Beliefs that go together change together: A longitudinal study of caregiver beliefs surrounding pediatric COVID vaccination

Kara Weisman (kara@foundry10.org)[1] & Dominic J. Gibson (dominic@foundry10.org)[1]
Conceptual Development Lab, foundry10

Abstract

Cognitive science offers powerful tools for addressing pressing public health needs. In the current line of work, 1700 US parents reported beliefs about a range of topics related to pediatric COVID vaccination. We deployed Bayesian network structure learning to develop a cognitive model of the relationships among these beliefs and their combined influence on vaccine endorsement. Nine months later, we re-measured these beliefs in a subset of returning participants (*n*=884) to examine how beliefs changed together over a particularly tumultuous period of time. In a simple comparison of observed co-changes to Time 1 correlations, as well as a comparison of co-changes to model-based simulations, the relationships among beliefs at Time 1 accurately predicted co-change over time (R^2 .84-.92 across models). This is consistent with the possibility that correlations across individuals at a single timepoint reflect an underlying belief network that also governs how beliefs change over time.

Keywords: intuitive theories; belief revision; cognitive modeling; Bayesian networks; vaccine attitudes; parenting; longitudinal analysis

How do beliefs change in light of new information? Can we understand and predict belief revision by understanding these changes in the context of wider networks of auxiliary beliefs, or more elaborate mental models? These are fundamental questions in cognitive science with direct implications for understanding and addressing pressing real-world challenges. Caregivers' beliefs about pediatric COVID vaccination present one especially timely case study with practical consequences for public health.

In this paper, we present a study of changes in US parents' beliefs related to pediatric COVID vaccination between December 2024 and September 2025. Toward the beginning of this nine-month period, a new presidential administration took office after an historic US election. The administration established new leadership at the Department of Health and Human Services, including at the CDC and FDA. These events and the ensuing changes in policy and recommendations regarding vaccination and other healthcare decisions were widely discussed in the media, but coverage was uneven across news sources and varied fairly dramatically in viewpoints (Basch et al., 2025). Meanwhile, several months into the study period the US recorded its highest number of new measles cases since the disease was declared eliminated in 2000 (Anderer, 2025), an event which was widely discussed in the media in the context of the decline in pediatric vaccination rates since the beginning of the COVID pandemic (Wang et al., 2025). At the same time, new scientific studies were published on the risks of COVID infection (e.g., O'Mahoney et al., 2025) and the benefits of COVID vaccines (e.g., Du et al., 2025; Grippin et al., 2025), while the availability of updated vaccines was the topic of substantial uncertainty (e.g., Young & Bartels, 2025).

The diversity of these events themselves, as well as the wide range of coverage in both traditional and social media, set the conditions for beliefs to shift in different directions across individuals, and perhaps even within one individual. At the same time, the increasing political polarization of these topics—not least among them, pediatric vaccination in general and the COVID vaccine in particular—meant that even individuals exposed to the same news story about the same event might interpret that information differently (Baxter-King et al., 2025). This created a unique opportunity to investigate the extent to which tools of cognitive science can help make sense of relationships among a complicated set of beliefs and predict how beliefs are likely to change over time in light of new information.

The current study is part of a larger project on caregiver beliefs about pediatric COVID vaccination. In the first study in this line of work, we adapted the approach pioneered by Powell, Weisman, & Markman (2023) to develop a cognitive model of the relationships among a variety of related beliefs. We used individual differences in these beliefs to infer via Bayesian network structure learning how they might work together to influence parents' vaccination choices for their children, and to simulate how interventions on these beliefs might spark belief revision (Weisman & Gibson, 2025).

In this paper, we describe the second study in this line of work. We administered the same belief measures to a subset of these same participants after a delay of approximately nine months. Our primary goal was to assess how closely the correlations between beliefs at the first timepoint would approximate the co-changes between beliefs over time. This is one way to assess whether cross-sectional explorations of individual differences provide reasonable approximations of the cognitive models that drive individuals' reasoning and belief revision—a fundamental tenet of this general approach to modeling intuitive theories and crafting educational interventions to induce conceptual change (Weisman & Markman, 2017; Powell et al., 2023).

Method

The sampling plan, inclusion/exclusion criteria, procedure, data processing, and modeling approach for Time 1 were pre-registered at osf.io/vgysz.

In G. Boleda, I. Dautriche, M. Maldonado, C. Saldana and T. Torrent (Eds.), *Proceedings of the 48th Annual Conference of the Cognitive Science Society*

Participants

As reported in Weisman & Gibson (2025), at Time 1 (December 2024), 1700 US adults were recruited through the online survey platform Prolific using the following inclusion criteria: participants must be at least 18y old, located in the US, and fluent in English; participants must have ≥1 child older than 6 months and younger than 18y at the time of the study; and participants must not have participated in any related pilot studies. We used survey administration options to achieve a balance between self-identified Democrats (36% of the final sample), Republicans (34%), or participants who identified with neither political party (29%). We excluded participants who completed the study in <5 min; failed ≥2 of the 3 attention checks embedded in the survey; indicated that they did not provide thoughtful answers to the questions in this survey; were flagged as engaging in fraudulent behavior by Qualtrics; or indicated that they did not have any children in the target age range. For modeling purposes, we further excluded any participant who declined to answer all of the questions within ≥1 belief scale, leaving us with a Time 1 sample of 1524 participants.

Approximately nine months later, we invited all Time 1 participants to a second study. We took care to re-recruit as many participants as possible by re-sending invitations roughly two weeks after the initial study post; messaging people who "timed out of" or "returned" the survey to give them more time to complete it; and leaving the survey open for four weeks (September 2-30, 2025). 993 returning participants participated at Time 2; after implementing our exclusion criteria, this left us with *n*=884 US parents for our longitudinal analyses (58% of our final Time 1 sample).

Participants ranged in age from 21-75y (median: 40y), and reported having between 1-6 children in the target age range (median number of children: 2; median child age: 9y; total number of children ≤19y: 1697). The demographics of returning participants were similar to the initial sample: The sample was predominantly mothers (66%), predominantly White (67%; 14% of participants identified as Black or African American, 9% as ≥1 race/ethnicity, and ≤5% as any other race/ethnicity), and generally highly educated (45% of participants reported having obtained a college degree or more advanced education, and an additional 32% reported having attended some college).

Procedure

Participants completed an online study titled "Parents' & Caregivers' Health Beliefs Follow-up Survey."

After providing some basic demographic information (e.g., number and ages of children), participants completed a series of 12 belief scales indexing a variety of beliefs relevant to pediatric COVID vaccination; this procedure was identical to Time 1. These scales assessed two abstract worldviews (*holistic balance*, McFadden et al., 2010; and *naturalism*, based on Powell et al., 2023); three general beliefs related to pediatric vaccination (general *vaccine hesitancy,* Shapiro et al., 2018; as well as *medical skepticism* and *parental expertise*, both based on Powell et al., 2023); and seven beliefs specific to COVID (*COVID conspiracy thinking*, *COVID rarity*, *COVID severity for children*, *COVID long-term risk for children*, *COVID vaccine effectiveness*, *COVID vaccine danger*, and *COVID vaccine community benefits*). Observed internal consistency was uniformly high at both timepoints (Time 1: all observed α>0.73, ω_T>.83; Time 2: all observed α>0.79, ω_T>.83). The high intercorrelations among belief scales (see Fig. 3) raise questions about discriminant validity that warrant further psychometric investigation; we leave this to future work.

The longitudinal data collection process provided an opportunity to further examine the reliability of these belief scales (many of which were newly developed for Time 1). Across scales, test-retest reliability was good: The median intraclass correlation (two-way, mixed effects, consistency, single measure) was .79, and only two scales fell below the traditional threshold of .70 (*COVID long-term risk for children*: ICC3=.61; *holistic balance*: ICC3=.63). Together with the internal consistency of these scales and their relationship to various proxies of vaccine decisions at Time 1 (Weisman & Gibson, 2025), this lends us confidence in treating these measures as psychometrically valid indices of credence in the beliefs of interest, and stable indicators of the relative differences in these beliefs across individuals.

As at Time 1, our primary outcome of interest was a *COVID vaccine endorsement* scale, which began with the following premise: "Imagine that a close friend or family member is a first-time parent and wants your advice to help them decide whether to get their child vaccinated against COVID. What would you advise them to do?" Participants rated their agreement with a variety of statements for and against COVID vaccination for children. We consider this to be an index of participants' attitudes toward pediatric COVID vaccines for children in general, a proxy of parents' decisions about COVID vaccines for their own families, and some indication of their influence on the vaccine decisions of other families in their social networks.

The 12 belief scales were presented in a random order for each participant (along with an older measure of COVID conspiracy thinking, which was not included in these analyses as per our preregistration). Directly after this, participants completed the *COVID vaccine endorsement* scale. Each scale was presented on a separate page of the survey, with items presented in a random order. For all items in all scales, participants responded on a 7-point Likert-type agreement scale from "Strongly disagree" to "Strongly agree" (with the additional option "I prefer not to say," coded as missing data).

At the end of the survey, participants reported on a variety of experiences, attitudes, and behaviors related to COVID vaccines, measles vaccines, and current events, as well as a variety of demographic measures.

Data preparation

At each timepoint, for each scale, we averaged each participant's responses (after reverse-coding and dropping

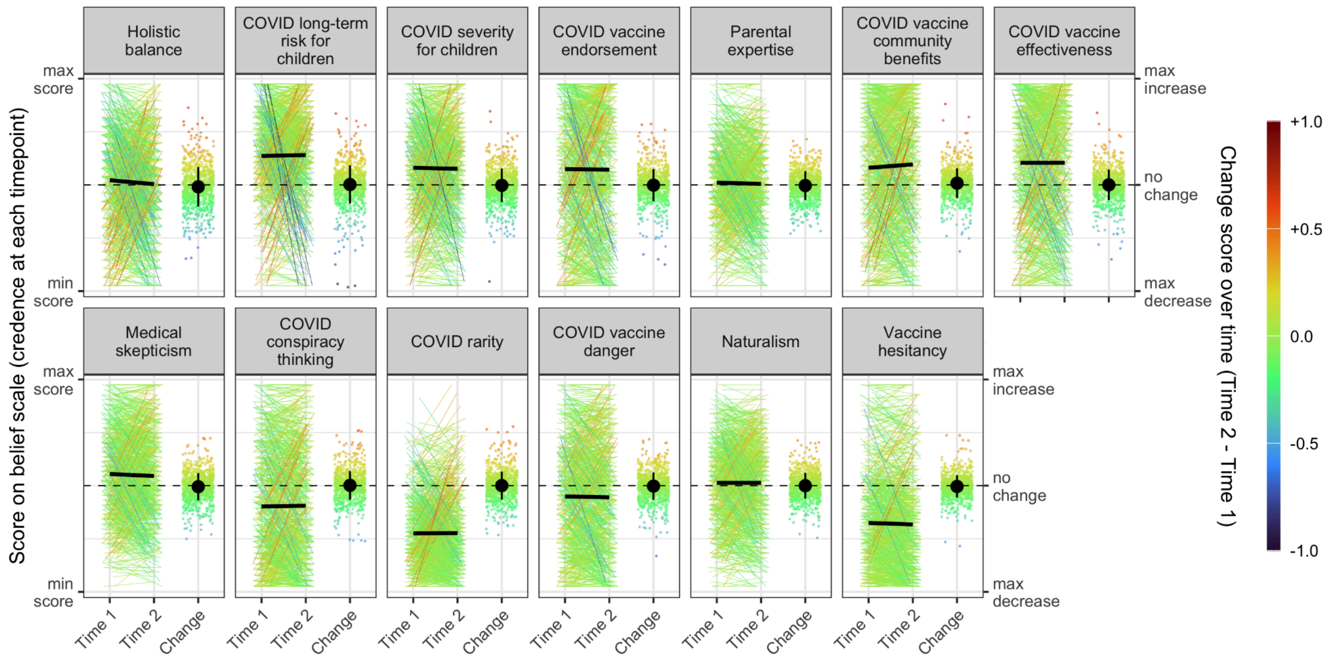

Figure 1: Belief change over time, by belief scale (ordered from most to least absolute change). The left side of each panel depicts scores at Time 1 (December 2024) and Time 2 (September 2025); the right side depicts change scores between timepoints. Small colorful lines and points correspond to individual participants, with increases in red and decreases in blue; larger black lines and points correspond to averages across the sample. Error bars for change scores are ±1 SD.

items with missing data). We rescaled these scores to range from 0-1, representing the participant's credence in each belief. As a measure of within-subjects change in credence over time, for each participant we calculated a change score for each scale by subtracting their Time 2 score from their Time 1 score (theoretical range: –1 to +1).

Results

Changes in beliefs over time

In contrast to previous observations of group-level changes in vaccine attitudes following specific real-world events (e.g., the 2018 measles outbreaks; Powell et al., 2023), there were no specific scales where we observed substantial group-level changes in belief scores; see Fig. 1. The largest and most reliable group-level changes were in beliefs about *holistic balance* (Δ=-0.018; t=-2.90, p=.004) and *COVID vaccine community benefits* (Δ=0.016, t=3.43, p<.001). Relative to a theoretical range of –1 to +1, these are very small shifts, corresponding to approximately one eighth of the distance between adjacent options on our 7-point response scale. (All other $|\Delta|$<.011, $|t|$<2.25, $|p|$>.024; corrections for multiple comparisons render these differences non-significant. T- and p-values are from a series of mixed effects models regressing scores onto timepoint with random intercepts by participant.)

However, this apparent group-level stability is deceptive: At an individual level, there were in fact substantial changes in beliefs during these nine months, but these changes ran in different directions for different people; see individual trajectories in Fig. 1 and absolute change scores in Fig. 2. Ignoring direction, average absolute change scores ranged from Δ_{abs}=.074 for *vaccine hesitancy* (~0.5 ticks up or down on our response scale) to Δ_{abs}=.141 for *holistic balance* (~1 tick up or down). In the aggregate, absolute changes were reliably non-zero for all scales (all t>28.10, all p<.001; t- and p-values are from a series of linear models regressing absolute change scores onto timepoint).

In documenting these differences between timepoints, we might well be charting small fluctuations in response tendencies over time. In addition to regression to the mean, this includes the possibility that we happened to catch people in different moments in their daily lives when they

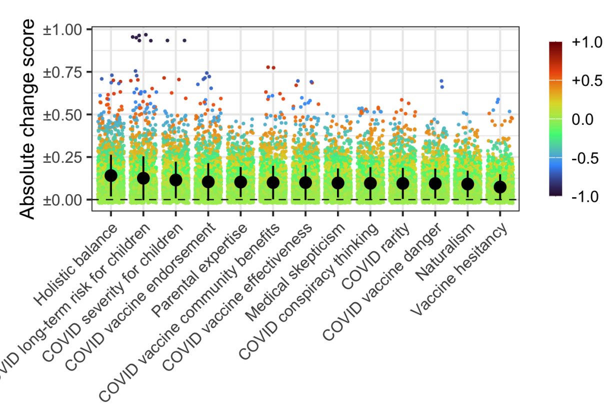

Figure 2: Absolute change scores. Error bars are ±1 SD.

were temporarily in a more trusting vs. skeptical frame of mind (or calm vs. anxious, community-minded vs. self-protective, and so on; Zaller & Feldman, 1992). At the same time—especially given that this was a particularly volatile period in the US—the diversity of belief change observed between these two timepoints might also reflect more robust and meaningful differences in individuals' awareness of different historical events, their exposure to different messages about these events, the social meaning and personal relevance of these events, and so on.

In any case—whether the changes in belief scores reflect relatively low-level fluctuations or higher-level belief revision—in the current analysis we are primarily interested in the extent to which these diverse examples of belief change were *coherent*: When an individual's reported credence in one belief changed, did other beliefs in the network change in internally consistent ways? How well did these co-changes align with the observed correlations in beliefs across individuals at the first timepoint?

Exploring the coherence of belief change under these conditions—with many plausible sources of influence on multiple beliefs in the network, and beliefs changing in different directions across individuals—requires a different approach than other natural experiments. In the longitudinal study featured in Powell et al. (2023), for example, the analysis plan hinged on the assumption that living through a time of major measles outbreaks would generally improve people's attitudes toward measles vaccines (albeit to different degrees), allowing the authors to model "time" as a direct intervention on vaccine attitudes. In the current study, there are too many countervailing forces exerting too many disparate kinds of influence across the sample to make any such assumption.

In the remainder of this paper, we present two alternative approaches to exploring whether belief revision proceeded coherently across historical time under these circumstances: (1) a simple approach rooted in direct comparisons of observed correlations; and (2) a model-based approach rooted in the belief network discovered via Bayesian network structure learning over the Time 1 correlations.

Analysis 1: Comparing single-timepoint correlations and co-changes over time

As a simple starting point, we compared single-timepoint correlations between beliefs to co-changes in those beliefs over time. The Pearson correlation between the 78 off-diagonal correlations at Time 1 (Fig. 3A) and the corresponding 78 co-changes between beliefs over time (Fig. 3B) was quite high, R=.92, indicating that beliefs that were the most strongly correlated across individuals at Time 1 were also the most tightly linked in their changes over time, while beliefs that were not strongly correlated at Time 1 were also less tightly linked in their changes over time; see Fig. 4A. Mantel permutation tests identified R=.41 as the 99th percentile of the null distribution, indicating that these two matrices are much more similar than we would expect by chance. This was not an artifact of limiting the Time 1 data to the subset of returning participants; the relationship with Time 1 correlations among *non*-returning participants (n=635) was just as strong (R=.94).

Fig. 4A reveals a distinctive S-shape to the relationship between Time 1 correlations and co-changes over time; indeed, a cubic model was a significantly better fit to the data (linear model: R^2=.84; cubic model: R^2=.92; F(2, 74)=41.32, p<.001). This cubic relationship may be rooted in the basic measurement properties of correlation coefficients: stronger correlations (closer to ±1) are measured more precisely than weaker correlations (closer to 0), and this relationship between correlation strength and measurement error is non-linear (i.e., as r approaches either +1 or –1, SE_r decreases quadratically). By extension, the observed cubic relationship might derive from differential attenuation of correlation estimates among weaker conceptual links (see Schmidt & Hunter, 1996). Alternatively, weaker correlations at Time 1 may be more likely to reflect confounds (e.g., individual response biases) that are washed out in change scores, creating an artificial inflation of r_{t1} at values close to 0. At a high level, these findings align with our basic intuition that two beliefs have more freedom to change independently over time if they are less fundamentally intertwined to begin with.

As illustrated in Fig. 4A, the biggest prediction errors were underestimations of the positive co-change between *COVID long-term risk for children* and *COVID severity for children* (residual from the cubic model: +.211) as well as the negative co-changes between these two beliefs and *COVID rarity* (residuals: -.158 and -.135, respectively). These three beliefs may be more tightly linked than they appeared at Time 1—or indeed may have become more tightly linked over the course of these nine months. For belief pairs involving *COVID vaccine endorsement*, our primary outcome of interest in this broader line of work, residuals were fairly small, (-.058 to +.092; median residual: +.007), suggesting that there were no obvious sources of systematic error in predicting how perturbations to the belief network would impact parents' attitudes toward pediatric COVID vaccines (our best proxy for their own vaccine decisions). Likewise, the two beliefs that were identified at Time 1 as the best candidates for interventions to increase *COVID vaccine endorsement*—*COVID vaccine effectiveness* and *COVID vaccine community benefits*—were also the beliefs that were most strongly linked to changes in endorsement over time.

Analysis 2: Bayesian network analysis

In the second analysis, we specify this belief network in the form of a generative cognitive model and compare observed co-changes to model-based simulations.

As detailed in Weisman & Gibson (2025), in previous work we used the Time 1 correlations among beliefs to conduct a bottom-up search for a network representing the relationships among these beliefs via Bayesian network structure learning. Following the procedure pioneered by Powell et al. (2023), we constrained our search to networks

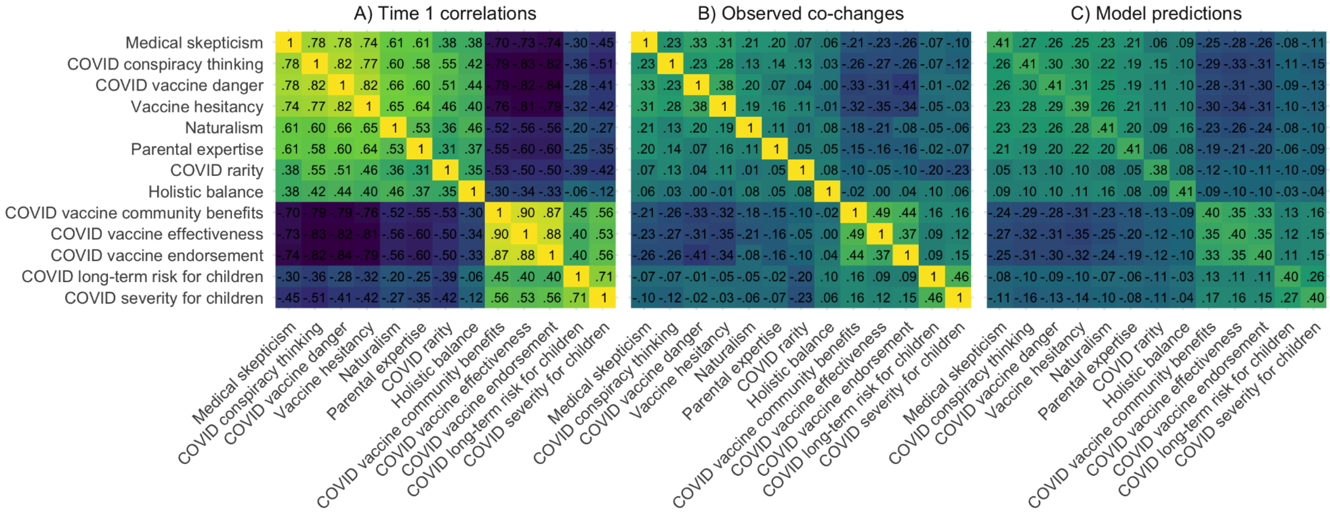

Figure 3: (A) Correlations between pairs of belief scores at Time 1, ordered via hierarchical clustering. (B) Correlations between change scores ("co-changes") over time. (C) Model predictions of the degree of coupling between beliefs; here, each column contains predictions for how evidence for that belief will propagate through the network to impact all other beliefs, and each row contains predictions for how sensitive a belief is to evidence for each of the other beliefs in the network.

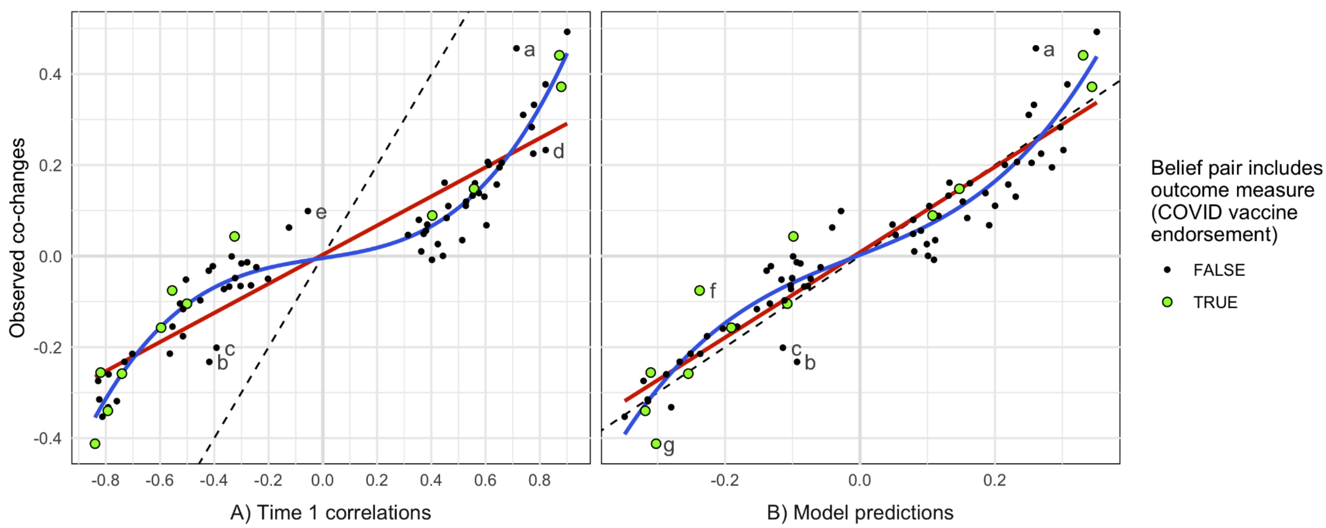

Figure 4: (A) The relationship between Time 1 correlations and observed co-changes. Each point represents one of the 78 possible pairs of beliefs (linear fit, in red: R^2=.84; cubic fit, in blue: R^2=.92). (B) The relationship between model predictions and observed co-changes (linear fit: R^2=.88; cubic fit: R^2=.91). Belief pairs including our outcome of interest are depicted in green, and belief pairs with the largest residuals are labeled, as follows: (a) *COVID long-term risk for children–COVID severity for children*; (b) *COVID rarity–COVID severity for children*; (c) *COVID long-term risk for children–COVID rarity*; (d) *COVID vaccine danger–COVID conspiracy thinking*; (e) *COVID long-term risk for children–Holistic balance*; (f) *COVID vaccine endorsement–Naturalism*; (g) *COVID vaccine danger–COVID vaccine endorsement*.

in which directed edges connecting nodes of beliefs flow from more general/abstract beliefs to more specific/concrete beliefs (a pre-registered "generative model principle"), and used the hill-climbing algorithm implemented in the "bnlearn" R package (Scutari, 2010) to identify the network structure that maximized the likelihood of the observed correlations between beliefs in a 70% training split (*n*=1238). In the original analysis, we estimated the parameters of this model as a Beta regression using maximum likelihood methods applied to the training split, and validated the model by testing its predictions on the held-out test split (R^2=.70). For the current analyses, we refitted model parameters using data from the subset of participants who returned for Time 2 (*n*=884). (The refitted parameters were very similar to the original Time 1 model.)

We used the resulting model to simulate a series of interventions to predict how beliefs should change together in response to evidence. For each of the 13 beliefs in the network, we introduced an evidence node as a child of that belief, conditioned on that evidence, propagated this evidence through the full network via approximate Bayesian inference, and recorded the posterior probability of each of the other 12 beliefs. We did this twice for each target belief: once introducing positive evidence (multiplying the odds by 10); and the other time introducing an equal amount of negative evidence (dividing the odds by 10). We combined the effects of these two simulations by subtracting posterior probabilities under negative evidence from posterior probabilities under positive evidence and dividing by two. The resulting sensitivity matrix J is presented in Fig. 3C, where each entry $J[i,j]$ represents the predicted change in i in response to evidence for j.

The correlation between the 78 off-diagonal predictions (upper triangle), and the corresponding observed co-changes was again quite high, R=.94; see Fig. 4B. As in Analysis 1, a cubic model was a significantly better fit, although in this case the difference was less pronounced (linear model: R^2=.88; cubic model: R^2=.91; nested model comparison: $F(2, 74)$=11.73, p<.001). Again, the biggest errors in model predictions were in the relationships among *COVID long-term risk for children*, *COVID severity for children*, and *COVID rarity* (all |residuals| > .132). For belief pairs involving *COVID vaccine endorsement*, residuals ranged from -.115 to +.118 (median residual: +.002). Finally, model predictions were equally good when we refitted model parameters on non-returning participants (R=.94), suggesting that the predictive power of the model is unlikely to be an artifact of non-random attrition between timepoints.

A limitation of this approach is that these predictions are based on simulations of interventions to one belief at a time, which is almost certainly an oversimplification of real-world events, especially in this time period. Nonetheless, the success of even this rough approximation provides some validation of the Time 1 model as a reasonably precise tool for predicting how beliefs change together over time.

Discussion

The current study is part of a larger project that aims to discover and specify with quantitative precision a network of beliefs that guides and constrains reasoning, decision making, belief revision, and behavior change in the domain of caregivers' decisions about pediatric vaccination. We documented changes over a 9-month period in a variety of beliefs related to pediatric COVID vaccination in a large sample of US parents, and found that beliefs that went together at Time 1 also tended to change together between timepoints—whether we operationalized "going together" using simple Time 1 correlations (Analysis 1) or a more sophisticated Bayesian modeling approach (Analysis 2). This is consistent with our theoretical assertion that cross-sectional correlations in beliefs represent both a summary of how different individuals think about pediatric COVID vaccines and a network of beliefs that captures the dynamics of belief change. (See von Klipstein et al., 2021, for a similar approach in the domain of psychopathology.)

We think of this network as a cognitive model that not only describes links between beliefs at the group level, but also guides reasoning and belief revision at the individual level. The results described here are consistent with this possibility—but they are also consistent with several alternative explanations. For example, an external influence could have directly affected several beliefs simultaneously (for everyone, or for some subgroup, or in different directions for different participants); this multipronged influence could have induced the observed "co-changes" among these beliefs without involving the propagation of evidence through a belief network. Likewise, over time participants may have moved toward or away from social identities associated with certain profiles of beliefs, without these beliefs being conceptually linked in their own mental models of the world (Cohen, 2003). Cross-lagged analyses over ≥2 timepoints and experimental interventions targeting one belief at a time offer two promising avenues for testing our claim that these correlations across beliefs reflect, at least in part, conceptual links in an underlying cognitive model. In turn, comparing meaningful subgroups of participants—for example, participants with different political identities, ethno-cultural identities, educational backgrounds, or medical histories—could shed light on the extent to which the structure of this (purported) cognitive model is held in common across people from a diversity of backgrounds, vs. emerging as an average over more distinct ways of reasoning about this domain.

Caregiver beliefs about pediatric vaccination are an interesting topic for many reasons. How do parents make choices about their children's health and well-being, when every option seems to carry its own set of risks, and the benefits are difficult to observe firsthand? Why do some parents end up making different choices than others, and how might the cognitive science of folk theories and belief revision be deployed to explain and support this reasoning process? Why did a global pandemic cause some people to become more hesitant about vaccines—and how might future events continue to shape public (and private) opinion about preventative care? Most urgently, from the perspective of public health: How can we encourage more parents to obtain life-saving vaccinations for their children? In the current study, cross-sectional correlations between beliefs generated fairly accurate predictions of how beliefs would change together over time, in turn suggesting that model-based simulations could be powerful tools for anticipating fluctuations in vaccination adherence in response to natural events, as well as planning educational interventions to support parents in making choices to promote their children's health and well-being.

Acknowledgements

Thanks to Allie Tung and Mikka Hoffman for assistance with survey design. We used ChatGPT and ClaudeAI to compare, critique, and refine analyses; to implement, improve, and test code; and to critique interpretations. The text of the paper was written by the human authors.

References

Anderer, S. (2025). US Measles Cases Hit 25-Year High. *JAMA*, *334*(8): 661. https://doi.org/10.1001/jama.2025.10974

Basch, C.H., Datuowei, E., Mewani, A., Jones, V., & Hillyer, G. C. (2025). Analysis of Google News coverage during the 2025 US measles outbreak. *American Journal of Infection Control*. https://doi.org/10.1016/j.ajic.2025.08.018

Baxter-King, R., Coppock, A., Straus, G., & Vavreck, L. (2025). Endorsements vs. information: Experimental evidence of backlash and parallel persuasion during the COVID-19 public health crisis. *PNAS Nexus*, *4*(6), pgaf185. https://doi.org/10.1093/pnasnexus/pgaf185

Cohen, G.L. (2003). Party over policy: The dominating impact of group influence on political beliefs. *Journal of Personality and Social Psychology*, *85*(5), 808–822. https://doi.org/10.1037/0022-3514.85.5.808

Du, Y., Paritala, S., Xu, Y., Maloney, P., & Lin, D. Y. (2025). Durability of 2024-2025 COVID-19 vaccines against JN. 1 subvariants. *JAMA Internal Medicine*, *185*(12), 1501-1504. doi.org/10.1001/jamainternmed.2025.5465

von Klipstein, L., Borsboom, D., Arntz, A. (2021). The exploratory value of cross-sectional partial correlation networks: Predicting relationships between change trajectories in borderline personality disorder. *PLoS ONE 16*(7): e0254496. https://doi.org/10.1371/journal.pone.0254496

McFadden, K.L., Hernández, T.D., & Ito, T.A. (2010). Attitudes toward complementary and alternative medicine influence its use. *EXPLORE*, *6*(6), 380–388. https://doi.org/10.1016/j.explore.2010.08.004

Grippin, A.J., Marconi, C., Copling, S., Li, N, Braun, C., Woody, C., … & Lin, S.H. (2025). SARS-CoV-2 mRNA vaccines sensitize tumours to immune checkpoint blockade. *Nature*, *647*, 488–497. https://doi.org/10.1038/s41586-025-09655-y

O'Mahoney, L.L., Routen, A., Gillies, C., Jenkins, S.A., Almaqhawi, A., Ayoubkhani, D., … & Khunti, K. (2025). The risk of Long Covid symptoms: a systematic review and meta-analysis of controlled studies. *Nature Communications*, *16*(1), 4249. https://doi.org/10.1038/s41467-025-59012-w

Powell, D., Weisman, K., & Markman, E.M. (2023). Modeling and leveraging intuitive theories to improve vaccine attitudes. *Journal of Experimental Psychology: General*, *152*(5), 1379–1395. https://doi.org/10.1037/xge0001324

Schmidt, F.L., & Hunter, J.E. (1996). Measurement error in psychological research: Lessons from 26 research scenarios. *Psychological Methods*, *1*(2), 199–223. https://doi.org/10.1037/1082-989x.1.2.199

Scutari, M. (2010). Learning Bayesian Networks with the bnlearn R Package. *Journal of Statistical Software*, *35*(3), 1-22. (package version: 4.8.3) https://doi.org/10.18637/jss.v035.i03

Shapiro, G.K., Tatar, O., Dube, E., Amsel, R., Knauper, B., Naz, A., ... & Rosberger, Z. (2018). The vaccine hesitancy scale: Psychometric properties and validation. *Vaccine*, *36*(5), 660-667. https://doi.org/10.1016/j.vaccine.2017.12.043

Wang, W., Wang, Q., Zhao, S., Zhou, X., … & Yu, H. (2025). COVID-19 pandemic and waning immunity disrupted measles population immunity and strategies to close immunity gaps. *Nature Communications*, *16*, 10966. https://doi.org/10.1038/s41467-025-65955-x

Weisman, K., & Gibson, D.J. (2025). Modeling a network of beliefs surrounding parents' endorsement of COVID vaccines for children. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *47*. Retrieved from https://escholarship.org/uc/item/8dq8m80z

Weisman, K., & Markman, E.M. (2017). Theory-based explanation as intervention. *Psychonomic Bulletin and Review*, *24*(5), 1555–1562. https://doi.org/10.3758/s13423-016-1207-2

Young, L.J., & Bartels, M. (2025, June 17). Can you still get a COVID vaccine this fall? Here's what to know. *Scientific American*. Retrieved from https://www.scientificamerican.com/article/can-you-still-get-a-covid-vaccine-this-fall-heres-what-to-know/

Zaller, J., & Feldman, S. (1992). A simple theory of the survey response: Answering questions versus revealing preferences. *American Journal of Political Science*, *36*(3), 579–616. https://doi.org/10.2307/2111583