beyond increasing the salience of statistical information. Early work in this area found that when a causal explanation was supplied for a statistic, people were more likely to incorporate that statistic into their intuitive probability judgments (e.g., Ajzen, 1977; Bar-Hillel, 1980; Tversky & Kahneman, 1980). Such effects were explained by suggesting that

the causal information led people to pay more attention to statistics when encoding a problem and deriving a solution. This attentional salience account seems insufficient to explain our results. It is not clear how paying more attention to the false positive rate in the causal condition would produce more stable probability estimates over repeated observations.

It is worth considering whether the current results can be explained by other approaches that make no appeal to causal Bayes net principles. One popular explanation of errors on problems like the mammogram task is a failure to appreciate implied nested set relations (Barbey & Sloman, 2007; Evans et al., 2000; Reyna & Brainerd, 2008; Sloman, Over, Slovak & Stibel, 2003). In the problem shown in Figure 1, for example, women with a positive mammogram represent a subset of the population, and only a further subset of these women has cancer. Instructing people about these set relations leads to a reduction in base rate neglect (e.g., Sloman et al., 2003). Our causal instructions could be interpreted as promoting certain beliefs about set structure (e.g., that there are two non-overlapping sets of women with positive mammograms, those with cancer and those with a cyst). However it is not clear why invoking this set structure would lead to the greater stability in estimates across observations found in our causal conditions.

One question that arises from these results is whether the causal stability effect requires a specific alternate cause for observed data (e.g., a cyst) or whether the same effect could be found by just pointing out that alternative causes exist. We suspect that simply reminding a participant of the existence of alternative causes will not have the same impact on their causal interpretation of judgment problems as offering them a specific alternate cause. This conjecture is supported by the findings of Beyth-Marom and Fischhoff (1983) who asked participants to judge the likelihood that an individual belonged to a given social category and examined the extent to which participants considered information about plausible alternative categories when making this judgment. Consideration of alternatives

increased when a specific alternative category was mentioned but not when it was simply noted that the target individual could have belonged to other groups (also see McKenzie, 1998 for related findings).

4.1 Implications for causal Bayes net models of cognition

The current work represents a novel extension of causal Bayes net models. Most previous causal Bayes net models have been developed to represent the causal relationships within a set of independent observations (e.g., inferring the causal relations between depression and insomnia by observing multiple individuals with different combinations of relevant symptoms). Rottman and Keil (2012) refer to such cases as “static” causal systems. In contrast, the causal condition in the current experiments can be seen as an example of a “dynamic” Bayesian system where there are dependencies between the causes of events across successive observations.⁷ A key implication of our findings is that events that lack a common causal explanation are likely to be viewed as statistically independent across multiple observations. An understanding of the relevant causal mechanisms however can lead the same events to be viewed as the common effects of competing causes.

We have argued that when a probabilistic causal relation is established between two variables (e.g., a cyst can cause a positive mammogram with a particular probability) this relationship is often expected to remain stable across observations. Of course this is only true as long as there is no direct change to or intervention on the causal system (e.g., the cyst shrinks over time). We also acknowledge that the observable effects of some causal mechanisms exhibit marked change over time (e.g., seasonal variations in climate) and may be less likely to lead to beliefs in stable outcomes across observations (cf. Rehder & Martin, 2011). However we suspect that the default assumption for many everyday causal mechanisms is stability across observations, at least over the medium-term. Returning to our initial example, if you believe that your internet account is paid up and your browser launches

without difficulty today then there is good reason for you to expect the same cause-effect relation to hold tomorrow.

4.2 Conclusions

We see the current results as having two important implications. First, they represent further evidence that a causal Bayes net approach can be useful in understanding how people interpret and solve problems involving judgment under uncertainty (cf. Krynski & Tenenbaum, 2007). Second, they show how provision of causal information can radically change the way people interpret statistical information across multiple observations. Overall these studies suggest that a better understanding of the judgment process will arise from a careful consideration of how people construct causal models of specific problems and how they interpret key statistics in the light of these models (cf. Newell, 2013).

Acknowledgements

This research was supported by Australian Research Council Discovery Grant DP120100266 to the first and third authors. The authors would like to thank Ann Martin and Kelly Jones for their assistance in data collection.

References

- Ahn, W., & Kim, N. S. (2001). The causal status effect in categorization: An overview. In D. L. Medin (Ed.) *The psychology of learning and motivation*. (Vol. 40, pp. 23–65). San Diego, CA: Academic Press.
- Ajzen, I. (1977). Intuitive theories of events and the effects of base-rate information on prediction. *Journal of Personality and Social Psychology*, *35*, 303–314.
- Barbey, A. K., & Sloman, S. A. (2007). Base rate respect: From ecological rationality to dual processes. *Behavioral and Brain Sciences*, *30*, 241–297.
- Bar-Hillel, M. (1980). The base-rate fallacy in probability judgments. *Acta Psychologica*, *44*, 211–233.
- Beyth-Marom, R., & Fischhoff, B. (1983). Diagnosticity and pseudodiagnosticity. *Journal of Personality and Social Psychology*, *45*, 1185-1195.
- Bes, B., Sloman, S., A., Lucas, C. G., & Raufaste, E. (2012). Non-Bayesian inference: Causal structure trumps correlation. *Cognitive Science*, *36*, 1178–1203.
- Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, *104*(2), 367.
- Cokely, E. T., Galesic, M., Schulz, E., Ghazal, S., & Garcia-Retamero, R. (2012). Measuring risk literacy: The Berlin numeracy test. *Judgment and Decision Making*, *7*, 25-47.
- Eddy, D. M. (1982). Probabilistic reasoning in clinical medicine. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases*. (pp. 249–267). Cambridge, UK: CUP.
- Evans, J. St. B. T., Handley, S. J., Perham, N., Over, D. E., & Thompson, V. A. (2000). Frequency versus probability formats in statistical word problems. *Cognition*, *77*, 197–213.

- Fenton, N., Berger, D., Lagnado, D., Neil, M., & Hsu, A. (2013). When 'neutral' evidence still has probative value (with implications from the Barry George Case). *Science and Justice*.
- Fernbach, P. M., Darlow, A., & Sloman, S. A. (2011). Asymmetries in predictive and diagnostic reasoning. *Journal of Experimental Psychology: General*, *140*, 168-185.
- Gigerenzer, G., & Hoffrage, U. (1995). How to improve Bayesian reasoning without instruction: Frequency formats. *Psychological Review*, *102*, 684–704.
- Gopnik, A., Glymour, C., Sobel, D. M., Schulz, L. E., & Kushnir, T. (2004). A theory of causal learning in children: Causal maps and Bayes nets. *Psychological Review*, *111*, 3-23.
- Griffiths, T. L., & Tenenbaum, J. B. (2005). Structure and strength in causal induction. *Cognitive Psychology*, *51*, 334–384.
- Hagmayer, Y., & Sloman, S. A. (2009). Decision makers conceive of themselves as interveners. *Journal of Experimental Psychology: General*, *128*, 22-38.
- Hayes, B. K., Newell, B. R., & Hawkins, G. E. (2013). Causal model and sampling approaches to reducing base rate neglect. In M. Knauff, M. Pauen, N. Sebanz, & I. Wachsmuth (Eds.) *Proceedings of the 35th Annual Conference of the Cognitive Science Society*, (pp. 567-572). Austin, TX: Cognitive Science Society
- Hayes, B. K. & Rehder, B. (2012). The development of causal categorization. *Cognitive Science*, *36*, 1102–1128.
- Holyoak, K. J., Lee, J. S., & Lu, H. (2010). Analogical and category-based inferences: A theoretical integration with Bayesian causal models. *Journal of Experimental Psychology: General*, *139*, 702-727.
- Kemp, C., & Tenenbaum, J. B. (2009). Structured statistical models of inductive reasoning. *Psychological Review*, *116*, 20–58.

- Kemp, C., Goodman, N. D. & Tenenbaum, J. B. (2010). Learning to learn causal models. *Cognitive Science*, *34*, 1185–1243.
- Kemp, C., Shafto, P., & Tenenbaum, J. B. (2012). An integrated account of generalization across objects and features. *Cognitive Psychology*, *64*, 35-73.
- Koehler, J. J. (1996). The base rate fallacy reconsidered: Descriptive, normative, and methodological challenges. *Behavioral and Brain Sciences*, *19*, 1–53.
- Krynski, T. R., & Tenenbaum, J. B. (2007). The role of causality in judgment under uncertainty. *Journal of Experimental Psychology: General*, *136*, 430–450.
- McKenzie, C. R. M. (1998). Taking into account the strength of an alternative hypothesis. *Journal of Experimental Psychology: Learning, Memory and Cognition*, *24*, 771-792.
- McNair, S., & Feeney, A. (2013). When does information about causal structure improve statistical reasoning? *Quarterly Journal of Experimental Psychology*. Advance online publication. doi: 10.1080/17470218.2013.821709
- Neapolitan, R. E. (2004). *Learning Bayesian networks*. Upper Saddle River, NJ: Pearson Prentice Hall.
- Newell, B. R. (2013). Judgment under uncertainty. In D. Reisberg (Ed.) *Oxford Handbook of Cognitive Psychology*. (pp. 602-615). New York, NY: Oxford University Press.
- Rehder, B. (2006). When similarity and causality compete in category-based property generalization. *Memory & Cognition*, *34*, 3–16.
- Rehder, B. (2010). Causal-based classification: A review. In B. Ross, (Ed.), *The Psychology of Learning and Motivation* (52), pp. 39-116. San Diego, CA: Elsevier Academic Press.
- Rehder, B., & Kim, S. (2010). Causal status and coherence in causal-based categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *36*, 1171

- Rehder, B. & Martin, J. B. (2011). A generative model of causal cycles. In L. Carlson, C. Hölscher & T. Shipley (Eds.) *Proceedings of the 33rd Annual Conference of the Cognitive Science Society*, (pp. 2944-2949). Austin, TX: Cognitive Science Society
- Reyna, V. F., & Brainerd, C. J. (2008). Numeracy, ratio bias, and denominator neglect in judgments of risk and probability. *Learning and Individual Differences*, *18*, 89–107.
- Rottman, B. J., & Keil, F. C. (2012). Causal structure learning over time: Observations and interventions. *Cognitive Psychology*, *64*, 93-125.
- Rouder, J. N., Speckman, P. L., Sun, D., Morey, R., & Iverson, G. (2009). Bayesian *t* tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review*, *16*, 225-237.
- Shank, R. C., & Abelson, R. P. (1995). Knowledge and memory: The real story. In R. S. Wyer (Ed.), *Advances in Social Cognition*, (Vol. 8, pp. 1–85). Hillsdale, NJ: Erlbaum.
- Sloman, S. A. (2005). *Causal models: How people think about the world and its alternatives*. New York, NY: Oxford University Press.
- Sloman, S. A., & Lagnado, D. A. (2005). Do we “do”? *Cognitive Science*, *29*, 5–39.
- Sloman, S. A., Love, B. C., & Ahn, W. (1998). Feature centrality and conceptual coherence. *Cognitive Science*, *22*, 189-228.
- Sloman, S. A., Over, D., Slovak, L., & Stibel, J. M. (2003). Frequency illusions and other fallacies. *Organizational Behavior and Human Decision Processes*, *91*, 296–309.
- Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: Statistics, structure, and abstraction. *Science*, *331*, 1279–1285.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, *185*, 1124–1131.
- Tversky, A., & Kahneman, D. (1980). Causal schemas in judgments under uncertainty. In M. Fishbein (Ed.), *Progress in social psychology* (pp. 49–72). Hillsdale, NJ: Erlbaum.

Waldmann, M. R. (1996). Knowledge-based causal induction. In D. R. Shanks, K. J. Holyoak & D. L. Medin (Eds.), *The psychology of learning and motivation, Vol. 34: Causal learning* (Vol. 34, pp. 47-88). San Diego: Academic.

Waldmann, M. R., Hagmayer, Y., & Blaisdell, A. P. (2006). Beyond the information given: Causal models in learning and reasoning. *Current Directions in Psychological Science, 15*, 307–311.

Footnotes

1. Despite making different assumptions about the details of Bayesian updating, Krynski and Tenenbaum (2007) make qualitative predictions that are similar to ours, i.e. after observing two positive mammograms they predict a large increase in estimates of the probability of cancer in the non-causal condition but little change in estimates in the causal condition.
2. Note that these predictions hold for individuals making a single estimate based on two positive mammograms (as in Experiment 1) as well as those making make two sequential predictions (Experiment 2). The appropriate derivations are available from the authors.
3. We note that in an unreported study where instructions could be interpreted such that the second mammogram *was* conditional on the outcome of the first mammogram, probability estimates increased across tests in both causal and non-causal conditions.
4. This interaction remained robust when the covariate was removed, $F(1,96) = 4.95$, $p = .03$.
5. To check whether this result depended on the specific values that defined “no change”, we ran chi-square analyses with definitions that varied across the range of 0 to +/- 10% difference between the first and second probability estimates. The pattern of results remained largely unchanged and remained statistically reliable across this range (all p 's < .04).
6. This representation of the causal condition assumes that subjects interpreted “Of those with a benign cyst, 50% received a positive test on the mammogram” as referring to women known not to have cancer. This constraint can be relaxed by assuming that some of the positive tests amongst those with a cyst were due to cancer (i.e. that the two causes were not mutually exclusive). Given that the probability of a positive test *without* cyst, $p(O_j|\overline{C}y_j)$, is $c_c m_c = (.02)(.765) = .015$ (again, assuming no other causes of a positive test other than cancer and cyst), then $m_{c_y} = [p(O_j|Cy_j) - p(O_j|\overline{C}y_j)] / [1 - p(O_j|\overline{C}y_j)] = [.50 - .015] / (1 - .015) = .492$, and the results of the subsequent analyses are virtually identical to those in the Appendix.

7. It might be argued that static Bayes nets are sufficient to explain the results of the Experiment 1 since the outcomes of the two mammogram tests could be viewed simultaneously. This argument however, does not apply to Experiment 2 where the test outcomes were viewed sequentially. In this case a dynamic Bayes net is certainly a more appropriate causal model.

Appendix

The joint distribution of the variables $\{V_1, V_2, \dots, V_n\}$ of a causal graphical model can be factorized as

$$p(V_1, V_2, \dots, V_n) = \prod_{j=1}^n p(V_j | Pa(V_j))$$

where $Pa(V_j)$ denotes the parents of V_j . Throughout we assume that the variables are binary (i.e., $V_j = 0$ [absent] or 1 [present]), that the causal relations are generative, and that multiple causes of an effect combine according to a noisy-or function,

$$p(V_j | Pa(V_j)) = 1 - (1 - b_j) \prod_{i \in Pa(V_j)} (1 - m_{ij})^{ind(i)}$$

where m_{ij} is the strength or *causal power* of the generative link between V_i and V_j , b_j represents alternative causal influences unaccounted for by other nodes in the network, and $ind(i)$ is an indicator function that returns 1 when V_j 's i^{th} parent is present and 0 otherwise. When V_j is a root node (i.e., has no parents), its probability is defined as c_j .

These definitions are now applied to the causal networks of Figure 2. For the network of Figure 2a, the probability of positive test j (O_j), given the states of cancer and cyst for that test is (C_j and Cy_j),

$$p(O_j | C_j, Cy_j) = 1 - (1 - b_U)(1 - m_C)^{ind(C_j)}(1 - m_{Cy})^{ind(Cy_j)}$$

whereas for the non-causal network in Figure 2b the probability of O_j given C_j is,

$$p(O_j | C_j) = 1 - (1 - b_U)(1 - m_C)^{ind(C)}$$

These expressions were used to generate the joint probability distributions of each network in Figure 2. The causal models' parameters were set to match those in the cover stories. First, the base rates of root variables C_1 and Cy_1 were $c_C = .02$ and $c_{Cy} = .30$, respectively. Second, the power of the $C_j \rightarrow O_j$ relationships was derived from the stated

values of $p(O_j|C_j)$ and $p(O_j|\bar{C}_j)$ according to Cheng's (1997) power-PC formula: $m_C = [p(O_j|C_j) - p(O_j|\bar{C}_j)]/[1 - p(O_j|\bar{C}_j)] = [.80 - .15]/[1 - .15] = .765$. In the non-causal condition, subjects were told "Of those [women] without breast cancer, 15% received a positive result on the mammogram." and so the strength of alternative causes of a positive test (b_U) was set to .15. In the causal condition, we assume that the propensity of cyst to cause a positive test, m_{Cy} , was .50 (which, when multiplied by cyst's base rate of .30, yields an alternative causal influence equal to that in the non-causal condition, .15) and that there were no other causes of a positive test, $b_U = 0$.⁶

The joint distributions of the causal and non-causal networks for these parameterizations are shown in Table A1. Note that because the two instantiations of cancer C_1 and C_2 are defined to be perfectly correlated, in the joint those variables have been collapsed them into a single variable C ; Cy_1 and Cy_2 are similarly collapsed into Cy . We derived the key conditional probabilities from these distributions, namely, the probability of cancer given one, $p(C = 1|O_1 = 1)$, or two, $p(C = 1|O_1 = 1, O_2 = 1)$, positive mammograms, also shown in Table A1. In the causal condition, the probability of cancer given a positive mammogram increases moderately from one positive test (.098) to two (.148). In the non-causal condition, in contrast, the increase in the probability of cancer is much larger, to .367. This larger increase in the non-causal condition obtains because of the improbability of the conjunction of two relatively rare events, that is, the improbability that two positive tests are *both* not due to cancer (two false positives). Indeed, in the non-causal condition the probability of two successive positive mammograms is low (.035). In contrast, in the causal condition the probability of them being due to a common alternative cause (cyst) is higher (.086).

Table A1

Joint probability distribution for alternative causes of a positive mammogram at O1 and O2.

C	C_y	O_1	O_2	Joint Probability	
				Causal	Non-causal
0	0	0	0	0.6860	0.4956
0	0	0	1	0.0000	0.0875
0	0	1	0	0.0000	0.0875
0	0	1	1	0.0000	0.0154
0	1	0	0	0.0735	0.2124
0	1	0	1	0.0735	0.0375
0	1	1	0	0.0735	0.0375
0	1	1	1	0.0735	0.0066
1	0	0	0	0.0008	0.0006
1	0	0	1	0.0025	0.0022
1	0	1	0	0.0025	0.0022
1	0	1	1	0.0082	0.0090
1	1	0	0	0.0001	0.0002
1	1	0	1	0.0006	0.0010
1	1	1	0	0.0006	0.0010
1	1	1	1	0.0047	0.0038
$p(C = 1)$.020	.020
$p(C = 1 O_1 = 1)$.098	.098
$p(C = 1 O_1 = 1, O_2 = 1)$.149	.367