Modeling a network of beliefs surrounding parents' endorsement of COVID vaccines for children

Kara Weisman (kgweisman@gmail.com) & Dominic J. Gibson (dominic@foundry10.org)
Conceptual Development Lab, foundry10, 4244 University Way NE, #85538 Seattle, WA 98145

Abstract

Cognitive science offers powerful tools for addressing pressing public health needs. Here we apply the cognitive science of intuitive theories and the tools of Bayesian networks to shed light on why so few children in the US have received COVID vaccines and how we might encourage caregivers to seek out these and other life-saving vaccines for their children. 1700 US parents completed 13 belief scales on a range of topics likely to influence their endorsement of pediatric COVID vaccines. We deployed structure learning techniques to develop a cognitive model of the relationships among beliefs and their influence on endorsement of vaccines for children, which accounted for 70% of the variance in participants' beliefs in a held-out testing split. Model-based simulations suggested that educational interventions focused on the effectiveness of COVID vaccines for supporting individual and community health may be most effective in increasing uptake of the COVID vaccine for children.

Keywords: intuitive theories; belief revision; cognitive modeling; Bayesian networks; Bayesian network structure learning; vaccine attitudes; parenting

The Centers for Disease Control and Prevention estimates that only around half of US children have ever received a COVID vaccine, and only 11.3% received this season's updated shot (CDC 2024a; 2025). In addition to children needlessly getting ill (and in a small number of cases severely so), this lack of uptake for pediatric COVID vaccines allows COVID to continue to spread and mutate in the community (Chen, 2023). The broader implications of this anti-vaccination sentiment are even more disturbing. US parents and caregivers have increasingly opted out of even the standard course of vaccinations for their children in the years since the COVID pandemic began (CDC, 2024b), leading to the return of diseases like measles (CDC, 2024c) and polio (Broadfoot, 2022)—diseases which were eliminated in the US in the 20th century, and are known to have devastating impacts on unvaccinated children. Meanwhile, medical centers and global health organizations have issued increasingly alarming warnings regarding the imminent threat of new pandemics with the potential to cause severe disease and high fatality rates (e.g., avian influenza). It has never been more important to understand vaccine hesitancy among parents and caregivers.

In recent years, many cognitive scientists have suggested that cognitive science offers both theories and tools that are well-suited to address such pressing real-world issues (Orchinik et al., 2024; Priniski & Holyoak, 2022; Weisman & Markman, 2017). For example, Powell, Weisman, & Markman (2023) used Bayesian network structure learning techniques to develop a cognitive model of the relationships among a variety of beliefs related to the standard course of pediatric vaccinations against "childhood diseases" like measles. The resulting model demonstrated strong out-of-sample predictions; explained participants' belief revision in response to educational interventions and to external events; and led to the development of a new intervention that improved vaccine attitudes.

In this paper, we use the approach pioneered by Powell et al. (2023) to develop a cognitive model of the relationships among a variety of beliefs related to pediatric COVID vaccination, working with a large sample of US parents. We deploy Bayesian network structure learning to infer how these beliefs might work together to influence caregivers' vaccination decisions, and use the resulting cognitive model to simulate how hypothetical interventions on these beliefs might spark belief revision. In addition to addressing the ongoing threat of COVID itself, we consider this a useful case study of why caregivers might opt in or out of vaccinating their children against any number of diseases—and, by extension, how researchers and medical professionals might intervene to encourage vaccination.

Method

Sampling plan, inclusion/exclusion criteria, procedure, data processing, and modeling approach were preregistered at https://osf.io/vgysz.

Participants

1700 US adults were recruited through the online survey platform Prolific in December 2024, using the following inclusion criteria: participants must be ≥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 Democrats (36% of our final sample), Republicans (34%), and participants who identified with neither political party (29%). We excluded participants who completed the study in <5 min (*n*=8); failed ≥2 of the 3 attention checks embedded in the survey

In D. Barner, N.R. Bramley, A. Ruggeri and C.M. Walker (Eds.), *Proceedings of the 47th Annual Conference of the Cognitive Science Society*

(n=8); reported that they did not provide thoughtful answers to the questions in this survey (n=0); were flagged by Qualtrics as engaging in fraudulent behavior (n=132); or indicated that they did not have any children in the target age range (n=24).[1] To build our model, we further excluded any participant who declined to answer all of the questions included in any of our belief scales (n=4).

This left us with a final sample of n=1524 US parents, who reported having a total of n=2809 children in the target age range. Participants ranged in age from 19-74y (median: 38y), and reported having between 1-9 children in the target age range (median number of children: 2; median child age: 8y). Participants included more mothers (63%) than fathers (36%); an additional 1% of participants did not identify as women or men. The sample was predominantly White (65%); an additional 19% of participants identified as Black or African American, 7% as ≥1 race/ethnicity, and ≤5% as as any other race/ethnicity. 60% of participants reported having obtained at least a college degree, and an additional 28% reported having attended some college.

Procedure

Participants completed an online study titled "Parents' and Caregivers' Health Beliefs Survey."

After providing some basic demographic information (e.g., number and ages of children), participants completed a series of 13 scales indexing a variety of beliefs relevant to vaccinating children against COVID; see Table 1.

This included five existing measures: Shapiro et al.'s (2018) *vaccine hesitancy* scale, which was developed prior to the COVID pandemic and focuses on the standard course of childhood vaccinations; extended versions of Powell et al.'s (2023) *naturalism*, *medical skepticism*, and *parental expertise* scales; and McFadden et al.'s (2010) *holistic balance* subscale (also used in Powell et al., 2023).

The study also featured seven newly developed measures: a *COVID conspiracy thinking* scale, reflecting popular conspiracies about the COVID pandemic; scales assessing beliefs about *COVID rarity*, *COVID severity for children*, *COVID long-term risk for children*, *COVID vaccine effectiveness*, and *COVID vaccine danger*, all loosely based on scales from Powell et al. (2023) but tailored to focus on COVID rather than measles and other "childhood diseases"; and a *COVID vaccine community benefits* scale, probing beliefs about the impact of COVID vaccines at the community level. These new scales, and our extensions of the *naturalism*, *medical skepticism*, and *parental expertise* scales, were based on ethnographic observations, personal experiences, correspondence with parents and medical professionals, and posts in online parenting forums. They were refined over the course of two pilot studies to achieve equally high internal consistency across scales; for any measures with observed Cronbach's α<0.80 in the first pilot study (n=97), we developed new items and re-piloted with new participants in a second pilot (n=51). For each scale, we retained a small number of items that increased internal consistency while maintaining a high degree of face validity (all observed αs>0.73 in the current study; see Table 1).

In addition to these 12 scales, participants also completed Enders et al.'s (2022) "COVID-19 Conspiracy Beliefs" scale. As preregistered, this was intended to validate our novel measure of *COVID conspiracy thinking* (observed correlation between scales: r=0.82), and was not included in the Bayesian network analysis described here.

These 13 belief scales were presented in a random order for each participant, directly followed by our primary outcome of interest: a novel index of *COVID vaccine endorsement*. This measure 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. This measure was intended to be an index of participants' current attitudes toward pediatric COVID vaccines in general and a proxy of vaccine decisions for their own families.

Each of these measures 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 infections, COVID vaccines, and vaccination more broadly, as well as a variety of demographic measures. In the current manuscript, we report simple correlations between the *COVID vaccine endorsement* scale and a subset of vaccine-related behaviors as an index of the validity of this novel outcome measure, leaving further analysis of these questions to future manuscripts.

Data preparation and modeling approach

Our approach followed Powell et al. (2023), Study 1.

For all scales, each participant's responses were averaged together after reverse-coding, dropping any items with missing data. We then rescaled these scores to range from 0-1, representing the participant's credence in each belief.

We used the observed correlations among these scores to conduct a bottom-up search for a Bayesian network representing the relationships among these beliefs.

Following Powell et al. (2023), our search incorporated a "generative model principle" stipulating that general and abstract states of affairs cause or otherwise generate more specific and concrete states of affairs. To do this, we sorted our belief scales into 4 preregistered tiers of abstractness (see Table 1) and constrained our search to networks in which directed edges connecting beliefs flow from more

[1] In addition to using reCaptcha score for bot detection, as preregistered (n=8 excluded), we further excluded anyone flagged for any type of fraud by Qualtrics. We have not yet implemented one of our preregistered exclusion criteria: excluding participants who provided nonsensical responses to a question about what they did in the study.

Table 1: Belief measures, ordered by tier of abstractness (preregistered). Example items are those with the strongest item-whole correlations. The last three columns indicate the number of items per scale (#), observed internal consistency (Cronbach's α), and correlations with *COVID vaccine endorsement* scores (Pearson's *r*).

Tier		Scale name	Example item	#	α	*r*
1	Abstract worldviews	*Holistic balance*	Health and disease are a reflection of balance between positive life enhancing forces and negative destructive forces.	5	.83	-0.25
		Naturalism	It is best to avoid eating genetically modified foods.	7	.77	-0.52
2	General beliefs	*General vaccine hesitancy*	Childhood vaccines are important for my child's health. (reversed)	10	.94	-0.78
		Medical skepticism	Pharmaceutical companies put pressure on the FDA and CDC to suppress negative findings.	7	.85	-0.71
		Parental expertise	Parents have insights into their children's health and well-being that no health professional can match.	5	.74	-0.57
3	COVID-specific beliefs	*COVID conspiracy thinking*	The government aimed to benefit by spreading lies about COVID.	6	.91	-0.78
		COVID long-term risk for children	Even a mild case of COVID can cause long-term problems for a child.	6	.90	0.35
		COVID rarity	People don't really get infected with COVID anymore.	8	.89	-0.33
		COVID severity for children	COVID is a serious disease for children.	7	.87	0.51
		COVID vaccine community benefits	When enough people in a community are vaccinated against COVID, rates of COVID-related hospitalizations and deaths in that community are lower.	7	.93	0.85
		COVID vaccine danger for children	mRNA vaccines (like the Pfizer and Moderna COVID vaccines) are not safe for children.	8	.91	-0.81
		COVID vaccine effectiveness	Vaccines have played a large role in reducing the number of deaths from COVID.	6	.93	0.85
4	Outcome	*COVID vaccine endorsement*	I would advise them to get their child vaccinated against COVID.	8	.94	1.00

general and abstract beliefs to more specific and concrete beliefs, and not in reverse. Edges between beliefs in the same tier were allowed to flow in either direction.

We 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 (training split: n=1238). We estimated the parameters of this model as a Beta regression, and validated the model by testing its predictions on the held-out testing data (n=310).

Finally, we used this model to simulate how *COVID vaccine endorsement* would be influenced by hypothetical interventions targeting the other beliefs in the network.

Results

Validation of outcome measure

The *COVID vaccine endorsement* (*CVE*) scale was highly correlated with all of our indices of COVID vaccine status and attitudes. Across the full sample (n=1524), 73% of participants reported having received ≥1 COVID vaccine since they first became available, and these reports were positively correlated with *CVE* scores (r_{pb}=0.58), as were reports of receiving the 2023-24 shot (36% vaccinated, r_{pb}=0.54); reports of receiving the 2024-25 updated shot (24% vaccinated, r_{pb}=0.41); and, among the remaining 76% of participants who had not yet received the 2024-25 shot, judgments of how likely they were to do so (on a scale of 0-10; r=0.63). We observed similarly strong relationships between *CVE* scores and reports of children's COVID vaccination status, both overall (49% of n=2809 children vaccinated; correlation with parent's *CVE* score: r_{pb}=0.61) and in recent seasons (2023-24: 30% vaccinated, r_{pb}=0.51; 2024-25: 20% vaccinated, r_{pb}=0.40; correlation with judged likelihood of obtaining this year's shot among the remaining 80% of children who had not been vaccinated this season: r=0.68). Finally, *CVE* scores were highly correlated with responses to the following question: "Imagine that your child/children's school is considering hosting a COVID vaccine clinic for a few weeks, where children and families could receive free COVID vaccines/boosters. How would you feel about this idea?" (7-point response scale, ranging from "Strongly against this idea" to "Strongly in favor of this idea"; correlation with *CVE* score: r=0.84). Taken together, we consider these correlations to be promising indications of the *CVE* scale's external validity as a proxy of parents' decisions about COVID vaccines for their own families, and perhaps as a more general indicator of whether an individual serves as an amplifier or detractor of efforts to increase COVID vaccination in their community.

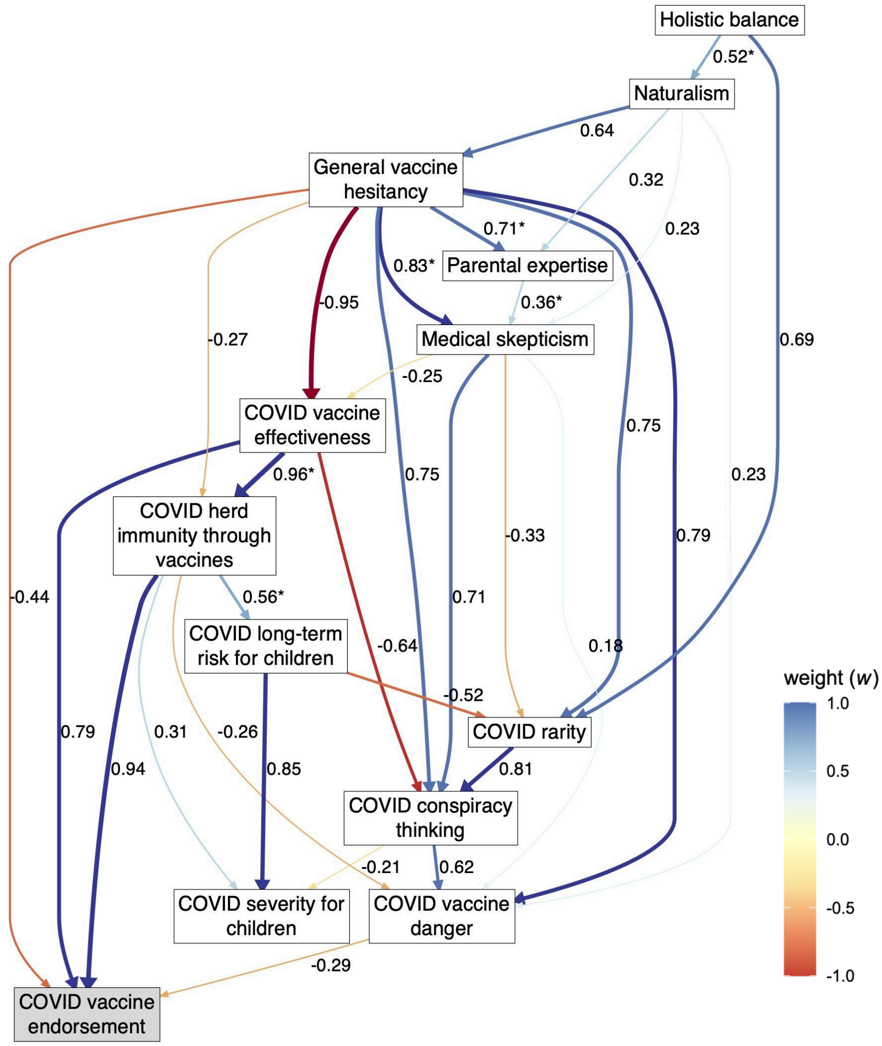

Figure 1: Bayesian network model of an intuitive theory (training split: n=1238). Our primary DV, *COVID vaccine endorsement*, is highlighted in gray. Edge color and size indicate the weight of the coefficient for each link between beliefs: generative links are depicted in blue, and preventive links in red. The direction of 6 edges, marked with asterisks, was not determinable from the combination of our data and modeling constraints and was selected randomly in the final model.

Bayesian network model of beliefs

Figure 1 presents the Bayesian network that maximizes the likelihood of the observed correlations between beliefs in the training split (n=1238), using the hill-climbing algorithm and the constraints specified by our generative model principle. This model has 32 edges connecting the 13 nodes in the network, including both generative links (depicted in blue) and preventative links (in red) of varying strength (ws ranged from -0.95 to +0.96).

To assess model performance, for each participant in the held-out test split (n=310), we generated predictions for scores on each target belief by conditioning the cognitive model network on observations of the remaining 12 beliefs, and compared these predictions to actual belief scores; see Figure 2. Pooling across beliefs, the average correlation between predicted and observed scores was high (r=0.84), accounting for 70% of the variance in observed scores. This out-of-sample performance suggests this model can usefully predict people's beliefs.

The cognitive model presented in Fig. 1 should be interpreted as a set of generative and preventative links between beliefs. For example, this model suggests that credence in *COVID vaccine community benefits* exerts a strong generative influence on *COVID vaccine endorsement*: Participants who believe that COVID vaccines promote public health at a population level (e.g., that high vaccination rates in a community result in fewer hospitalizations in that community) are in turn more likely to endorse pediatric COVID vaccines. As another example, this model suggests that *general vaccine hesitancy* exerts a strong negative influence on beliefs about *COVID vaccine effectiveness*: Participants who express general concerns about childhood vaccination are in turn less likely to believe that COVID vaccines effectively reduce the risk or severity of COVID infections. Following Powell et al. (2023), we consider this interconnected network of links between beliefs to be something like an intuitive theory supporting parents' reasoning about pediatric COVID vaccination.

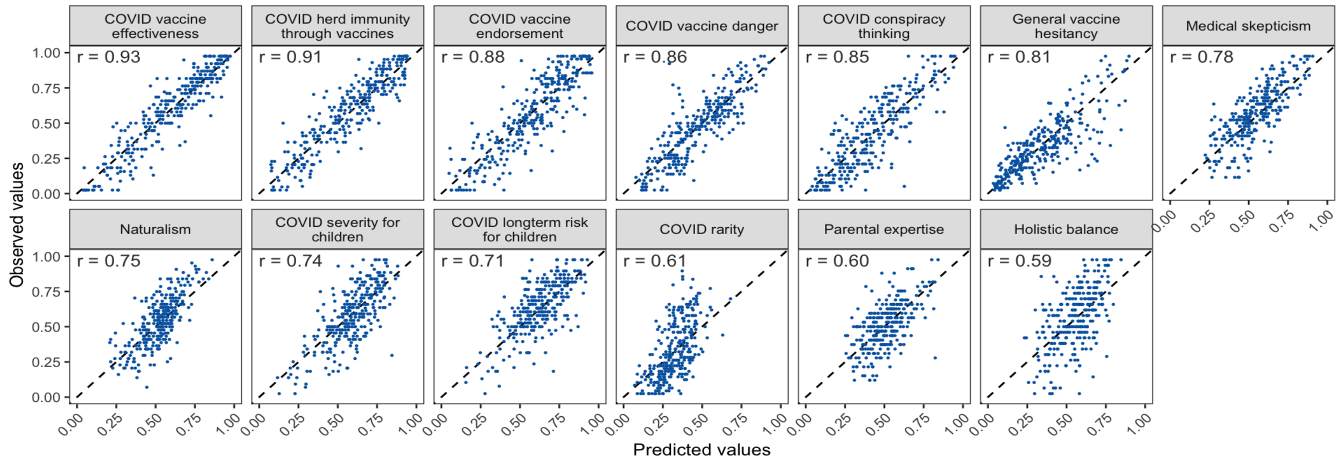

Figure 2: Model predictions versus observed values for 13 target beliefs in the held-out test split (*n*=310). Belief scales are ordered from strongest to weakest correlation; dashed lines represent perfect fits (*r*=1.00). This figure also illustrates the distribution of scores for each belief scale in the test split (along the vertical axis).

The nature of these links likely varies across the network, and the dataset itself cannot distinguish among these different interpretations of the probabilistic relationships between beliefs (see Powell et al., 2023; Williamson, 2001). In many cases, we can imagine a causal story, although drawing causal conclusions would of course require additional work: the link from *COVID vaccine community benefits* → *COVID vaccine endorsement* might represent the possibility that parents endorse pediatric COVID vaccines because they believe such vaccines can promote community health. Other links might be better interpreted as logical implication (e.g., *COVID long-term risk to children* → *COVID severity for children* might signify an understanding that if COVID carries a substantial risk of long-term consequences for children it should be considered "severe" by definition), or set membership (e.g., *COVID conspiracy thinking* → *COVID vaccine danger* might reflect the fact that narratives about the creation of dangerous vaccines are one part of, or variety of, COVID conspiracy theories).

There are six edges in the network (marked with asterisks in Fig. 1) whose direction was not determined from the combination of our training data and the constraints of our generative model principle; following Powell et al., 2023, these directions were set randomly. This indeterminate directionality could stem from limitations with our sample or scales, but might also reflect deeper complexities not captured by the current model.

For example, we echo Powell et al. (2023) in speculating that there may be unmeasured common causes of credences in *holistic balance* and *naturalism*, or credences in *medical skepticism* and *parental expertise*—edges whose directionality was also indeterminate in Powell et al.'s (2023) model. We might extend this kind of explanation to two closely related cases: the edges connecting *medical skepticism* and *parental expertise* to *general vaccine hesitancy* (a scale that was not included in Powell et al.'s study). Alternatively, more accurate representations of the connections among these beliefs might include undirected or bidirectional edges, types of influence relations which are not possible to represent in the current modeling context.

In the fifth case of indeterminate directionality—the link between *COVID vaccine effectiveness* and *COVID vaccine community benefits*—we speculate that participants (or at least some subset of participants) may not draw a strong distinction between vaccine effectiveness at the level of the individual vs. benefits at the community level. Further analysis, as well as additional empirical work, could shed light on whether it would make sense to collapse these two nodes into a single index of vaccine effectiveness—or, conversely, whether sharpening this distinction for participants could in turn increase vaccine endorsement.

We find the last case of indeterminate directionality—the link between *COVID vaccine community benefits* and *COVID long-term risk for children*—to be the most puzzling, both in its directionality and in the very presence of this edge in the network. Further empirical work might be required to understand how beliefs about the long-term risk of COVID for children relate to the community-level impact of COVID vaccines, or what other explanations might help explain this feature of the current model (e.g., unmeasured latent constructs). This might include direct experimental manipulations of beliefs, as well as qualitative studies of how participants interpret these scales and explain the relationship between these beliefs.

Such limitations notwithstanding, we find the model that emerged from this study to be intelligible and generally in line with our own intuitions about the intuitive theories that guide parents' vaccine decisions for their children. For example, the four beliefs with direct connections to *COVID vaccine endorsement* are all concrete beliefs about vaccines: *COVID vaccine community benefits*, *COVID vaccine effectiveness*, *general vaccine hesitancy*, and *COVID vaccine danger*. The identification of these factors as most proximal to vaccine decisions provides some evidence that the model is capturing important and sensible relationships between beliefs. Likewise, the fact that a structure-learning process that adheres to our a priori tiers of abstractness produces a model with accurate out-of-sample predictions is further evidence that this is a plausible approximation of the intuitive theories guiding reasoning and decision-making about pediatric COVID vaccines.

Simulating interventions

As a first step toward applying this model to design effective educational interventions, we used the model to simulate how hypothetical educational interventions on each of these target beliefs might cause parents to revise their endorsement of pediatric COVID vaccines. Estimates of the relative effectiveness of these hypothetical interventions, operationalized as the amount of pro-vaccine change in *CVE* scores assuming a fixed amount of pro-vaccine change in each targeted belief, are presented in Figure 3.

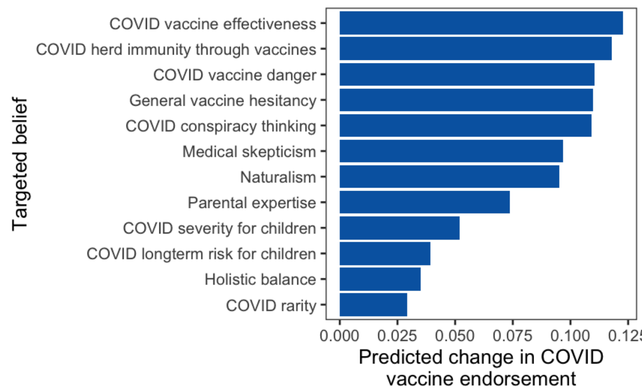

Figure 3: Model-based predictions of how hypothetical pro-vaccine interventions targeting each belief would increase endorsement of pediatric COVID vaccines.

Discussion

In a large-scale study of US parents, we applied the cognitive science of intuitive theories and belief revision, together with the tools of Bayesian networks and structure learning algorithms, to shed light on a pressing public health concern: low rates of COVID vaccination among US children. This bottom-up approach allowed us to uncover a network of beliefs feeding into caregivers' decisions about their children's health, without knowing in advance which beliefs might be most influential or how these beliefs might be interrelated. Specifying this network as a cognitive model allowed for quantitative precision in our theory-building and facilitated formal simulations of how changes in a single belief might propagate throughout the network. In addition to shedding light on some of the psychological factors that might account for low COVID vaccination rates among children, this work lays the foundation for a streamlined search for effective educational interventions to encourage caregivers to obtain COVID vaccines for their children.

First, our model suggests that increasing pediatric COVID vaccine uptake might require addressing caregivers' beliefs about COVID vaccines specifically, rather than focusing on the importance of vaccines for children in general. Although our data suggest that *general vaccine hesitancy* may exert some direct preventative influence on vaccine endorsement, the indirect paths from *general vaccine hesitancy* through *COVID vaccine effectiveness*, *COVID vaccine community benefits*, and (to a lesser degree) *COVID vaccine danger* appear to be much stronger.

Second, the model suggests that it is beliefs about the effectiveness of COVID vaccines, more than beliefs about the danger of COVID vaccines or the danger of COVID infections, that exert the strongest influence on endorsement of pediatric COVID vaccines. Indeed, *COVID long-term risk for children* and *COVID severity for children* played fairly minimal roles in the model, and were less strongly correlated with COVID vaccine endorsement than many other belief scales (see Table 1). This contrasts with Powell et al.'s (2023) findings regarding measles and other "childhood diseases," in which beliefs about vaccine danger and disease severity were strong predictors of vaccine endorsement and effective targets of pro-vaccine interventions. We speculate that these differences might reflect participants' awareness of the actual effectiveness of COVID vs. measles vaccines at preventing infection in the vaccinated individual; gaps in participants' understanding of the community-level benefits of vaccination; and differences in participants' confidence about the risks to children of COVID vs. other infections. More broadly, these findings raise the interesting possibilities that vaccine hesitancy might be instantiated differently across diseases or that the nature of vaccine hesitancy might be shifting over time.

In line with these observations, our model-based simulations suggest that one powerful way to increase uptake of pediatric COVID vaccines in this population could be to educate caregivers about the documented benefits of COVID vaccines, both at the individual level and at the community level. Following Weisman & Markman (2017), we emphasize that such interventions would be most effective if they targeted specific gaps and misconceptions and provided compelling explanations of how and why COVID vaccines are beneficial—e.g., acknowledging that people might still contract COVID even if they have been vaccinated, explaining that preventing infection is not the only goal of vaccination, and providing an alternative concept of "benefit" that includes reducing the severity of infections and reducing rates of transmission to others in addition to preventing infection for the vaccinated person.

These simulations could also be understood as predictions about the effects of real-world events on COVID vaccine uptake. Just as receiving an educational intervention that countered general vaccine hesitancy would, per these simulations, have a strong positive influence on a caregiver's endorsement of COVID vaccines (Table 1), real-world events that increase general vaccine hesitancy—such as the spread of misinformation about vaccines—could have an equally strong negative impact on caregivers' decisions about whether to vaccinate their children against COVID, not to mention diseases that pose even more direct and imminent threats to children themselves, such as measles, polio, and influenza.

Acknowledgements

Thanks to Allie Tung and Mikka Hoffman for their assistance with survey design and data collection, and to Lisa Castaneda and Sam Bindman for their role in the project's conceptualization. The code used to analyze these data was closely based on the code developed by Derek Powell for the studies described in Powell, Weisman, & Markman (2023).

References

Broadfoot, M. (2022, December 15). The basics of polio and why it's making a comeback. *Mayo Clinic News Network*. https://newsnetwork.mayoclinic.org/discussion/the -basics -of-polio-and-why-its-making-a-comeback/

Enders, A.M., Uscinski, J., Klofstad, C., & Stoler, J. (2022). On the relationship between conspiracy theory beliefs, misinformation, and vaccine hesitancy. *PLOS One*, *17*(10), e0276082.

Centers for Disease Control and Prevention. (2024, November 15). COVID-19 vaccination coverage and vaccine confidence among children. https://www.cdc.gov/covidvaxview/interactive/children.html

Centers for Disease Control and Prevention. (2024, October 2). Vaccination coverage and exemptions among kindergartners. https://www.cdc.gov/schoolvaxview/data/index.html

Centers for Disease Control and Prevention. (2024, November 14). Measles cases surge worldwide, infecting 10.3 million people in 2023. https://www.cdc.gov/media/releases/2024/p1114-measles-cases.html

Centers for Disease Control and Prevention. (2025, January 29). COVID-19 Vaccination coverage and intent for vaccination, children 6 months through 17 years, United States. https://www.cdc.gov/covidvaxview/weekly-dashboard/child-coverage-vaccination.html

Chen Y.T. (2023). Effect of vaccination patterns and vaccination rates on the spread and mortality of the COVID-19 pandemic. *Health policy and technology*, *12*(1), 100699.

McFadden, K.L., Hernández, T.D., & Ito, T.A. (2010). Attitudes toward complementary and alternative medicine influence its use. *EXPLORE*, *6*(6), 380–388.

Orchinik, R., Dubey, R., Gershman, S.J., Powell, D.M., Bhui, R. (2024). Learning from and about scientists: Consensus messaging shapes perceptions of climate change and climate scientists. *PNAS Nexus*, *3*(11), 485.

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.

Scutari, M. (2010). Learning Bayesian Networks with the bnlearn R Package. *Journal of Statistical Software*, *35*(3), 1-22. (package version: 4.8.3)

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.

Weisman, K., & Markman, E.M. (2017). Theory-based explanation as intervention. *Psychonomic Bulletin and Review*, *24*(5), 1555–1562.

Williamson, J. (2001). Bayesian networks for logical reasoning. In *Proceedings of the AAAI Fall Symposium on using Uncertainty within Computation* (pp. 136-143).