

Tell me something I don't know: How perceived knowledge influences the use of information during decision making

Samantha Kleinberg (samantha.kleinberg@stevens.edu)

Department of Computer Science, 1 Castle Point on Hudson
Hoboken, NJ 07030 USA

Jessecae K. Marsh (jessecae.marsh@lehigh.edu)

Department of Psychology, 17 Memorial Drive East
Bethlehem, PA 18015 USA

Abstract

We are often confronted with new causal information about the world, such as what causes a disease. What we think we know may influence if and how we choose to use this new information. Yet as prior work has shown, we are not always successful at evaluating our own knowledge. We explored how helping people better understand what they know about a domain can influence their ability to use new causal information in a decision-making context. Participants self-assessed their knowledge (Experiment 1) or completed an objective assessment of their knowledge (Experiment 2) of diabetes, before making diabetes-related decisions, either with or without new causal information. Without a knowledge assessment, participants were less accurate with new causal information compared to without such information, replicating previous work. However, reassessing their knowledge increased participants' decision-making accuracy with causal information. We discuss why helping people realize the limits of their causal understanding may make them better supplement it with new information.

Keywords: decision making; illusion of explanatory depth; causality; diabetes

Introduction

How does what we think we know influence how we use information? Imagine you are playing a game where you are a zombie and must decide which brains to eat to fuel your zombie lifestyle. As you, reader, are likely not a zombie, you probably do not have prior experience eating brains. As such, you will have to rely solely on information provided to you by the game of how different brains create different nutritional outcomes. If instead you are at a breakfast buffet deciding whether to have a pastry or a fruit salad, would information on how these foods affect your health longterm or fuel your workout later in the morning be as helpful? Any new information will exist alongside your prior experience with breakfast foods, beliefs about the relationship between nutrition and health, and maybe even experience in eating each of these foods before a workout. Importantly, the beliefs you bring to the breakfast buffet may be incorrect or conflict with the new information you see, which in turn may affect how much you trust this new information. Much prior work has shown people can successfully learn about and use causal information in novel scenarios like the zombie one. Yet little is known about how people's beliefs (e.g. carbs are unhealthy) interact with new causal information (e.g. carbs fuel our muscles and can improve performance) and influence their choices.

Prior work on decision-making has shown that people can successfully learn causal models (Lagnado, Waldmann, Hagmayer, & Sloman, 2007; Waldmann, Hagmayer, & Blaisdell, 2006) and use these models to make decisions (Hagmayer & Sloman, 2009; Sloman & Hagmayer, 2006). An essential feature of these studies is that the only information people can use to make decisions is the causal model provided by the experimenters. As in our zombie example, people must use the information provided in these studies as it is the only way to answer the study questions. However, it is not clear whether people will use information the same way in familiar scenarios. That is, when people have existing causal beliefs about a decision domain, how do they view new causal information meant to guide that decision?

Decision making in familiar domains brings added complications in that people may believe they are knowledgeable about the domain, even when they are not. People have been shown to overestimate their causal understanding of the world (Shulman, 2015) and more expertise does not lead to better estimates (Fisher & Keil, 2016). A classic demonstration of this is the illusion of explanatory depth (IOED). In the IOED task of Rozenblit and Keil (2002), people estimate how well they understand how an everyday item like a faucet works, generate a causal explanation of how that item works, and finally re-assess their understanding. The illusion is that people believe they understand how things work much better than they actually do – until they have to explain the process in gory detail. This overestimation of causal understanding has been shown in a multitude of domains including physical systems (Rozenblit & Keil, 2002), politics (Fernbach, Rogers, Fox, & Sloman, 2013; Vitriol & Marsh, 2018), and mental health (Zeveney & Marsh, 2016). Situations where people do not have causal beliefs about a domain (regardless of the correctness of those beliefs) may be rare. Thus, it is critical to understand how people's perception of their knowledge may influence how they use the new information that they constantly receive.

If our perceptions are wrong we may fail to make use of information that could assist us, like calorie counts on menus or the causal models being output by many machine learning methods. In fact, our recent work has shown that people may struggle to use new causal information to make decisions in familiar domains. Zheng, Marsh, Nickerson, and Kleinberg (2020) asked participants to make a real-world decision

(e.g., food intake in weight management), with and without the aid of causal information, whose content was based on established guidelines. Participants who were presented with information that should have been useful made worse choices compared to participants who were given no supplemental information. Yet when the same participants made decisions in novel domains, they successfully used the very same type of causal information to make decisions.

Causal information hindering decision making is a troubling finding for the many fields trying to uncover causal relationships from data to help people make everyday decisions. Health officials provide information on how diet influences disease risk; economists provide indicators for recessions; and machine learning has made it possible to uncover this information in many more domains. But will anyone listen? People's inability to understand the limits of their knowledge without being confronted by its inadequacy (as in the IOED task or when taking an exam) may influence whether and how they use this information. Yet before choosing to eat the brains of a psychologist instead of a computer scientist, or whether to eat a croissant or an apple, we do not regularly sit down and interrogate our beliefs on nutrition. We propose that helping people recognize the gaps in their causal knowledge may allow them to better use new information when making decisions. People may be both more receptive to new causal information that fills those gaps and may be more able to integrate it since they have already identified holes in their knowledge.

We conducted two experiments to test how exposing gaps in individuals' knowledge influences their use of causal information during decision making. Since diagrams are a common output of methods to learn causal models from data, and commonly used for studies on cognition, we specifically use causal diagrams for conveying new information. We investigate both subjective and objective methods of manipulating perceived knowledge. In Experiment 1 we test whether a subjective reassessment of knowledge, using an IOED task, leads to better use of causal information at decision time. While the IOED has been demonstrated in many domains, it is not yet known whether exposing gaps in one's knowledge can improve information use during decision making.¹ In Experiment 2 we test whether objective assessments of knowledge have the same effect as subjective self-reflection.

Experiment 1

What we think we know may affect our willingness or ability to use new information. In particular, people may not consider new information if they think they already have relevant knowledge. If people do not use new causal information or do not use it effectively due to illusions of understanding, then exposing the lack of depth people have in their causal under-

¹While Fernbach et al. (2013) found that completing an IOED task on political issues could attenuate extreme political attitudes and in turn make people less likely to donate to relevant advocacy groups, prior work has not addressed whether the IOED could be used to influence decision making accuracy.

standing prior to decision time may make people more willing to use new causal information. We test this using decisions surrounding prevention and management of Type 2 diabetes (T2D), with self-perceived understanding manipulated using an IOED task prior to the decision-making task.

Method

Participants We recruited participants from Amazon Mechanical Turk (MTurk) using Turkprime panels. Because over 9% of people in the US have T2D (Centers for Disease Control and Prevention, 2017), focusing on decision making in this domain allowed us to recruit a sizable population of people both with and without personal experience managing T2D. Thus we can test whether perceived knowledge due to personal experience with a domain is more or less malleable than perceived knowledge gained outside of personal experience. We aimed to recruit 200 people with T2D and 200 people without diabetes (PWOD). To prevent temporal differences in data collection between groups, we used a stopping rule of running on Mturk for 5 consecutive business days, or when a group hit 200, whichever came first. Of the target 200 PWOD, 171 completed the study and reported they did not have T2D and were not a caregiver for someone with diabetes. We recruited a total of 103 participants with T2D in our timeframe. We screened the IOED responses to ensure participants were engaging in the explanation task as a manipulation check. Eleven participants (6 PWOD and 5 T2D) were dropped for copying from other sources (i.e., websites) or for not providing an explanation. All participants were U.S. residents aged 18-64, and were compensated \$2.75 for participation.

Materials We developed five questions about T2D that spanned a range of types of causal thinking. Two questions related to managing diabetes (management questions), with one requiring participants to choose which single causal factor would be most effective for managing T2D, while the other required choosing between single and interactive causes for the best way to manage T2D. There were three questions targeting risk reduction for T2D (prevention questions). One involved choosing which of a series of single preventative factors would be most effective for reducing risk of T2D. In another, participants chose between single and interactive causes for the best way to reduce risk for T2D. In the last, participants chose between causes that had direct, indirect, or both direct and indirect effects on T2D risk. Each question has a single target answer, which is either the only answer where a cause is present, or the answer that dominates others by making the effect most likely. One question was:

Tyler is a graduate student who is nervous about his upcoming exams. He usually skips breakfast and goes straight to the library in the morning to study. He spends most of his time on schoolwork, and eats most meals in the cafeteria or fast casual restaurants. He read a pamphlet on diabetes at the health center and realizes he's at

risk, but is too focused on his exams to figure out what to do next.

What change will MOST reduce Tyler’s risk of diabetes?

- A. Have oatmeal with a banana in the morning and walk to school
- B. Play basketball on the weekends
- C. Cook at home more often
- D. Manage his stress by listening to calming music

To complement the questions, we created two causal diagrams that represent scientific understanding of important factors for T2D management and prevention, as shown in Figure 1 (Noble, Mathur, Dent, Meads, & Greenhalgh, 2011). As in our prior work (Zheng et al., 2020), these graphs contain information that can be specifically used to help answer the decision-making questions. To keep the figures from being overwhelming, we limited the total number of factors in each diagram to 6. Each graph depicted generative (represented with a +) and preventative (represented with a -) causal relationships.

Procedure Participants began the experiment by completing an IOED task, based on the protocol of Rozenblit and Keil

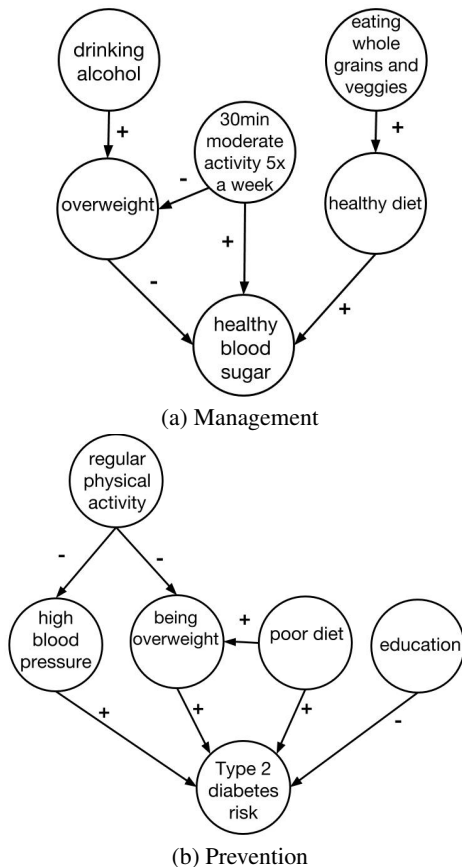


Figure 1: Causal diagrams provided with decision-making questions in diagram condition.

	PWOD		T2D	
	IOED	control	IOED	control
No info	43	40	17	27
Diagram	44	44	32	27

Table 1: Number of subjects per condition in Experiment 1.

(2002). Participants rated their understanding of a stimulus, then explained the causal relationships between features of that stimulus, and finally re-rated their understanding of the stimulus. Participants were randomly assigned to either the explanation condition (IOED; PWOD: $n = 87$; T2D: $n = 49$), or to a control description condition (control; PWOD: $n = 84$; T2D: $n = 54$). Participants first read the introductory materials of Rozenblit and Keil (2002) that introduced the 1 (low understanding) to 7 (high understanding) rating scale. Participants in the IOED condition then rated their understanding of 1) how people develop diabetes later in life, and 2) how lifestyle choices and treatments manage diabetes symptoms (Time 1 [T1] rating). The control condition prompts were rephrased to avoid triggering causal thinking while still asking participants to reflect on their knowledge (Zeveney & Marsh, 2016). The prompts were: 1) diabetes later in life, and 2) lifestyle choices and treatments used for diabetes. Both sets of prompts map to the decision-making questions, which focused on either prevention or management of T2D. After T1 ratings, IOED participants created a causal explanation that included all the steps in each process in detail. In the control condition the causal explanation task was replaced with a description task where participants were asked to list all characteristics of each prompt. Since the critical feature for exposing illusions in the IOED is the act of generating a causal explanation, these descriptive prompts should not engage the same causal understanding introspection (Zeveney & Marsh, 2016). Finally, participants re-rated their understanding for both prompts (Time 2 [T2] rating).

After the IOED task, all participants completed the 5 multiple choice decision-making questions about diabetes. Participants in both IOED and control conditions were randomly assigned to receive either a diagram with each decision-making question (diagram) or no information beyond the question text (no info). Participants received instructions on what the nodes, arrows, and plus and minus signs indicated in the causal diagrams. In prior work (Zheng et al., 2020) we found both that participants were able to use similar diagrams and that providing more detailed instructions than we did here did not change participant responding. The number of subjects per condition are shown in detail in Table 1. The order of the decision-making questions and of the answer choices were randomized for every subject.

Results

Demonstration of IOED We first tested whether our manipulation produced an IOED in diabetes knowledge, since we aim to examine effects of the IOED. Since there were no




Figure 2: Decision-making accuracy for each group in Experiment 1. Bars denote standard error.

significant interactions involving experience² we pooled all participants and then ran a 2 (Time: T1 vs. T2; within) x 2 (IOED manipulation: IOED vs. control; between) mixed ANOVA over judgments. We found a significant main effect of Time ($F(1,261) = 13.3, p < .001, \eta_p^2 = .049$) and a marginal main effect of IOED manipulation, $p = .080$. Importantly, we found the expected interaction between IOED manipulation and time, $F(1,261) = 4.08, p = 0.044, \eta_p^2 = .015$. We further analyzed the interaction through Sidak-corrected t -tests. In the IOED condition, T1 ratings ($M = 4.63$) were significantly higher than T2 ratings ($M = 4.24; p < .001$). In the control condition, there was no difference between ratings at T1 ($M = 4.84$) and T2 ($M = 4.73; p = .248$). These findings support that the IOED manipulation worked as intended and that the control condition did not also result in an IOED exposure.

Effect of IOED on decision accuracy We now explore how the IOED manipulation affected accuracy on the decision making questions. Our major question of interest is whether there is an interaction between IOED condition and information condition. That is, does self-assessing knowledge with the IOED task improve use of causal diagrams? To test this, we used a multilevel model (MLM) analysis to account for the repeated nature of our data and correlations between our measures (3 prevention and 2 management questions). We entered our factors of IOED manipulation (IOED vs. control; between), information condition (diagram vs. no info; between), and experience (T2D vs. PWOD; between) as fixed effects into the MLM model, which allows testing the significance of main effects and their interactions through F -tests. There is good reason to separate the T2D and PWOD groups in analysis given the differences in their personal experience in the domain of diabetes and our interest in seeing whether personal experience influences how likely peo-

ple are to update their perceptions. The separation is supported by the three-way interaction being marginally significant, $F(1,255) = 3.422, p = .066$. Thus we conduct further analyses separately for the PWOD and T2D groups as MLM models of IOED manipulation (IOED vs. control; between) by information condition (diagram vs. no info; between).

As seen in Figure 2a, self-assessment of knowledge did influence the use of causal diagrams for PWOD. We found a significant main effect of information condition ($F(1,161) = 6.03, p = .015$), no main effect of IOED manipulation ($p = .272$), and the predicted significant interaction ($F(1,161) = 5.09, p = .025$). Exploring the interaction, we replicated previous findings (Zheng et al., 2020) that presenting a causal diagram can interfere with decision making: control participants who did not complete the IOED task showed higher accuracy in the no info condition ($M = .748$) than control participants in the diagram condition ($M = .601; p = .001$). Our key comparison is whether completing the IOED procedure would improve decision making. We found that IOED participants did show better accuracy in the diagram condition ($M = .706$) than control participants in the diagram condition, $p = .017$. This suggests that an IOED procedure can help overcome some of the difficulties of trying to integrate a diagram with previous knowledge. However, for our main comparison of diagram to no info, there was not a significant difference in accuracy for IOED participants in the diagram and no info conditions ($M = .712; p = .887$). In other words, while the IOED procedure removed the previously observed negative effect of diagrams, it alone was not enough to make diagrams better than no info.

As shown in Figure 2b, the IOED manipulation did not have the same effect for T2D participants. There was not a significant main effect of IOED manipulation, $p = .280$, or information condition, $p = .266$. The interaction was also not significant, $p = .650$.

Discussion

We find that reflecting on causal knowledge can influence how people use causal information at decision time, but this may work differently depending on people's experience in the decision domain. For individuals without personal domain

²We conducted a 2 (Time: T1 vs. T2; within) x 2 (IOED manipulation: IOED vs. control; between) x 2 (experience: T2D vs. PWOD; between) mixed ANOVA on participants' understanding judgments averaged across the two diabetes prompts. There was a main effect of experience ($F(1,259) = 90.3, p < 0.001$), but no significant two-way interactions of experience with Time or IOED manipulation ($ps > .23$) and no three-way interaction ($p = .732$).

experience (PWOD in this experiment), re-evaluating knowledge removed the detrimental effect of causal information observed in the control group. For individuals with domain experience (T2D), we did not observe this improvement. Why is this? One possibility is that having domain experience may result in stronger causal beliefs. We have some evidence for this in that our main effect of experience for IOED ratings in Experiment 1 (see footnote 2) is because of higher self-assessment ratings overall for T2D participants than PWOD participants. Thus for individuals with personal domain experience, who as a result are highly confident in their knowledge, beliefs may be stickier and the IOED task alone may not be enough to make people truly reconsider what they know. Understanding how expertise in a domain and strong confidence in one's beliefs may influence belief change is an important avenue for future researchers to explore.³

Experiment 2

Our first experiment showed that subjective reflection on knowledge in a domain can positively influence the use of new information during decision making. In Experiment 2 we test whether an objective measure of knowledge has the same effect as self-reflection. If participants becoming aware of gaps in their knowledge is responsible for inoculation against potential detrimental effects of causal information, we expect that being exposed to objective assessments of their knowledge will have a similar effect as the IOED. Given the differences in results between PWOD and T2D, we now focus solely on individuals without personal domain experience.

Method

Participants As in Experiment 1, we recruited participants through Turkprime who were U.S. residents aged 18-64 and we compensated them \$2.75 for participation. We excluded all participants who did not complete the study ($n = 3$), who report having diabetes of any type ($n = 13$), or who identified as caregivers for someone with diabetes ($n = 20$), resulting in a final sample of 261 people without diabetes (PWOD). We also screened participants to ensure that they were spending time on our critical quiz manipulation. We excluded participants who took less than 60 seconds to complete the 24-question quiz ($n = 7$) and subjects who took longer than 3

³To test whether participants were simply ignoring the diagram, leading to similar results with and without it in the IOED condition, we examined time spent on the decision making questions. If participants attended to the diagram, then we would predict participants would take longer in the diagram condition than in the no info condition. We ran a 2 (IOED manipulation) x 2 (information condition) x 2 (experience) between-subjects ANOVA on completion times summed across the 5 decision making questions. We found a significant main effect of information condition, $F(1, 255) = 6.60$, $p = .011$, with participants taking significantly longer in the diagram condition compared to the no info condition, suggesting participants were not ignoring the extra information. There was also a significant main effect of IOED manipulation, $F(1, 255) = 4.71$, $p = .031$, and a marginal main effect of experience, $F(1, 255) = 3.63$, $p = .058$. There were no significant interactions, $ps > .15$, suggesting that participants across the other manipulations took longer with the diagrams.

	PWOD		
	feedback	no feedback	no quiz
No info	43	41	45
Diagram	45	42	45

Table 2: Number of subjects per condition in Experiment 2.

standard deviations to complete the quiz ($n = 4$).

Materials We used the five decision-making questions of Experiment 1. To assess knowledge of T2D we use the DKQ-24 (Garcia, Kouzekanani, Villagomez, Hanis, & Brown, 2001), which is a set of 24 true/false questions that have been validated for assessing diabetes knowledge.

Procedure We used the same procedure as Experiment 1, with the difference being the replacement of the IOED task with the DKQ-24. Consistent with Experiment 1, the DKQ-24 quiz was presented immediately prior to the decision-making questions. Participants were randomly assigned to complete the quiz without feedback ($n=83$, no feedback), to complete the quiz and see their score ($n=88$, feedback), or to not do the quiz ($n = 90$, no quiz). For the feedback condition, the percentage of correct responses was shown immediately after the quiz, and prior to the decision-making questions. We included the no feedback condition to determine whether just doing a quiz could prompt participants to reflect on their knowledge, as they may realize solely from taking the quiz that they are guessing on many questions. Within each condition, half the participants saw diagrams along with the decision-making questions and half did not. The number of subjects per condition is shown in Table 2. The order of decision-making questions and answer choices for each question was randomized for each participant.

Results

Effect of objective knowledge assessment on diagram use

We use the same MLM approach as in Experiment 1 to analyze performance on our five decision making questions. We are specifically interested in whether an objective knowledge assessment can increase accuracy with causal information in the same way a subjective one did. We found no difference in performance for the feedback and no feedback conditions, so we collapse those two participant groups into one quiz group.⁴ We entered our factors of quiz manipulation (quiz vs. no quiz; between) and information condition (diagram vs. no info; between) as fixed effects into the MLM model.

As shown in Figure 3, we observe the same patterns as in Experiment 1: accuracy followed a similar boost after an objective knowledge test as after a subjective perceived knowledge test. We found a significant main effect of quiz manipulation ($F(1, 246) = 6.86$, $p = .009$), and a significant main

⁴Entering quiz format (feedback vs. no feedback; between) and information condition (diagram vs. no info; between) as fixed effects into a MLM model did not find significant main effects or a significant interaction, $ps > .22$.

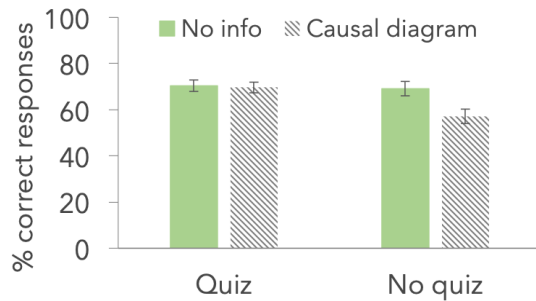


Figure 3: Decision-making accuracy for Experiment 2. Quiz condition represents results for participants who did the quiz with or without feedback. Error bars indicate standard error.

effect of information condition ($F(1, 246) = 7.51, p = .007$). We also found a significant interaction ($F(1, 246) = 5.73, p = .017$). Sidak-corrected comparisons exploring the interaction found the same pattern of results as in Experiment 1. No quiz participants showed higher accuracy in the no info condition ($M = .721$) than no quiz participants in the diagram condition ($M = .586; p = .001$). This replicates the basic detrimental effect causal diagrams can have. In the diagram condition, quiz participants did show better accuracy ($M = .718$) than no quiz participants, $p < .001$. Much like the IOED procedure of Experiment 1, an objective knowledge quiz can help overcome some of the difficulties experienced when trying to use a diagram. However, as also seen in Experiment 1, performance in the quiz condition was not better with the diagram than in the no info condition ($M = .727, p = .875$). In sum, the objective knowledge test influenced accuracy in the same way as the subjective test of Experiment 1.⁵

Relation between knowledge and accuracy Unlike the IOED task, where we only capture what participants think they know, with the DKQ it is possible to examine the effect of people’s actual knowledge. We tested whether there was a relationship between knowledge-level, measured with the DKQ-24, and decision-making accuracy. Since DKQ scores are not available for the no quiz condition, we tested for correlations in the quiz with feedback and quiz without feedback conditions. There was no correlation between quiz score and accuracy regardless of feedback (overall: $r = -.045, p = .576$; feedback: $r = -.024, p = .828$; no feed-

⁵As in Experiment 1, we analyzed completion times to see if participants were attending to the extra information in the diagram condition. We conducted a 2 (quiz manipulation: quiz vs. no quiz) x 2 (information condition: diagram vs. no info) between-subjects ANOVA over the completion times summed across the 5 decision making questions. As in Experiment 1, we found a main effect of information condition, $F(1, 246) = 13.4, p < .001$, suggesting again that participants are attending to the extra information. The main effect of quiz manipulation was not significant, $p = .282$. The interaction was significant, $F(1, 246) = 4.49, p = .035$. Follow-up Sidak corrected comparisons found that in the no quiz condition, participants took significantly longer in the diagram than the no info condition, $p < .001$. While the means were in the same direction for the quiz participants, there was not a significant difference between the diagram and no info conditions, $p = .202$

back: $r = -.041, p = .726$). Thus we find that even when participants have relevant domain knowledge, this does not necessarily aid them in making decisions in the domain. This is true overall for both the no info ($r = -.120, p = 0.304$) and diagram condition ($r = .004, p = 0.970$).

Discussion

Experiment 2 builds on the previous experiment’s findings on subjective assessment by showing that objective assessments of an individual’s knowledge also improve use of causal diagrams. This effect was not due to a general improvement in decision making, as we found that neither the IOED nor the DKQ had any effect in the no info condition. At the same time, we do not find that any of these interventions lead to significantly better accuracy with causal information than no information – they simply remove the detrimental effect of causal information, or potentially inoculate participants against negative effects of information.

General Discussion

Across two experiments, we demonstrate that reflecting on one’s knowledge in a domain can alter how new causal information is used. We previously showed that causal diagrams that provide new information can result in lower decision making accuracy (Zheng et al., 2020). We replicate these findings in our control conditions that did not involve knowledge assessment (i.e., control and no quiz conditions). We also show that inspecting one’s knowledge can fully remove this effect. While knowledge inspection did improve accuracy with diagrams relative to no inspection, it did not lead to higher accuracy with diagrams compared to when participants relied solely on what they already knew.

An open question in our experiments is the mechanism by which inspecting beliefs leads to better performance with causal diagrams. One possibility is that individuals realize their previous knowledge is incomplete and then opt to use the newly presented information instead. Relying solely on the diagrams presented in our experiments should have produced excellent performance. However, our participants did not perform better than in the no information condition. Rather than wholesale ejecting previous knowledge, people may try to fill in gaps exposed by knowledge inspection with the newly provided information. In this way, the diagram is actively integrated with prior knowledge. The integration process may be taxing, perhaps explaining no gain in accuracy compared to attempting to use prior knowledge alone. Alternatively, after inspecting their own knowledge participants may have thought they were capable of answering the question on their own and choose to ignore the diagram, resulting in performance that is more similar to the no info condition. It is a question for future research to determine how exactly knowledge revision may be happening with new causal diagrams.

Understanding how people incorporate new causal information into decision making is important given the constant onslaught of new scientific information about old problems

(e.g., what foods cause heart disease, what investments cause an increase in one's financial portfolio). Our results suggest that our beliefs could keep us from making use of information that could assist us, like calorie counts on menus or the causal models being output by many machine learning methods. Further work is needed to better understand how we can best tell people something they do not know, and then actually get them to use it.

Acknowledgments

This work was supported in part by the James S. McDonnell Foundation and the NSF under award number 1915182. We thank Kaitlyn Sharo, Tara Xu, and Min Zheng for assistance.

References

- Centers for Disease Control and Prevention. (2017). *National diabetes statistics report, 2017*. Atlanta, GA: Centers for Disease Control and Prevention.
- Fernbach, P. M., Rogers, T., Fox, C. R., & Sloman, S. A. (2013). Political extremism is supported by an illusion of understanding. *Psychological Science, 24*(6), 939–946.
- Fisher, M., & Keil, F. C. (2016). The curse of expertise: When more knowledge leads to miscalibrated explanatory insight. *Cognitive Science, 40*(5), 1251–1269.
- Garcia, A. A., Kouzekanani, K., Villagomez, E. T., Hanis, C. L., & Brown, S. A. (2001). The starr county diabetes education study. *Diabetes Care, 24*(1), 16–21.
- Hagmayer, Y., & Sloman, S. A. (2009). Decision makers conceive of their choices as interventions. *Journal of Experimental Psychology: General, 138*(1), 22–38.
- Lagnado, D. A., Waldmann, M. R., Hagmayer, Y., & Sloman, S. A. (2007). Beyond covariation: Cues to causal structure. In A. Gopnik & L. E. Schultz (Eds.), *Causal learning: Psychology, philosophy, and computation* (p. 154–172). Oxford University Press.
- Noble, D., Mathur, R., Dent, T., Meads, C., & Greenhalgh, T. (2011). Risk models and scores for type 2 diabetes: systematic review. *BMJ, 343*.
- Rozenblit, L., & Keil, F. (2002). The misunderstood limits of folk science: An illusion of explanatory depth. *Cognitive Science, 26*(5), 521–562.
- Shtulman, A. (2015). How lay cognition constrains scientific cognition. *Philosophy Compass, 10*(11), 785–798.
- Sloman, S. A., & Hagmayer, Y. (2006). The causal psychology of choice. *Trends in Cognitive Sciences, 10*(9), 407–412.
- Vitriol, J. A., & Marsh, J. K. (2018). The illusion of explanatory depth and endorsement of conspiracy beliefs. *European Journal of Social Psychology, 48*(7), 955–969.
- 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*(6), 307–311.
- Zeveney, A., & Marsh, J. K. (2016). The illusion of explanatory depth in a misunderstood field: The IOED in mental disorders. In A. Pagafragou, D. Grodner, D. Mirman,

& J. C. Trueswell (Eds.), *Proceedings of the 38th Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.

Zheng, M., Marsh, J. K., Nickerson, J. V., & Kleinberg, S. (2020). How causal information affects decisions. *Cognitive Research: Principles and Implications, 5*(1).