COGNITIVE SCIENCE
A Multidisciplinary Journal

Cognitive Science 39 (2015) 1912–1924

ISSN: 0364-0213 print / 1551-6709 online
DOI: 10.1111/cogs.12210

Effects of Causal Structure on Decisions About Where to Intervene on Causal Systems

Brian J. Edwards,[a,b] Russell C. Burnett,[a] Frank C. Keil[a]

[a]*Department of Psychology, Yale University*
[b]*Department of Psychology, Northwestern University*

Received 14 June 2013; received in revised form 14 July 2014; accepted 23 July 2014

Abstract

We investigated how people design interventions to affect the outcomes of causal systems. We propose that the abstract structural properties of a causal system, in addition to people's content and mechanism knowledge, influence decisions about how to intervene. In Experiment 1, participants preferred to intervene at specific locations (immediate causes, root causes) in a causal chain regardless of which content variables occupied those positions. In Experiment 2, participants were more likely to intervene on root causes versus immediate causes when they were presented with a long-term goal versus a short-term goal. These results show that the structural properties of a causal system can guide the design of interventions.

Keywords: Causal reasoning; Intervention

1. Introduction

Causal knowledge is essential for understanding how the world works. Our knowledge of the causal relationships between events not only allows us to make predictions about how the occurrence of one event will influence another (e.g., if I do not study hard, then I will do poorly on the test), it also enables us to act on the environments in which we live. We can use this causal information to *intervene* on causal systems of events to achieve more desirable outcomes or avoid undesirable ones. In this example, I could increase the amount of time I spend studying to improve my performance on tests. More precisely, we define an *intervention* as an action that manipulates the value of a particular variable in a causal system (e.g., the amount of time I spend studying).[1] Designing interventions is central to our interaction with the world, and understanding how to intervene

Correspondence should be sent to Brian J. Edwards, Department of Psychology, Northwestern University, 2029 Sheridan Rd., 102 Swift Hall, Evanston, IL 60208. E-mail: brian.edwards@aya.yale.edu

on a system is arguably the most fundamental component of causal reasoning (Woodward, 2003; see also Woodward, 2007).

For example, we can use our causal knowledge to evaluate the effectiveness of various strategies for improving the quality of health in the United States. The following causal chain is a simple model of how education might lead to good health: Education causes access to good health information, which causes people to practice healthy habits, which causes an overall high level of health. Causal knowledge like this sheds light on potential interventions such as increasing education funding to improve the quality of education and launching campaigns to promote healthy habits.

Several factors influence our reasoning about interventions. One class of factors concerns our knowledge of a specific content domain and the mechanisms by which interventions lead to outcomes. Some mechanisms may be more reliable than others; for example, one might believe that the mechanism by which healthy habits lead to good health is more reliable than mechanisms linking other possible causes to good health. This may be a reason for preferring to intervene on people's habits. Pragmatic factors, such as the cost (in money, time, or effort) and feasibility of various interventions, might also affect intervention decisions. It is possible that a campaign to promote healthy habits would cost less than the investment in education needed to achieve a comparable increase in the quality of health.

The factors just mentioned involve specific knowledge of the causal system at hand (hereafter "content factors"). Here, we focus on a second type of factor: the abstract *structure* of the causal system. By *structure*, we mean information about what has causal influence on what, without any further detail about the content of the variables, the mechanisms by which these causal relationships operate, or any pragmatic factors that might be involved. Although content knowledge and structural knowledge are not necessarily distinct and can interact—content knowledge can suggest a causal structure that governs how a particular system works—structural considerations may also influence interventions above and beyond mechanistic and pragmatic factors. People may choose an intervention targeting healthy habits in part because it is the *immediate (proximal) cause* of good health, or, alternatively, intervene on the quality of education because it is the *root (distal) cause* of good health (in this particular chain).

Explaining causal reasoning and related forms of cognition at least partly in terms of abstract structure has supporting precedents. For example, causal structures can explain why some objects make better members of categories than others (for a review, see Rehder, 2010). Ahn's (1998) causal status hypothesis predicts that features of objects are more important to category membership to the degree that they are deep in causal structure; that is, the degree to which they cause or enable other features associated with the category (see also Rehder & Hastie, 2001). Similarly, the degree to which symptoms are further apart in a branching tree of causes (i.e., the degree of diversity of symptoms) can influence inductions about the likelihood of a root cause being present (Kim & Keil, 2003; Kim, Yopchick, & de Kwaadsteniet, 2008). More generally, the structure of causal relationships has become a major factor in making behavioral predictions about causal learning and reasoning from the perspective of causal Bayes net theory (for a review, see

Gopnik et al., 2004). Although the present experiments were not motivated by causal Bayes net theory and have implications for a much wider range of approaches, in light of our results, the General Discussion does consider possible implications for causal Bayes nets.

The present studies explore how people's interventions on causal systems might depend on structural factors. Perhaps people prefer proximal interventions—that is, interventions on variables that are close to the outcome of interest in causal structure—because people deem these interventions to be more reliable or more efficacious. White (1997) proposed a "dissipation effect," whereby people judge that the effects of interventions become weaker as they propagate through a causal network (see also White, 2000). This theory predicts a preference for proximal interventions since such interventions have less distance over which to dissipate.

Proximal interventions may also be preferred if causal relationships are construed as probabilistic, or "noisy." A probabilistic causal relationship is one that is not perfectly reliable; it operates with a probability less than 1. When you throw a light switch, the light does not always turn on; perhaps the switch has broken, the bulb has burned out, or there has been a power failure. Every probabilistic link in a causal chain is an opportunity for failure. If a link fails to operate, then the flow of causation is stopped, and downstream effects will fail to appear. Thus, if one regards a given causal chain as a sequence of probabilistic links, one might prefer a proximal intervention over a distal one because proximal interventions travel through fewer links in the chain and therefore have lower probabilities of failure.

However, sometimes more distal interventions may be preferred. Perhaps root-cause interventions are seen as especially effective because they address the underlying causes of a problem and provide a stable and lasting solution. In contrast, an intervention on the immediate cause may be perceived as a "quick fix" that can be easily undone. For example, an effective strategy for fighting terrorism might consist not only of capturing and killing as many terrorists as possible, but also of eliminating the conditions that cause people to become terrorists in the first place. Thus, we predict that when people are interested in achieving long-term solutions (vs. short-term solutions), they will prefer interventions that target root causes. Experiment 2 investigates this hypothesis.

In Experiment 1, we ask whether there is a tendency to intervene on variables that occupy specific positions in a causal chain, such as the immediate or root causes of an outcome. We predict that structural factors, in addition to content-based factors such as pragmatic considerations and knowledge of particular mechanisms, will affect intervention decisions. We therefore designed our experiments to allow for these nonstructural effects and yet isolate the unique effects of abstract causal structure.

2. Experiment 1

We manipulated the positions of the specific content variables in a causal chain (e.g., $A \rightarrow B \rightarrow C \rightarrow D$ vs. $C \rightarrow B \rightarrow A \rightarrow D$) and asked participants how they would intervene to affect the outcome. If structural knowledge affects intervention decisions, participants should be more likely to intervene on each variable when it appears in some positions than in others.

2.1. Method

2.1.1. Participants

We tested 144 adults who were recruited at busy locations on the Yale University campus or through the psychology department subject pool. Participants received a snack or course credit for their participation.

2.1.2. Materials

The stimuli were four causal chains covering a range of real-world phenomena: why a person enjoys playing poker, a child's interest in becoming a veterinarian, why a student does poorly in school, and factors that cause the stock market to go down. We selected a broad range to see whether structural influences would be at work across very different content domains. For each stimulus, we presented participants with a causal chain followed by a list of interventions that could be used to affect the outcome. Participants were asked to choose the best intervention (e.g., encouraging a student to do his homework) to achieve some goal (e.g., improving a student's performance on tests). The variables in each chain and the results are shown in Table 1; full versions of sample stimuli for each experiment are included in the Appendices and full versions of all stimuli are available in the supplementary material.

Each causal chain had four variables and was constructed such that the positions of the first three variables could be interchanged and still form a plausible causal chain. Thus, each stimulus item could be presented in $3! = 6$ possible orders. For each item, each participant randomly received one of the six orders. The order of the stimuli was counterbalanced.

2.1.3. Procedure

The experimenter handed each participant a pencil-and-paper questionnaire containing the four items described above. The questionnaire took about 5 min to complete.

2.2. Results and discussion

All interventions were considered plausible. No intervention was chosen less frequently than 21% of the time or more frequently than 44% of the time. We then analyzed the distribution of participants' responses to determine whether causal structure influenced their choice of interventions. If people's interventions are completely determined by their content knowledge, then for each variable (e.g., whether Mark does his homework), people's interventions should be uniformly distributed over the three locations in the causal chain (root, middle, and immediate). This is because each variable appeared in each position one-third of the time. However, if causal structure influences participants' interventions, then they should be more likely to intervene at some positions than at others.

Table 1
Number of interventions on each variable and position in the causal chain in Experiment 1

Stimulus	Variable	Root	Middle	Immediate	Total
Poker	Frank's poker ability	20	19	23	62
	Amount of time Frank spends playing poker	13	9	17	39
	Amount of money Frank wins at poker	18	12	13	43
	How much Frank enjoys playing poker	(Outcome of causal chain)			
	Total	51	40	53	144
Veterinarian	Whether Sally watches animals at the zoo	17	6	16	39
	Whether Sally talks about animals	15	5	21	41
	Whether Sally volunteers at an animal center	24	14	26	64
	Whether Sally becomes a veterinarian	(Outcome of causal chain)			
	Total	56	25	63	144
School performance	Whether Mark does his homework	38	11	4	53
	Whether Mark understands a class lesson	39	9	13	61
	Whether Mark pays attention in class	25	2	3	30
	Whether Mark does well on a test	(Outcome of causal chain)			
	Total	102	22	20	144
Economy	Consumer confidence about the economy	22	6	9	37
	Amount of money consumers spend	31	11	14	56
	Number of companies laying off workers	32	11	8	51
	Stock market performance	(Outcome of causal chain)			
	Total	85	28	31	144
All	Total	294	115	167	576

For each item, we used a chi-square test of independence to find out whether the distribution of people's interventions over the three locations (root, middle, immediate) deviated from a uniform distribution. We found that causal structure had a significant effect on participants' interventions for three of four stimuli (becoming a veterinarian: $\chi^2(2) = 17.0$, $p < .001$, improving school performance: $\chi^2(2) = 91.2$, $p < .001$, and improving the economy: $\chi^2(2) = 42.9$, $p < .001$).

Overall, participants were most likely to intervene on the root cause, selecting the variable in this position 294 of 576 times (51%). Furthermore, they were more likely to intervene on the immediate cause, choosing this variable 167 times (29%), than on the middle cause, which was chosen 115 times (20%). (Not all of the stimuli showed this pattern, and we discuss this variability below.) A series of t-tests confirmed these findings. For each participant, we calculated the proportion of interventions that fell on the root variable as opposed to any of the other variables. On average, these proportions were reliably greater than 1/3, indicating a significant preference for intervening on the root cause, $t(143) = 7.30$, $p < .001$. A similar analysis revealed that when one of the nonroot variables was chosen, people had a statistically significant tendency to intervene on the immediate cause as opposed to the middle cause. On average, these proportions were reliably greater than 1/2, $t(124) = 2.15$, $p < .05$, suggesting a preference to intervene on an immediate cause and avoid interventions on intermediate variables.

Despite these trends, participants' responses varied considerably across the four items. For the "school performance" and "economy" items, participants had a strong tendency to intervene on root causes. In contrast, a plurality of participants chose immediate-cause interventions when reasoning about the "poker" and "veterinarian" items. Perhaps for the "school performance" and "economy" items, participants cared more about influencing multiple variables in the chain. For example, in a chain in which companies laying off workers is presented as the root cause of stock market performance, a person may want to perform an intervention that decreases layoffs not just because it will have a positive effect on the stock market's performance, but also because it will increase consumer spending and be beneficial to the workers who keep their jobs.

Not surprisingly, participants' content knowledge also affected their choice of interventions for three of four stimuli. Content effects were shown by employing the same form of chi-square analysis but counting interventions not over locations in causal structure but over specific manipulations such as encouraging Mark to do his homework. Chi-square tests showed reliable deviations from uniform distributions for poker ($\chi^2(2) = 6.29$, $p < .05$), becoming a veterinarian ($\chi^2(2) = 8.04$, $p < .05$), and school performance ($\chi^2(2) = 10.8$, $p < .01$), meaning that some specific manipulations were selected more frequently than others, regardless of causal structure. For the fourth item, improving the economy, participants probably did not ignore content and instead may have found each intervention to be equally effective or disagreed on which intervention was best.

We note that when deciding how to intervene, participants likely considered their prior knowledge about the relationships between variables and efficacy of the proposed interventions along with the specific causal chains they were given, which in some cases may have been at odds. Although we did not ask participants about the degree to which they relied on their prior knowledge versus the specific causal chains, if participants relied exclusively on their prior knowledge, there should be no observed effects of causal structure because locations of content in the chains were counterbalanced.

However, the distribution of participants' interventions across the different orders shows that both content and structural knowledge affected their choice of interventions. Participants' choices of interventions appeared to be more dependent on causal structure in some stimulus sets than in others, presumably because beliefs about other effects of particular interventions, expectations about the ease or reliability of particular interventions, and so on—are more important in some domains and causal systems than in others. The present experiment was not designed to assess the relative magnitudes of effects of structure and other factors across stimuli or across domains. Rather, it simply tested whether there *are* systematic effects of causal structure across domains, and indeed the results are consistent with this hypothesis.

3. Experiment 2

Experiment 2 tested whether framing a problem in a short-term or long-term context would influence whether people intervene proximally (near the immediate cause) or distally (near the root cause). In the Introduction, we suggested that interventions on root causes might be perceived as achieving more stable and lasting effects than interventions on immediate causes, which might be regarded as "quick fixes" with little long-term efficacy. Thus, we hypothesize that people presented with problems in a long-term context will intervene more distally than people presented with problems in a short-term context.

3.1. Method

3.1.1. Participants

We tested 41 adults. Participant recruitment and compensation were identical to Experiment 1. Each participant was randomly assigned to the short-term (n = 21) or long-term condition (n = 20).[2]

3.1.2. Materials

Our stimuli were seven causal chains covering a range of real-world and artificial phenomena. We included artificial stimuli because participants' content knowledge might influence how they intervene on real-world systems. The real-world stimuli were as follows: preventing a heart attack, improving the quality of health, caring for one's car (adapted from Rehder & Hastie, 2001), and preventing terrorism. The artificial stimuli were as follows: an alien mind-reading game (adapted from Steyvers, Tenenbaum, Wagenmakers, & Blum, 2003), transmission of a made-up disease, and casting a spell. The stimuli were presented in the above order, with the real-world stimuli appearing first. For each stimulus, we presented participants with a causal chain and provided them with a list of interventions that could be used to affect the outcome.

For the "health" and "terrorism" stimuli, we explicitly asked participants to select the best short-term or long-term intervention. For the "heart attack" and "car" stimuli, we used more subtle experimental manipulations. For the "heart attack" stimulus, the short-term and long-term manipulations were preventing a heart attack in a 70-year-old and a 30-year-old, respectively, and for the "car" stimulus, the experimental manipulations were preventing muffler damage to an old car and a new car, respectively. The purpose of the subtle manipulations was to test whether an effect of short-term vs. long-term context depends on explicitly invoking a short-term or long-term goal. The manipulations for the artificial stimuli varied across items. The real-world causal chains and results are shown in Table 2; the variables are shown in the order in which they formed the chain.

3.1.3. Procedure

The experimenter handed each participant a pencil and paper questionnaire containing the seven stimulus items described above. The questionnaire took approximately 10 min to complete.

3.2. Results and discussion

We analyzed participants' responses based on the distance between the variable they intervened on and the outcome. Because the lengths of the causal chains varied by item,

Table 2
Number of interventions on each variable in Experiment 2

	Condition	
	Short	Long
Unhealthy Diet	13	19
High Cholesterol	2	0
Blocked Arteries	2	0
Insufficient Oxygen to Heart	3	1
Heart Attack		
Better Education	4	13
Better Health Information	5	5
Healthy Habits	12	2
Better Health		
Budgeting Little Money	2	2
Buying Butane-Laden Fuel	6	14
Fuel-Filter Gasket Corrosion	4	2
Engine Runs Hot	4	0
Carbon Monoxide in Exhaust	5	2
Muffler Rusts		
U.S. Economic Meddling	5	5
Muslim Poverty	6	8
Wounded Pride	2	2
Anti-Americanism	4	4
People Join Al-Qaeda	3	1
Terrorism		

we transformed the intervention distance to a linear [0,1] scale with 0 being an intervention on the immediate cause and 1 being an intervention on the root cause.

For the four real-world stimuli, we performed a 2 × 4 repeated measures ANOVA on intervention distance with condition (short term or long term) as a between-subjects variable and causal system as a within-subjects variable to determine whether participants in one condition intervened significantly farther away from the outcome. Consistent with our hypothesis, the average intervention distance was significantly greater in the long-term condition than in the short-term condition, $F(1, 37) = 14.1$, $p < .01$ (see Table 2).

For the three artificial stimuli, the intervention distance did not differ across the short-term and long-term conditions, $F(1, 36) < 1$, *n.s.* The reason for the absence of an effect for these items remains unclear. One possibility is that the "short-term" and "long-term" manipulations may have had little meaning for these systems. Alternatively, participants may have construed these systems as deterministic, in which case interventions on immediate and root causes should be equally reliable. Interestingly, for these items, 98% of interventions were on immediate or root causes, suggesting that in the absence of content knowledge, people consider these variables the best places to intervene.

Although the real-world stimuli data suggest that people intervene more distally when presented with a long-term intervention goal than with a short-term goal, it is possible that our stimuli were biased such that interventions on the relatively distal variables were better long-term solutions, whereas interventions on the relatively proximal variables were better short-term solutions. We performed an additional experiment to explore this possibility (see supplementary material), which showed that our stimuli were not biased in this way.

Framing a problem in a short-term versus a long-term context influenced where in a causal chain people intervened to manipulate the outcome. Consistent with our intuition that interventions on the root cause provide stable and permanent solutions to long-term problems by addressing the underlying cause and that interventions on the immediate cause provide more direct and fast solutions to short-term problems, participants asked for the best long-term intervention intervened significantly farther back in the causal chain than participants asked for the best short-term intervention. Thus, individuals clearly incorporated their knowledge of the system's causal structure into their decisions regarding where to intervene.

4. General discussion

Structural knowledge can influence people's decisions about how to intervene on causal systems. In Experiment 1, participants were most likely to intervene on root and immediate causes and were less likely to intervene on variables in the middle of the causal chain. However, the relative preference for interventions on root versus immediate causes varied across items, suggesting that nonstructural factors also influence intervention decisions. Experiment 2 identified one specific factor that affects preferences for root-cause versus immediate-cause interventions. Participants who were given a long-term goal intervened more distally than those who were given a short-term goal. When attempting to find a long-term solution, one is more likely to be concerned with choosing an intervention that produces a stable and permanent solution. Indeed, the very term "root cause" suggests an intervention that satisfies these properties.

We also proposed two considerations that support a preference for immediate-cause interventions. First, as an intervention propagates through a causal network, its effects tend to "dissipate" and weaken in strength (e.g., White, 1997). Second, if the causal links between variables are probabilistic, the farther away from the outcome that one intervenes, the more likely it is that one of these links will fail. While the prevalence of immediate-cause interventions provides some support for these hypotheses, future work is needed to tease these hypotheses apart.

Because Bayes net theory has been so widely used to frame research on causal cognition, and because some of our results do bear on the theory, it is useful to briefly relate our work to this theoretical framework even though our two studies were motivated by questions that apply equally well to many other approaches to causal reasoning.

Experiment 1 asked whether there are tendencies to intervene at certain locations in a causal chain, such as root or immediate causes. In general, without additional knowledge about the reliabilities of the links in a chain or the effectiveness of particular interventions, Bayes net theory does not predict a preference to intervene at one location or another. However, if more specific information about the reliability of links in the chain is available, then the theory *may* make predictions. In a case where every link in a causal chain is perfectly reliable (i.e., the causal chain is fully deterministic), the theory makes no prediction regarding a preference to intervene at any location over any other. This is intuitive; if every link is reliable, then the probability of a failure is zero regardless of how many links an intervention must "pass through" to affect an outcome. On the other hand, if the links in a chain are not perfectly reliable, then one should intervene as close as possible to the outcome of interest to maximize the probability of success. The more links the intervention must "pass through" to affect an outcome, the greater the probability of a failure somewhere along the way. Future work could examine whether people would show a stronger preference for immediate-cause (vs. root-cause) interventions when the causal links are probabilistic (vs. deterministic).

Thus, with limited exceptions described below, Bayes net theory can predict either no preference or a preference for proximal/immediate-cause interventions, depending on assumptions about the reliabilities of causal links. Our finding of a general preference for immediate causes over intermediate ones is at least consistent with, if not predicted by Bayes net theory. But without assuming additional causal pathways, hidden causes, or differences in the efficacies of interventions, the probability of the final outcome cannot increase by intervening farther upstream in a causal chain. Thus, Bayes net theory cannot generally predict a preference for more distal interventions over more proximal ones. However, if the intervention targeting a root-cause variable is deemed more efficacious at influencing its target variable than possible interventions on other variables, the root-cause intervention could be preferred. Nonetheless, Bayes net theory cannot predict that a particular intervention would be chosen more frequently when the target variable is listed as the root cause than when it is listed as an intermediate cause. Therefore, our finding in Experiment 1 of a preference for root-cause interventions over intermediate ones would seem to be unaccounted for by unmodified Bayes net theory.

Designing interventions is a fundamental component of causal cognition. Similar to Bayes net theory and other models of human causal cognition, our work suggests that structural knowledge, in addition to other factors such as content and mechanism knowledge, significantly affects how people evaluate the effectiveness of alternative interventions. Our specific findings, because they are not well predicted by existing theories, may help foster new additions to theories of the role of structure in causal reasoning.

Acknowledgments

We thank Kara Gaughen and Eric Smith for running subjects, Greg Laun for assistance with data analysis, and anonymous reviewers for valuable suggestions. The research was supported by an NSF graduate fellowship to Brian J. Edwards and NIH grant R37HD023922 to Frank C. Keil.

Notes

1. Our definition of *intervention* should not be confused with that of Pearl (2000) and Woodward (2003), who defined an intervention as an exogenous action that breaks the causal link between a variable and its usual causes, and sets the target variable to a specific value. We chose our definition because it fits better with an everyday notion of intervention.
2. A small number of participants were excluded from some analyses for providing blank or uncodable responses.

References

Ahn, W. (1998). Why are different features central for natural kinds and artifacts? The role of causal status in determining feature centrality. *Cognition*, *69*(2), 135–178.

Gopnik, A., Glymour, C., Sobel, D. M., Schulz, L. E., Kushnir, T., & Danks, D. (2004). A theory of causal learning in children: Causal maps and Bayes nets. *Psychological Review*, *111*(1), 1–30.

Kim, N. S., & Keil, F. C. (2003). From symptoms to causes: Diversity effects in diagnostic reasoning. *Memory and Cognition*, *31*(1), 155–165.

Kim, N. S., Yopchick, J. E., & de Kwaadsteniet, L. (2008). Causal diversity effects in information seeking. *Psychonomic Bulletin & Review*, *15*(1), 81–88.

Pearl, J. (2000). *Causality*. Cambridge, UK: Cambridge University Press.

Rehder, B. (2010). Causal-based categorization: A review. *Psychology of Learning and Motivation*, *52*, 39–116.

Rehder, B., & Hastie, R. (2001). Causal knowledge and categories: The effects of causal beliefs on categorization, induction, and similarity. *Journal of Experimental Psychology: General*, *130*(3), 323–360.

Steyvers, M., Tenenbaum, J. B., Wagenmakers, E., & Blum, B. (2003). Inferring causal networks from observations and interventions. *Cognitive Science*, *27*(3), 453–489.

White, P. A. (1997). Naive ecology: Causal judgments about a simple ecosystem. *British Journal of Psychology*, *88*(2), 219–233.

White, P. A. (2000). Naive analysis of food web dynamics: A study of causal judgment about complex physical systems. *Cognitive Science*, *24*(4), 605–650.

Woodward, J. (2003). *Making things happen: A theory of causal explanation*. New York: Oxford University Press.

Woodward, J. (2007). Interventionist theories of causation in a psychological perspective. In A. Gopnik & L. Schulz (Eds.), *Causal learning: Psychology, philosophy, and computation* (pp. 19–36). Cambridge, UK: Cambridge University Press.

Supporting Information

Additional Supporting Information may be found in the online version of this article on Wiley Online Library:

Appendix S1. Supplemental Materials.

Appendix A: Sample item for Experiment 1

Economy

The following causal chain explains why the stock market goes down.

Consumers do not feel confident about the economy → Consumers spend less money → Companies lay off workers → The stock market goes down

Consumers' not feeling confident about the economy causes consumers to spend less money.
Consumers' spending less money causes companies to lay off workers.
Companies' laying off workers causes the stock market to go down.

The U.S. government wants the stock market to go up. Which intervention should the U.S. government do to make the stock market go up?

1. Educate consumers about the economy to increase confidence.
2. Reduce taxes on the middle class so consumers will spend more money.
3. Provide economic incentives for companies to not lay off workers.

The U.S. government should do intervention _____ to make the stock market go up.

Appendix B: Sample items for Experiment 2

Two items (one real-world and one artificial) for Experiment 2 are shown here. The experimental manipulations (short term and long term) are shown in bold, with the short-term manipulation shown first and the long-term manipulation shown second in brackets.

Health

A government commission has identified the following causes of better health.

Better education causes people to have access to better health information.
Access to better health information causes people to adopt healthy habits.
Healthy habits cause better health.

Better education → Access to better health information → Healthy habits → Better health

The U.S. government is searching for the best **short-term [long-term]** policy to improve the quality of health in the United States. Where in the causal chain should the U.S. government intervene to improve the quality of health in the United States?

Possible interventions

1. Increase education funding to improve the quality of education.
2. Increase access to better health information.
3. Encourage healthy habits.

The U.S. government should do intervention _____ to improve the quality of health in the United States.

Alien

You are about to answer a question about a small group of space aliens. Within this group, some aliens have special mental connections with other aliens. For example, it could be that Alien A implants thoughts in the mind of Alien B, and Alien C implants thoughts in the minds of Alien D and Alien E. These special mental connections do not change over time; in other words, if Alien A implants thoughts in Alien B, then whatever Alien A is thinking will also be in the mind of Alien B, but it will never affect what other aliens are thinking. Sometimes these aliens play a game in which they intentionally think certain simple thoughts to send and receive them. In particular, they use a very small vocabulary of thoughts that consists of just two words: "DAX" and "BLICK." Each alien can only hold one thought at a time; for example, if an alien is thinking "BLICK," then it cannot also think "DAX" at the same time.

The following information describes a sequence of events that causes Alien 4 to think "DAX."

Alien 1 thinking "DAX" causes Alien 2 to think "DAX."
Alien 2 thinking "DAX" causes Alien 3 to think "DAX."
Alien 3 thinking "DAX" causes Alien 4 to think "DAX."

Alien 1 thinks "DAX" → Alien 2 thinks "DAX" → Alien 3 thinks "DAX" → Alien 4 thinks "DAX"

Bob has a mind zapper that can implant a thought in the mind of Alien 1, Alien 2, or Alien 3. Bob wants Alien 4 to think "DAX" at the end of the game. The game is very short; it lasts five minutes. [The game is very long; it lasts three hours.] Where in the causal chain should Bob intervene to cause Alien 4 to think "DAX?"

Possible interventions

1. Give Alien 1 the "DAX" thought.
2. Give Alien 2 the "DAX" thought.
3. Give Alien 3 the "DAX" thought.

Bob should do intervention _____ to cause Alien 4 to think "DAX."