

When Default Options Explain Away Preferences: A Causal Reasoning Account of Mental State Reasoning from Default Options

Xingyu (Shirley) Liu (xil086@ucsd.edu)

Department of Psychology
University of California, San Diego

Craig McKenzie (cmckenzie@ucsd.edu)

Rady School of Management & Department of Psychology
University of California, San Diego

Adena Schachner (adschachner@ucsd.edu)

Department of Psychology
University of California, San Diego

Abstract

People often infer that those who actively switch away from a default option have stronger preferences than those who passively accept it (termed asymmetric preference inferences). We test whether this classic effect reflects rational causal inference about how defaults provide alternative explanations for others' mental states. This account predicts that asymmetric inferences should occur only when accepting the default provides an alternative explanation for choice (e.g., following a recommendation), and that asymmetry should diminish or disappear when it does not (e.g., a default licensing indulgence in a preferable option). In a pre-registered study (N=120), participants showed this effect: They made asymmetric inferences only when the default provided an alternative explanation for preference, and made symmetrical inferences when it did not. These findings suggest this classic effect reflects rational causal inference, providing a framework for predicting when people make asymmetric preference inferences from defaults.

Keywords: Default Effects; Social inference, Preference; Causal Reasoning

Introduction

Imagine that you and a friend are at a restaurant choosing between fries and salad as a side dish – but by default, all orders come with salad. Your friend gets the salad. How much does your friend like salad compared to fries? What if the default had been fries instead? This story involves a default, or a pre-selected choice option, which is known to impact people's choice behaviors such that they are more likely to choose an option when it is the default than when it is not (Jachimowicz et al., 2019).

Defaults not only impact people's choices, but they can also influence the inference people make about other people's preferences based on their choices. People often infer that, for the same chosen option, a choice-maker who actively switches from a default has a stronger preference for that option than a choice-maker who passively accepts the default (Davidai et al., 2012; Leong et al., 2020; Lin et al., 2018). For example, people rated a person who ordered salad as caring more about healthy eating when the default dish was not the salad (i.e., fries) than when it was salad (Leong et al., 2020). Similarly, people perceived organ donors as preferring organ donation more, or finding organ donation more meaningful, in countries where the default is to not be an organ donor than

where it is to be an organ donor (Davidai et al., 2012; Lin et al., 2018). Thus, for the same chosen option, people often infer that a choice-maker has a stronger preference for that option if they made their choice by switching from rather than accepting a default (asymmetric inferences; Leong et al., 2020).

Why do these asymmetric inferences occur? In the current paper, we propose that these asymmetric inferences may reflect a structured and rational causal inference (Tenenbaum et al., 2006). We use a causal inference framework to predict when asymmetric inferences should occur, and when they should diminish. In particular, while people often infer that one's choices are caused by their preferences, a default can provide an alternative reason for one's choice, other than their preferences. For example, accepting a default can sometimes be due to following the default-setter's implicit recommendation of a certain option, even if this option differs from the choice-maker's true preference (McKenzie et al., 2006). In causal reasoning, people tend to favor a single parsimonious causal explanation (Lombrozo, 2007). A default can thus "explain away" the hypothesis that people's choices are caused by their preferences when they accept a default – but not when they switch from one, resulting in asymmetric preference inferences. This pattern of reasoning, in which one alternative explanation weakens the evidence of others, is a signature of causal inference (Gopnik & Sobel, 2000; Pearl, 2000; Tenenbaum et al., 2006).

This causal account makes novel predictions about when asymmetric inferences should occur versus diminish. More specifically, it predicts that asymmetric inferences should diminish or disappear when accepting the default does not provide an alternative explanation for one's choices — for example, when there is reason to believe that a default option aligns with one's preferences. In this study, we test this prediction, and in doing so test the extent to which reasoning about defaults is dependent on causal reasoning about others' behaviors and underlying mental states.

Default Options as Alternative Explanations

Default options themselves can convey choice-relevant social information, which may in turn provide alternative explanations for one's choice. For example, defaults can convey implicit recommendations of a certain option by the default-setter (e.g. a policy-maker selecting the default policy;

Mckenzie et al, 2006). Therefore, we can attribute the asymmetric inference to a kind of structured causal inference: In the case of accepting a default option, there is often more than one plausible explanation for a person's choice. The choice may reflect their preference for the default over the alternative, or their choice may reflect their wish to follow the default-setter's recommendations. Since people tend to favor parsimony during causal reasoning (Lombrozo, 2007), a causal reasoning account suggests that this alternative explanation should weaken the evidence for other explanations for a person's choice, in this case, the person's preference.

This account is useful because it can help us make predictions about when asymmetric inferences should occur, and when they should diminish or disappear. In particular, accepting the default should send a strong signal of following recommendations when the default option is socially desirable, or if it conveys normative expectations – a form of benign paternalism. For example, many real-world applications of defaults aim to promote healthy or responsible behavior like saving for retirement (Madrian & Shea, 2001) or green energy use (Ebeling & Lotz, 2015).

Previous research on default effects has only examined cases like these, in which accepting the default signals following social values or norms. For example, in Leong et al. (2020), participants read vignettes in which choice-makers chose salad by either accepting or switching from a default; here accepting the default may signal conforming to an implied norm of healthy eating. In studies on organ donation (Davidai et al., 2012; Lin et al., 2018), accepting the default of organ donation may signal following an implied norm of altruism (see Davidai et al., 2012, Exp. 3), while accepting a default of not donating organs may signal conformity to norms of conservatism or religious practice.

These cases likely serve as strong alternative explanations for choices. Here, the choice-maker has a plausible and strong reason to follow the default-setter's recommendations, sometimes even against their own preferences— to conform to the implied social value or norm. Thus, accepting the default is more likely explained by following the default-setter's recommendations than following one's own preferences, leading to asymmetric inferences when accepting versus switching from a default. The causal reasoning account predicts that asymmetric inferences should be found in these cases from previous research.

However, the causal reasoning account makes a distinct and novel prediction for cases in which a default option does not provide a strong alternative explanation for a person's choice—cases that past studies have not tested. In these contexts, accepting the default should remain better explained by one's own preferences than a desire to follow the default-setter's recommendations, resulting in diminished asymmetric inferences, or even completely symmetric inferences, based on accepting or switching from a default.

This type of diminished asymmetry may occur in multiple contexts. For example, if a default is seen as a sign of the status quo (e.g., the continuation of enrollment in a particular

plan or subscription), it sends a weaker signal of recommendation by the default-setter (Jachimowicz et al., 2019), and thus may suggest a weaker motivation to accept the default. If a choice-maker does not trust or value the default-setter, rendering their recommendations irrelevant (e.g., a default dish set by a random employee rather than a nutritionist), then this should reduce asymmetric inferences. Most relevant to the current investigation, the causal reasoning account predicts diminished asymmetry if the default option aligns with a choice-maker's existing preference (when this existing preference is expected or known). For instance, a default option can signal the licensing of indulgence — letting choice-makers choose what they want, flaunting some typical normative restraint. In this case, choosing the default remains consistent with the idea that choice indicates preference, rather than providing an alternative explanation for the choice. Consider the case of a restaurant setting the default side-dish as fries, rather than salad, to convey a license to indulge. Here, accepting the default fries may be seen as following one's preference, which happens to coincide with the default-setter's recommendations. The causal reasoning account thus predicts that choosing the default option will still be seen as a sign of strong preference, in this case.

The Current Study

In the current study, we test this key prediction of the causal reasoning account, asking if people make asymmetric inferences only when the default provides a plausible alternative explanation for accepting the default. If people are using causal reasoning, then this asymmetry should diminish in strength or even disappear when the default does not provide a plausible alternative explanation for preference – as in the case when it appears to license indulgence in a generally desirable option. The causal reasoning account predicts that people will infer a similarly-strong preference for choice-makers who switch from versus accept the default, in this case.

To test this, we asked participants to make inferences about others' preferences from their choices, in the context of one of two different default options. We tested adult participants; however, stimuli were designed to also be comprehensible for children, to allow for future developmental comparison. In particular, the situation presented was that of two child characters choosing what to have for snack: Either broccoli, or chocolate. Crucially, one of the snacks had been set as the default option by their parents. Participants observed one character accept the default as their choice (Default-Acceptor), and the other switch away from it (Switcher). Participants judged how strongly each character preferred their chosen snack.

Here, while a default option of broccoli conveys the social norm of healthy eating (benign paternalism), a default option of chocolate conveys a license to indulge in the less healthy but likely preferable option (licensed indulgence). We chose chocolate and broccoli based on the expectation that adults would expect that most children find one option (chocolate)

preferable to the other (broccoli). As a result, we are able to compare a situation where accepting the default provides a plausible alternative explanation, other than the child's own preference (the broccoli default) versus a situation where accepting the default does not provide an alternative explanation (the chocolate default).

The causal reasoning account makes predictions about the extent to which people will infer *different* (asymmetric) preference strengths from choosing the same option (e.g. Broccoli), when this involves accepting versus switching from the default. Across the two Default conditions, broccoli was chosen by two characters—one who has accepted it as the default option, and one who has switched to it from a default option of chocolate. In this case, both causal reasoning and prior literature on asymmetric inferences make the same prediction: Participants should infer that the Default-Acceptor has a weaker preference for broccoli over chocolate than the Switcher does. By the causal reasoning account, this is because when broccoli is the chosen snack, accepting the broccoli default has a clear alternative explanation (following the healthy recommendation). In contrast, switching from the chocolate default can only be explained by strong preference (if parents have licensed indulgence in chocolate, they have little reason to choose broccoli unless they actually prefer it to chocolate). This should lead to strongly asymmetric inferences from the choice of broccoli, under the two defaults.

The causal reasoning account predicts that this asymmetry should diminish, or even disappear, when comparing the two characters who chose chocolate. This is because when the default is chocolate, this default appears to license indulgence in a generally preferable option. Thus, choosing the default of chocolate does not provide an alternative explanation for preference, but instead is consistent with preference. When the default is broccoli, these inferences should also occur, as usual: Switching away from a default of broccoli to choose chocolate provides evidence of preference. Thus, if people use causal reasoning, this should lead to minimally asymmetric inferences from the choice of chocolate – people should infer a similarly-strong preference for chocolate under both defaults.

The experiments' design, sample size, and analysis plan were preregistered on OSF, available here: https://osf.io/prb84?view_only=3afb3cd59c74f60bbe8fca077e7c8a2.

Method

Participants

As preregistered, $N=120$ adults from the U.S. participated, recruited from www.prolific.com (Mean age = 37.29 years, Range 18-84; 44 men, 74 women, 2 non-binary gender). This sample size was chosen based on bootstrapping analyses from a pilot sample (in which we found similar results, available here: https://osf.io/qx4th/?view_only=069a1bbf35f149388661c7b7d1cca08f) which suggested that 120 participants would provide sufficient power to detect the predicted differences in

effect size when the characters chose broccoli versus chocolate. 16 additional participants were also tested, but excluded and replaced due to the pre-registered exclusion criterion of not passing one or both attention checks.

Design

To allow for future comparison of adults' reasoning to children's, all stimuli and measures were designed to be comprehensible to children, as well as adults. The experiment used a 2 (Default: Broccoli vs. Chocolate) by 2 (Choice: Broccoli vs. Chocolate) within-subject design, where each participant participated in four trials, and each trial was a unique condition. The four trials were blocked into pairs, based on Default condition (Broccoli Default; Chocolate Default); which Default condition was presented first was counterbalanced across participants. Each participant saw one pair of trials with girl characters, and one with boy characters. Whether the characters in each Default condition were both boys, or both girls, was also counterbalanced across participants.

Procedure and Stimuli

In the task, adult participants read and viewed illustrated vignettes about child characters at school (which were designed to be comprehensible to both adults and future child participants). Illustrations were created using Vyond (www.vyond.com) and Microsoft PowerPoint.

For each Default condition, participants saw an image of two child characters standing near a table with two bowls, one marked with a star (either Broccoli or Chocolate, see Fig. 1a). Participants read that the star showed which type of snack had been selected as the default choice by parents (*"Before snack time, the parents picked out which snack they think is a good choice for the kids. All the parents of the kids picked [broccoli or chocolate] for them – that's why there's a star on that bowl"*). It was also reinforced that the characters had the ability to freely choose, though they were aware of the default (*"Every kid at school gets to choose either a chocolate or a broccoli. The kids get to choose which snack they want. But they can see which one has the star!"*). In the Broccoli Default Condition, the star was on the broccoli bowl (see Fig 1). In the Chocolate Default Condition, the star was instead on the chocolate bowl.

Within each Default condition, on each of two subsequent trials, participants then saw one of the two characters choose a snack, and rated that characters' relative preference for Chocolate versus Broccoli (*"which snack does [the character] like more?"*, binary forced-choice). They were then asked about the strength of that characters' preference (*"do they like [that snack] a lot more, or a little more?"*, binary forced-choice). These two binary answers were combined to create a 4-point scale of participants' inferred preference strength:

"Like broccoli a lot more (1)", "Like broccoli a little more (2)", "Like chocolate a little more (3)", "Like chocolate a lot more (4)". This was our main dependent measure.




Figure 1: Method and Stimuli. The snack with a star was described as the default option (either Broccoli or Chocolate). On each trial, one character chose the default option (pictured); the other chose the non-default option.

Images are copyrighted by and used by permission of VYOND, trademark of GoAnimate, Inc.

More specifically, within each Default condition, participants first saw one character choose the default snack (Default-Acceptor trial, Fig 1), and rated that character's preference. Participants then saw the other character choose the alternative (Switcher trial), and rated that character's preference. Participants also completed one memory check within each Default condition: Before the second trial, participants were asked: "Can you remind me why there is a star on the bowl?" (free-response). Two researchers independently coded these responses for accuracy (e.g. "the parents picked it"; "it was what the parents suggested"). As pre-registered, any participant who failed one or both memory checks was excluded and replaced.

After completing two trials, participants then moved on to the other Default condition, completing two other trials that were identical except for having novel characters, and the opposite snack marked as the Default. After the last trial, participants were asked to explain why they gave the response that they did on the final trial (free response). As preregistered, these free-responses were coded to detect and exclude bots and non-fluent English speakers; this did not result in any exclusions in the current sample.

Results

To test our prediction that the nature of the Default will impact whether people make asymmetric preference inferences, we performed an ordinal logistic regression with Inferred Preference as the dependent variable, and with the predictors of Default, Choice, and their interaction (all fixed effects), and participant (as a random effect). As predicted, we found a significant interaction between Default and Choice ($b = -10.98, p < .001$; nested model comparison with vs. without the interaction term: $\chi^2(1) = 490.74, p < .001$).

In particular, in line with our predictions, participants made more asymmetric inferences when the broccoli was the chosen snack, than when chocolate was the chosen snack. Comparing the two characters who chose broccoli, participants judged that the Default-Acceptor liked broccoli only a little more than chocolate ($M = 2.05$). In contrast, participants judged that the Switcher (who had switched

away from a Default of chocolate) ($M = 1.43$) liked broccoli a lot more than chocolate. The inferred preference for broccoli was significantly stronger for the Switcher compared to the Default-Acceptor ($V = 2375.5, p < .001$; Wilcoxon's signed rank test) – in line with prior findings of asymmetric preference inferences in situations with defaults.

However, as predicted, this asymmetry was diminished when the chosen snack was chocolate. Comparing the two characters who chose chocolate, participants rated that both the Default-Acceptor ($M = 3.87$) and the Switcher (who had switched away from a default of broccoli) liked chocolate a lot more than broccoli ($M = 3.84$; see Fig. 2). The inferred preference strength for chocolate was not significantly different for the Switcher and the Default-Acceptor ($V = 72, p = .528$; Wilcoxon's signed rank test).

To compare the difference in the extent of asymmetric inferences for different chosen snacks, we performed an ordinal logistic regression with inferred preference for the snack the character chose as the dependent variable, and the predictors of Default, Choice, and their interaction (all fixed effects), and participant (as a random effect). We found a significant interaction between Default condition and Choice ($b = -4.18, p < .001$), indicating that the asymmetric inferred preferences diminished significantly when chocolate was the chosen snack than when broccoli was the chosen snack.

To provide a further comparison, we compared the effect size of the difference between the Default-Acceptor vs the Switcher, for each of the choice conditions. This effect size was larger when broccoli was the chosen snack ($r = .520, 95\% CI [.39, .64]$) than when chocolate was the chosen snack ($r = .045, 95\% CI [.00, .22]$). Using bootstrapping, we simulated the dataset 1000 times to generate two distributions of effect size—effect sizes of the difference in inferred preference when the chosen snack was broccoli vs. chocolate), respectively. We then compared the 95% Confidence Interval (CI) of the two distributions. There was no overlap in the 95% CIs, indicating that the two effect sizes are reliably different.

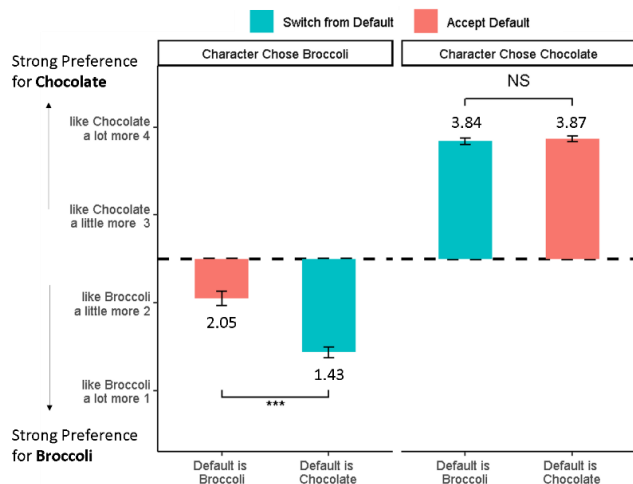


Figure 2: Results. Participants' average inferred preference in each condition. As predicted, participants made asymmetric inferences when broccoli was chosen, and this asymmetry diminished when chocolate was chosen.

Discussion

Overall, our findings provide strong evidence that asymmetric preference inferences from default options reflect a process of rational causal inference. In particular, we find that people make asymmetric inferences about others' preferences only when accepting the default provides a plausible alternative explanation for the choice. When participants observed characters choosing between broccoli and chocolate, they made strongly asymmetric inferences about preference strength when broccoli was chosen, but not when chocolate was chosen.

These results support the causal reasoning account in two key ways. First, we replicated the classic asymmetric inference effect (Davidai et al., 2012; Leong et al., 2020; Lin et al., 2018) in the case where accepting the default provided a clear alternative explanation. When broccoli was the default, participants inferred that someone who actively switched from an alternative to choose it had a stronger preference than someone who passively accepted it. This aligns with previous work showing that defaults can serve as implicit recommendations (McKenzie et al., 2006), providing an alternative explanation for accepting the default that weakens evidence for one's own preferences.

Second, and more critically, we found that this asymmetry disappeared when chocolate was the chosen snack. This novel finding supports our prediction that asymmetric inferences would not occur when accepting the default fails to provide an alternative explanation for the choice. When chocolate was the default, accepting it was consistent with acting the choice-maker's preferences rather than following a recommendation, since the default appeared to license indulgence in a generally desirable option. This finding suggests that people engage in structured causal reasoning about how defaults relate to preferences and choices: Accepting a default option only "explains away" preference as the explanation for one's choice if the default provides an alternative explanation for the choice (as in Pearl, 2000; Tenenbaum et al., 2006). These results provide a theoretical framework for understanding when default options will impact social inferences, and when they will not. In addition to the case examined here (defaults that license indulgence), this framework generates novel predictions about other contexts where asymmetric inferences might diminish. In particular, this should depend on the extent to which a default option provides a strong, plausible reason for a person to follow the default-setter's recommendation. Defaults are less likely to provide such a reason when defaults are randomly-selected rather than intentionally chosen, when the default is seen as a sign of the status quo (e.g., the continuation of enrollment in a particular subscription; Jachimowicz et al., 2019), or when the choice-maker actively mistrusts the default-setter (Tannenbaum et al., 2017). Future work may test whether in these cases, accepting versus switching from a default results in similarly-strong preference inferences.

The causal reasoning account also raises intriguing questions about the development of social inference from defaults. Prior work has shown that by ages 5-6, young

children can reason about how physical and epistemic constraints influence others' choices when inferring preferences (Pesowski et al., 2016). An open question is whether children can similarly reason about how social constraints, like defaults, provide alternative explanations for others' choices. Default options have been shown to impact young children's behavior (Zhao et al., 2023). Therefore, we are currently examining whether children also consider default options when inferring preferences, and whether children show adult-like patterns of asymmetric inferences. Interestingly, a sample of children aged 7-8 years showed the opposite pattern of responses than adults: they made strongly asymmetric inferences when chocolate was the chosen snack, but not when the broccoli was the chosen snack (reproducible code and results available here: https://osf.io/qrmfv/?view_only=c134b1f25eb4d00a5450294146a0a04). This raises the possibility that children may not see parents' recommendations as a strong alternative reason for potentially undesirable choices (e.g., broccoli). In ongoing work, we seek to explain why and how children's reasoning may be different than adults'.

One limitation in our current study is that the default option was operationalized more as a recommendation than a default option *per se*. A traditional default option is an option that the choice-maker automatically receives unless they opt-out, typically examined in a context where opting-out is not effortful or costly. People often infer that the default option reflects an implicit recommendation (McKenzie et al., 2006). In our study, the default option conveyed the parents' recommendation, but the character still needed to actively choose an option in the form of a free choice. We originally chose this design to carefully equate cost or effort for each of the choices; however, as a result, the current task more directly reveals causal reasoning about others' preference in the presence of *recommendations*. Since defaults often work because they convey implicit recommendations, we expect the causal reasoning process seen here to generalize to more traditional instantiations of defaults (where action is only required to opt-out). In ongoing work, we are currently testing this, using a more traditional operationalization of the default where a character automatically receives the default option unless they chose an alternative (e.g., the default lunch option at a cafeteria).

Another potential limitation comes from possible ceiling effects. In our dataset, the inferred preference for chocolate when it was chosen was extremely high, for both accepting and switching from the default. On the one hand, the similarly high inferred preferences in these two conditions supports our prediction that asymmetry would be reduced, and that participants see the chocolate default as aligning with the character's existing preference—a licensed indulgence. However, this raises a concern that ceiling effects may mask asymmetry in the inferred preference for chocolate. We reasoned that a more continuous dependent measure with greater range (rather than a 4-point scale) would give us more power to detect asymmetry in this condition, should it exist. We therefore ran a pre-registered study (N=120) with the

exact same design, but with a continuous measure of inferred preference (a scale from -100 to 100, where -100 is “like broccoli a lot more than chocolate” and 100 is “like chocolate a lot more than broccoli”). We replicated our current findings with this measure, even when results remained below ceiling in all conditions. Specifically, participants made strongly asymmetric inferences when broccoli was the chosen snack ($M=-25.82$ when default was broccoli, vs. -70.22 when default was chocolate); this asymmetry diminished significantly when chocolate was the chosen snack, in spite of remaining below ceiling ($M=85.54$ vs. 81.40 ; reproducible code and results available here: https://osf.io/5jt32/?view_only=a7e68505cb034d8f81bc6404b18304bc).

Finally, our focus on food choices and parental defaults opens up additional questions about how different domains and institutional contexts might influence preference inferences. Many real-world defaults are set by institutions rather than individuals - for instance, governments setting organ-donation policies, energy companies setting energy-use defaults, or workplaces setting retirement-savings defaults (Madrian & Shea, 2001; Ebeling & Lotz, 2015). The situation tested (snacks at school) allows for a straightforward comparison of defaults set by larger institutions (policy-makers setting default food choices for school menus), versus defaults set by directly-known individual social partners (such as parents). This comparison may be particularly informative for children, who slowly develop an understanding of the nature of institutions over middle childhood (Noyes et al., 2020; Noyes et al., 2022).

In conclusion, this work suggests that a classic effect in judgment and decision-making—asymmetric preference inferences from defaults—reflects a process of rational causal inference. People appear to reason about how defaults provide explanations for others’ choices, such that choices from defaults leads to asymmetric inferences only when the default provides a plausible alternative explanation. This framework provides new insights into the cognitive processes underlying social inference from choice architecture, allowing for novel predictions about when and where default options will result in asymmetric preference inferences across a large range of contexts.

Acknowledgements

We are very thankful to Lim Leong for his work that inspired this project, to Jennifer Lozano and Alan Lopez for their amazing job collecting data with the children’s sample, and to the UCSD MaD LAB and the Decision-Making lab for their generous help and support.

References

Davidai, S., Gilovich, T., & Ross, L. D. (2012). The meaning of default options for potential organ donors. *Proceedings of the National Academy of Sciences*, *109*(38), 15201–15205. <https://doi.org/10.1073/pnas.1211695109>

Ebeling, F., & Lotz, S. (2015). Domestic uptake of green energy promoted by opt-out tariffs. *Nature Climate*

Change, 5(9), Article 9. <https://doi.org/10.1038/nclimate2681>

Gopnik, A., & Sobel, D. M. (2000). Detectingblickets: How young children use information about novel causal powers in categorization and induction. *Child Development*, *71*(5), 1205–1222. <https://doi.org/10.1111/1467-8624.00224>

Jachimowicz, J. M., Duncan, S., Weber, E. U., & Johnson, E. J. (2019). When and why defaults influence decisions: A meta-analysis of default effects. *Behavioural Public Policy*, *3*(2), 159–186. <https://doi.org/10.1017/bpp.2018.43>

Leong, L. M., Yin, Y., & McKenzie, C. R. M. (2020). Exploiting asymmetric signals from choices through default selection. *Psychonomic Bulletin & Review*, *27*(1), 162–169. <https://doi.org/10.3758/s13423-019-01699-2>

Lin, Y., Osman, M., Harris, A. J. L., & Read, D. (2018). Underlying wishes and nudged choices. *Journal of Experimental Psychology: Applied*, *24*(4), 459–475. <https://doi.org/10.1037/xap0000183>

Lombrozo, T. (2007). Simplicity and probability in causal explanation. *Cognitive Psychology*, *55*(3), 232–257. <https://doi.org/10.1016/j.cogpsych.2006.09.006>

Madrian, B. C., & Shea, D. F. (2001). The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior. *The Quarterly Journal of Economics*, *116*(4), 1149–1187.

McKenzie, C. R. M., Liersch, M. J., & Finkelstein, S. R. (2006). Recommendations Implicit in Policy Defaults. *Psychological Science*, *17*(5), 414–420. <https://doi.org/10.1111/j.1467-9280.2006.01721.x>

Pearl, J. (2000). Causality: Models, reasoning, and inference (pp. xvi, 384). *Cambridge University Press*.

Noyes, A., Dunham, Y., & Keil, F. C. (2022). Subjectivity and social constitution: Contrasting conceptions of institutions, artifacts, and animals. *Developmental Psychology*, *59* (2), 377–389. DOI: 10.1037/dev0001499

Noyes, A., Keil, F. C., Dunham, Y. (2020). Institutional actors: Children’s emerging beliefs about the causal structure of social roles. *Developmental Psychology*. *56*(1), 70–80. <https://doi.org/10.1037/dev0000847>

Pesowski, M. L., Denison, S., & Friedman, O. (2016). Young children infer preferences from a single action, but not if it is constrained. *Cognition*, *155*, 168–175. <https://doi.org/10.1016/j.cognition.2016.07.004>

Tannenbaum, D., Fox, C. and Rogers, T. (2017), On the misplaced politics of behavioral policy interventions, *Nature Human Behavior*, *1*(7): 0130. <https://doi.org/10.1038/s41562-017-0130>

Tenenbaum, J. B., Griffiths, T. L., & Kemp, C. (2006). Theory-based Bayesian models of inductive learning and reasoning. *Trends in Cognitive Sciences*, *10*(7), 309–318. <https://doi.org/10.1016/j.tics.2006.05.009>

Zhao, L., Mao, H., Zheng, J., Fu, G., Compton, B. J., Heyman, G. D., & Lee, K. (2023). Default settings affect children’s decisions about whether to be honest. *Cognition*, *235*, 105390.