Modeling and leveraging intuitive theories to improve vaccine attitudes

Derek Powell¹, Kara Weisman², and & Ellen M. Markman³

¹ Arizona State University

² University of California - Riverside

³ Stanford University

Author Note

Supplementary Materials, including all data and analysis code for this project are available at <https://osf.io/sdfpj/>. Preregistration information for Studies 3a and 3b is available at <https://osf.io/uemq4>. Earlier modeling and analyses of data from Study 1 were presented at the 40th Annual Conference of the Cognitive Science Society. A six-page paper also appeared in the proceedings (Powell et al., 2018).

The authors made the following contributions. Derek Powell: Conceptualization, Investigation, Methodology, Data collection, Formal Analysis, Visualization, Writing; Kara Weisman: Conceptualization 6t, Investigation, Data collection, Formal Analysis, Visualization, Writing; Ellen M. Markman: Conceptualization, Methodology, Writing.

Correspondence concerning this article should be addressed to Derek Powell, 4701 W. Thunderbird Road, Glendale, AZ 85306. E-mail: dmpowell@asu.edu

Abstract

Much of the richness of human thought is supported by people's intuitive theories—mental frameworks capturing the perceived structure of the world. But intuitive theories can also contain and reinforce dangerous misconceptions. In this paper, we take up the case of misconceptions about vaccine safety that discourage vaccination. These misconceptions constitute a major public health risk that predates the coronavirus pandemic but that has become all the more dire in recent years. We argue that addressing such misconceptions requires awareness of the broader conceptual contexts in which they are embedded. To build this understanding, we examined the structure and revision of people's intuitive theories of vaccination in five large survey studies (total $N = 3196$). Based on these data, we present a cognitive model of the intuitive theory surrounding people's decisions about whether to vaccinate young children against diseases like measles, mumps, and rubella (MMR). Using this model, we were able to make accurate predictions about how people's beliefs would be revised in light of educational interventions, design an effective new intervention encouraging vaccination, and understand how these beliefs were affected by real-world events (the measles outbreaks of 2019). In addition to presenting a promising way forward for promoting the MMR vaccine, this approach has clear implications for encouraging the uptake of COVID-19 vaccines, especially among parents of young children. At the same time, this work provides the foundation for richer understandings of intuitive theories and belief revision more broadly.

Keywords: Intuitive theories; Belief revision; Vaccine attitudes; Cognitive modeling

Word count: 11853

Modeling and leveraging intuitive theories to improve vaccine attitudes

One remarkable feature of human psychology is our ability to form rich mental representations of what is true, how things work, and how best to navigate the world around us. Much of this type of thinking can be understood as deriving from people’s *intuitive theories*—explanatory frameworks, analogous to scientific theories, that capture the causal and logical structure of the world within a domain. From early in life, people are driven to seek out such causal explanations and to use their mental representations of these causal systems to explain, predict, and intervene in the world (see Carey, 2009; Gelman & Legare, 2011; Gerstenberg et al., 2021; Gopnik & Wellman, 1994; Keil, 1994; Murphy & Medin, 1985; H. M. Wellman & Gelman, 1992 for influential treatments of this topic). Intuitive theories inform reasoning in a wide variety of domains, including reasoning about physical systems (e.g., Baillargeon, 2004; Battaglia et al., 2013; Spelke, 1994; Ullman et al., 2017), biological processes (e.g., Carey, 1985, 2009; Medin & Atran, 1999) and the contents of other people’s minds (e.g., Baker et al., 2017; Gopnik & Wellman, 1994; H. Wellman, 2015).

Because intuitive theories are such powerful tools for making sense of the world, when these theories are mistaken there can be serious consequences. In this paper, we focus on an important and timely example: misconceptions about vaccine safety.

Even before the coronavirus pandemic, misconceptions about vaccine safety drove public health crises around the world. Undervaccination resulted in outbreaks of measles affecting over one thousand children across the United States in 2019, and over 80,000 children across Europe in 2018 (Thornton, 2019). Misconceptions about childhood vaccination—particularly the fraudulent claim that vaccines cause autism or other dangerous side-effects—have proved remarkably difficult to correct (e.g., Betsch & Sachse, 2013; Horne et al., 2015; Lewandowsky et al., 2012). Sadly, only a handful of educational interventions have been found to improve vaccine attitudes (e.g., Dubé et al., 2015; Horne

et al., 2015; Vannice et al., 2011; Wallace et al., 2006; and for some other effective approaches see e.g. Gagneur et al., 2018; Opel et al., 2013). This is especially sobering in the context of the coronavirus pandemic: widespread reluctance to be vaccinated against COVID-19 threatens our ability to end a pandemic that has already claimed more than five million lives around the world.

The difficulty of convincing people to vaccinate themselves and their children illustrates a general pattern of findings showing that people often do not respond to evidence in ways that educators might wish (e.g. Lewandowsky et al., 2012). Part of the challenge in correcting misconceptions is that they are situated within larger, internally coherent belief systems—intuitive theories—that guide how people interpret and respond to evidence (e.g., Cook & Lewandowsky, 2016; Gershman, 2019; Jern et al., 2014; Lewandowsky et al., 2012). In particular, prior work suggests the “attitude roots” of vaccine skepticism may spread quite wide, with vaccine misconceptions correlated with quite distant beliefs about the world (Hornsey et al., 2018).

Consider some of the beliefs people who are skeptical of childhood vaccination expressed when we asked them to explain their views in an open-ended qualitative survey (conducted pre-pandemic; $n = 16$). Several participants saw vaccines as “unnatural” and expressed concerns about toxic additives in vaccines, such as heavy metals and mercury. One said, “Vaccines are unnatural additions to the body. The human body should naturally fight off the disease and we should not be controlling how nature plays its role on mankind.” Some doubted whether certain diseases were worth worrying over, while others feared that the contents of vaccines “are too strong and too concentrated for a baby’s body to properly handle.” A number of respondents were skeptical of the medical community’s objectivity in testing the safety of vaccines. One respondent wrote that vaccines were “A BIG MONEY MAKING SCAM BY THE BIG PHARMS” (emphasis theirs); another thought that health professionals “are brainwashed by the medical media and the government.”

To people holding these beliefs, it could well seem unreasonable to vaccinate a 2-month-old baby against five or more diseases at once, as the CDC recommends. Moreover, if skeptics believe that the medical community is unduly influenced by pharmaceutical companies, they might doubt the validity of medical findings that appear to favor these companies' interests. Altogether, these beliefs might sustain misconceptions about vaccine safety, even in the face of counter-evidence: CDC guidelines, clinical studies, and even the advice of the family's pediatrician could be interpreted not as evidence that vaccines are safer than the person thought, but instead as further proof of the widespread corruption in the medical community.

The way in which evidence affects any particular belief is unavoidably shaped by other interrelated beliefs. Duhem (1954) and Quine (1951) famously argued that scientific theories necessarily incorporate *auxiliary hypotheses* that influence how observations are interpreted: when an observation is inconsistent with a hypothesis under certain auxiliary assumptions, a scientist cannot know whether to revise the hypothesis or the auxiliary assumptions. Similarly, we argue that intuitive theories incorporate auxiliary beliefs that affect how other beliefs are revised in light of evidence: when an observation is inconsistent with a person's current belief under certain auxiliary assumptions, the person cannot know whether to revise the belief or the auxiliary assumptions (also see Gershman, 2019).

Here we present a line of theoretical and empirical work that seeks to understand misconceptions and belief revision in the context of people's intuitive theories. We focus on reducing hesitancy toward childhood vaccination by identifying intuitive theories surrounding these vaccination decisions. Through this case study we hope to shed light on why such misconceptions persist and how interventions might address misconceptions more effectively across domains.

We argue that correcting misconceptions requires presenting evidence in a way that is sensitive to and persuasive within the broader conceptual context in which those

misconceptions reside (see Cook & Lewandowsky, 2016; Gershman, 2019; Hornsey et al., 2018; Weisman & Markman, 2017). Understanding the revision of any one specific belief (e.g., a misconception like *vaccines are dangerous*) demands a holistic understanding of the wider set of interrelated beliefs that make up people’s intuitive theories.

Modeling intuitive theories

Quine argued that all of one’s beliefs are intertwined in a “web of belief” (Quine, 1951; Quine & Ullian, 1978)—an apt metaphor for intuitive theories. Intuitive theories consist of relations between beliefs that specify how states of affairs causally, logically, or otherwise interconnect with one another. Specifically, we consider intuitive theories to be generative mental models that entail probabilistic dependencies among possible states of affairs. Given this definition, we argue that any pair of beliefs related by a intuitive theory can be expected to exhibit several qualitative features:

1. When two beliefs are related by an intuitive theory, evidence directly affecting one of those beliefs will also affect the other in accordance with their relationship in the intuitive theory (e.g. evidence for B will increase credence in A, and vice versa). This is a clear and direct consequence of the laws of probability and Bayes’ rule.
2. Where beliefs about two states of affairs, A and B, are related through an intuitive theory (e.g., A causes B), those beliefs will be systematically correlated across individuals in proportion to the average conditional probabilities (e.g. people who believe A will also tend to believe B); This can be seen through an application of the law of total probability, which implies a linear relationship between the two beliefs across individuals.
3. Thus, we can often reasonably expect there will be some systematic relationship between the average change in each belief following observation of some evidence and the correlation among those beliefs across individuals. We provide a more formal explication of these last two points in Supplementary Materials.

Here, we sought to formalize this understanding of intuitive theories by representing an intuitive theory as a Bayesian network. Originally designed as tools for artificial intelligence (Pearl, 1988), Bayesian networks are a class of probabilistic graphical models that have been deployed widely across cognitive science (e.g. Gopnik et al., 2004; Griffiths & Tenenbaum, 2005; Jern et al., 2014). The heart of a Bayesian network is a graphical structure—not unlike a “web of belief”—that represents variables as nodes connected by edges, where those edges represent probabilistic or statistical relations.

Bayesian networks

Figure 1 presents several examples of Bayesian networks. Each of the networks is a *directed acyclic graph* (DAG), a type of graphical model. Graphical models represent variables as nodes connected by edges. In a DAG, these edges are directed, flowing from parent to child, and are constrained so that they cannot form cycles or loops. In a Bayesian network, these directed edges encode conditional dependencies among variables, so that the conditional probability distribution for a given variable is defined only in terms of its direct parents.

In Figure 1a, the edge from *fire* $\rightarrow$ *smoke* indicates that the probability of smoke is increased by fire. So, conditional on fire, smoke becomes more likely. The Bayesian network captures how conditioning on any one node affects the probability of all the other nodes in the network. The directed edges in a Bayesian network capture *asymmetric* relationships: different edge directions imply different conditional dependencies and independencies. A criterion called “d-separation” provides a formal means for deciding which nodes in a network conditionally depend on one another (Pearl, 1988). To illustrate, consider some relationships depicted in Figure 1c. Intuitively, it should be clear that observing (conditioning on) lightning will affect the probability of fire and all of its descendents. Indeed, these nodes are “d-connected” in the network. But, observing (conditioning on) lightning should not affect the probability of accidents (these nodes are

Figure 1

Several examples illustrating different features of Bayesian Networks.

A) A graphical representation of influence relations: Fire causes smoke. This influence is represented by the conditional dependencies captured in a conditional probability table (CPT).

B) A “parent” node can be linked to multiple “children” and a “child” node can be linked to multiple “parents”: A child and its parents together compose a “family”. This graph has three families: {Accidents, lightning, drought, fire}, {fire, smoke}, and {fire, ash}.

C) Not all influence relations correspond to causation: we could say a storm consists of rain and lightning (marked as dashed arrows), which in turn causes thunder. Causes can be generative or preventive, e.g. rain prevents drought and fire (red arrows). Some nodes may exert countervailing influences on others: A storm consists of rain, which can prevent fires, but a storm also consists of lightning, which can cause fires. In addition, parents may influence their children independently or interactively. Lightning and accidents could be seen as independent causes of fire—an accident may cause a fire whether or not there is lightning, and vice versa. However, drought conditions interact with both of these other parent nodes, making them more likely to produce a fire. In the present work, we will consider only the role of independent influence relations.

d-separated in the network).

It can be tempting to interpret the edges in a Bayesian network as representing causal relationships. Strictly speaking however, the Bayesian network is simply an encoding of the probabilistic relationships among variables. These dependencies can be interpreted most generally as representing “influence relations” (Williamson, 2001). As modelers, we might sometimes give these influence relationships more specific interpretations: causal relationships and logical implications (Williamson, 2001) are two prominent sorts of influence relations. In other cases, influence relations might be interpreted as representing “constituency” (e.g. see Figure 1c) or an entity’s having a particular property (e.g. an edge from “being an apple” to “being red” might represent that redness is a common property of apples). However, in all these cases these interpretations are external to the Bayesian network model itself, where the edges signify only probabilistic relationships.

Bayesian networks like those depicted in Figure 1 can be made to capture each of the qualitative features of intuitive theories laid out above. First, a Bayesian network captures people’s reasoning about these states of affairs given observations: observing smoke would lead someone to infer that there is likely to be a fire, and observing fire would lead someone to predict smoke will soon follow (Feature 1). Using the conditional probability values in Figure 1a, observing smoke would shift someone with a prior probability of fire $p(\text{fire}) = .2$ to a posterior probability of fire $p(\text{fire}|\text{smoke}) = .83$. And for this same person, observing fire would shift their credence in smoke from $p(\text{smoke}) = .238$ to $p(\text{smoke}|\text{fire}) = .99$.

This model can also capture how these beliefs might covary across individuals: if two people hold different degrees of belief in *fire*, but share the same understanding of the relationship between *fire* and *smoke*, this model specifies the degree of belief each should have in *smoke* (Feature 2). If Alice holds that $p(\text{fire}) = .4$ and Bob holds that $p(\text{fire}) = .1$, then Alice should belief $p(\text{smoke}) = (.99)(.4) + (.05)(.6) = .426$ and Bob

should believe $p(\textit{smoke}) = (.99)(.1) + (.05)(.9) = .144$.

We can leverage both of these facts to derive a model of people’s intuitive theory relating these affairs based on the correlations among these beliefs. Because the conditional probabilities of *smoke* given *fire* and $\neg\textit{fire}$ crucially determine both the change in beliefs given observations and the correlation of beliefs across individuals, average changes in beliefs following evidence will generally be related to the initial correlation of those beliefs across individuals (Feature 3; see Supplementary Materials for further explanation).

Formulating these features in a formal model, rather than a verbal theory, allows us to specify not just the presence but the quantitative strength of these relations, and to simulate the cumulative influences of many factors in more complex networks (for example, incorporating beliefs about the frequency of wildfires, recent weather conditions, etc.).

Other network-based approaches have also been applied to model systematic relationships among beliefs and attitudes, such as neural network models (e.g. see Holyoak & Powell, 2016; Holyoak & Simon, 1999; Simon et al., 2004; Spellman et al., 1993; Thagard, 1989). In these models, beliefs or attitudes are represented by nodes (neurons) connected by undirected or bi-directed edges—in contrast to the directed edges used in Bayesian Networks. One application of neural-network models has been the modeling of “coherence-shifts”: patterns of belief revision wherein the beliefs or attitudes that should putatively serve as inputs to an evaluation are themselves modified during the process of evaluation. However, the directed edges of Bayesian networks are similarly capable of capturing these coherent reasoning dynamics: When conditioning on a node within a network, information flows in both directions along the edges. To illustrate, consider how observing smoke will lead a reasoner to infer that there is a fire, just as readily as observing fire will lead them to predict there will be smoke.¹

¹ With this in mind, Bayesian Networks also offer potential models of coherence-based reasoning effects based on rational belief updating mechanisms (Jern et al., 2014).

Where the structure of an intuitive theory is known, Bayesian networks thus offer powerful tools for predicting and understanding how evidence will affect people’s beliefs. However, there are many cases where the structure of the intuitive theory is itself the very question under investigation. Consider vaccine hesitancy: Earlier, we speculated that beliefs about the infant immune system and beliefs about the trustworthiness of the medical community (among many others) might affect people’s beliefs about vaccines. But the structure of the intuitive theory that might relate these beliefs is not at all clear.

Fortunately, the formalism of Bayesian networks provides tools for navigating these challenges. Roughly, when the relationships among beliefs are unknown, we can use the qualitative features of intuitive theories to *infer* the structure of an intuitive theory from observed correlations among beliefs. Structure learning algorithms (see Scutari & Denis, 2021) provide a means for searching through the (super-exponentially large) possible space of Bayesian networks to find the structure most likely to have generated some observed data. These approaches allow us to scale the Bayesian network modeling framework beyond the smaller models characteristic of early computational work on intuitive theories (e.g., Griffiths & Tenenbaum, 2005).

Current studies

Here, we used Bayesian network structure learning algorithms to construct a model of the intuitive theory underlying adults’ decisions about whether or not to vaccinate children against diseases like measles, mumps, and rubella.

In Study 1, we used Bayesian network structure learning algorithms to develop a cognitive model of the relationships among 14 beliefs related to vaccine hesitancy in a large sample of U.S. adults. In Study 2, we replicated the success of an existing intervention emphasizing the dangers of childhood diseases (Horne et al., 2015) and used our model to provide a more detailed understanding of *how and why* this intervention increases vaccine intentions and its effects on related beliefs. In Study 3a and 3b, we drew inspiration from

our model to create and validate the success of a novel intervention aimed at addressing concerns about toxic additives in vaccines. Finally, in Study 4 we tested whether people’s beliefs changed coherently following relevant real-world events, by examining beliefs before and after serious outbreaks of measles in the United States in early 2019.

Taken together, this work illustrates the promise of understanding belief revision in terms of intuitive theories and the utility of discovering and formalizing intuitive theories via computational models. Of particular interest in the current moment, these findings lay a foundation for empirically-validated, theory-based educational interventions that could encourage people to vaccinate their children and themselves against dangerous diseases like measles and COVID-19.

Study 1: Assessing and modeling intuitive theories of vaccination

Across all our studies, our primary outcome of interest was participants’ intentions to vaccinate their children against diseases like measles, mumps, and rubella (which we refer to henceforth as *vaccination intentions*). Drawing on a variety of sources, including research on anti-vaccine skepticism (e.g., Dubé et al., 2015; Larson et al., 2015; Salmon et al., 2015; Williams, 2014), anti-vaccine websites, and our own qualitative work with vaccine skeptics (see Supplementary Materials), we generated a list of 13 additional beliefs that, though surely not exhaustive, might jointly offer a reasonable means of understanding and predicting vaccine intentions.

We translated these qualitative insights into 14 psychometrically-valid scales. In addition to 1) *vaccination intentions* (e.g. “If I had a young baby I would have him or her vaccinated.”), these included a variety of beliefs about vaccines, including beliefs about 2) *vaccine danger*, the overall danger of vaccines (e.g., “Many children have severe adverse reactions to the MMR vaccine”); 3) *toxic additives in vaccines* (e.g., “Exposure to certain additives in vaccines can cause major health problems.”); and 4) *vaccine effectiveness*, how effective vaccines are in preventing disease (e.g., “Your chances of getting a disease after

being vaccinated against it are incredibly low”); as well as a variety of specific claims about childhood diseases like measles, mumps, and rubella, including beliefs about 5) *disease rarity* (e.g., “Diseases like measles and whooping cough are so rare that there is no real need to worry about them”) and 6) *disease severity* (e.g., “Diseases like measles and whooping cough are extremely dangerous for young children”). Beyond this, we theorized that participants’ understanding of the infant immune system might be relevant, including beliefs that 7) the infant immune system is weak (*IIS: weakness*; e.g., “Babies’ immune systems are weak”); 8) the infant immune system is limited in its capacity and can be easily overwhelmed (*IIS: limited capacity*; e.g., “Babies’ immune systems are severely limited in how much they can cope with at one time.”); and 9) vaccines strain the infant immune system (*IIS: vaccines strain*; e.g., “Vaccines strain the immune system much like the actual disease”). We also included general theories about parenting and medicine: 10) general *parental protectiveness* (e.g., “Parents should strive to control their baby’s environments as much as possible”); 11) *parental expertise*, namely the belief that parents usually know more about their children’s health than medical experts (e.g., “Every individual is different and a parent knows their own child best”); and 12) *medical skepticism*, including concerns about pharmaceutical companies and corruption in the medical community (e.g., “Pharmaceutical companies put pressure on the FDA and CDC to suppress negative findings”); and two broad worldviews: 13) *naturalism*, a general preference for natural over artificial things (e.g., “Natural things are always better than synthetic alternatives”); and 14) *holistic balance*, one important aspect of attitudes toward alternative medicine (e.g., “The body is essentially self-healing and the task of a health care provider is to assist the healing process”; (McFadden et al., 2010)).

Each scale measured agreement with 4-6 statements, including at least one reverse-coded item, and was highly reliable (observed reliability, Study 1: all Cronbach’s $\alpha \geq .73$). We developed 13 of these scales, and included one preexisting scale (the “holistic balance” subscale from (McFadden et al., 2010). See Supplementary Materials for all items

included in all scales.

In Study 1, we administered these 14 belief scales to participants in an online survey with 1129 participants. We then used these intercorrelations to build a cognitive model of the relations among these beliefs via Bayesian network structure learning techniques.

Method

Participants. A total of 1202 people recruited from Amazon Mechanical Turk (mTurk) participated in this study conducted in August of 2017. Participants were paid \$1.60 for about 8 minutes of their time. 73 participants (6%) failed at least one of two attention check questions and were excluded from further analysis, leaving a final sample of $n = 1129$. See Supplementary Materials for demographic information about this sample.

Procedure. Participants (final $n = 1129$) responded to the 14 belief scales. Each scale was presented on a separate survey page, and participants rated a series of statements on a 7-point scale from “Strongly disagree” (coded as -3) to “Strongly agree” (+3). Scales, and questions within each scale, were presented in a random order for each participant. Attention checks (e.g., “Please select somewhat agree”) were embedded within two to four randomly chosen scales and participants who failed any of these checks were excluded from analyses. We scored each belief scale by averaging a participant’s responses to each item (after reverse-coding where appropriate), then rescaling these values to fall between 0 and 1. We considered this score to represent the participant’s credence in the belief.

The 14 belief scales were followed by demographics questions and a handful of questions about whether participants were or were expecting to be a parent. Participants were then debriefed with a brief statement noting that vaccines are safe and pointing them toward a WHO website for further information.

Results

The goal of this study was to use these data to conduct a bottom-up search for a Bayesian network that would represent the intuitive theory surrounding vaccination decisions as a network of related beliefs. There were strong intercorrelations among these 14 beliefs confirming that these beliefs are, in fact, closely related to each other and to intentions to vaccinate children (see SM). Based on the observed associations among participants’ beliefs about each of these issues, we sought to infer how these states of affairs (e.g., that vaccines are dangerous, that diseases are severe) are connected in people’s intuitive theories—how people understand these states of affairs to cause, imply, or otherwise influence one another.

Bayesian network parameterization. A Bayesian network is a directed acyclic graph that represents variables as nodes connected by directed edges. Each set of “child” and “parent” nodes in the larger network is a “family” (for some examples, see Figure 1b). The network’s directed edges encode conditional dependencies among variables, and many parameterizations for the individual families are possible. In our model, the probability of some “child” belief node y can be expressed as a noisy-logical function of the influence of one or more “parent” belief nodes x :

$$P(y = 1) = \left(1 - (1 - w_0) \prod_i (1 - w_i)^{x_i}\right) \prod_j (1 - w_j)^{x_j} \quad (1)$$

Here, each parent node x is either a generative (i) or preventive (j) influence, that generates or prevents the child node with some probability weight (w_i or w_j). The weight w_0 represents the “leak” probability, or the baseline probability that the child occurs without any other generating influences. The first term represents the probability that y is generated (or, put another way, that not all of the possible generating forces fail to generate it). The second term represents the probability that it is not then prevented [i.e., that all the preventing forces fail; see Maaskant and Druzdzel (2008)].

This parameterization is a substantial simplification compared with the full flexibility possible within a Bayesian network. A cognitive model composed of these types of relations is capable of capturing a causal system of non-interactive generative and preventive causes. But it bears noting that this simplification comes with limitations. For instance, returning to the example network in Figure 1c, this parameterization would not permit capturing the relationships between dryness, lightning, and fire. Dryness is an enabling condition that both increases the probability of fire but also increases the probability that lightning would cause a fire. So the influences of these variables are not independent and would need a more complex parameterization to be properly represented.

Finally, having transformed participants' scores on each belief scale to fall between 0 and 1, we estimated the parameters of this model as a Beta regression using maximum likelihood methods (see Supplementary Materials for details).

A cognitive model underlying the data model. We aimed to represent not only the relationships among variables in the data but the *cognitive model* that gave rise to these observed relationships, i.e., the relationships among beliefs in the mind. We actually estimated two models sharing the same network structure as depicted in Figure 2: First, we estimated a data model assuming that participants possess a shared mental model from which they generate their responses to the belief scales. In this model, the values of $p(y = 1)$ and x in the equation above corresponded to participants' continuous responses to the belief scales (e.g. a participant's score for the *vaccine danger* scale). Then, from the parameters estimated in the data model we created a cognitive model of participants' reasoning, where the x are binary, indicating whether that state of affairs obtains (e.g., whether "vaccines are dangerous").

We used this cognitive model to simulate participants' degrees of beliefs given other beliefs, as well as how beliefs would change following exposure to evidence.

Structure learning, constrained by a “generative model principle.”

Rather than specifying a priori all of the connections among the 14 beliefs included in this study, we deployed Bayesian network structure learning to search for the most likely network structure given the observed correlations between beliefs—with the constraint that this structure conform to some of our strongest high-level intuitions about what an intuitive theory should look like in this domain.

In particular, we applied a constraint we called the “generative model principle,” based on the assumption that the intuitive theory is a generative model, so that these influence relations (“edges”) should flow from causes or other “generating” states of affairs toward effects. For instance, we stipulated that vaccine intentions—our outcome of interest—would be a “child node,” with edges only flowing inwards and not outwards (since one’s intention to vaccinate cannot “generate” the facts about vaccinations safety, effectiveness, etc). This generative model principle instills a common core assumption in research on intuitive theories and causal explanatory reasoning (Cheng, 1997; e.g. Gerstenberg et al., 2021; Gopnik et al., 2004).

To apply our generative model principle, we sorted the 14 measured beliefs into four tiers based on how general or abstract each belief was, reasoning that general and abstract states of affairs should be understood to cause (or otherwise generate) more specific and concrete states of affairs, so edges should flow from the general and abstract toward the specific and concrete (e.g., if natural things are better than artificial things, this should set an expectation about whether a specific class of artificial things—vaccines—is safe). We considered *holistic balance* and *naturalism* to be the most general and abstract beliefs among those we measured; followed by theories about parenting and medicine; then specific beliefs about vaccines, diseases, and the infant immune system; and finally our outcome of interest, *vaccine intentions*—the most concrete measurement of a specific intended action.

We used this sorting to induce a partial ordering of these 14 beliefs, stipulating that

directed edges that violated this ordering must not appear in the final network of beliefs (e.g., that there should be no edge from *vaccine intentions* to *naturalism*, although there could be an edge from *naturalism* to *vaccine intentions*).

To conduct an efficient search through the super-exponential space of graphs, we tested several learning algorithms implemented in the “bnlearn” R package (Scutari, 2010) and evaluated their performance through cross validation. To avoid overfitting, we used an 80% split of the data ($n = 904$) for training, and reserved the remaining 20% for evaluating model performance. We compared performance of different structure learning algorithms using k-fold cross validation within the training split. Graphs were scored by Bayesian Information Criteria, an asymptotic estimate of the likelihood of the data under a given graph structure. Based on these results, we chose to apply a score-based “hill-climbing” algorithm using our generative model constraints (see Supplementary Materials for details). Applying these techniques resulted in an estimate of the network shown in Figure 2.

We evaluated the resulting cognitive model’s performance by testing its accuracy in predicting responses in the 20% testing split ($n = 226$). We generated predictions for held-out participants’ responses for each variable by conditioning the cognitive model network on the remaining 13 observed variables. Figure 3 compares the model’s predictions with participants’ actual responses. Collapsing across all variables, the average correlation between predicted and observed responses was $r = 0.67$, accounting for 45% of the variance in observed responses.

Our goal in developing this cognitive model was to create a psychologically plausible and intelligible model of the intuitive theory surrounding vaccination decisions. However, it is also reasonable to ask whether other less-constrained models might be more predictive. To address this, we compared the Bayesian network model’s predictions to a suite of 14 naive “maximal” regression models, each using values on all 13 other belief scales as predictors for each belief. The model presented in Figure 2—which we take to be simpler,

Figure 2

*Bayesian network cognitive model resulting from application of hill-climbing algorithm with our generative-model constraints applied to data in Study 1 training split ($n = 904$). The network has 29 edges connecting these 14 variables. Edge polarity and weights are indicated by color, with red hues indicating negative or preventive influences, and blue hues indicating positive or generating influences. The direction of the edge was not determinable from the data and modeling constraints for two edges: *Naturalism* $\rightarrow$ *Holistic Balance* and *Medical Skepticism* $\rightarrow$ *Parental Expertise*; see *Supplementary Materials*.*

more psychologically plausible, and more intelligible—performed as strongly as these “maximal” regression models (see Supplementary Materials).

Figure 3

Observed responses and cognitive model predictions for each belief among participants in the Study 1 testing split. Correlations between observed and predicted values ranged from .380 to .898 across the different belief scales. In general, the model showed greater predictive accuracy for more central beliefs (e.g., vaccine danger) than for more peripheral beliefs (e.g., parental protectiveness; see Supplementary Materials for details).

What does this cognitive model reveal about the intuitive theory underlying vaccination decisions? The three nodes with direct connections to *vaccine intentions* are beliefs about *vaccine effectiveness*, *vaccine danger*, and the severity of childhood diseases (*disease severity*)—concrete beliefs that are, at face value, especially closely related to vaccination decisions. This gives us confidence that the model is capturing important aspects of the intuitive theory of interest.

Other parts of the model shed new light on the role of lay theories in vaccination decisions. For example, *naturalism*—the general view that natural things are better than

artificial things —appears to be strongly related to *medical skepticism* and *parental expertise*; and all three of these abstract beliefs are related to concrete beliefs that, in turn, feed into participants’ vaccination intentions.

This model structure also helps explain some initially surprising findings. For example, we originally speculated that dispelling misconceptions about the capacity of the infant immune system could promote positive attitudes toward vaccination. We were disappointed to observe the weak first-order correlation between this belief and *vaccine intentions* in our behavioral data ($r = -.09$ in our training split). The model sheds light on this surprising (lack of) relationship: Although the belief that the infant immune system is limited in capacity is positively related to the belief that vaccines strain the immune system—discouraging vaccination, as we had assumed—it also seems to promote the belief that childhood diseases have severe consequences for young children (*disease severity*), which might, in turn, encourage vaccination. These countervailing forces—which, on average, seem to have cancelled one-another out in this sample—dissuaded us from our plan to develop an intervention aimed at dispelling misconceptions about limited capacity in order to encourage childhood vaccination.

It is worth pointing out that not all of the relationships represented in the model should be interpreted as *causal*. Due to Bayesian networks’ tight connections with structural causal networks, it can be tempting to give the edges this interpretation, but in some cases these relationships are better thought of as probabilistic dependencies or “influence” relations (Williamson, 2001). For instance, the edge running from vaccine danger to toxic additives in vaccines clearly does not represent a causal influence: vaccines being dangerous does not cause them to contain toxins. Instead, a more plausible interpretation of this edge is that it reflects something like a set membership relation: that toxins are one type of danger that might be present in vaccines.

In other cases, our data and explicit constraints do not perfectly identify the

directions of the edges. This is the case for two edges (Naturalism $\rightarrow$ Holistic Balance and Medical Skepticism $\rightarrow$ Parental Expertise; see Supplementary Materials for details). In both cases, we do not feel it is intuitive which direction the edge should flow in, nor even that it should be directed at all. Such a counterintuitive result might be taken as indicating a limitation of Bayesian network models, which are composed solely of directed, acyclic relationships. We note that prior work has successfully modeled coherent belief revision by leveraging models with undirected or bidirected relations (e.g. Holyoak & Simon, 1999; Simon et al., 2004; Spellman et al., 1993; Thagard, 1989). Alternatively, we might instead hypothesize there is an additional unmeasured variable that connects these beliefs that could more intelligibly account for their association within the context of a directed network. For instance, a more abstract epistemic belief might influence both medical skepticism and parental expertise beliefs. To speculate, such an abstract belief might be in one’s own expertise or in how we come to have knowledge (e.g. through direct experience versus science).

More generally, as they operate solely on correlations among variables, structure learning algorithms cannot be guaranteed to correctly identify causal directionality. Looking forward to future applications of these methods, researchers might seek other ways to resolve ambiguities left by structure learning algorithms. The simplest approach would be to set edges manually based on top-down theorizing, or even to set stronger constraints in advance of fitting models. Another potential future direction would be to probe people’s mental causal models more directly, either by asking them directly what they think about the causal relationships among variables or by asking them to reason through hypothetical interventions on variables. These more direct responses might then be used to constrain or augment the structure learning algorithm’s outputs.

Summary

From a set of potentially relevant beliefs and a general principle guiding the likely relationships among these beliefs, we generated a cognitive model of an intuitive theory of this domain. In Studies 2-4, we use this same model to predict how this intuitive theory guides reasoning about new information.

Study 2: Targeting disease severity beliefs

Only a handful of educational interventions have been shown to improve attitudes toward childhood vaccination. One of the successful educational interventions in this domain was tested by Horne and colleagues (2015), who found that emphasizing the dangers of preventable diseases like mumps, measles, and rubella improved participants' attitudes toward vaccines. Their intervention included descriptions of the symptoms of mumps, measles, and rubella; vivid photos of children suffering from these diseases; and a brief story from the perspective of a parent whose child was hospitalized with measles (see Supplementary Materials). This "disease risk" intervention improved attitudes toward vaccines compared to an irrelevant control essay and compared to an intervention aimed at debunking concerns about vaccines causing autism. In Study 2, we aimed to replicate this success and leverage the cognitive model developed in Study 1 to gain insight into how and why such an intervention successfully induces belief revision in this domain.

In Study 2, we used the cognitive model established in Study 1 to predict how Horne and colleagues' (2015) intervention would affect the beliefs in this network, including people's intentions to vaccinate their children. Participants completed the 14 belief scales from Study 1, then returned two days later to complete either the disease risk intervention or a no-intervention control, followed by the same 14 belief scales.

Method

Participants. We recruited 860 people to participate in a two-day study via mTurk in April 2018. Two days later, participants who completed Phase 1 and passed attention checks were invited via email to participate in Phase 2 of the study, with 559 participants returning (726 invited, retaining 77% of the Phase 1 sample). Participants were paid \$1.35 for their participation in each phase of the study. 54 participants (9.70%) failed at least one attention check on Phase 2 and were excluded from further analysis. We also excluded from analyses 12 additional participants who accessed the second survey more than once. This left a final sample of 493 participants.

Procedure. Procedures for Phase 1 of Study 2 were nearly identical to those of Study 1, but the debriefing statement was moved to the end of Phase 2. In Phase 2 (approximately two days later), participants completed a study very similar to Horne and colleagues' (2015) original intervention study. Participants were randomly assigned to either the disease risk condition or a no-intervention control condition, with the constraint that the distributions of participants' Phase 1 *vaccine intentions* scores should be similar across the two conditions; we accomplished this by stratifying participants into three groups based on their pretest vaccine intentions, and randomly assigning them to conditions in equal numbers within each group.

Participants ($n = 493$ passed attention checks in total) in the intervention conditions were presented with a version of Horne and colleagues' "disease risk" intervention, modified to present subparts in a fixed (rather than randomized) order, with slight changes to wording to improve clarity, and with a final statement summarizing the take-home message of the intervention; see Supplementary Materials. They then completed the same 14 belief scales and demographic questions that they had completed in Phase 1.

Participants in the no-intervention condition completed the belief scales and demographics questions without reading any material first.

Results

Figure 4 shows the change in participants’ attitudes across all 14 belief scales, from pretest to posttest. As expected, the largest change was observed for *disease severity* beliefs—the belief directly targeted by the intervention. Most importantly, the disease risk intervention significantly increased participant’s *vaccination intentions*, conceptually replicating earlier findings regarding this intervention (Horne et al., 2015).

We conducted a mixed-effects Bayesian ordinal regression to assess the effects of the disease risk intervention on responses to the *Vaccination Intentions* scale items, using the “brms” package for R (Bürkner, 2017). We regressed responses on phase (Phase 1 vs. 2, dummy-coded with Phase 1 as the baseline), condition (dummy-coded with the disease risk condition as the baseline, so as to assess difference scores in our condition of primary interest), and the interaction between phase and condition. We included random intercepts by subject and scale item. In the lme4 syntax, the model was:

```
response ~ phase * condition + (1|subject) + (1|item)
```

In this model, the unique effect of the intervention on vaccination intention scores is assessed by the interaction term (indicating the unique effect of Phase 2 for intervention versus no-intervention participants). The disease risk intervention increased participants’ vaccination intention ratings (interaction effect: $b = 0.36$, 95% Credible Interval (CI): [0.12, 0.60]; see Figure 4, panel A). We consider this to be a full replication of Horne and colleagues’ (2015) results.

Visual inspection of Figure 4, panel A, further reveals that changes were also observed across the belief network identified in Study 1. Not only did participants exposed to the intervention come to see childhood diseases as more severe, they also came to see vaccines as less dangerous, and demonstrated less medical skepticism as well as less commitment to “holistic balance”. These results suggest that targeted educational

interventions could have wide-ranging effects, from influencing who people turn to for future information, to shaping the more abstract worldviews they bring to bear in understanding health and science.

Figure 4

Belief changes following the intervention in Study 2. Left, mean difference scores from pretest to posttest for the 14 belief scales across experimental conditions. Scales are arranged in order of their predicted changes according to the cognitive model. Right, mean observed and predicted change scores from the cognitive model for the disease risk intervention condition. In both plots, error bars represent 95% confidence intervals. In both plots, changes were calculated after transforming responses to lie between 0 and 1.

Modeling belief revision. How did this intervention—which focused solely on the dangers of childhood diseases—affect beliefs about *vaccine danger*, *medical skepticism*, *holistic balance*, and so on? One explanation is that this intervention provides direct evidence for *disease severity* beliefs and that other beliefs are revised as a result of their relationships with *disease severity* as captured in the intuitive theory. To examine this, we used the same model developed in Study 1 to simulate how the 13 other beliefs in the

intuitive theory would change as a result of the intervention’s direct effect on *disease severity* beliefs.

To model the effects of the intervention presented in Study 2 on the intuitive theory surfaced in Study 1, we augmented the cognitive model’s network representing the intuitive theory with an additional node representing the information about measles, mumps, and rubella presented in the intervention. This evidence node was added as a child of the *disease severity* node, allowing the intervention to act as virtual evidence for *disease severity* (following Pearl, 1988).

To compute the conditional probability table (CPT) for this new node as a function of *disease severity* beliefs, we estimated an evidence ratio based on participants’ pretest and posttest credences in *disease severity*. To so do, we made use of the log-odds expression of Bayes rule, with the following regression equation:

$$\text{logit}(\text{posttest}) = \text{logit}(\text{pretest}) + \ln(ER_0) + x \ln(ER_1) \quad (2)$$

Where x was a binary variable representing the presence (1) or absence (0) of evidence at posttest, and participants’ prior and posterior (pretest and posttest) credences were transformed into log odds. The variables ER_0 and ER_1 represent measures of the evidence seen by participants in the control and intervention conditions. These are defined as:

$$ER_i = \ln \frac{P(d_i|h = 1))}{P(d_i|h = 0))} \quad (3)$$

In equation 3, ER_0 represents whatever background “evidence” participants would see between phases outside the experimental context, that would be shared by participants in both conditions. ER_1 represents the evidence provided by the intervention information, which is assumed to be independent from ER_0 . These log evidence ratios were both

estimated via regression. Then, we constructed a conditional probability table consistent with ER_1 and augmented the cognitive model with this evidence node (see Supplementary Materials for details).

Using this augmented model, we are able to “condition” our cognitive model on this evidence node, and thereby predict how observing evidence will affect all of the beliefs in the network. As shown in Figure 4, panel B, this simulation is a reasonable approximation of what we observed in Study 2: The model’s predictions correlate strongly with observed changes ($r = 0.870$).

Given that our cognitive model is parameterized via noisy-logical functions wherein each parent independently influences each child node, it should be noted that the marginal predicted influence of each belief is closely related to its first-order correlation with disease severity. This is in keeping with the qualitative features of intuitive theories we laid out in the introduction, and we should often expect this to hold true on average for any plausible model of an intuitive theory. However, future models with more complex or nuanced parameterizations could potentially make additional predictions not evident from first-order correlations, such as individualized predictions for how individuals or groups would respond to evidence.

Study 3a and 3b: Targeting beliefs about vaccine toxicity

The intervention developed by Horne and colleagues (2015) and employed in Study 2 is notable for being one of the few educational interventions shown to decrease vaccine hesitancy—but is disease severity the only (or most) effective target for educational interventions in this domain?

One advantage of having developed a formal, computational representation of a network of beliefs supporting vaccination decisions is that it can guide our search for effective interventions, allowing us to compare how interventions targeting each of these other beliefs might potentially change people’s vaccination intentions (see Figure 5).

Countering myths about toxic additives is one approach the cognitive model developed in Study 1 predicts should be effective.

Figure 5
Predicted change in vaccination intentions based on cognitive model following an equivalent hypothetical change in each of the other 13 beliefs.

We developed a novel educational intervention that addresses fears about toxic additives in vaccines. This “vaccine ingredients” intervention explained the common components of vaccines, emphasized the minute quantities of all vaccine ingredients, and explained the roles that various vaccine additives play in ensuring the stability and safety of vaccines, such as antibiotics, adjuvants, and stabilizers. We gave special attention to defusing concerns about the use of formaldehyde in vaccines because our assessment of anti-vaccination websites suggested that this was a particular concern for many vaccine skeptics. The intervention acknowledged that some vaccines contain trace amounts of formaldehyde, but pointed out that formaldehyde is a metabolic byproduct that occurs naturally in all living things. The intervention highlighted that formaldehyde is already present in children’s bodies (and even in foods like pears) in far larger quantities than

contained in any vaccine. The full intervention text is presented below.

“Vaccine ingredients” intervention text

Vaccines are one of the most important medical advances in history, but some parents worry that vaccines contain toxins that could harm their babies. In fact, all of the ingredients in a vaccine serve important purposes that help ensure the safety and effectiveness of the vaccine. Decades of research have shown vaccine ingredients to be safe, even in much larger quantities than those found in vaccines.

What’s in a vaccine?

Active ingredients

The key ingredient in all vaccines is one or more active ingredients, whose purpose is to train the immune system to recognize and combat the pathogens that cause diseases. To do this, certain molecules from the pathogen are introduced into the body to trigger an immune response. These molecules are called antigens, and they are present on all viruses and bacteria. By injecting these antigens into the body, the immune system can safely learn to recognize them as hostile invaders, produce antibodies, and remember them for the future.

Other Ingredients

A typical dose of an injected vaccine is just 0.5 millilitres of liquid, which is just a few drops. Apart from active ingredients, the main ingredient in vaccines is water. All other ingredients weigh a few thousandths of a gram or even less. All of the ingredients in vaccines have been proven to be safe for children in these minute quantities.

All of the ingredients in a vaccine serve important purposes, and are present to ensure that the vaccine is safe and effective. Some of these added ingredients are adjuvants, added to enhance the immune system response; others are antibiotics, to prevent contamination during the manufacturing process; and

still other ingredients act as preservatives and stabilizers.

To take one example, consider one of the scariest things you may have heard about the ingredients in vaccines: that some vaccines contain formaldehyde. It may surprise you to learn that formaldehyde is actually produced by the human body as part of normal metabolic processes. Human bodies—even babies’ bodies—naturally produce, process, and safely excrete formaldehyde. The amount of formaldehyde in vaccines is extremely small compared to what our bodies produce on their own: the amount of formaldehyde naturally present in a 2-month-old infant’s blood (around 1.1 milligrams) is ten times greater than the amount found in any vaccine (less than 0.1 milligrams). In addition, many healthy foods—such as apples, grapes, bananas, and pears—naturally contain formaldehyde. A pear, for example, contains around 50 times more formaldehyde than is found in any vaccine.

Methods

We conducted two preregistered experiments² (Study 3a and Study 3b) testing the effects of this novel intervention aimed at dispelling concerns about toxic additives in vaccines using the same pretest-posttest design as Study 2. After analyzing data from the first of these studies, we were surprised to observe that in several cases participants’ beliefs changed significantly from pretest to posttest among participants in the no-intervention control condition. For this reason, two months later we conducted a direct replication (Study 3b) in the hopes that any transient factors that might have caused these changes in beliefs might have dissipated, but we obtained very similar results (see Supplementary Materials).

These studies followed similar design procedures as Study 2, with Phase 2 conducted 1-2 days after Phase 1. In both Study 3a (April 2019) and Study 3b (June 2019;

² <https://osf.io/csnez/registrations>

combined final $n = 1026$) the vaccine ingredients intervention was compared against no-intervention control. Unlike Study 2, we did not stratify participants before random assignment to conditions.

Participants. We recruited 1491 people to participate in a two-day study via mTurk in April (Experiment 3a) and June (Experiment 3b) of 2019. One day after their initial participation, participants who completed part 1 and passed attention checks were invited via email to participate in Day 2 of the study, with 1077 participants returning (1379 invited, retaining 78.10% of the Day 1 sample). Participants were paid \$1.35 for their participation in each phase of the study and were allowed to participate in part 2 up to two days after their initial participation in part 1. 51 participants (4.70%) failed at least one attention check on Day 2 and were excluded from further analysis. This left a final sample of 1026 participants. see Supplementary Materials for demographic information about this sample.

Procedure. Procedures for Studies 3a and 3b closely followed those of Study 2, but these studies tested the effects of a novel intervention aimed at dispelling concerns about toxic additives in vaccines. In both studies, participants were randomly assigned to either a no-intervention control condition or to the novel “vaccine ingredients” intervention.

Results

Regression analyses. As they shared the same design, we combined data from Studies 3a and 3b for analysis (but see Supplementary Materials for separate analyses by study). As in Study 2, we conducted a mixed-effects Bayesian ordinal regression to assess the effects of the vaccine ingredients intervention on participants’ *vaccination intentions* scores.

As in Study 2, Study 3a and 3b used a difference-in-difference design. Thus the effect of the vaccine intervention is measured by testing the interaction between condition and phase. The vaccine ingredients intervention improved participants’ vaccination

intention scores relative to the control conditions (interaction effect: $b = 0.32$, 95% CI: [0.15, 0.49]); see Figure 6, panel A. It is worth noting that reliable changes were also observed from pretest to posttest in the control conditions (phase effect: $b = -0.18$, 95% CI: [-0.30, -0.06]), suggesting that these attitudes were also being influenced by other factors during the course of the study. Given that these studies were randomized controlled trial designs, these findings do not impair our ability to conclude that the intervention was causally efficacious in affecting vaccine intentions. It is important to emphasize that the test of the intervention’s effectiveness is a test of an interaction—an interaction demonstrating that the improvements in intentions in the intervention condition were greater than those in the control condition.

Modeling belief revision. To model the effect of the vaccine ingredients intervention, we followed the same procedure as in Study 2, but this time augmented the cognitive model with a node representing the vaccine ingredients intervention as a child of the vaccine toxins belief node (for details, see Supplementary Materials). This allowed us to predict how other beliefs in the network would change in light of the observed changes in participant’s beliefs about whether vaccines contain dangerous toxins; see Figure 6, panel B.

Collapsing across these two studies, Figure 6 shows the change in participants’ ($n = 1026$) attitudes across all 14 belief scales from pretest to posttest. Most importantly, the vaccine ingredients intervention significantly improved participants’ vaccination intentions compared to the no-intervention control condition. Changes were also evident across many of the 14 beliefs, with the largest changes being a reduction in vaccine toxicity beliefs.

As in Study 2, we used our model to predict how the 13 other beliefs in the intuitive theory should be expected to change given the average magnitude of the change observed in participants’ vaccine toxicity beliefs following the intervention. As shown in Figure 6, the model’s predictions are significantly related to the observed changes, but the correlations are relatively weaker ($r = 0.696$) than those observed in our other studies.

Figure 6

Belief changes following the intervention in studies 3a and 3b. Left, mean difference scores from pretest to posttest for the 14 belief scales across experimental conditions. Scales are arranged in order of their predicted changes according to the cognitive model. Right, mean observed and predicted change scores from the cognitive model for the disease risk intervention condition. In both plots, changes were calculated after combining data from both studies and transforming responses to lie between 0 and 1 and error bars represent 95% confidence intervals.

We speculate that changes from pretest to posttest among control participants in both experiments were related to media coverage of measles outbreaks in the U.S. during the time these studies were conducted (April and June of 2019), so that participants' beliefs were not only affected by our intervention between pretest and posttest, but also potentially by additional news stories or social media content related to the outbreaks. As we employed a pretest-posttest design and randomly assigned participants to conditions, these changes do not impair our ability to draw causal inferences as to the efficacy of our intervention. However, these additional unmodeled influences may be negatively impacting

performance of the model’s predictions.

Nevertheless, the success of this intervention in changing people’s responses to our *vaccine intentions* scale highlights the potential of this computational, theory-based approach to understanding belief revision. This model allowed us to design a novel educational intervention focusing on vaccine ingredients that successfully changed participants’ beliefs in a domain where many previously-researched educational approaches have failed—while sparing the time and resources we might have otherwise expended on interventions which our model predicts are unlikely to be successful.

Study 4: Natural experiment

The early months of 2019 saw the worst U.S. measles outbreak in decades, with over 1000 cases reported by June, mostly in the state of New York. These outbreaks were covered in every major newspaper and by every major news program. Would learning about these tragic events lead people to revise their beliefs about vaccines?

Methods

To examine this possibility, we invited participants from Study 1 (conducted in August-September 2017) to participate in a new study measuring each of those 14 beliefs again (conducted from May 1st to 3rd, 2019; approximately 18 months later). Of the 1129 participants in Study 1, a total of 549 participants returned and 530 successfully completed the study, including passing attention checks. Approximately 76% of participants reported recently seeing, hearing, or reading about measles outbreaks. See Supplementary Materials for demographic information about this sample.

Results and discussion

Figure 7, panel A, shows the average change in participants’ beliefs following the outbreaks of 2019. Among participants who were exposed to relevant news covered, we observed changes in a number of beliefs, with notable increases in intentions to vaccinate,

Figure 7

Participants' belief changes from participation in Study 1 to Study 4. Left, mean difference scores comparing beliefs reported in August-September 2017 and May 2019 for the 14 belief scales, across participants who did and did not report exposure to news of the ongoing New York measles outbreak. Scales are arranged in order of their predicted changes according to the cognitive model. Right, mean observed and predicted change scores from the cognitive model for those participants who did report exposure to this news. Belief changes were modeled assuming that news of the outbreak directly impacted participant's vaccine intentions, the most central of the 14 beliefs, in order to evaluate the overall coherence of these belief changes. In both plots, changes were calculated after transforming responses to lie between 0 and 1 and error bars represent 95% confidence intervals.

perceptions of vaccine efficacy, and perceptions of disease severity, as well as notable decreases in concern about vaccine danger. In contrast, most beliefs remained unchanged for those participants who did not report exposure to media, although there were modest changes in their beliefs about disease severity and disease rarity. These outbreaks appear to have significantly impacted people's beliefs surrounding vaccination.

Mirroring our analyses of interventions, we regressed responses to the *vaccine intentions* scale items on predictors for phase (August-September 2017 vs. May 2019, dummy-coded), news exposure (dummy-coded with non-exposure as the baseline), and an interaction between phase and news exposure onto participants' scores, using a Bayesian mixed-effects ordinal regression including random intercepts by subject and item. To assess the effects of recent news exposure on vaccine intentions, we examined the interaction between phase and news exposure. Participants' vaccine intentions were increased by recent news exposure compared to their responses 18 months earlier (interaction effect: $b = 0.48$, 95% CI: [0.22, 0.74]); see Figure 7, panel A.

Modeling belief revision. Were changes in participants' beliefs following exposure to news of the measles outbreak coherent with respect to the intuitive theory specified by the cognitive model developed in Study 1? Unlike in the intervention studies, where a specific belief was quite precisely targeted by the intervention, it is likely that many beliefs were addressed by media coverage of these outbreaks. Given that we cannot measure the exact information that participants were exposed to, it is not feasible to model the precise locus of its impacts. However, we can leverage our model to assess whether the changes in participants' beliefs were *coherent*.

Throughout our prior studies, we have focused primarily on *vaccine intentions* as the most crucial and proximal belief for determining vaccination decisions. So, we sought to examine whether the changes across all the measured beliefs were coherent with respect to the changes in vaccination intentions. Using our model, we generated predictions for how the 13 other beliefs in the intuitive theory should be expected to change given the

average magnitude of the change observed in vaccine intentions. To accomplish this, we followed the same procedure as in previous studies, but here augmented the model with a node representing media exposure as a child of the *vaccine intentions* node. This provides an evaluation of coherence similar to prior work on coherence-based reasoning, where the coherence of belief revision is evaluated with respect to some emerging decision, evaluation, or change in one specific belief or attitude (e.g. Holyoak & Simon, 1999; Simon et al., 2004; Spellman et al., 1993).

As shown in Figure 7, panel B, the changes in participants' beliefs are very similar to the model's predictions ($r = 0.955$), suggesting that the changes in their beliefs were highly coherent with respect to their intuitive theories of the domain. Modeling these beliefs as united by an intuitive theory not only captures how these beliefs are revised immediately in light of evidence, but also how they are changed by real-world events.

Discussion

Across five studies, we examined how understanding people's intuitive theories surrounding vaccination decisions could help us formulate educational interventions that would counteract misconceptions about vaccine safety and encourage vaccination. We identified 14 beliefs related to vaccination decisions and captured the relationships among these beliefs in a generative probabilistic cognitive model. Using this model we were able to predict how these beliefs would be affected by controlled educational interventions as well as real-world events. Inspired by our findings and model, we developed an intervention targeting beliefs about toxins in vaccines and found it successfully increased people's intentions to vaccinate their children—adding to the very short list of effective interventions in this domain.

Our findings indicate that beliefs related to vaccination decisions can be understood as situated within a coherent intuitive theory: these beliefs were systematically related across individuals (Study 1), and these relationships were predictive of changes across these

beliefs following evidence, whether that evidence took the form of educational interventions (Studies 2, 3a, 3b) or real-world events (Study 4). Modeling this intuitive theory formally as a Bayesian network allowed us to create a cognitive model that demonstrates these key features of intuitive theories quantitatively. Looking forward, this formal representation—though simplified and imperfect—lays a foundation for richer and more complex descriptions in the future.

We tested two educational interventions: one emphasizing the dangers of childhood diseases and another debunking concerns about toxins in vaccines. Both were effective in increasing people’s intentions to vaccinate. How might these educational interventions affect vaccine uptake if deployed in the real world? We are somewhat cautious in our optimism: In conventional terms of standardized effect sizes, both intervention effects would be considered quite modest and their impacts on actual future decisions remain unknown. Still, we nevertheless see both of these interventions as promising foundations for real-world educational efforts. First, as they are quite brief, these interventions present a very low cost for delivery. In addition, their brevity supports opportunities for deploying them with strategic timing and from trusted sources. For instance, these arguments could be deployed in the context of visits with family physicians, including immediately before vaccination decisions are to be made. Delivering these messages at the point-of-decision could maximize their impacts. And, we predict their efficacy would be improved if they were delivered by a trusted family physician, rather than coming from a relatively anonymous online survey.

We credit our development of a quantitative model of the intuitive theory surrounding vaccination decisions as supporting our work as interventionists. By providing a graphical representation of the structure of this intuitive theory, this approach helped us to understand relationships between beliefs that were puzzling when considered only as correlations (e.g., the lack of correlation between beliefs about the infant immune system and vaccine intentions), and allowed us to “spot” opportunities to develop novel

interventions. One exciting future direction for these and other brief interventions would be to tailor these interventions to an individual's beliefs: Models with richer parameterizations could simulate belief change for an individual given a certain set of prior beliefs, paving the way for determining which specific intervention or combination of interventions would have the most impact for a particular decision-maker.

More broadly, the approaches we have introduced here may shed light on long-standing theoretical questions about belief revision. Many researchers have argued that people's belief updating is subject to various biases, or is skewed by motivational factors (Flynn et al., 2017; e.g. Kahan, 2013; Kunda, 1990). Yet many of the phenomena traditionally taken as evidence for these theories can also be explained by normative models of belief revision, so long as we are willing to posit that other auxiliary beliefs can affect the relationships between a belief and evidence (e.g., Gershman, 2019; Jern et al., 2014). Adjudicating between competing views of belief revision will require us to move beyond minimal models and illustrative simulations by assessing whether people actually do hold these auxiliary beliefs, and whether their broader belief systems behave according to the rules of probabilistic inference. The Bayesian network structure learning techniques applied here are one tool we might use to build these richer models.

Implications for and insights from the COVID-19 pandemic

Though these studies were conducted before the onset of the COVID-19 pandemic and focused on childhood diseases, many similar concerns are likely to animate hesitancy toward COVID-19 vaccines, and our findings may provide some leverage into addressing the current surge of vaccine hesitancy.

First, in the spirit of Horne and colleagues' (2015) disease risk intervention, we suspect that education focused on how severely ill many people become and the prevalence of long-term health problems following even mild cases of COVID-19 might encourage people to accept coronavirus vaccines. Indeed, controlled intervention studies have shown

that messages focused on personal risk increase vaccination intentions (Motta et al., 2021). However, the severity of COVID-19 is highly variable across individuals and across variants of the virus, potentially rendering the formation, updating, and use of beliefs about severity much more nuanced in the case of COVID-19 than in the case of diseases like measles.

Our findings also suggest that beliefs about vaccine efficacy are central to determining vaccination intentions. The high prevalence and mutation rate of COVID-19 makes the risk of infection substantial even for the vaccinated, although these vaccines remain highly effective at preventing severe illness and death. The rising prevalence of breakthrough infections and the need for boosters relatively soon after initial doses are likely to lead people to doubt the effectiveness of vaccines—but we see promise in educational interventions that clarify the many ways in which a vaccine can be “effective” (e.g., by preventing infection vs. by preventing severe illness, preventing death, or lowering transmissibility).

Looking forward, living through the COVID-19 pandemic is likely to have deepened the public’s understanding of infectious diseases and public health. For example, the concept of “herd immunity” has likely never had greater public visibility and appreciation than now. This understanding seems like a prime target for future interventions aimed at increasing the uptake of vaccines for both COVID and other diseases.

Finally, like all pandemics, the COVID-19 pandemic is inescapably social. In the U.S. and abroad, COVID-19 vaccines and mitigation measures have become deeply politicized. According to Gallup polling in the U.S., COVID-19 vaccination rates have at times been nearly twice as high among Democrats as among Republicans (Gallup, 2021). Messages from political elites often imply concrete beliefs about the world that could be incorporated into a model of an intuitive theory: for instance, that COVID-19 is “like the flu” or that infection can be treated with off-label pharmaceuticals—or, conversely, that COVID-19 vaccines are safe and effective. Still, these affiliations are also associated with

social commitments and personal values that may behave differently than the beliefs examined here. Whereas beliefs about “matters of fact” might be driven by an underlying intuitive theory and thereby sensitive to evidence, these sorts of identity-constituting “beliefs” may operate by different principles (Van Leeuwen, 2014), such as being driven primarily by people’s affinity toward their social groups rather than by concrete external evidence (Kahan, 2016).

More broadly, people from different social groups have different experiences and knowledge that are likely to shape their intuitive theories of vaccination and the medical establishment in important ways. For example, White vs. Black Americans are likely to have different experiences with medical racism. Future work should explore how differences in people’s wider intuitive theories might explain vaccination behaviors, as well as how a more nuanced understanding of these theories might support messages tailored to different individuals and communities.

Broader implications and applications

The current studies add to a growing body of work demonstrating how leveraging intuitive theories can spur belief revision and behavior change in the domain of physical health (cf. Weisman & Markman, 2017). It also introduces a new approach to modeling and understanding intuitive theories more broadly. In prior work, researchers focused on identifying a small set of beliefs or concepts that were core to the intuitive theory and specific attitudinal or behavioral outcomes. For example, in previous work on vaccination attitudes, researchers focused on how people balanced perceptions of the risk of vaccine side effects and the risk of childhood diseases to inform their intentions to vaccinate their children (Horne et al., 2015). The approach we introduced here expands the set of beliefs considered beyond a small core to a wider set of auxiliary beliefs that may influence decisions. This is made feasible by leveraging a bottom-up approach that aims to discover the network structure connecting those wide-ranging beliefs.

Here we speculate about how this kind of approach might be applied to the topics of climate change and racial justice: complex domains in which variability in people's beliefs leads to serious societal consequences, and in which widespread misconceptions demand intervention. In both cases, we hypothesize that people's behaviors are shaped by a wide-ranging network of interconnected beliefs that might be better understood through approaches similar to those developed here.

Consider first the domain of climate change: Without swift and drastic efforts to reduce CO₂ emissions, global climate change threatens to increase extreme weather events, disrupt agriculture, destroy ecosystems, and displace millions of people over the next few decades. Climate change is a complex geophysical, economic, political, and social phenomenon. Many changes, both individual and structural, are likely needed to effectively reduce CO₂ emissions and mitigate the worst effects of climate change. Our work suggests that efforts to understand and encourage any of these behaviors might be best-served by seeking to understand the broader web of beliefs that sustain and support them.

Small models of beliefs related to climate change with Bayesian Networks have already yielded some promising results. Cook and Lewandowsky (2016) developed a model of how a core free-market "worldview" affects beliefs and belief revision about whether human-caused climate change is really happening. Schotsch and Powell (2022) explored people's causal mental models of how CO₂ emissions cause global warming and how warming in turn causes different outcomes such as rising sea levels and extreme weather. Elaborating on these small models to discover and model a wider network of related beliefs has great potential for crafting effective interventions in this domain.

We might ask: How do people understand the science connecting CO₂ emissions to warming and warming to other negative consequences and do they trust this science? How do people weigh the efficacy of their own individual behaviors versus government regulations? How much do they trust regulators or specific regulations to be effective and

how concerned are they about their economic side-effects? How do they expect climate-driven extreme weather, crop losses, economic disruptions, and mass migration will affect them personally? What obligations do they feel toward future generations or even their own children for the kind of world they leave behind? What value do they place on preserving the natural world and its rich ecosystems and biodiversity?

The complexity of this domain will likely present a challenge to cataloging even a wide auxiliary set of beliefs that could fully circumscribe it. Thus the goal need not be to identify an authoritative and all-encompassing “intuitive theory” of climate change. Rather, interventionists might be better served by developing more constrained models of intuitive theories that bear directly on a specific target behavior of interest.

Intuitive theories may also inform people’s reasoning about social and political issues. To highlight one important example, following the police killing of George Floyd in 2020, protests for racial and social justice organized by groups like Black Lives Matter spread across the U.S. and elsewhere. Across different people and groups of people these events spurred a wide range of reactions: Some called for police departments to be defunded, others for police funding to be increased. Some sought to educate themselves by purchasing anti-racist books, others began pressuring school boards to eliminate purported elements of “critical race theory” from K-12 curricula. Some were freshly or further awakened to and disquieted by racial injustice in American society, others were persuaded by the white supremacist “great replacement” conspiracy theory.

One way to try to understand these vastly different reactions is to imagine that they were shaped by large, complex networks of beliefs about the nature of race; about American history and politics; about policing; about justice, equity, and other moral principles; and so on. The approach outlined in the current studies could provide a useful tool for navigating this complexity, yielding insights into the interrelations among these beliefs.

To draw a few examples from the vast literature on American racism: How does

understanding of the historical impacts of slavery, reconstruction, the post reconstruction era and the civil rights era relate to beliefs about the racial inequalities that persist in the present day? How does this historical awareness in turn relate to whether people conceptualize racism as residing in beliefs, attitudes, and behaviors of individuals or as embedded in societal structures, institutions, and practices? How do construals of race as a biological trait versus as a social construct impact thoughts on these issues? And how does the conceptualization of racism as an individual or structural problem encourage people to address racism actively or passively ignore it? It is our hope that mapping out connections among these and other beliefs could reveal new avenues for correcting people's misconceptions and promoting mutual understanding that could spur anti-racist actions and social change.

Conclusions

Not all human reasoning need be governed by intuitive theories, but intuitive theories are likely to be crucial wherever people must integrate information to evaluate, explain, or make decisions. Intuitive theories are powerful cognitive tools, for better and for worse. Coherent networks of beliefs allow people to understand, predict, and navigate the world, but they can also sustain misconceptions and lead to suboptimal—even dangerous—decisions and behaviors. The present work on understanding vaccine hesitancy suggests a general theoretical approach to understanding belief revision more broadly: by mapping out a set of beliefs relevant to a given domain, researchers and educators might identify optimal avenues for persuasion and education. Further, by formalizing intuitive theories into quantitative models, our approach supports the use of bottom-up approaches to discover structure among a wider set of auxiliary beliefs and provide more precise tests of theories of belief revision. We hope the studies presented here have shed light on how people might come to form the beliefs and misconceptions that they do, why those misconceptions can be so hard to combat, and how to develop educational interventions

that are persuasive and effective in encouraging and sustaining positive behaviors.

References

- Baillargeon, R. (2004). Infants' Physical World. *Current Directions in Psychological Science*, 13(3), 89–94. <https://doi.org/10.1111/j.0963-7214.2004.00281.x>
- Baker, C. L., Jara-Ettinger, J., Saxe, R., & Tenenbaum, J. B. (2017). Rational quantitative attribution of beliefs, desires and percepts in human mentalizing. *Nature Human Behaviour*, 1(4), 0064. <https://doi.org/10.1038/s41562-017-0064>
- Battaglia, P. W., Hamrick, J. B., & Tenenbaum, J. B. (2013). Simulation as an engine of physical scene understanding. *Proceedings of the National Academy of Sciences*, 110(45), 18327–18332. <https://doi.org/10.1073/pnas.1306572110>
- Betsch, C., & Sachse, K. (2013). Debunking vaccination myths: Strong risk negations can increase perceived vaccination risks. *Health Psychology*, 32(2), 146–155. <https://doi.org/10.1037/a0027387>
- Bürkner, P.-C. (2017). Brms: An R Package for Bayesian Multilevel Models Using Stan. *Journal of Statistical Software*, 80(1), 1–28. <https://doi.org/10.18637/jss.v080.i01>
- Carey, S. (1985). *Conceptual change in childhood*. MIT Press.
- Carey, S. (2009). The Origin of Concepts. In *The Origin of Concepts*. Oxford University Press.
- Cheng, P. W. (1997). From covariation to causation: A causal power theory. *Psychological Review*, 104(2), 367–405. <https://doi.org/10.1037/0033-295X.104.2.367>
- Cook, J., & Lewandowsky, S. (2016). Rational Irrationality: Modeling Climate Change Belief Polarization Using Bayesian Networks. *Topics in Cognitive Science*, 8(1), 160–179. <https://doi.org/10.1111/tops.12186>
- Dubé, E., Vivion, M., & MacDonald, N. E. (2015). Vaccine hesitancy, vaccine refusal and the anti-vaccine movement: Influence, impact and implications. *Expert Review of Vaccines*, 14(1), 99–117.

<https://doi.org/10.1586/14760584.2015.964212>

Duhem, P. M. M. (1954). *The aim and structure of physical theory*. Princeton University Press.

Flynn, D. J., Nyhan, B., & Reifler, J. (2017). The Nature and Origins of Misperceptions: Understanding False and Unsupported Beliefs About Politics: Nature and Origins of Misperceptions. *Political Psychology*, 38, 127–150. <https://doi.org/10.1111/pops.12394>

Gagneur, A., Gosselin, V., & Dubé, È. (2018). Motivational interviewing: A promising tool to address vaccine hesitancy. *Vaccine*, 36(44), 6553–6555. <https://doi.org/10.1016/j.vaccine.2017.10.049>

Gallup. (2021). *More in U.S. Vaccinated After Delta Surge, FDA Decision*. Gallup. <https://news.gallup.com/poll/355073/vaccinated-delta-surge-fda-decision.aspx>

Gelman, S. A., & Legare, C. H. (2011). Concepts and Folk Theories. *Annual Review of Anthropology*, 40(1), 379–398. <https://doi.org/10.1146/annurev-anthro-081309-145822>

Gershman, S. J. (2019). How to never be wrong. *Psychonomic Bulletin & Review*, 26(1), 13–28. <https://doi.org/10.3758/s13423-018-1488-8>

Gerstenberg, T., Goodman, N. D., Lagnado, D. A., & Tenenbaum, J. B. (2021). A counterfactual simulation model of causal judgments for physical events. *Psychological Review*, 128(5), 936–975. <https://doi.org/10.1037/rev0000281>

Gopnik, A., Glymour, C., Sobel, D. M., Schulz, L. E., Kushnir, T., & Danks, D. (2004). A Theory of Causal Learning in Children: Causal Maps and Bayes Nets. *Psychological Review*, 111(1), 3–32. <https://doi.org/10.1037/0033-295X.111.1.3>

Gopnik, A., & Wellman, H. M. (1994). The theory theory. In L. A. Hirschfeld & S. A. Gelman (Eds.), *Mapping the Mind* (First, pp. 257–293). Cambridge University Press. <https://doi.org/10.1017/CBO9780511752902.011>

Griffiths, T. L., & Tenenbaum, J. B. (2005). Structure and strength in causal

- induction. *Cognitive Psychology*, 51(4), 334–384.
<https://doi.org/10.1016/j.cogpsych.2005.05.004>
- Holyoak, K. J., & Powell, D. (2016). Deontological coherence: A framework for commonsense moral reasoning. *Psychological Bulletin*, 142(11), 1179–1203.
<https://doi.org/10.1037/bul0000075>
- Holyoak, K. J., & Simon, D. (1999). Bidirectional Reasoning in Decision Making by Constraint Satisfaction. *Journal of Experimental Psychology: General*, 29.
- Horne, Z., Powell, D., Hummel, J. E., & Holyoak, K. J. (2015). Countering antivaccination attitudes. *Proceedings of the National Academy of Sciences*, 112(33), 10321–10324. <https://doi.org/10.1073/pnas.1504019112>
- Hornsey, M. J., Harris, E. A., & Fielding, K. S. (2018). The psychological roots of anti-vaccination attitudes: A 24-nation investigation. *Health Psychology*, 37(4), 307–315. <https://doi.org/10.1037/hea0000586>
- Jern, A., Chang, K. K., & Kemp, C. (2014). Belief polarization is not always irrational. *Psychological Review*, 121(2), 206–224.
<https://doi.org/10.1037/a0035941>
- Kahan, D. M. (2016). The Politically Motivated Reasoning Paradigm, Part 1: What Politically Motivated Reasoning Is and How to Measure It. In R. A. Scott & S. M. Kosslyn (Eds.), *Emerging Trends in the Social and Behavioral Sciences* (First, pp. 1–16). Wiley. <https://doi.org/10.1002/9781118900772.etrds0417>
- Kahan, D. M. (2013). Ideology, motivated reasoning, and cognitive reflection. *Judgment and Decision Making*, 8(4), 19.
- Keil, F. C. (1994). The birth and nurturance of concepts by domains: The origins of concepts of living things. In L. A. Hirschfeld & S. A. Gelman (Eds.), *Mapping the Mind* (First, pp. 234–254). Cambridge University Press.
<https://doi.org/10.1017/CBO9780511752902.010>
- Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin*, 108(3),

- 480–498. <https://doi.org/10.1037/0033-2909.108.3.480>
- Larson, H. J., Jarrett, C., Schulz, W. S., Chaudhuri, M., Zhou, Y., Dube, E., Schuster, M., MacDonald, N. E., & Wilson, R. (2015). Measuring vaccine hesitancy: The development of a survey tool. *Vaccine*, *33*(34), 4165–4175. <https://doi.org/10.1016/j.vaccine.2015.04.037>
- Lewandowsky, S., Ecker, U. K. H., Seifert, C. M., Schwarz, N., & Cook, J. (2012). Misinformation and Its Correction: Continued Influence and Successful Debiasing. *Psychological Science in the Public Interest*, *13*(3), 106–131. <https://doi.org/10.1177/1529100612451018>
- Maaskant, P. P., & Druzdzel, M. J. (2008). An independence of causal interactions model for opposing influences. *Proc. 4th European Workshop on Probabilistic Graphical Models*, 185–192.
- 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>
- Medin, D. L., & Atran, S. (Eds.). (1999). *Folkbiology*. MIT Press.
- Motta, M., Sylvester, S., Callaghan, T., & Lunz-Trujillo, K. (2021). Encouraging COVID-19 Vaccine Uptake Through Effective Health Communication. *Frontiers in Political Science*, *3*. <https://www.frontiersin.org/article/10.3389/fpos.2021.630133>
- Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, *92*(3), 289–316. <https://doi.org/10.1037/0033-295X.92.3.289>
- Opel, D. J., Heritage, J., Taylor, J. A., Mangione-Smith, R., Salas, H. S., DeVere, V., Zhou, C., & Robinson, J. D. (2013). The Architecture of Provider-Parent Vaccine Discussions at Health Supervision Visits. *PEDIATRICS*, *132*(6), 1037–1046. <https://doi.org/10.1542/peds.2013-2037>

- Pearl, J. (1988). *Probabilistic reasoning in intelligent systems*. Morgan Kaufmann.
- Powell, D., Weisman, K., & Markman, E. M. (2018). *Articulating lay theories through graphical models: A study of beliefs surrounding vaccination decisions*. 6.
- Quine, W. V. (1951). Two Dogmas of Empiricism. *The Philosophical Review*, 60(1), 20–43. <https://doi.org/10.2307/2181906>
- Quine, W. V., & Ullian, J. S. (1978). *The web of belief* (Second edition). Random House.
- Salmon, D. A., Dudley, M. Z., Glanz, J. M., & Omer, S. B. (2015). Vaccine Hesitancy. *American Journal of Preventive Medicine*, 49(6), S391–S398. <https://doi.org/10.1016/j.amepre.2015.06.009>
- Scutari, M. (2010). Learning Bayesian Networks with the bnlearn R Package. *arXiv:0908.3817 [Stat]*. <https://arxiv.org/abs/0908.3817>
- Scutari, M., & Denis, J.-B. (2021). *Bayesian networks: With examples in R* (Second edition). CRC Press.
- Simon, D., Snow, C. J., & Read, S. J. (2004). The Redux of Cognitive Consistency Theories: Evidence Judgments by Constraint Satisfaction. *Journal of Personality and Social Psychology*, 86(6), 814–837. <https://doi.org/10.1037/0022-3514.86.6.814>
- Spelke, E. (1994). Initial knowledge: Six suggestions. *Cognition*, 50(1), 431–445. [https://doi.org/10.1016/0010-0277\(94\)90039-6](https://doi.org/10.1016/0010-0277(94)90039-6)
- Spellman, B. A., Ullman, J. B., & Holyoak, K. J. (1993). A Coherence Model of Cognitive Consistency: Dynamics of Attitude Change During the Persian Gulf War. *Journal of Social Issues*, 49(4), 147–165. <https://doi.org/10.1111/j.1540-4560.1993.tb01185.x>
- Thagard, P. (1989). Explanatory coherence. *Behavioral and Brain Sciences*, 12(3), 435–467. <https://doi.org/10.1017/S0140525X00057046>
- Thornton, J. (2019). Measles cases in Europe tripled from 2017 to 2018. *BMJ*, 1634.

- <https://doi.org/10.1136/bmj.l634>
- Ullman, T. D., Spelke, E., Battaglia, P., & Tenenbaum, J. B. (2017). Mind Games: Game Engines as an Architecture for Intuitive Physics. *Trends in Cognitive Sciences*, 21(9), 649–665. <https://doi.org/10.1016/j.tics.2017.05.012>
- Van Leeuwen, N. (2014). Religious credence is not factual belief. *Cognition*, 133(3), 698–715. <https://doi.org/10.1016/j.cognition.2014.08.015>
- Vannice, K. S., Salmon, D. A., Shui, I., Omer, S. B., Kissner, J., Edwards, K. M., Sparks, R., Dekker, C. L., Klein, N. P., & Gust, D. A. (2011). Attitudes and Beliefs of Parents Concerned About Vaccines: Impact of Timing of Immunization Information. *Pediatrics*, 127(Supplement 1), S120–S126. <https://doi.org/10.1542/peds.2010-1722R>
- Wallace, C., Leask, J., & Trevena, L. J. (2006). Effects of a web based decision aid on parental attitudes to MMR vaccination: A before and after study. *BMJ*, 332(7534), 146–149. <https://doi.org/10.1136/bmj.38678.681840.68>
- Weisman, K., & Markman, E. M. (2017). Theory-based explanation as intervention. *Psychonomic Bulletin & Review*, 24(5), 1555–1562. <https://doi.org/10.3758/s13423-016-1207-2>
- Wellman, H. (2015). Theory of Mind. In *Emerging Trends in the Social and Behavioral Sciences* (pp. 1–17). American Cancer Society. <https://doi.org/10.1002/9781118900772.etrds0360>
- Wellman, H. M., & Gelman, S. A. (1992). Cognitive Development: Foundational Theories of Core Domains. *Annual Review of Psychology*, 43(1), 337–375. <https://doi.org/10.1146/annurev.ps.43.020192.002005>
- Williams, S. E. (2014). What are the factors that contribute to parental vaccine-hesitancy and what can we do about it? *Human Vaccines & Immunotherapeutics*, 10(9), 2584–2596. <https://doi.org/10.4161/hv.28596>
- Williamson, J. (2001). Bayesian Networks for Logical Reasoning. In C. Gomez & T.

Walsh (Eds.), *Proceedings of the AAAI Fall Symposium on using Uncertainty within Computation* (pp. 136–143).