Title
It’s not so simple: morphosyntactically simpler languages are not always easier to learn

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 43(43)

ISSN
1069-7977

Author
Semenuks, Arturs

Publication Date
2021

Peer reviewed

It's Complicated: Improving Decisions on Causally Complex Topics

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 make frequent decisions about how to manage our health, yet do so with information that is highly complex or received piecemeal. Causal models can provide guidance about how components of a complex system interact, yet models that provide a complete causal story may be more complex than people can reason about. Prior work has provided mixed insights into our ability to make decisions with causal models, showing that people can use them in novel domains but that they may impede decisions in familiar ones. We examine how tailoring causal information to the question at hand may aid decision making, using simple diagrams with only the relevant causal paths (Experiment 1) or those paths highlighted within a complex causal model (Experiment 2). We find that diagrams tailored to a choice improve decision accuracy over complex diagrams or prior knowledge, providing new evidence for how causal models can aid decisions.

Keywords: decision making; causal reasoning; complexity

Introduction

Decision making in health is highly complex, and we receive voluminous guidance on how to manage it. We are told to eat certain foods to prevent disease, avoid others because they cause illness, learn about new findings on the best way to exercise to stay healthy, and ultimately translate all of this into daily action. Whether implicit or explicit, these guidelines all rely on causal relationships. However, we are more likely to receive piecemeal guidance about specific causal relationships (e.g., stress causes weight gain) than a complete causal model (e.g., full guidance on all of the factors that result in weight gain). Models that bring together all causes of a health phenomenon and depict how they are interrelated have been manually created for some topics and computational methods have made them possible to learn for many others (Pearl, 2000; Kleinberg, 2012). Figure 1 shows a smaller version of the Obesity System Atlas developed by the UK government. The full diagram encompasses psychological, environmental, social, and other factors influencing obesity and has over 100 nodes with many more causal connections between them.[1] Given the complexity of real computationally created causal models, those who want to use these models to help people make decisions must choose between providing the fullest accounting of the causal structure (which may overwhelm people) and providing simplified versions (that leave out relevant detail). Little is known about how this trade-off affects decision making.

[1] https://www.gov.uk/government/publications/reducing-obesity-obesity-system-map We have adapted the diagram to be suitable for non-medical audiences.

Prior work suggests that complex causal diagrams may be difficult to learn from. Work on cognitive load theory (CLT) suggests that when learning about complex phenomena people perform better when the information is split into more manageable chunks (van Merriënboer & Sweller, 2005). For decision making with a causal model like the complex obesity diagram, we could consider social systems one chunk and psychological ones another, focusing on subsets of the model that are of a size we can handle. However, factors interact across these systems so ignoring their connections may lead to the wrong conclusions.

Further, we do not necessarily prefer simplified information. Korman and Khemlani (2020) found that people perceive a single integrated causal model as more complete than a system that is split into multiple models. People also find some complex causal explanations more satisfying than ones that include fewer causes (Zemla et al., 2017). While some work has shown a preference for simpler explanations (Lombrozo, 2007), such preferences depend on both what type of system is being explained (Johnson et al., 2019) and its perceived complexity (Lim & Oppenheimer, 2020). Taken together, this work suggests that if people are presented with a complex diagram, they may judge it as more complete than a set of diagrams created from its components. On the other hand, without any way to easily chunk it, prior work suggests the diagram may be hard to learn from.

While people may prefer a comprehensive model that more fully accounts for outcomes, it is not yet known whether people can successfully use complex models in decision making. Prior work has primarily focused on people's ability to learn and make inferences from smaller structures (Rottman & Hastie, 2014) and in scenarios unrelated to prior experience and knowledge. We are focused on decision making on familiar topics, such as health, where people bring their own expectations and beliefs (whether right or wrong) to the problem. Thus, simple explanations could be considered incomplete if people perceive a gap based on their own knowledge.

Our recent work on decision making with causal models in real-life scenarios (e.g. maintaining weight, managing diabetes) found that simple models can lead to worse decisions than when people rely on what they already know (Zheng et al., 2020). In that study, we observed this effect only when the questions were about familiar topics, rather than novel scenarios. We later found that manipulating perceived knowledge removed the detrimental effect of causal models, yet did

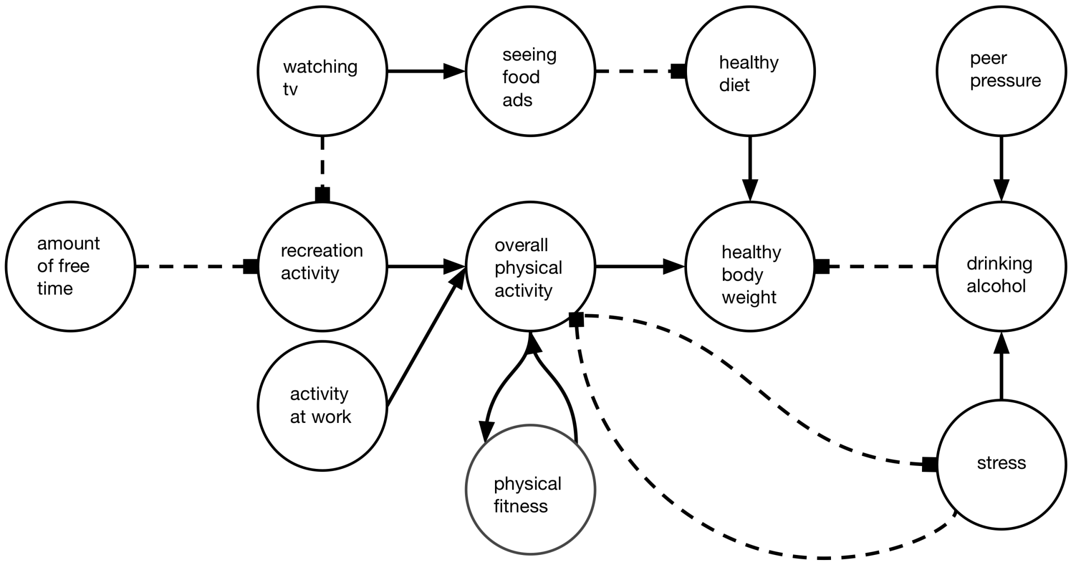

Figure 1: Complex causal diagram on managing bodyweight.

not improve decision making over just prior knowledge alone (Kleinberg & Marsh, 2020). It is an open question as to why exactly causal models lead to worse decisions in these studies. It may be difficult to integrate causal models with prior knowledge or understand which aspects of a model are most important for a choice. Alternately, the simplicity of the models used in previous research could be at the root of the challenges, as the models may leave out details people expect to see.

We aim to advance understanding of when and which causal models are useful for decision making. If we provide a complex comprehensive model that captures more causes, will people be able to use it? Alternatively, if we provide a simpler model but make it directly relevant to the question at hand, without any broader causal information, will this too improve decisions? We conduct two experiments designed to shed light on the trade-off between simplicity and complexity in decision making. We use both a familiar topic that many individuals have experience with (decisions surrounding bodyweight) and an unfamiliar domain as a control topic (alien dance-off). In Experiment 1 we compare decision making using simple diagrams containing only information relevant to each choice versus a complex diagram that covers information beyond each specific question. In Experiment 2 we test whether drawing people's attention to the relevant part of a complex diagram has the same effect as presenting only that component in a simple form.

Experiment 1

In this experiment we look at how information detail influences decision-making accuracy. We focus on decision making around bodyweight, as this is a topic for which people often make decisions and receive information. We test whether simple diagrams tailored directly to each decision can improve accuracy over a complex diagram that includes information from all the simple diagrams and other information not needed for the decision at hand.

Method

Participants We recruited 300 U.S. residents aged 18-64 from Prolific, with 299 completing. Participants were compensated $3.00. We excluded participants who failed our attention check or submitted unusual and duplicate responses (n = 36). Thus 263 participants remain in the analysis.

Materials Inspired by the complex diagram shown in Figure 1, we created a set of questions that each target a decision related to healthy bodyweight and that can be directly aided by using the diagram. We chose four different types of causal pathways contained in the complex diagram and created a corresponding simple diagram (Figure 2) that provides only the subset of the complex diagram that is relevant to answering that question. These represent (a) a direct preventative relationship (prevent), (b) a common effect structure (two causes), (c) a mix of positive and negative relationships (mixed causes), and (d) a causal chain (direct cause). The first three diagrams represent causal pathways that increase in complexity (from 2 to 4 nodes) and the correct answers require participants to choose a direct preventative cause, a combination of two positive causes, and to activate a positive cause while deactivating a negative one. These simple diagrams are tailored to each corresponding question and contain solely the causal pathways relevant to the target answer. The direct cause diagram is different in that we now include information beyond that needed for the target answer to test whether simple diagrams influence participants' likelihood of choosing the most direct cause. Causes in many areas such as health are rarely deterministic, and from a probabilistic perspective, the proximal, direct, cause is most likely to bring

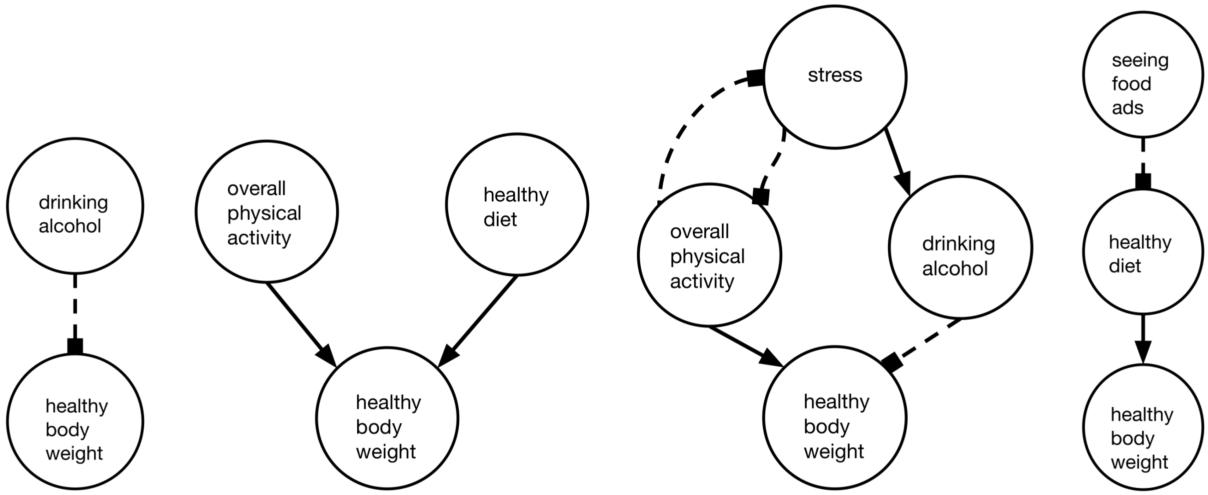

(a) Prevent (b) Two causes (c) Mixed causes (d) Direct

Figure 2: Simple diagrams used in Experiment 1.

about the effect. Despite this, people have been shown to prefer root nodes in making interventions (Lagnado & Sloman, 2006). We analyze this question separately since it is an open question if participants will choose the direct or distal cause and how diagram complexity will influence this choice.[2]

We use the same diagram format (using solid arrows to indicate generative/positive causes, and dashed lines ending in a solid box to indicate preventative/negative causes) as in the original obesity system map.[3] Different questions were developed for each diagram, with all having a similar structure. For example, the question paired with the direct diagram was:

Daniel has been overweight for most of his life. He lives with his brother and sister, both of whom also would like to lose weight. They don't have much free time during the work week, so on weekends they like to relax by watching TV and cooking from new recipes. Daniel is concerned about his weight but doesn't know what he could do differently.

Which of the following is the BEST suggestion for Daniel?
A. Don't do anything, weight is genetic
B. Fast forward through TV commercials
C. Get takeout pizza instead of cooking
D. Add more vegetables to his weekend recipes

For questions paired with the prevent, two causes, and mixed causes diagrams, there was one answer choice that we designate as "correct" given the causal diagram. This answer represented intervening on a node in the diagram that maximized producing the desired effect. The other answer choices either represented causal relationships that were not in the diagrams, ineffective relationships that contradicted the diagrams (e.g., an option to increase alcohol intake when the diagram suggests reducing it). For the direct cause question, one answer choice corresponded to the direct cause and one to the distal root cause. The remaining two choices were incorrect in the same way as the other models' questions.

We created four control questions that targeted the same causal pathways and used the same simple and complex diagram structures. However we replaced the node labels and question text with a scenario participants should not be familiar with, namely an alien dance off. In the alien dance off, aliens can cause others to dance faster (positive causal relationship) or dance slower (negative causal relationship).

Procedure Participants were instructed on the meaning of nodes and solid/dashed lines, as well as the rules of alien dance offs for the control questions. Participants were randomly assigned to one of three conditions for the bodyweight question: simple diagram (simple), complex diagram (complex), or no diagram (no diagram). The simple diagrams provide information directly tailored to the question at hand, while the complex diagrams encompass all of the relationships shown across the simple diagrams. In the no diagram condition, participants did not receive a visual aid and had to answer the questions using their background knowledge alone. For the control questions, simple and complex condition participants received the opposite diagram of what they received in the main experimental questions (simple received complex and complex received simple) to make it less obvious the control questions targeted the same pathways for which they had previously answered questions. Participants in the no diagram condition were randomized to receive either simple or complex diagrams. The order of experimental and

[2] We tested a fifth question that depicted a feedback loop, but it proved difficult to interpret, so we do not discuss it further here.

[3] We previously (Zheng et al., 2020) used plus and minus signs to indicate generative and preventative causes. We reran our prior weight management question with the current solid and dashed line format. We replicated the earlier main effect, namely decreased accuracy with the diagram compared to no diagram. Thus it does not appear that this presentation style influences accuracy.

control questions and their answer choices were randomized for each participant. After completing all questions participants completed two free-text response questions designed to evaluate their understanding of the control question set up and the meaning of solid and dashed lines. We excluded participants who failed this attention check.

Results

Influence of complexity on decision accuracy To determine if participants could use the diagrams, we first explored mean accuracy on the control questions. We collapsed all participants who received a simple control diagram (complex condition plus simple half of the no diagram condition) and all participants who received a complex control diagram (simple condition plus complex half of the no diagram condition). We calculated mean accuracy across the three main control questions (prevent, two causes, and mixed causes). Using a one-way ANOVA with control diagram condition (simple vs. complex) as a between-subjects factor we found a significant main effect ($F(1, 261) = 27.2, p < .001, \eta_p^2 = .094$) indicating that performance was significantly better with a simple diagram ($M = .814$, $SE = .023$) than a complex diagram ($M = .626$, $SE = .028$). These means suggest participants are able to use simple diagrams to make accurate choices.

Our control analysis established that targeted diagrams can be more useful to answering our decision making questions than complex diagrams that capture more information. We then tested whether in real world questions where prior knowledge may interact with the provided information, simple diagrams still have an edge. We calculated mean accuracy across the three main bodyweight questions (prevent, two causes, and mixed causes). Using a one-way ANOVA with diagram condition (no diagram, simple, complex) as a between-subjects factor we found a main effect, $F(2, 260) = 9.65$, $p < .001$, $\eta_p^2 = .069$. Sidak corrected t-tests showed that the significant main effect was driven by the simple diagram condition: performance in the simple condition ($M = .841$, $SE = .030$) was significantly better than the no diagram ($M = .683$, $SE = .031$; $p = .001$) and the complex ($M = .681$, $SE = .028$; $p < .001$) conditions, while the no diagram and complex conditions did not differ from each other, $p = 1$.

Influence of simplicity on selection of direct causes We now compare performance using the simple diagram of Figure 2d against the complex one of Figure 1, to test whether the simpler diagram may trigger a preference for root causes. We compared the percentage of people who chose the answer corresponding to a direct proximal cause (healthy diet) versus the answer corresponding to a more distal root cause (seeing food ads). In the no diagram condition, the vast majority of participants chose the answer representing the most direct intervention (96.4%; $SE = 2.02$; see Figure 3), and only one chose the answer corresponding to a distal cause (1.19%; $SE = 1.18$). This establishes a baseline measure of how likely people are to choose these two options based on their existing knowledge. We find that diagrams can significantly alter this choice. When provided with the simple diagram fewer participants chose the direct cause (44.3%; $SE = 5.30$) and significantly more selected the distal (root in this diagram) cause (54.5%; $SE = 5.31$). Yet we did not find this effect with the complex diagram, where most participants chose the proximal direct cause (84.6%; $SE = 3.78$) and few chose the more distal cause (12.1%; $SE = 3.42$). Comparing frequency of choosing the direct cause across conditions using a Pearson's Chi squared test with N-1 correction, significantly more participants chose the direct cause in the complex than the simple condition ($p < .001$, χ^2=19.0), and in the no diagram compared to the simple condition ($p < .001$, χ^2=39.8). The difference between the no diagram and complex conditions was also significant ($p = .009$, χ^2=6.88) Thus, when diagrams highlight a cause further back in the chain, that option is more popular. However, in a complex diagram, this preference for distal causes diminishes.

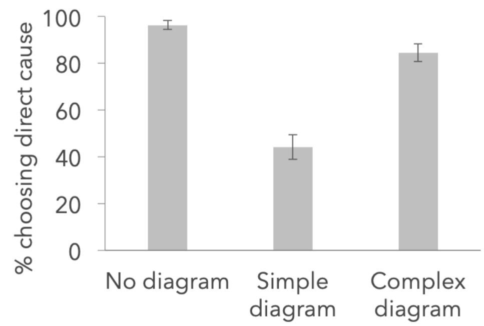

Figure 3: Selection of direct cause in Experiment 1.

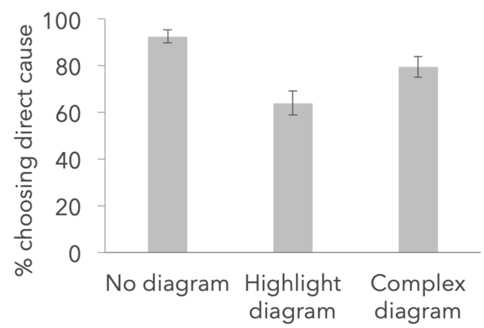

Figure 4: Selection of direct cause in Experiment 2.

We also analyzed the preference for direct causes in the corresponding control question. In the simple condition, fewer people chose the direct proximal cause (37.4%; $SE = 4.23$) compared to the distal root cause (49.6%; $SE = 4.37$), while the majority of participants in the complex condition chose the direct cause (56.1%; $SE = 4.32$), with fewer people choosing the distal cause (25.6%; $SE = 3.81$). Again, a simple diagram makes people more likely to choose a more distal cause than a complex diagram does ($p = 0.002$, χ^2=9.20).

Discussion

Prior work has uncovered a conflict: people are able to learn about and use causal models (Rottman & Hastie, 2014) and yet visual depictions of such models can lead to worse decisions when combined with people's existing knowledge (Zheng et al., 2020). Our findings begin to resolve this conflict by illuminating where models can help in everyday decisions and what models are most helpful. First, we can reject the hypothesis that simple models impede decisions because they do not provide a full accounting of a causal structure. As our experiments on both real-world (weight management) and control (alien dance off) questions find, simpler models can outperform complex ones. Second, a critical difference between the simple models used in this experiment and those used in prior work is that our diagrams contain only information needed to successfully answer the question. Indeed, we replicated the detrimental effect of diagrams from prior work when we tested our solid/dashed format with previously used diagrams that are not tailored to a specific question and include extra information beyond the causal paths pertaining to the right answer. Thus, this provides a first positive step: causal models can be helpful even when people have prior knowledge if they are tailored to the specific decision at hand.

Lastly, using our question that pits direct against indirect causes, we find that presenting a subset of a model versus the more complete one can strongly influence people's choices. Prior work has found people have a preference for intervening on root nodes in causal networks (Hagmayer & Sloman, 2009; Lagnado & Sloman, 2006; Yopchick & Kim, 2009). It appears we can trigger this preference when presenting simple diagrams (where a node that is not normally a root appears to be one). Thus it is important to consider potential negative effects of how causal models are presented, whether that is steering people toward more distal causes (which may not be desirable) or including information beyond that which is necessary (which can impede decision making).

Experiment 2

Our first experiment showed that diagrams tailored to a decision can aid decision making compared to relying on prior knowledge or a complex and comprehensive diagram. We now investigate whether similar results can be obtained by highlighting what information is relevant to a decision in a complex model. That is, rather than creating different diagrams for each decision, we now test whether directing people's attention to relevant aspects of a diagram is sufficient to help them obtain the benefits of this information while enabling them to ignore irrelevant information (and thus not having the negative effects seen in prior work).

Method

Participants We recruited 300 U.S. residents aged 18-64 from Prolific, with 290 completing. Participants were compensated $3.00. As before we dropped participants who failed our attention check or submitted nonsense answers ($n = 37$). A total of 253 participants remain in the analysis.

Materials We use the same study protocol and materials as in Experiment 1. The key difference now is that instead of simple diagrams, we highlight the nodes and edges from each simple diagram within the complex diagram. For example, instead of the two node prevention diagram, those nodes and the edge between them are depicted in orange within the complex diagram. Orange was chosen to grab people's attention without prompting any associations with positive or negative effects. Thus participants in the highlighted diagram condition (highlight) see different parts of the complex diagram highlighted for each question, while the no diagram and complex groups remain the same as in Experiment 1. We follow the same approach for the control questions. The control complex diagram has nodes rotated and moved from their positions in the bodyweight diagram, so we do not expect participants to realize they have the same causal structure.

Procedure The procedure was the same as Experiment 1, except we added an explanation to the instructions of the meaning of the highlighting.

Results

Effect of highlighted causal paths on decision accuracy We again begin by testing whether the diagram format made a difference for accuracy on the control questions. We collapsed all participants who saw the highlighted version of the diagram into one group and similarly all participants who saw the complex version without highlighting into another group. Using a one-way ANOVA with diagram condition (highlighted vs. complex) as a between-subjects factor we did not find a significant main effect, $p = .157$, indicating that performance was not significantly better with the highlighted control diagram ($M = .620$, $SE = .0263$) compared to the complex diagram ($M = .562$, $SE = .029$).

While there was no difference in the control condition, we now examine whether highlighting the relevant paths makes a difference to the experimental questions. Our hypothesis was that highlighting would have a similar effect to the simple diagram. Using a one-way ANOVA with diagram condition (no diagram, highlighted, complex) as a between-subjects factor we now do find a main effect of diagram condition, $F(2, 250) = 5.19$, $p = .006$, $\eta_p^2 = .040$. Sidak corrected follow up t-tests show that performance in the highlighted condition ($M = .757$, $SE = .033$) was marginally better than the no diagram condition ($M = .654$, $SE = .031$; $p = .079$) and significantly better than the complex condition ($M = .615$, $SE = .034$; $p = .006$). The no diagram and complex conditions did not differ from each other, $p = .778$. Overall, the pattern of means found here replicate Experiment 1, and suggest that focusing attention on the relevant aspects of a diagram may be sufficient for making them useful for a decision.

Effect of highlighted paths on preference for direct causes In Experiment 1 we found that a simplified diagram increased the preference for distal causes. We now test whether highlighting the same paths within a complex diagram yields the same effect. The vast majority of participants in the no diagram condition again chose the answer corresponding to the most direct cause (92.6%; $SE = 2.91$; see Figure 4), while only two participants chose the indirect one (2.47%; $SE = 1.72$). Highlighting decreased the percentage of people choosing the direct cause (64.0%; $SE = 5.09$; $p < .001$, $\chi^2 = 19.855$ compared with no diagram) and led a large percentage of people to chose the distal cause (32.6%; $SE = 4.97$). As in Experiment 1, most participants in the complex condition chose the direct cause (79.5%; $SE = 4.43$) and few chose the more distal cause (13.3%; $SE = 3.72$). While significantly more participants in the complex condition chose the direct cause compared to in the highlighted condition ($p = 0.025$, $\chi^2 = 5.03$) we observed a significant drop compared to the no diagram condition ($p = 0.016$, $\chi^2 = 5.78$).

In the control questions we do not see this pattern. Instead, similar percentages of participants chose the direct cause (highlighted: 48.3%; $SE = 4.37$; complex: 51.9%, $SE = 4.52$, $p = 0.568$, $\chi^2 = 0.326$) and distal cause (highlighted: 27.8%; $SE = 3.92$; complex: 24.4%, $SE = 3.89$) across both conditions and the direct cause was by far the most frequent response in both (highlighted: $p = 0.0007$, $\chi^2 = 11.5$; complex: $p < 0.001$, $\chi^2 = 19.8$).

Discussion

This experiment builds on our findings in Experiment 1 by identifying another way information can aid decision making. We find that highlighting the relevant causal paths elicits the same decision-making behavior as providing only those paths with our real-world decision questions. Together these experiments show that tailoring information or directing people's focus to only the paths relevant to a decision can improve their performance, and can make diagrams that are not otherwise improving performance, useful and usable. Interestingly, we do not show this same effect in the control condition of this experiment. We hypothesize that the alien nature of an alien dance off may make it more challenging for participants to understand if more distal causes have the same downstream influences as causes do in real life. It is an interesting question for future work to determine why highlighting of complex diagrams in an unfamiliar domain may be less likely to shift people to focusing on distal causes.

General Discussion

Given the prevalence of causal information we receive (particularly about health) and share with others, it is vital to understand whether such information actually improves our decisions. Prior work has shown both that people are adept at learning causal models and using them to intervene on systems (Rottman & Hastie, 2014) and that such models can also lead to worse decisions when they interact with people's existing beliefs and knowledge (Zheng et al., 2020; Kleinberg & Marsh, 2020). We move toward reconciling this view through experiments comparing causal information tailored to a decision and comprehensive causal information. We find that information tailored to a decision that does not include or draw attention to causal paths not relevant for the choice led to better choices in our real-world decision-making questions. This suggests one way that causal models in prior work may have been interfering with decisions - by prompting individuals to consider alternative causes. By either removing these or focusing attention on the relevant ones, we reduce the likelihood participants will select them.

Our results also provide insight into how we can improve people's reasoning with complex information. For some real world decisions, such as choosing a federal economic policy or restructuring the healthcare system, the complex inner workings of the relevant causal system may be important to represent in making a decision. In this case, highlighting a path relevant to a given decision may benefit the reasoner while still allowing the full complexity of the system to be represented. More generally, while prior work has shown that that complex models can be overwhelming and that people have a hard time ignoring irrelevant information (Bastardi & Shafir, 1998; Hall, Ariss, & Todorov, 2007), we have found ways to mitigate this complexity and help people use complex information to make an accurate decision.

While participants in our study successfully used the tailored diagrams, future work is needed to probe how such interventions influence model trust. That is, given prior work on people's preferences and beliefs about the completeness of models, it is important to determine whether simplification could have negative consequences by making people less likely to trust the model. This may also provide reason to prefer highlighting, reassuring individuals about the completeness of the model while allowing useless information to be ignored. These are important questions to be asked in future research. Overall, our two experiments provide new insights into how people may tangle with complexity in diagrams and decisions in the real world.

We found that providing people with diagrams tailored to a specific decision-making context could aid decision making. This suggests that instead of generic advice pertaining to a variety of choices, people may need highly focused guidance for the specific decision at hand. While this may seem impractical, many aspects of life are already personalized to our interests and needs (e.g. book and movie recommendations). Future work is needed to bridge cognitive science and machine learning (which can be used to create this personalized guidance), and ultimately help translate knowledge of what people should to do guidance they can follow as they make real-world choices.

Acknowledgments

This work was supported in part by the James S. McDonnell Foundation and the NSF under award number 1915182.

References

Bastardi, A., & Shafir, E. (1998). On the pursuit and misuse of useless information. *Journal of Personality and Social Psychology*, *75*(1), 19–32.

Hagmayer, Y., & Sloman, S. A. (2009). Decision makers conceive of their choices as interventions. *Journal of Experimental Psychology: General*, *138*, 22–38.

Hall, C. C., Ariss, L., & Todorov, A. (2007). The illusion of knowledge: When more information reduces accuracy and increases confidence. *Organizational Behavior and Human Decision Processes*, *103*(2), 277–290.

Johnson, S. G., Valenti, J., & Keil, F. C. (2019). Simplicity and complexity preferences in causal explanation: An opponent heuristic account. *Cognitive Psychology*, *113*, 101222.

Kleinberg, S. (2012). *Causality, Probability, and Time*. New York: Cambridge University Press.

Kleinberg, S., & Marsh, J. K. (2020). Tell me something I don't know: How perceived knowledge influences the use of information during decision making. In S. Denison, M. Mack, Y. Xu, & B. Armstrong (Eds.), *Proceedings of the 42nd Annual Meeting of the Cognitive Science Society* (pp. 1849–1855). Cognitive Science Society.

Korman, J., & Khemlani, S. (2020). Explanatory completeness. *Acta Psychologica*, *209*, 103139.

Lagnado, D. A., & Sloman, S. A. (2006). Time as a guide to cause. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *32*(3), 451.

Lim, J. B., & Oppenheimer, D. M. (2020). Explanatory preferences for complexity matching. *PloS One*, *15*(4), e0230929.

Lombrozo, T. (2007). Simplicity and probability in causal explanation. *Cognitive Psychology*, *55*(3), 232–257.

Pearl, J. (2000). *Causality: Models, Reasoning, and Inference*. Cambridge University Press.

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

van Merriënboer, J. J., & Sweller, J. (2005). Cognitive load theory and complex learning: Recent developments and future directions. *Educational Psychology Review*, *17*(2), 147–177.

Yopchick, J. E., & Kim, N. S. (2009). The influence of causal information on judgments of treatment efficacy. *Memory & Cognition*, *37*, 29–41.

Zemla, J. C., Sloman, S., Bechlivanidis, C., & Lagnado, D. A. (2017). Evaluating everyday explanations. *Psychonomic Bulletin & Review*, *24*(5), 1488–1500.

Zheng, M., Marsh, J. K., Nickerson, J. V., & Kleinberg, S. (2020). How causal information affects decisions. *Cognitive Research: Principles and Implications*, *5*(6), 1–24.