Competing Narratives in Action: An Empirical Analysis of Model Adoption Dynamics*

Marco Angrisani† Anya Samek‡ Ricardo Serrano-Padial§

May 17, 2023

Abstract

We use a longitudinal dataset measuring beliefs and behaviors to study the dynamics of model – or *narrative* – adoption during the Covid-19 pandemic. We show that individuals switch beliefs about the effectiveness of preventive behaviors following changes in perceived risk. The adoption of models promoting preventive behaviors is procyclical and model switching is influenced by exposure to conflicting information. We explain the data using a heterogeneous-agent model of competing narratives in which agents exhibit motivated beliefs. Adopting misspecified models increases infection rates, highlighting the importance of promoting accurate beliefs to guide behavior in the presence of novel risks.

*We thank Aislinn Bohren, Antonio Cabrales, Juan Carlos Escanciano, William Fuchs, Nagore Iriberri, Juuso Toikka and audiences at UPenn, University of the Basque Country, Universidad Carlos III, and the 2022 Stony Brook Game Theory Conference for helpful comments and suggestions.

†University of Southern California

‡University of California San Diego & NBER

§Drexel University

1 Introduction

Novel risks are becoming increasingly common. From climate change to global pandemics, these risks have far-reaching effects on individuals and on entire economies. A growing theoretical literature describes how individuals adopt models of the world – or *narratives* – to guide their behavior in this uncertain landscape. These narratives provide a causal link between behaviors and outcomes in the presence of changing information and evolving risks. By shaping individual behavior, narratives also play a critical role in the evolution of the risks themselves.

Despite the growing theoretical literature on narrative selection, there is no empirical evidence on how individuals actually adopt narratives in the real world. This is partly because it is challenging to obtain longitudinal data on beliefs. In this paper, we aim to bridge this gap by studying narrative adoption dynamics empirically. To do so, we use the Covid-19 pandemic as a case study. We take advantage of a longitudinal survey that collected data from a representative panel of individuals in the United States (U.S.) on behaviors they engaged in to avoid Covid-19 infection, as well as their beliefs about the effectiveness of these behaviors.

The Covid-19 pandemic provides a compelling case study of the challenges that individuals face when choosing which narratives to adopt under emergent risks. During the pandemic, individuals were exposed to conflicting information as well as fluctuating prevalence of infection. From the early stages of the pandemic, there was a debate over the effectiveness of preventive behaviors such as wearing masks and social distancing. The Centers for Disease Control (CDC) in the U.S. later admitted that public health guidance was "confusing and overwhelming."[1] Further, the proliferation of conflicting perspectives across media outlets made it difficult for people to identify and comply with effective behaviors for curbing the spread of the virus (Bursztyn et al., 2020).

In this paper, we make three main contributions. First, we document the dynamics of adopting narratives about the effectiveness of preventive behaviors against novel risks. Specifically, we observe substantial narrative switching over time driven by changes in perceived risk and exposure to conflicting information. Individuals switch between correct beliefs that preventive behaviors curb infection and incorrect beliefs that preventive behaviors are ineffective. Such narrative switching cannot easily be explained by standard learning and belief updating models. Therefore, our second contribution is to propose and calibrate a heterogeneous-agent model of competing narratives based on Eliaz and Spiegler (2020). Our model fits the data remarkably well. Third, we quantify the welfare impact of incorrect beliefs by estimating the counterfactual drop in infection rates were all agents to adopt the correct narrative.

[1] See, for example, the New York Times article https://www.nytimes.com/2022/08/17/us/politics/cdc-rochelle-walensky-covid.html.

For our analyses, we rely on data from the Understanding Coronavirus in America Study (UCAS). This is a unique longitudinal survey administered throughout the pandemic within the Understanding America Study (UAS), a nationally-representative online panel of adults in the U.S. The UCAS surveyed nearly 9,000 individuals every 2-4 weeks between March 2020 and July 2021. Crucially, the survey period covered the first three waves of the pandemic, thus providing rich, cyclical variation in infection risk over time. The UCAS questionnaire repeatedly asked individuals about their perceived risk of Covid-19 infection and their beliefs about the effectiveness of different preventive behaviors.

We find that a majority of individuals switch back and forth between believing that a preventive behavior (e.g., avoiding restaurants and bars) is effective in reducing infection risk and believing that the preventive behavior is ineffective. Belief in the effectiveness of the behavior is procyclical – as infection risk goes up (down), the share of individuals who believe in the effectiveness of the behavior also goes up (down). Adherence to preventive behaviors is also procyclical and is driven to a large extent by the belief in their effectiveness.

Exposure to conflicting information partly drives narrative switching. The propensity to switch narratives is higher for individuals who were exposed both to a conservative news outlet such as Fox News (which promoted the notion that preventive behaviors were ineffective) and public health officials like the CDC and the World Health Organization (which promoted the effectiveness of preventive behaviors) than for individuals who were exposed to only one type of news source. Narrative switching is also higher for individuals who identify as Republicans than those who identify as Democrats.

Under standard learning and belief updating models, we would expect beliefs about the effectiveness of preventive behaviors to converge as evidence accumulates over time. In contrast, our data show that changes in perceived risk induce narrative switching. To explain this empirical pattern, we propose a theoretical model in which agents have motivated beliefs and adopt narratives based partly on their incentives. In our model, agents choose between two competing narratives to maximize their anticipated utility. According to the "effective" narrative, both the prevalence of the virus and the adoption of preventive behaviors affect the probability of becoming infected. According to the "ineffective" narrative, behaviors do not affect infection risk and, thus, should not be adopted. In this context, an agent would adhere to a preventive behavior if the benefit from reducing infection risk is higher than the cost of adopting the behavior. Both infection risk and the cost of behavior adoption are heterogeneous in the model. We show that the fraction of agents choosing the effective narrative goes up with virus prevalence as long as the distribution of costs is sufficiently dispersed.

We calibrate our model to assess its ability to generate the belief and behavior dynamics observed in the data. To do so, we prove how to identify the unobserved cost distribution from data on perceived risks. We find that the model reproduces both the level and cycling of narrative and behavior adoption in the data. We formally test its goodness of fit by showing that the model-generated time series is cointegrated with the data and explains above 60% of its variation.

We present a counterfactual welfare analysis of the impact of competing narratives on infection rates. This analysis aims to quantify the change in infection rates if individuals who believe preventive behaviors are ineffective had adopted the effective narrative. For this purpose, we develop an econometric framework based on our theoretical model to estimate the probability of adhering to the preventive behavior after the counterfactual narrative switch. We find that the drop in infection rates would have been between 2.5% and 4.6% (significant at the 1% level). To put this magnitude in context, a 16-month mandate imposing full adherence to the eight behaviors we study would have led to a drop in infection rates of 24%. Hence, the successful promotion of a unified public health narrative would have led to a reduction in infection rates equivalent to 10-20% of the mandate.

Our paper offers several takeaways. Our empirical analysis shows that narrative switching is prevalent and can have significant welfare consequences. Our theoretical model illustrates the need to incorporate motivated beliefs and exposure to conflicting information to generate these dynamics. Methodologically, we show how to run counterfactual analysis by identifying unobserved preference parameters using data on individual beliefs and behaviors.

2 Related Literature

The paper is related to the small but growing literature studying how agents choose among competing models of the world or narratives. We contribute to this literature by providing the first empirical evidence of narrative switching and documenting the impact of competing narratives on welfare.

Existing theoretical research on narrative selection follows two main approaches. The first approach – which we adopt – assumes that narrative choice is influenced by the agent's preferences, subject to fitting some aspects of the observed data. In the 'competing narratives' models of Eliaz and Spiegler (2020) and Eliaz et al. (2022), agents choose the narrative that yields the highest anticipated utility given the data.[2] The second approach – which is less consistent with our data – assumes that agents' choice of narrative is driven only by some goodness of fit to the observed data (Olea et al., 2019; Galperti, 2019; Schwartzstein and Sunderam, 2021; Ba, 2022).

We contribute to the theory literature by extending the competing narratives model of Eliaz and Spiegler (2020) to include heterogeneous agents. Under this framework,

[2]Relatedly, He and Libgober (2020) study the evolutionary selection of models by looking at whether agents endowed with a given model enjoy higher payoffs over time than agents using alternative models.

agents have motivated beliefs and select narratives that are in their best interest (Kunda, 1990; Bénabou and Tirole, 2016). In our model, agents form beliefs about the causal link between costly behaviors and their consequences. Misspecification arises when agents omit relevant causal relationships.[3] Our model differs from much of the related theoretical literature in two important respects that make it suitable for empirical analysis. First, narratives focus on actions and outcomes at the individual level, rather than dealing with policies and aggregate outcomes. Second, agents face heterogeneous costs and risks.

We contribute to the empirical literature by documenting the dynamics of narrative selection. No work that we are aware of has presented direct evidence of narrative switching. Prior empirical work has identified variation in forecasters' confidence in inflation predictions over time, which could be interpreted as an indirect measure of narrative switching (Giacomini et al., 2020). Existing empirical work has also found evidence for motivated beliefs. For example, experiments show that individuals have motivated beliefs (Saccardo and Serra-Garcia, 2023) and adjust their beliefs in a self-serving manner when faced with positive versus negative feedback (Zimmermann, 2020). However, the work on motivated beliefs has not explored the dynamics of belief cycles over time as we do.

The political economy literature has considered theories of competing narratives in the context of political ideology and the mobilization of political opinion (Eliaz et al., 2022). We contribute to this literature by providing evidence at the individual level of political ideology driving narrative switching. This complements existing work showing the influence of politics on pandemic outcomes at the aggregate level (Wallace et al., 2022; Krieger et al., 2022).

3 Data

We conduct our empirical analysis using data from the UAS, a nationally representative online panel of adults aged 18 and older residing in the U.S. Features of the UAS ensure a high data quality dataset, comparable to that obtainable from more traditional survey modes, such as in-person or phone interviews (Angrisani et al., 2019).[4] The UAS is the only nationally representative panel in the U.S. that continually assessed individuals' experiences during the Covid-19 pandemic.

Nearly 9,000 UAS members participated in the Understanding Coronavirus in America Study (UCAS), a longitudinal study featuring surveys every 2-4 weeks between March

[3] This is related to theoretical work that defines model misspecification as the omission of relevant variables (Cowell et al., 2007; Pearl, 2009; Eyster and Piccione, 2013; Spiegler, 2016; Mailath and Samuelson, 2020; Levy and Razin, 2021; Levy et al., 2022).

[4] These features include the following: 1) panel members are recruited exclusively through Address Based Sampling; 2) if members do not have access to the Internet, they receive a tablet and broadband Internet access (and related training); 3) members are regularly invited to complete two surveys per month and receive compensation at the rate of $20 per 30 minutes of survey time. Since 2014, the attrition rate in the UAS has been about 5% per year.

2020 and July 2021.[5] Crucially, this period covers the first three major waves of the Covid-19 pandemic. The first round of the UCAS was fielded to the entire UAS panel on March 10, 2020. The survey was administered every two weeks, from April 1, 2020, to February 16, 2021. After February 16, 2021, and until July 21, 2021, the survey's frequency was changed to every four weeks. Accordingly, the UCAS consists of 29 rounds of data collection.[6] Out of the 9,927 UAS members invited to participate in the UCAS, 8,628 (87%) answered the survey at least once over the observation period. Among those who answered at least once, 35% completed all 29 rounds, 55% completed at least 25 rounds, and 75% completed at least 12 rounds.

The UCAS survey contains longitudinal information about individuals' infection status, perceived risk of Covid-19 infection, beliefs about effectiveness of various preventive behaviors and adoption of such behaviors, and sources of information about Covid-19. The survey also includes rich socio-demographic information on participants. We next describe each of these variables in detail. The survey questions eliciting the main variables used in the empirical analysis are reported verbatim in Appendix E.

Infection Status. In each round, participants were asked to report whether they were diagnosed with Covid-19, tested positive for Covid-19, or thought that they were infected with Covid-19. Since Covid-19 test availability was not widespread in the first months of the pandemic and exhibited significant differences across population segments, for our analysis we combine this information into a self-reported infection indicator. The indicator takes the value 1 if the participant answered yes to any of the above questions and 0 otherwise. Figure 1 compares the prevalence of Covid-19 infections based on our indicator and the number of new Covid-19 cases per 10,000 inhabitants in UCAS participants' counties of residence.[7] As can be seen, our self-reported measure of infection status follows the three initial pandemic waves (March-April 2020, July 2020, and November 2020-January 2021) closely and fluctuates with official county-level data over the observation period. Our indicator shows a higher prevalence of infection relative to the official count, which is to be expected given that it accounts for both actual and presumed infection.

Perceived Infection Risk. Participants were asked to report their subjective probability of being infected with Covid-19 in the next three months by answering the question "*On a scale of 0 to 100 percent, what is the chance that you will get the coronavirus in*

[5] Table A.1 in Appendix A summarizes the demographic characteristics of survey participants.

[6] While the first UCAS survey was in the field, UAS members were were invited to continue participating every two weeks. Those who provided consent to participate were randomly assigned a number between one and fourteen, determining the day they were asked to answer the survey in each bi-weekly cycle. Upon invitation, participants had two weeks to complete their survey.

[7] To generate Figure 1, we merged restricted residential information for each UCAS participant with county-level caseloads in each round.

Figure 1: Prevalence Rates: Survey Self-Reports vs. Official County-Level Data

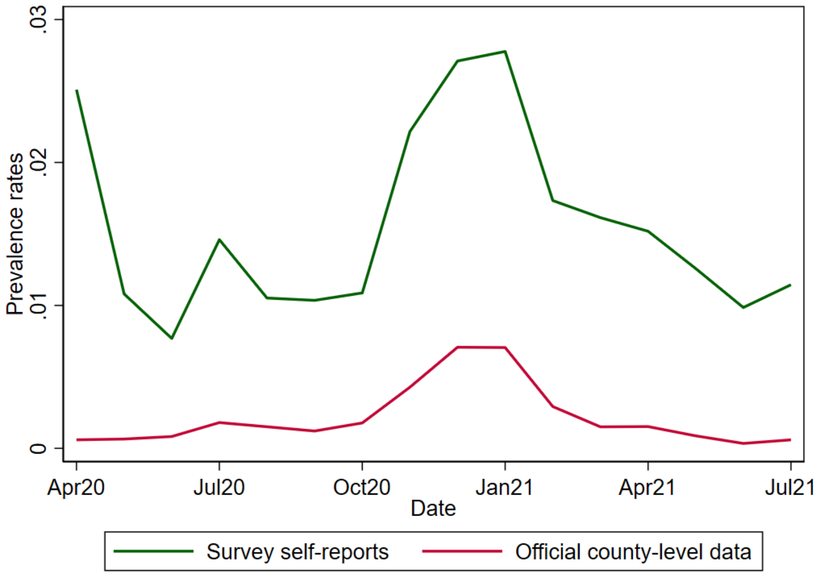

Note: This figure shows the prevalence of Covid-19 infection over time 1) using the indicator generated from self-reports and 2) official county-level data.

the next three months? If you're not sure, please give your best guess." Figure 2 shows the evolution of individuals' risk perception throughout the pandemic. The left panel displays the overall perceived infection risk, which closely tracks the initial pandemic waves through January 2021, but then declines sharply after the first vaccination campaign in the spring of 2021. The right panel splits the sample by our infection status indicator in the next round of data collection, showing that the current perceived infection risk is systematically higher among those who reported that they were infected in the following round.

Figure 2: Perceived Infection Risk, Overall (left) and by Future Infection Status (right)

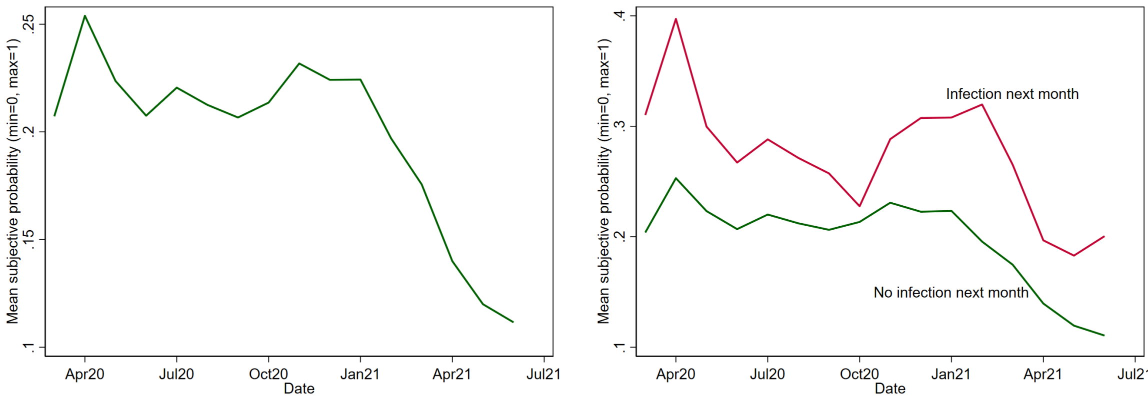

Note: This figure shows self-reported subjective probability of being infected with Covid-19 in the next 3 months over time. The right panel splits the sample by future infection status.

Beliefs about Effectiveness of Preventive Behaviors. In each round, the survey elicited individuals' beliefs about the effectiveness of different behaviors to prevent Covid-19 infection. Specifically, participants were shown a list of behaviors in a table and asked to answer the question *"How effective are the following actions for keeping you safe from coronavirus?"* for each behavior, using a 5-point Likert scale (from *Extremely Ineffective* to *Extremely Effective*). In our analysis, we consider the following behaviors: 1) wearing a mask, 2) avoiding restaurants and bars, 3) avoiding public places, 4) avoiding clinics or hospitals, 5) avoiding planes, 6) avoiding public transit, 7) avoiding high-risk people and 8) washing hands.[8] We create a binary measure of effectiveness that takes the value of 1 if the participant indicated that the behavior was *Extremely Effective* or *Somewhat Effective* and 0 otherwise.

Adoption of Behaviors. In each round, participants were asked to report whether they engaged in any of the preventive behaviors listed above in the last seven days.

Sources of Information. In rounds 1, 7, and 20 through 29, participants were asked about their sources of information about Covid-19 risks. They were presented with a list of sources that included the following: 1) public health officials (CDC, WHO, HHS, local public health officials), 2) television (ABC, CBS, CNN, NBC, MSNBC, Fox News, local television), and 3) social networks (friends, family, coworkers and social media).

Individual Characteristics. An advantage of using UCAS data is the availability of a wide range of background variables for each survey participant. In our analysis, we use socio-demographic information (age, educational attainment, household income and gender), political preferences (identifying as Republican, Democrat, or Independent/Other) and urbanicity (living in a rural, urban or mixed area).

4 Empirical Findings

We first analyze the dynamics of beliefs about effectiveness of preventive behaviors and adoption of preventive behaviors. To streamline the exposition, we illustrate the dynamics using the behavior "avoid restaurants and bars," given that it was one of the key public health recommendations during the pandemic. The dynamics for other behaviors follow a similar pattern and are shown in Appendix B.

We find that a significant share of respondents switch between believing that the preventive behavior is effective and that the behavior is ineffective. Switching is driven

[8] Participants were also asked about the effectiveness of praying or seeing a doctor if infected/exposed. We do not include them in our analysis because (i) praying is not considered a protective behavior from a publich health perspective, and (ii) there is some redundancy between avoiding hospitals/clinics and seeing a doctor.

by the pandemic waves. As infection risk goes up, people tend to switch towards believing that the behavior is effective. As infection risk goes down, people tend to switch towards believing that the behavior is ineffective. Since the belief in the behavior's effectiveness is associated with a higher likelihood of adopting the behavior, this amplifies the procyclicality of adoption rates.

4.1 Belief Dynamics

Figure 3 presents the fraction of respondents who believe that avoiding restaurants and bars is effective by date. As can be seen, this fraction fluctuates with perceived risk (which closely matches fluctuations in county prevalence rates as reported in Figure 1). Sixty-five percent of respondents switch beliefs at least once, with an average of 3.6 switches per individual over the observation period. In Appendix B, we show that respondents are statistically significantly more likely to switch toward believing the behavior is effective following an increase in perceived risk.[9] We also find that the trend changes characterizing these belief cycles are statistically significant by regressing the fraction of agents believing in the effectiveness of each behavior on a linear time trend for each of the phases of the pandemic waves (see Table B.2 in Appendix B).

Figure 3: Share of Individuals who Think Avoiding Restaurants is Effective (left axis) and Perceived Infection Risk (right axis)

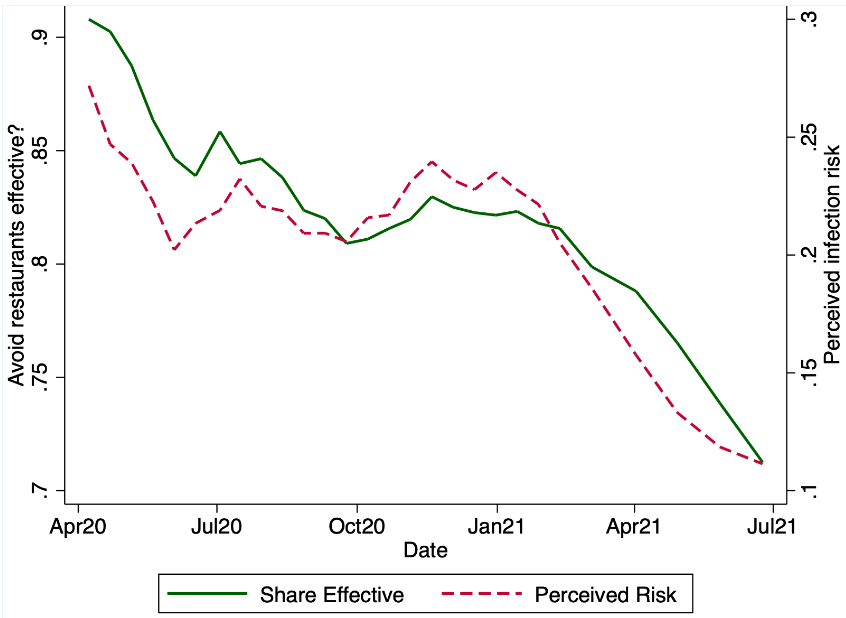

Note: This figure graphs proportion who think avoiding restaurants is effective against perceived infection risk over time.

[9]We run fixed effect regressions of the change in perceived infection risk on belief switches. Regression results and statistics on the number of switches by behavior are reported in Table B.3 in Appendix B.

We interpret these results as evidence that people choose between different models – or *narratives* – linking actions to consequences. To provide evidence for this interpretation, we show that individuals who were exposed to multiple information sources promoting competing narratives were more likely to switch beliefs than individuals exposed to information sources that promoted the same narrative, after controlling for changes in perceived risk and sociodemographic characteristics. We use data on the use of information sources from our survey and group sources into three categories: public health officials, mainstream television and Fox News. Of these sources, the first two promoted the effective narrative versus Fox News, which also promoted the ineffective narrative.[10]

In Table 1, we regress indicators of whether a respondent believes "avoid restaurants" is effective (Column 1) and whether the respondent switched beliefs in the next survey round (Column 2) on binary variables indicating which information source the individual was exposed to in the current round (using rounds 1, 7, and 20-29 in which information sources were available). As can be seen in Column 1, exposure to public health officials and mainstream television is associated with a statistically significantly higher likelihood of believing that the preventive behavior is effective whereas exposure to Fox News is associated with a statistically significantly lower likelihood of believing that the preventive behavior is effective.

To test for our hypothesis that exposure to competing narratives increases the likelihood of switching beliefs, we include interaction terms for Fox News with public health officials and Fox News with mainstream television. The regression output is displayed as odds ratios, with Column 2 showing that individuals who were exposed to Fox News and public health officials or Fox News and mainstream television were significantly more likely to switch beliefs than those not simultaneously exposed to both sources. To account for other sources of variation in exposure to information, we also include controls for whether the respondent lives in an urban or rural area, but this does not have significant effects on beliefs.

The regression also shows that Republicans were less likely to believe the behavior was effective and more likely to switch beliefs than Democrats. This is consistent with the fact that some prominent figures of the Republican Party promoted the ineffective narrative some of the time.[11] On the other hand, prominent figures of the Democratic

[10]Examples of promotion of ineffective narratives by Fox News hosts include Tucker Carlson's statement on October 23, 2020 that "*almost everyone — 85% — who got the coronavirus in July was wearing a mask, and they were infected anyway. So clearly (wearing a mask) doesn't work the way they tell us it works*" (https://www.politifact.com/factchecks/2020/oct/15/tucker-carlson/tucker-carlson-distorts-new-cdc-report-makes-false/), and Laura Ingraham's statement on December 8, 2020 that "*You know what the biggest lie is, is that restaurants are spreaders of Covid.*" (https://www.politifact.com/factchecks/2020/dec/10/laura-ingraham/ingraham-wrongly-claims-no-science-suggests-restau/).

[11]For example, in an interview from July of 2020, Republican President Donald Trump stated, *"And I don't agree with the statement that if everybody would wear a mask, everything disappears."* NBC news

Table 1: Logit regressions on the belief that avoiding restaurants is effective and on the propensity to switch beliefs in the next period.

	Pr(effective)[a]	$Pr(switch)$[b]
Information Sources		
Public Health Officials	1.583***	0.836***
Mainstream TV	1.734***	0.747***
Fox News	0.877***	1.030
Fox News *and* Public Health Officials	0.800***	1.307***
Fox News *and* Mainstream TV	0.938	1.183***
Friends, Coworkers and Social Media	0.980	1.049*
Party		
Republican	0.350***	1.623***
Independent/Other	0.503***	1.333***
Urban Category		
Mixed	1.021	1.030
Urban	0.972	1.044
Age		
30-39	0.945	0.995
40-49	1.013	0.907**
50-59	1.026	0.850***
60+	1.377***	0.796***
Education		
Some college	0.954*	0.997
Bachelor	1.298***	0.730***
Graduate Studies	1.419***	0.642***
Income		
30,000-59,999	1.085***	0.867***
60,000-99,999	1.132***	0.810***
100,000+	0.946*	0.882***
Female	1.109***	0.989
Perceived Risk	1.011***	0.995***
Change in Perceived Risk		1.003***
Observations	69,997	63,572

Note: Default categories are Democrat, rural, 18-29 years old, high school or less, income<$30,000, male. p-values computed using robust standard errors; $^{*}\ p < 0.10$, $^{**}\ p < 0.05$, $^{***}\ p < 0.01$.
[a]Odds ratios from logit regression of belief that behavior is effective.
[b]Odds ratios from logit regression of changes in belief from previous wave.

reported that Trump's messaging on mask wearing was "muddied." (https://www.nbcnews.com/politics/donald-trump/trump-says-i-don-t-agree-cdc-director-s-mask-n1234253.)

party were consistent in promoting the effective narrative.

In terms of socio-demographics, greater age, higher income, and higher educational attainment are all associated with a reduction in the propensity to switch beliefs.

In Table C.4 in Appendix C, we conduct similar regressions for the other preventive behaviors that we considered, including mask wearing, avoiding other public spaces and washing hands. There we find similar results.

4.2 Behavior Dynamics

We also observe significant fluctuations in adoption of preventive behaviors. Figure 4 presents the fraction of respondents who avoid restaurants and bars over time. It shows that the fraction fluctuates with the state of the pandemic, closely matching the three initial pandemic waves.

Figure 4: Share of Participants who Avoid Restaurants (left axis) and Average Perceived Risk (right axis)

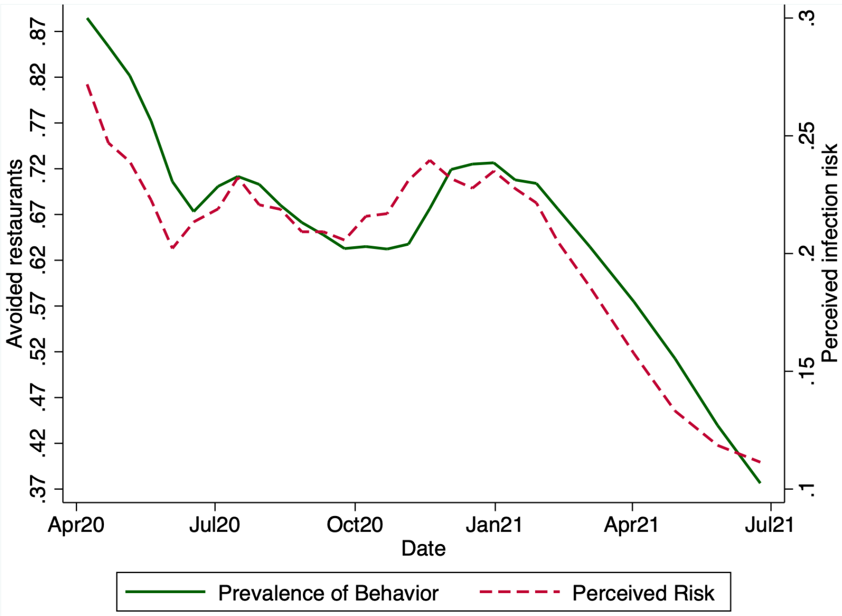

Note: This figure shows the proportions of survey respondents who avoid restaurants and the perceived risk of contracting Covid-19 in the next 3 months over time.

Table 2 displays the adoption of each preventive behavior separately by whether respondents believe the behavior is effective or ineffective. Not surprisingly, beliefs about the effectiveness of a preventive behavior are strongly associated with the take-up of the behavior. For example, among respondents who believe avoiding restaurants is effective, 80.6% report that they avoided restaurants. Among respondents who believe avoiding restaurants is ineffective, only 26.2% report that they avoided restaurants. Fixed effect logit regressions of the likelihood of adopting the behavior on the belief about its effectiveness all show positive and statistically significant coefficients.

Table 2: Prevalence of Behavior by Narrative

Behavior	Effective (%)	Ineffective (%)	All (%)	FE Logit[a] (odds ratio)
wear mask	90.1	60.0	85.6	4.93***
avoid…				
restaurants	80.6	26.2	66.3	6.09***
public places	78.7	34.6	74.5	3.08***
clinics/hospitals	92.1	90.5	91.8	1.33***
plane	98.2	96.7	98.0	1.30***
public transit	96.9	95.5	96.7	1.22***
high-risk people	84.8	46.4	81.6	2.97***
wash hands	94.5	67.2	93.0	3.00***

Note: Data from rounds 1-29 except for avoid hospitals (2-29) avoid planes and public transit (8-29).

[a]Conditional logit regressions of each behavior on the belief in its effectiveness; $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$.

Furthermore, we identify which behaviors led to reductions in infection rates over the subsequent three months. Table 3 presents the fraction of respondents reporting an infection over the next three months by whether they adopted the behavior or not in the last seven days. The last column shows the odds ratio estimates of a logit regression of reported infection status in the next three months on all behaviors.[12] Most (7 out of 8) behaviors are associated with a reduction in infection risk and half (4 out of 8) are associated with a statistically significant reduction in infection risk. As expected, avoiding restaurants, public places, clinics and hospitals, and avoiding public transit are all statistically significantly associated with reducing infection risk. While wearing a mask, avoiding planes, and avoiding high-risk people are also negatively associated with infection risk, their estimated coefficients are not statistically significant. It is possible that mask-wearing is not significantly associated with reduced infection risk because once individuals avoid most public places, wearing a mask becomes less important. Further, since few people took planes or interacted with high-risk people during this period, we may not expect large differences in infection risks by adoption of these behaviors.

[12]The regression uses data starting from round 8, which is the first wave in which questions about effectiveness of avoiding planes and public transit were introduced. We do not use fixed effects because very few individuals report infections in any given round, limiting our ability to obtain precise estimates.

Table 3: Infection Rates by Behavior

Behavior	Infection Rate (%) did not adopt behavior	Infection Rate (%) adopted behavior	Logit[a] (odds ratio)
wear mask	6.1	5.3	0.95
avoid...			
restaurants	6.3	4.9	0.86^{**}
public places	6.5	4.9	0.79^{***}
clinics/hospitals	7.0	5.2	0.72^{***}
plane	7.7	5.5	0.85
public transit	8.2	5.4	0.76^{**}
high-risk people	6.4	5.1	0.95
wash hands	5.9	5.3	1.21^{*}

Note: Data from rounds 8-29 ($N = 88,216$). p-values computed using clustered standard errors at the individual level; $^{*}\, p < 0.10$, $^{**}\, p < 0.05$, $^{***}\, p < 0.01$.

[a]Regression of infection status in the next three months on the set of eight behavior dummies.

5 Theory

In this section, we present a theoretical model of narrative adoption and show theoretically and quantitatively that the model can reproduce the observed narrative and behavior cycles during the pandemic. All proofs are relegated to Appendix D.

There is a continuum of agents indexed by $i \in [0, 1]$. Agents are exposed to a binary risk $I \in \{0, 1\}$ and can take a costly action $a \in \{0, 1\}$ to lower the probability that the risk is realized ($I = 0$). Risk exposure is heterogeneous in the population. Specifically, each agent belongs to a risk category $n \in \{1, \cdots, N\}$ determining her risk distribution $p_n(I|a, \theta)$ conditional on her action a and the state of the world $\theta \in [0, 1]$. The state of the world represents aggregate factors affecting individual risks such as infection prevalence during the pandemic. We assume that $p_n(0|a, \theta)$ is increasing and differentiable in θ for all n and a, i.e., higher θ is associated with higher risks. Taking action 1 lowers the risk for all agents in all possible states of the world, i.e., $p_n(0|1, \theta) < p_n(0|0, \theta)$ for all n, θ. Let $g_n(\theta) = p_n(0|0, \theta) - p_n(0|1, \theta)$ denote the gains in risk reduction from taking the action. We assume that $g_n(\cdot)$ is uniformly bounded away from zero and strictly increasing for all n. This is a natural assumption in the context of the pandemic, since the benefit from taking the preventive action is larger at higher prevalence rates.

The cost of taking action $a = 1$ for agent i is given by c_i, which is distributed according to F in the population. F is continuous with density f and has full support on $(0, \bar{c})$, with $\bar{c} > \max_{\theta,n} g_n(\theta)$.[13] An agent's type is thus given by her risk category n_i and her cost

[13]This implies that there is a mass of agents with costs higher than the expected reduction in risk.

c_i. We assume that c_i and n_i are independent. Each agent observes the realization of the state of the world θ before choosing a. The payoff of agent i is given by $u_i(I, a) = I - c_i a$.

Agents need to model the impact of a on their risk probability in order to inform their choices. They are exposed to alternative narratives about the relationship between variables θ, a and I in the population. We follow Spiegler (2016) and define a narrative as a subjective directed acyclic graph (DAG), where each edge represents a causal link between two variables. Specifically, agents can choose between the *effective* narrative ($\mathcal{E}$), which states that both θ and a cause I, and the *ineffective* narrative ($\mathcal{N}$) which assumes that only θ causes I. The associated DAGs are:

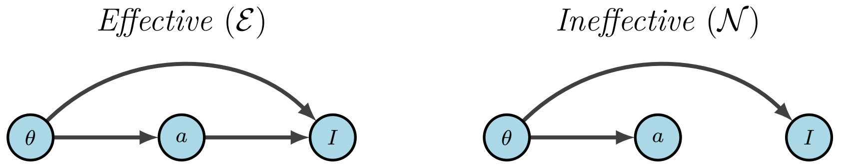

The edge from θ to a means that agents condition their action on the observed state of the world. The DAGs represent conditional independence restrictions among the variables: while narrative $\mathcal{E}$ allows for I to be correlated with a conditional on θ, narrative $\mathcal{N}$ asserts that I and a are conditionally independent (but not unconditionally so due to their mutual dependence on θ). The choice of narrative is driven by motivated beliefs as in the model of Eliaz and Spiegler (2020): each agent adopts the narrative that yields the highest anticipated utility, i.e., she maximizes her subjective expected payoff given the optimal action under each narrative.

We impose the restriction that agents' perceived risks must be derived from the true risk distributions $\{p_n\}_{n=1}^N$. The interpretation is that, while agents choose which variables to pay attention to by adopting different narratives, they use population frequencies to form their beliefs. For instance, an agent of risk type n adopting narrative $\mathcal{N}$ believes that the probability of being infected is $p_n(0|\theta)$ for all a, whereas the perceived risk under $\mathcal{E}$ is $p_n(0|a, \theta)$, with both probabilities coming from the true risk distribution.

We assume that there is a fraction $\gamma \in (0, 1)$ of rational agents, who always adopt the effective narrative, while the remaining agents can choose a narrative from $\{\mathcal{E}, \mathcal{N}\}$. Being rational does not depend on risk category or cost.

The timing is as follows. First, the state of the world θ is realized. After learning θ, each agent adopts a narrative and chooses an action to maximize her expected utility. Finally, I is realized according to $p_{n_i}(I = 0|a, \theta)$.

There exist two additional narratives, depicted below, which are both missing the edge from θ to I. We assume that agents do not consider these narratives for two reasons. First, they imply that agents' perceived risks are independent of the state of the world, which is at odds with the data. Second, they may lead to equilibrium non-existence for some parameter values.

5.1 Agent's Problem

We solve the agent's problem backwards. We first pin down the optimal action a_i^r under each narrative $r \in \{\mathcal{E}, \mathcal{N}\}$ and then we characterize the choice of narrative. Let p_i^r denote the perceived risk distribution of agent i under each narrative.

The optimal action a_i^r given narrative r solves

$$\max_{a \in \{0,1\}} p_i^r(1|a, \theta) - c_i a. \tag{1}$$

If agent i adopts narrative $\mathcal{E}$ she chooses $a = 1$ if $p_{n_i}(1|1, \theta) - c_i \geq p_{n_i}(1|0, \theta)$. That is, since $p_{n_i}(1|1, \theta) - p_{n_i}(1|0, \theta) = g_{n_i}(\theta)$, her optimal choice is given by whether gains in risk reduction are greater than the costs of taking the action:

$$a_i^{\mathcal{E}}(\theta) = \begin{cases} 1 & g_{n_i}(\theta) \geq c_i \\ 0 & \text{otherwise.} \end{cases} \tag{2}$$

Accordingly, her anticipated utility is

$$V_i^{\mathcal{E}}(\theta) = \begin{cases} p_{n_i}(1|1, \theta) - c_i & g_{n_i}(\theta) \geq c_i \\ p_{n_i}(1|0, \theta) & \text{otherwise.} \end{cases} \tag{3}$$

If instead the agent adopts $\mathcal{N}$ then she believes that $p_i^{\mathcal{N}}(0|1, \theta) = p_i^{\mathcal{N}}(0|0, \theta) = p_{n_i}(0|\theta)$ and chooses $a_i^{\mathcal{N}} = 0$ to avoid cost c_i. Her anticipated utility is thus given by

$$V_i^{\mathcal{N}}(\theta) = p_{n_i}(1|\theta). \tag{4}$$

The agent chooses narrative $\mathcal{E}$ if $V_i^{\mathcal{E}}(\theta) \geq V_i^{\mathcal{N}}(\theta)$. Such a choice involves comparing $p_{n_i}(1|a, \theta)$ to the unconditional probability $p_{n_i}(1|\theta)$. Let $\alpha_{n_i} := Pr_{n_i}(a = 1|\theta)$ denote the fraction of agents of risk category n_i choosing $a = 1$. Then $p_{n_i}(1|\theta)$ satisfies

$$p_{n_i}(1|\theta) = p_{n_i}(1|1, \theta)\alpha_{n_i} + p_{n_i}(1|0, \theta)(1 - \alpha_{n_i}).$$

First note that an agent maximizing anticipated utility would never choose narrative $\mathcal{E}$ and action 0 whenever $\alpha_{n_i} > 0$, since $p_{n_i}(1|\theta) > p_{n_i}(1|0, \theta)$. Accordingly the agent chooses $r_i = \mathcal{E}$ and $a_i^{\mathcal{E}} = 1$ if

$$p_{n_i}(1|1, \theta) - c_i \geq p_{n_i}(1|\theta). \tag{5}$$

Otherwise, she chooses $r_i = \mathcal{N}$ and $a_i^{\mathcal{N}} = 0$. Rearranging this condition, we obtain

$$c_i \leq p_{n_i}(1|1,\theta) - p_{n_i}(1|\theta) = (p_{n_i}(1|1,\theta) - p_{n_i}(1|0,\theta))(1 - \alpha_{n_i}),$$

yielding the following cutoff rule for optimal narrative adoption:

$$r_i = \begin{cases} \mathcal{E} & c_i \leq g_{n_i}(\theta)(1 - \alpha_{n_i}) \\ \mathcal{N} & \text{otherwise.} \end{cases} \tag{6}$$

5.2 Equilibrium

In equilibrium, the fraction of agents taking $a = 1$ in risk category n, denoted by α_n^*, must be consistent with agents optimally choosing the narrative that maximizes their anticipated utility given α_n^*. The share of agents choosing $a = 1$ consists of both rational agents with $a_i^{\mathcal{E}} = 1$ and agents with motivated beliefs opting for narrative $\mathcal{E}$. The former are those with $c_i < g_n(\theta)$ while the latter satisfy $c_i \leq g_n(\theta)(1 - \alpha_n)$. Accordingly, α_n^* is given by

$$\alpha_n^* = \gamma F(g_n(\theta)) + (1 - \gamma)F(g_n(\theta)(1 - \alpha_n^*)). \tag{7}$$

Let $\sigma_n^* \in [0, 1]$ be the share of agents with motivated beliefs who adopt narrative $\mathcal{E}$ in equilibrium, which is equal to $F(g_n(\theta)(1 - \alpha_n^*))$. The next result establishes the existence of a unique equilibrium, up to a measure zero set of agents who are indifferent between narratives and therefore also indifferent between actions.

Proposition 1. *An essentially unique equilibrium exists and exhibits a mix of narratives in each risk category, i.e.,* $0 < \sigma_n^* < 1$ *for all* $n = 1, \cdots, N$.

The intuition behind the presence of both narratives in equilibrium is straightforward. As more agents adopt $\mathcal{E}$ and choose $a = 1$, the unconditional probability $p_n(1|\theta)$ gets closer to the conditional probability $p_n(1|1,\theta)$. In this context, agents have a strong incentive to switch to $\mathcal{N}$ since their perceived risk is similar across narratives and they can avoid the costs of taking the action. The opposite forces are at play as more agents adopt $\mathcal{N}$: the gap between $p_n(1|1,\theta)$ and $p_n(1|\theta)$ widens, providing incentives for agents with low cost c_i to switch to $\mathcal{E}$ and take action $a = 1$.

5.3 Comparative Statics

We next identify the conditions under which the fraction of agents adopting the effective narrative is increasing in the state of the world θ. Define the elasticity of the decumulative cost distribution $1-F$ by $e_F(x) := \frac{xf(x)}{1-F(x)}$. It measures the relative increase in the fraction of agents with costs lower than x following a given relative increase in x. Abusing notation, let $\alpha_n(\theta)$ and $\sigma_n(\theta)$ respectively denote the equilibrium values of α_n^* and σ_n^* when the state of the world is θ.

Proposition 2. *The function $\alpha_n(\cdot)$ is strictly increasing for all n. Moreover, there exists $K(\theta) > \dfrac{1}{\gamma}$ such that $\sigma_n(\cdot)$ is strictly increasing at θ if and only if $e_F(g_n(\theta)) < K(\theta)$.*

The intuition behind this result is as follows. A higher θ leads to higher gains $g(\theta)$, making the difference $V_i^{\mathcal{E}} - V_i^{\mathcal{N}}$ go up for all i. In addition, more rational agents have gains greater than costs. Accordingly, a higher θ leads to a higher adoption rate α_n. However, while such an increase does not affect the beliefs under narrative $\mathcal{E}$, it reduces the difference between $p_n(1|1,\theta)$ and $p_n(1|\theta)$, making narrative $\mathcal{N}$ more attractive relative to $\mathcal{E}$. Therefore, whether $V_i^{\mathcal{E}} - V_i^{\mathcal{N}}$ goes up or down depends on whether the larger gains from taking the action are not completely offset by the smaller difference in perceived risk. For this to be the case, the adoption rate must grow slowly so that the gap between $p_n(1|1,\theta)$ and $p_n(1|\theta)$ does not shrink too fast, which happens when $F(g_n(\theta))$ does not increase quickly. That is, as long as the elasticity of the cost distribution is not too high. The next corollary provides a sufficient condition for this to be the case.

Corollary 1. *If $e_F(x) \leq \dfrac{1}{\gamma}$ for all $x \in [\min_n g_n(0), \max_n g_n(1)]$ then $\sigma_n(\cdot)$ is strictly increasing for all $\theta \in [0,1]$ and all $n = 1, \cdots, N$.*

6 Quantitative Analysis

We calibrate the model to study its quantitative implications. We are interested in the model's ability to generate cycles in equilibrium action frequency α_n^* and narrative prevalence σ_n^* similar to those in the data. For illustrative purposes, we focus on "avoiding restaurants and bars" as the preventive action and assume that there is a single risk category.

Our approach is as follows. We first calibrate the parameters of the model, namely, the fraction of rational agents γ and the cost distribution F. We separately estimate the gains $g(\theta)$ from taking the preventive action in each period and obtain the times series of gains $\{g_t\}$. We then input the time series into equilibrium equation (7) of the calibrated model to obtain the series of equilibrium adoption rates $\{\alpha_{nt}^*\}$ and narrative prevalence $\{\sigma_{nt}^*\}$. We use survey weights to estimate both the model parameters and the gains to ensure that our results are representative of the U.S. adult population.

We estimate the fraction of agents who are rational by calculating the share of agents in the data who always think that taking the action is effective (i.e., $r_{it} = \mathcal{E}$ for all t).

We do not directly observe costs in our data. However, we are able to estimate the cost distribution from individual data on risk perceptions by making use of the following key result:

Lemma 1. *For any rational agent i there exist $\underline{p}_i \leq \overline{p}_i$ such that the following conditions are equivalent:*

(i) $g_{n_i}(\theta) > c_i$

(ii) $p_{n_i}(0|1,\theta) > \underline{p}_i$

(iii) $p_{n_i}(0|0,\theta) > \overline{p}_i$.

Moreover, if $c_i \in [g(0), g(1)]$ then $c_i = \overline{p}_i - \underline{p}_i$ with $\overline{p}_i, \underline{p}_i$ given by

$$\underline{p}_i = \min_{\theta}\{p_{n_i}(0|1,\theta) : a_i^{\mathcal{E}}(\theta) = 1\}, \quad \overline{p}_i = \max_{\theta}\{p_{n_i}(0|0,\theta) : a_i^{\mathcal{E}}(\theta) = 0\}. \tag{8}$$

The first part of the lemma directly follows from the assumption that $g_n(\theta)$ and $p_n(0|a,\theta)$ are strictly increasing in θ for $a = 0, 1$. The second part is due to the continuity of g_n.

Lemma 1 states that, as long as most costs fall into $[g_{n_i}(0), g_{n_i}(1)]$, F can be approximated by the distribution of $\overline{p}_i - \underline{p}_i$. Thus, if we are able to identify $\overline{p}_i - \underline{p}_i$ from subjects' perceived risks we could estimate F. Since $\underline{p}_i$ represents the lowest perceived risk of rational agent i when it takes the preventive action we can approximate it using $\min_t\{p_{it} : a_{it} = 1, r_{it'} = \mathcal{E} \text{ for all } t'\}$. Similarly, $\max_t\{p_{it} : a_{it} = 0, r_{it'} = \mathcal{E} \text{ for all } t'\}$ is the empirical counterpart of $\overline{p}_i$. Accordingly, we can estimate F it using the sample of individual belief differences given by $\Big(\max_t\{p_{it} : a_{it} = 0, r_{it'} = \mathcal{E} \text{ for all } t'\} - \min_t\{p_{it} : a_{it} = 1, r_{it'} = \mathcal{E} \text{ for all } t'\}\Big)_{i=1}^{I}$, where I is the number of individuals in the dataset. To do so, we assume that F belongs to the Gamma family and estimate its shape (α) and rate (β) parameters.

We derive the gains from taking the action by estimating for each t the difference in future infection prevalence between those choosing $a_{it} = 0$ and those choosing $a_{it} = 1$. We define future infection using the same three-month window as the window for perceived risk p_{it}.[14]

Table 4 presents the model parameter values and summary statistics for gains, while Figure 5 shows our estimated cost distribution and gains over time.

[14] To avoid poorly measured infection rates due to significant shortage of testing at the beginning of the pandemic, we excluded initial survey rounds for which the number of tests performed did not reach 100 per 100,000 people. According to data from the Covid Tracking Project (https://covidtracking.com), such figure was reached in mid May 2020, coinciding with round 5, so we omitted rounds 1 to 4 from our analysis. Since the survey frequency was changed to every 4 weeks after wave 24, we estimate risk gains for each month by pooling the data for waves 5-6, 7-8,$\cdots$, 23-24. To convert 3-month estimates of differences in infection prevalence into gains $g_n(\theta)$ we assume that agents perceive gains as they are realized over time. Specifically, we define gains at the beginning of month t as differences in infection prevalence in months $\{t-2, t-1, t\}$. This assumption reflects a compromise between agents' need to observe enough data to estimate gains given low prevalence rates while avoiding outdated beliefs since the data arrives continuously. This gives us 12 monthly data points, from August 2020 to July 2021.

Table 4: Model Parameters

Parameter	Value
Fraction Rational (γ)	35%
Cost Distribution (F)	$Gamma[0.155, 1.511]$
Gains ($g_n(\theta)$)	
Mean	1.69%
Std. Deviation	0.21%

Note: F =$Gamma[\alpha, \beta]$ with α =shape parameter and β = rate parameter. Gains = p.p. difference in Covid prevalence over next 3-months between those who did not avoid restaurants and those who did.

Figure 5: Cost distribution (left) and gains from avoiding restaurants (right)

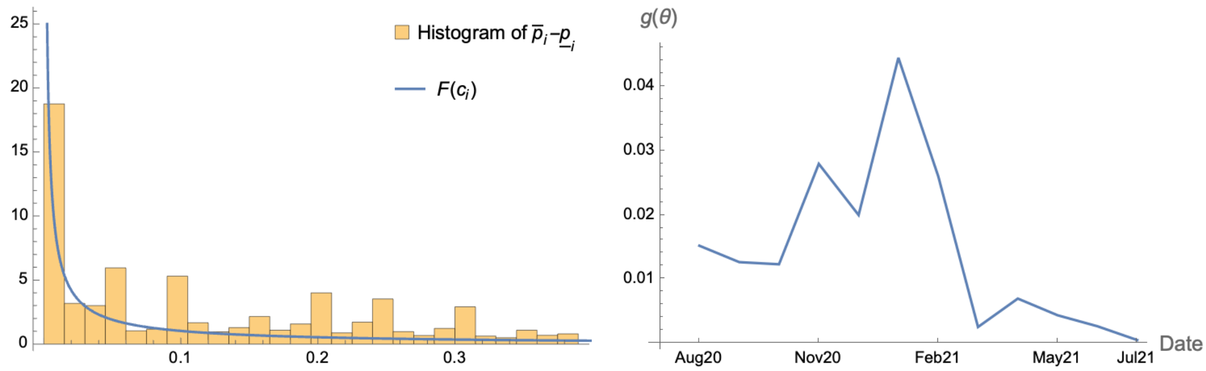

Note: This figure shows a histogram of the estimated cost distribution (left) and the estimated gains over time (right) for the action 'avoiding restaurants.'

Figure 6 compares the action frequency and narrative prevalence generated by the model with those in the data, both in absolute levels and in deviations from the mean. The model falls short of matching action frequency levels and the prevalence of the effective narrative. This is due to the fact that average perceived risks are much higher than infection rates, leading to a mismatch between costs and estimated gains. This mismatch could be due to a variety of factors, including survey participants overestimating risks and underestimating infection rates due to the high prevalence of asymptomatic Covid cases. One way to account for this mismatch is to normalize gains using the ratio of average perceived risk to average gains, where each average is taken both across individuals and across survey rounds. The top panel of Figure 6 presents the action and narrative levels generated by the original and the normalized models. The latter closely matches the average levels in the data. More importantly, as shown in the bottom panel of the figure, the model reproduces quite well the cycles associated with pandemic waves as well as the downward trend in both action and narrative prevalence.

Overall, the model does a very good job reproducing both behavior and belief cycles, despite calibrating each of its components independently rather than trying to fit

Figure 6: Comparison of action frequency and narrative prevalence between the model and the data.

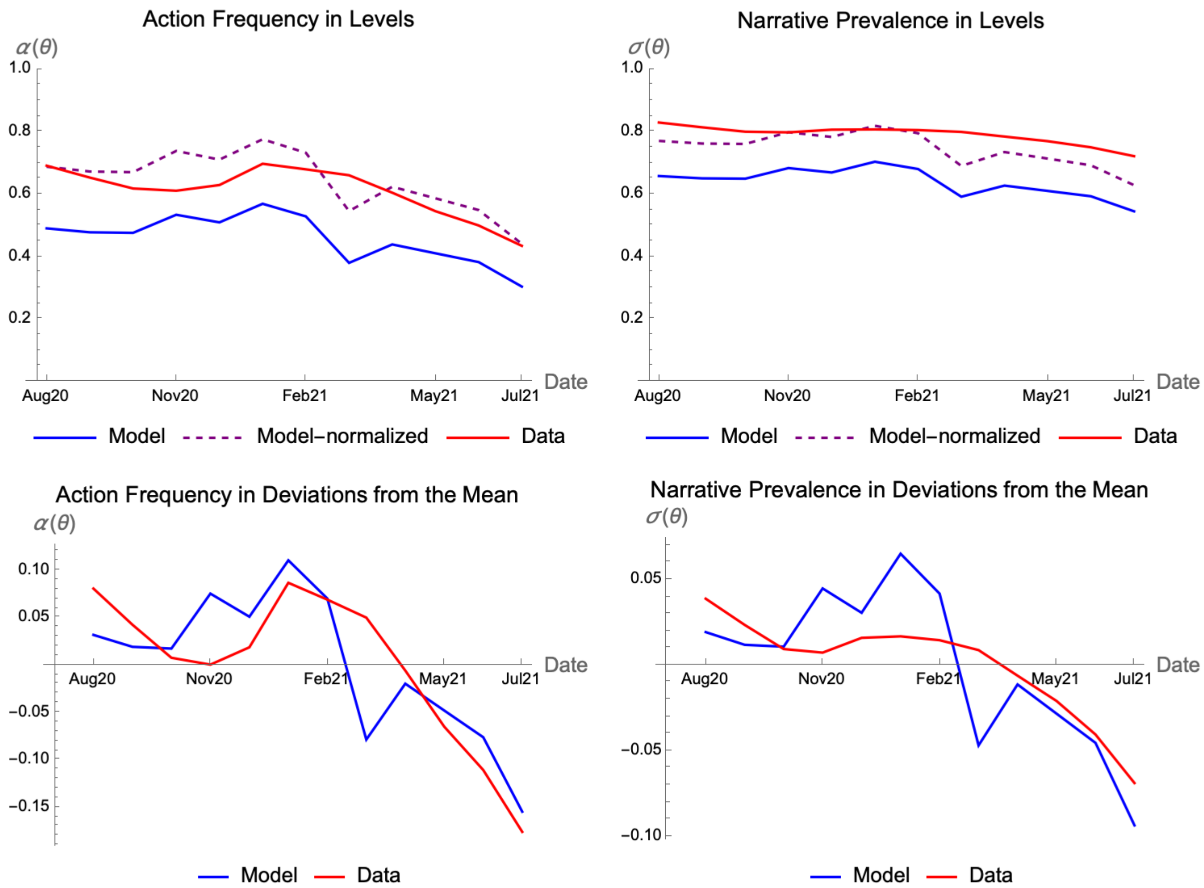

the model to the data. To formally assess the model's goodness-of-fit, we perform the Engle–Granger cointegration test using the time series generated by the model on action prevalence. If the two series are cointegrated, then the model fully captures the dynamics of the data, i.e., the data is a linear combination of the model plus stationary noise. In addition, a high R-squared and a coefficient close to one in the OLS regression of the data on the model-generated series would mean that the model mimics the variation in the data to a high degree. Table 5 presents the results from the test. It includes as a comparison the "all rational" benchmark in which all agents are assumed to believe in the effectiveness of the preventive behavior ($\gamma = 1$).

The results of the Engle-Grainger test are encouraging. First, as a pre-requisite for the test, both the data and the model-generated series exhibit unit roots. Second, we fail to reject that the residuals from the OLS regression $\text{data}_t = \alpha + \beta \text{model}_t + u_t$ are autocorrelated both using the Ljung-box test and selecting the best-fitting ARIMA model according to the Akaike information criterion (an MA(0) process). Finally, the estimated β is close to one (0.816) and the R-squared is quite high (0.616). As a comparison, the

Table 5: Cointegration test for time series of action frequency

Step	Model ($\gamma = 0.35$)	All Rational ($\gamma = 1$)
1. Unit Root Test of Series[a]		
Data	0.57	
Model	0.56	0.55
2. OLS Regression		
Coefficient	0.820	0.716
p-value	0.002	0.002
R^2	0.622	0.617
3. White Noise Test of Residuals		
ARIMA Process[b]	MA(0)	MA(0)
Ljung-Box Test p-values		
$t-1$	0.190	0.186
$t-2$	0.339	0.335
$t-3$	0.072	0.073

[a] Dickey-Fuller test p-value. A high p-value fails to reject the presence of a unit root
[b] Best fitting ARIMA process according to the Akaike Information Criterion

all-rational benchmark, apart from being unable to generate narrative cycles, exhibits a lower OLS coefficient (0.711) suggesting a poorer fit to the behavior prevalence data.

7 Counterfactual Analysis

In this section, we evaluate the impact of adopting the ineffective narrative on welfare. To do so, we develop an econometric framework based on our theoretical model to estimate individual behavior in the counterfactual scenario in which all agents are endowed with the effective narrative.

Our approach consists of two steps. First, for each preventive behavior a and for each agent i who reports adopting narrative $r_{it} = \mathcal{N}$ and choosing $a_{it} = 0$ in period t, we compute the probability that she would take action $a = 1$ under narrative $\mathcal{E}$. Second, we use these probability estimates to compute the counterfactual change in aggregate infection rates.[15]

We face two main challenges, one in each step. In the first step, we observe neither agents' perceived risk under narrative $\mathcal{E}$, since agents only report perceived risk given their chosen narrative and behavior, nor the cost of taking the preventive action. We overcome this lack of observability by using Lemma 1 to derive a counterfactual choice

[15] As in the prior section, we use survey weights to ensure that our estimates apply to the US population.

rule that only relies on reported infection risk. In the second step, since our data includes many preventive behaviors, we cannot tractably account for all the possible combinations of behavior changes contributing to the overall change in infection rates.[16] We tackle this issue by estimating the average marginal effect on risk associated with each behavior and summing these effects across behaviors.

7.1 Choice Rule

We derive the counterfactual choice rule for agents with $(r_{it}, a_{it}) = (\mathcal{N}, 0)$ and use it to identify and estimate the probability of choosing action $a = 1$. Recall that, according to (2), agent i would choose $a = 1$ under narrative $\mathcal{E}$ if $g_{n_i}(\theta) > c_i$.

As Lemma 1 states, we can express choice rule (2) in terms of cutoffs on conditional risk probabilities instead of (unobserved) individual costs. Specifically, we can use expression (8) to identify and estimate cutoffs $\underline{p}_i, \overline{p}_i$ and then apply rules (ii) and (iii) to pin down agent i's counterfactual action $\hat{a}_{it}$.[17] To do so, we need to deal with the fact that, since the agent's narrative in period t is $\mathcal{N}$, her reported infection risk p_{it} corresponds to $p_{it} = p_{n_i}(0|\theta_t)$ instead of $p_{n_i}(0|a_{it}, \theta_t)$. Nonetheless, we know that $p_{it} \in (p_{n_i}(0|1, \theta_t), p_{n_i}(0|0, \theta_t)))$ since the unconditional risk probability is a weighted average of the probabilities conditional on actions. Hence, if $p_{it} < \underline{p}_i$ then $p_{n_{it}}(0|1, \theta) < \underline{p}_i$ and thus $\hat{a}_{it} = 0$. Similarly, if $p_{it} > \overline{p}_i$ then $\hat{a}_{it} = 1$. Finally, we are unable to determine $\hat{a}_{it}$ if $p_{it} \in (\underline{p}_i, \overline{p}_i)$. Figure 7 illustrates the three possible scenarios depending on the value of p_{it}.

Figure 7: Counterfactual Action as a Function of Reported Infection Risk

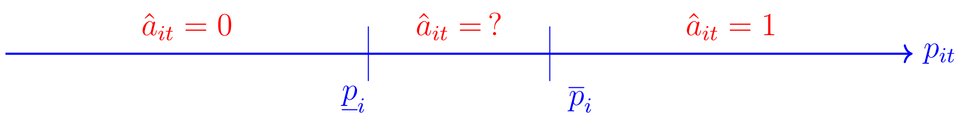

Turning to the estimation of cutoffs, let $\underline{H}(\cdot|Z_i)$ be the probability distribution of $\underline{p}_i$ in the population, where Z_i is a vector of sociodemographic characteristics representing different risk categories. Since $\underline{p}_i = \min_\theta\{p_{n_i}(0|1, \theta) : a_i^{\mathcal{E}}(\theta) = 1\}$ by (8), we can estimate $\underline{H}$ using the cross-section of observations $\left(\min_t\{p_{it} : a_{it} = 1, r_{it} = \mathcal{E}\}, Z_i\right)_{i=1}^{I}$. We assume a flexible functional form for $\underline{H}$, given by the zero-one inflated Beta distribution, which generalizes the Beta distribution to allow for the possibility that $\underline{p}_i \in \{0, 1\}$. Similarly,

[16]The data we analyze includes 8 behaviors. Analyzing all of these would yield 256 different combinations.

[17]Note that agents who switch narratives use cutoff rule $c_i < g(\theta)(1 - \alpha(\theta))$ instead of the cutoff rule $c_i < g(\theta)$ used by rational agents. Hence, c_i no longer coincides with $\overline{p}_i - \underline{p}_i$. However, as long as $g(\theta)(1 - \alpha(\theta))$ is increasing, we can still find cutoffs $\underline{p}_i, \overline{p}_i$ using (8) such that $c_i < g(\theta)(1 - \alpha(\theta))$ is equivalent to $p_{n_i}(0|1, \theta) > \underline{p}_i$ and $p_{n_i}(0|0, \theta) > \overline{p}_i$.

we estimate the cdf of $\overline{p}_i$, denoted by $\overline{H}$, by running a zero-one inflated Beta regression using data $\Big(\max_t\{p_{it} : a_{it} = 0, r_{it} = \mathcal{E}\}, Z_i\Big)_{i=1}^{I}$.

Equipped with the estimated distributions, we assign probabilities to each of the scenarios depicted in Figure 7 for each individual. By doing so, we can derive bounds on the probability that $\hat{a}_{it} = 1$ for each of the preventive behaviors included in the survey. Specifically, the probability that $\hat{a}_{it} = 1$ must be at least as high as $\overline{H}(p_{it}|Z_i)$ since $\underline{p}_i < p_{it}$ implies that $\hat{a}_{it} = 1$. Likewise, it can be no larger than the probability that $\underline{p}_i < p_{it}$, i.e., $\underline{H}(p_{it}|Z_i)$.

7.2 Change in Infection Rates

Our next step is to use the estimated probability of taking each preventive behavior $j = 1, \cdots, m$ to derive the expected difference in infection rates between the counterfactual and the actual data. To do so we need to estimate the probability of becoming infected for any given vector of actions $\mathbf{a} = (a^1, \cdots, a^m)$. Our analysis covers eight behaviors, making the estimation of infection risk conditional on every possible realization of $\mathbf{a}$ intractable.[18] Instead, we regress infection rates on the vector of actions $\mathbf{a}$ for each $t = 1, \cdots, T$ and compute the average marginal effect of each action a^j.[19] Under this approach, the overall change in infection risk for any agent i with counterfactual action probability $Pr(\hat{a}_{it}^j = 1)$ is given by

$$\Delta Pr(I_{it} = 0) = \sum_j Pr(\hat{a}_{it}^j = 1)[p_{n_i}(0|a^j = 1, \bar{\mathbf{a}}_{-jt}, \theta_t) - p_{n_i}(0|a^j = 0, \bar{\mathbf{a}}_{-jt}, \theta_t)], \quad (9)$$

where $\bar{\mathbf{a}}_{-jt}$ is the vector of average frequencies of all behaviors but j in period t.

Table 6 presents the estimated percent reduction in infection rate implied by individual behaviors in the counterfactual scenario in which all agents are endowed with the effective narrative. We compute bootstrap standard errors clustered at the individual level, where at each iteration we implement the econometric procedure described above. We provide estimates using both the lower and the upper bounds on the probability of taking $a = 1$.

As Table 6 illustrates, infection rates would have been between 2.5% and 4.6% lower if all agents in the population had adopted the effective narrative. The estimated reduction is statistically significant and quite large, considering that we do not account at all for any spillover or dynamic effects. Assuming a proportionate reduction in mortality, the adoption of the effective narrative would have prevented roughly between 15,000 and

[18] We include the following behaviors: avoid restaurants, wear a mask, avoid high-risk individuals, avoid public places and gatherings, avoid clinics and hospitals, avoid travel by plane, avoid public transit and wash your hands. We do not have data for the initial seven waves for travel related behaviors so our estimate do not fully reflect the effect of a higher adherence to those behaviors.

[19] We omit the first 4 waves of data due to the limited availability of covid tests prior to May of 2020. See footnote 14.

Table 6: Counterfactual Change in Infection Rates and Behavior

Behaviors	Lower bound (p-value)	Upper bound (p-value)	Mandate (p-value)
All	2.47%	4.59%	24.11%
	(< 0.001)	(< 0.001)	(< 0.001)

p-values based on bootstrap standard errors clustered at the individual level (200 replications).

28,000 deaths from a total of about 605,000 deaths by July 2021 in the U.S.

To put these estimates into context, in the last column of Table 6 we present the hypothetical impact of a mandate that imposes all eight behaviors in the population throughout the whole survey timeline. Under such a mandate, the overall infection rate would have been about 24% lower. Accordingly, eliminating the adoption of the ineffective narrative while keeping the adoption of preventive behaviors voluntary would have led to a reduction in infection rates 10 to 20% as large as the reduction due to a 16-month-long mandate.

8 Conclusions

The 2017 American Economic Association Presidential Address by Robert Shiller included a call to expand the field of economics to "include serious quantitative study of changing popular narratives." In this paper, we respond to this call by conducting the first evaluation of narrative adoption using empirical data. To do so, we use the Covid-19 pandemic as a case study for evaluating the adoption of models or *narratives* about the effectiveness of risk mitigating measures such as avoiding restaurants or wearing a facemask.

In our analysis, we use longitudinal data from a nationally representative survey covering the most intense and critical waves of the Covid-19 pandemic, from March 2020 to July 2021. We document substantial narrative switching driven by changes to perceived risk and exposure to conflicting narratives. These belief cycles pose a challenge to standard notions of belief updating. Therefore, we propose and calibrate a model of narrative adoption driven by motivated reasoning and show that it fits the data remarkably well.

Our welfare analysis highlights the importance of promoting accurate beliefs. We find that the prevalence of the ineffective narrative had substantial negative impacts on public health, leading to roughly 15,000-28,000 excessive deaths from a total of about 605,000 deaths by July 2021 in the U.S. Such non-trivial reductions point to the important consequences of failures in communication in times of crisis. Some years into the pandemic, public health authorities admitted that public communication failed to build consensus about the effectiveness of preventive behaviors. As the former National Institutes of Health Director Francis Collins later stated, "the big thing that I didn't do, and I don't think a lot of the communicators did, was to say this is an evolving crisis, this is going to change every time we made a recommendation, whether it was about social distancing or mask wearing or receiving vaccines."[20]

[20]As reported in STAT News, see https://www.statnews.com/2022/09/19/francis-collins-trust-science-covid-communication-failures/.

References

Angrisani, M., B. Finley, and A. Kapteyn (2019). Can internet match high quality traditional surveys? Comparing the health and retirement study and its online version. In G. Tripath, D. Jacho-Chavez, and K. Huynh (Eds.), *Advances in Econometrics, Vol. 39*. Bingley, UK: Emerald Publishing Limited.

Ba, C. (2022). Robust model misspecification and paradigm shifts. *arXiv preprint arXiv:2106.12727*.

Bénabou, R. and J. Tirole (2016). Mindful economics: The production, consumption, and value of beliefs. *Journal of Economic Perspectives 30*(3), 141–64.

Bursztyn, L., A. Rao, C. P. Roth, and D. H. Yanagizawa-Drott (2020). Misinformation during a pandemic. Technical report, National Bureau of Economic Research.

Cowell, R. G., P. Dawid, S. L. Lauritzen, and D. J. Spiegelhalter (2007). *Probabilistic networks and expert systems: Exact computational methods for Bayesian networks.* Springer Science & Business Media.

Eliaz, K., S. Galperti, and R. Spiegler (2022). False narratives and political mobilization. *arXiv preprint arXiv:2206.12621*.

Eliaz, K. and R. Spiegler (2020). A model of competing narratives. *American Economic Review 110*(12), 3786–3816.

Eyster, E. and M. Piccione (2013). An approach to asset pricing under incomplete and diverse perceptions. *Econometrica 81*(4), 1483–1506.

Galperti, S. (2019). Persuasion: The art of changing worldviews. *American Economic Review 109*(3), 996–1031.

Giacomini, R., V. Skreta, and J. Turen (2020). Heterogeneity, inattention, and bayesian updates. *American Economic Journal: Macroeconomics 12*(1), 282–309.

He, K. and J. Libgober (2020). Evolutionarily stable (mis) specifications: Theory and applications. *arXiv preprint arXiv:2012.15007*.

Krieger, N., C. Testa, J. Chen, W. P. Hanage, and A. J. McGregor (2022). Relationship of political ideology of us federal and state elected officials and key covid pandemic outcomes during the vaccine era: April 2021-march 2022. *SSRN 4131161*.

Kunda, Z. (1990). The case for motivated reasoning. *Psychological bulletin 108*(3), 480.

Levy, G. and R. Razin (2021). A maximum likelihood approach to combining forecasts. *Theoretical Economics 16*(1), 49–71.

Levy, G., R. Razin, and A. Young (2022). Misspecified politics and the recurrence of populism. *American Economic Review 112*(3), 928–62.

Mailath, G. J. and L. Samuelson (2020). Learning under diverse world views: Model-based inference. *American Economic Review 110*(5), 1464–1501.

Olea, J. L. M., P. Ortoleva, M. M. Pai, and A. Prat (2019). Competing models. *arXiv preprint arXiv:1907.03809*.

Pearl, J. (2009). *Causality: Models, Reasoning and Inference*. Cambridge university press.

Saccardo, S. and M. Serra-Garcia (2023). Enabling or limiting cognitive flexibility? evidence of demand for moral commitment. *American Economic Review 113*(2), 396–429.

Schwartzstein, J. and A. Sunderam (2021). Using models to persuade. *American Economic Review 111*(1), 276–323.

Spiegler, R. (2016). Bayesian networks and boundedly rational expectations. *The Quarterly Journal of Economics 131*(3), 1243–1290.

Wallace, J., P. Goldsmith-Pinkham, and J. L. Schwartz (2022). Excess death rates for republicans and democrats during the covid-19 pandemic. Technical report, National Bureau of Economic Research.

Zimmermann, F. (2020). The dynamics of motivated beliefs. *American Economic Review 110*(2), 337–363.

Appendix

A Sample Characteristics

Table A.1: Demographic Characteristics: UCAS Participants and Non-Participants

	In UCAS	Not In UCAS	H_0: In=Not In (p-val)
Male	41.0	40.8	0.906
White	64.3	52.2	0.000
Black	8.8	8.1	0.436
Other Race	10.5	12.2	0.086
Hispanic	17.0	26.8	0.000
Age 18-29	12.9	19.6	0.000
Age 30-39	19.6	22.7	0.013
Age 40-49	18.0	19.4	0.286
Age 50-59	19.1	15.1	0.001
Age 60+	30.3	23.2	0.000
High School or Less	22.2	23.7	0.253
Some College	37.0	37.0	0.995
Bachelor	24.2	23.2	0.468
Graduate Studies	16.6	16.0	0.648
Married	54.0	48.1	0.000
Separated/Divorced/Widowed	20.5	19.5	0.441
Never Married	25.5	32.4	0.000
Working	60.3	64.4	0.009
HH Income<\$30,000	24.8	28.3	0.011
\$30,000 ≤HH Income<\$60,000	25.4	24.1	0.368
\$60,000 ≤HH Income<\$100,000	23.9	21.4	0.060
HH Income≥\$100,000	25.9	26.1	0.855
Rural	16.8	10.5	0.000
Mixed	45.1	39.4	0.000
Urban	38.1	50.0	0.855
Democrat	46.5	53.8	0.455
Republican	37.6	30.8	0.471
Independent/Other	15.8	15.4	0.855
N	8,628	1,299	

B Narrative Dynamics: All Behaviors

Table B.2 presents the statistical significance and direction of estimates of the linear trend in the share of participants believing that a behavior is effective. We first group the survey periods into different time periods, each one associated with an upward or a downward phase in average perceived risk (we find similar phases if we look at infection prevalence). This results into five periods: an initial downward phase from the first pandemic wave ($t = 2, \cdots, 6$), and the upward and downward phases of waves two (resp. $t = 6, \cdots, 9$ and $t = 9, \cdots, 14$) and three (resp. $t = 14, \cdots, 21$ and $t = 22, \cdots, 29$). For each behavior and period, we estimate the regression $\sigma_t = \alpha + \beta t + \varepsilon_t$, where σ_t is the share believing the behavior is effective. Figure B.1 depicts the belief dynamics across all behaviors.

Table B.2: Trend regressions for narrative adoption across behaviors

	wave 1	wave 2		wave 3	
Average perceived risk	$\downarrow$	$\uparrow$	$\downarrow$	$\uparrow$	$\downarrow$
Share believing that is effective to...					
wear mask	$\rightarrow$	$\uparrow$	$\downarrow$	$\uparrow$	$\downarrow$
avoid...					
restaurants	$\downarrow$	$\rightarrow$	$\downarrow$	$\uparrow$	$\downarrow$
public places	$\downarrow$	$\rightarrow$	$\downarrow$	$\uparrow$	$\downarrow$
clinics/hospitals	$\downarrow$	$\rightarrow$	$\downarrow$	$\uparrow$	$\downarrow$
travel	$\rightarrow$	$\rightarrow$	$\downarrow$	$\uparrow$	$\downarrow$
high-risk people	$\searrow$	$\uparrow$	$\searrow$	$\rightarrow$	$\downarrow$
wash hands	$\rightarrow$	$\uparrow$	$\rightarrow$	$\rightarrow$	$\downarrow$
no. obs	5	4	6	8	8

$\downarrow, \uparrow$: p-value$< .05$; $\searrow$: p-value$< .1$; $\rightarrow$: p-value$> .1$

Figure B.1: Share of agents who think the behavior (left axis) is effective and average perceived risk (right axis)

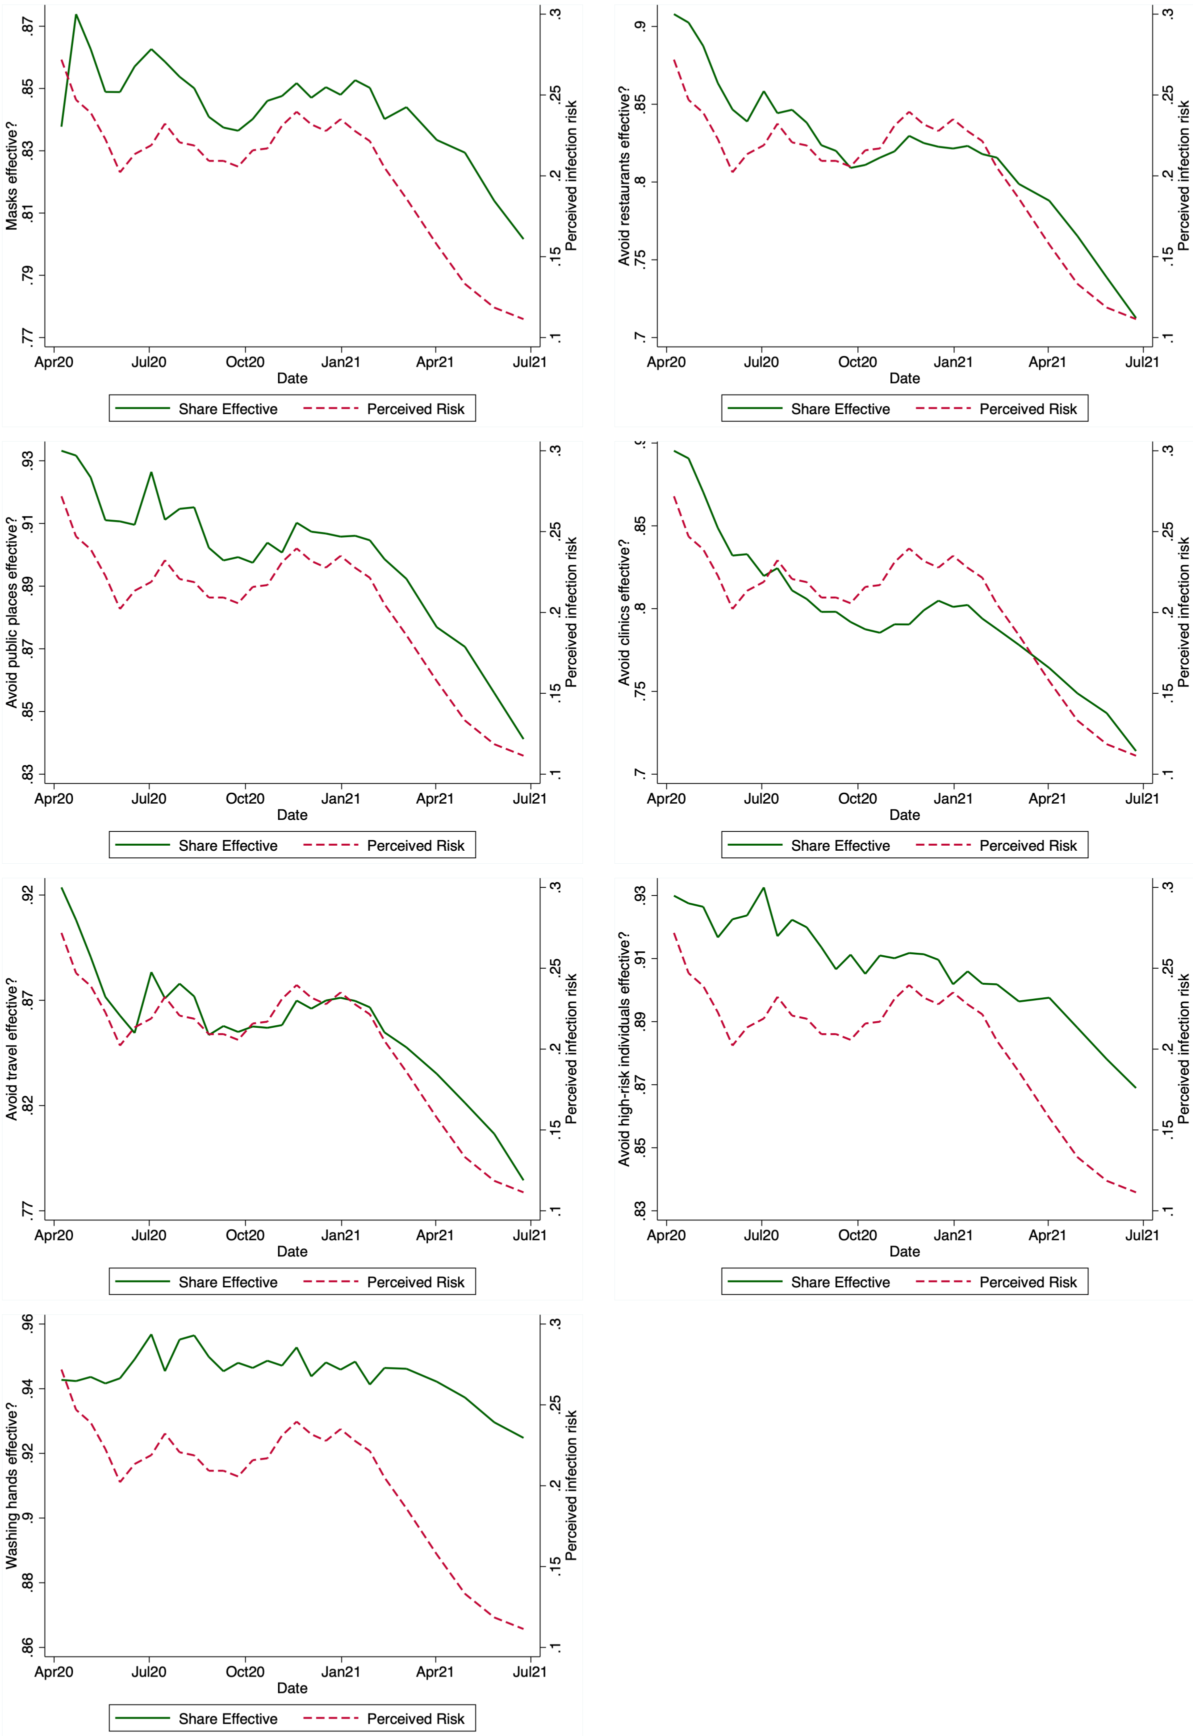

Table B.3 presents the interquartile range of the number of individual belief switches between 'behavior is effective' and 'behavior is not effective' (column 2), and the fraction of participants that switched beliefs at least once (column 3). Column 4 shows, using fixed-effect linear regressions with robust std. errors, whether the increase in perceived risk is associated with an increase in the probability of switching to 'behavior is effective'. The dependent variable equals -1 if the agent switched from effective to not effective, 0 if no switch took place and 1 if the switch was from not effective to effective. Similar results were obtained by estimating conditional ordered logit regressions.

Table B.3: Number of belief switches across behaviors

Belief about	IQR #switches	%Switchers	Effect of Δp_{it}
wear mask	$0-4$	73	$\uparrow$
avoid...			
restaurants	$0-4$	65	$\uparrow$
public places	$0-3$	49	$\uparrow$
clinics/hospitals	$0-6$	73	$\uparrow$
travel	$0-4$	60	$\uparrow$
high-risk people	$0-3$	50	$\nearrow$
wash hands	$0-2$	38	$\rightarrow$
no. obs			$\sim 168K$

p-values computed using robust standard errors; $\uparrow$: $p < .01$; $\nearrow$: $p < .05$; $\rightarrow$: $p > .05$

C Belief Switching and Information Sources

Table C.4: Logit Regressions on the probability to switch beliefs: Other behaviors

	Mask	Public Places	Hospitals	High-risk indiv.	Travel	Wash hands
Information Source						
Public Health Officials	0.788***	0.686***	0.947**	0.725***	0.781***	0.647***
Mainstream TV	0.753***	0.686***	0.904***	0.788***	0.729***	0.813***
Fox News	1.192***	1.109*	1.081	0.933	1.137**	0.954
Fox News *and* Public Health Officials	1.109	1.325***	1.040	1.305***	1.291***	1.498***
Fox News *and* Mainstream TV	1.141*	1.215***	1.078	1.271***	1.082	1.130
Friends, Coworkers and Social Media	1.099***	1.078**	0.967	0.996	1.070***	0.935*
Party						
Republican	1.807***	1.949***	1.003	1.567***	1.863***	1.224***
Independent/Other	1.521***	1.442***	0.999	1.514***	1.494***	1.401***
Perceived Risk	0.998***	0.994***	0.997***	0.996***	0.994***	0.999
Change in Perceived Risk	1.003***	1.004***	1.002***	1.003***	1.004***	1.002**
Observations	63,574	63,574	63,573	63,574	63,574	63,575

Note: Odds ratios from logit regression of changes in belief from previous wave.
All regressions include the following set of demographic controls: urban category, age, household income, education and gender.
p-values computed using robust standard errors; * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

D Omitted Proofs

Proof of Proposition 1. We show existence and uniqueness up to a measure zero subset of agents by arguing that eq. (7) has a unique solution. First, note that the left hand side (LHS) of eq. (7) is increasing and continuous, while the right hand side (RHS) is decreasing and continuous. Second, the RHS is equal to $F(g_n(\theta))$ when $\alpha_n^* = 0$, which is strictly positive given that F has full support. Finally, the RHS is equal to $\gamma F(g_n(\theta)) < 1$ when $\alpha_n^* = 1$. Accordingly, the RHS is strictly higher than the left hand side at $\alpha_n^* = 0$ and strictly lower than the LHS at $\alpha_n^* = 0$, implying that they must cross exactly once at $\alpha_n^* \in (0, 1)$. This also means that $\sigma_n^* > 0$, since $F(g_n(\theta)(1 - \alpha_n^*)) > 0$ by F having full support on $(0, \bar{c})$.

To prove that $\sigma_n^* < 1$ note that, by assumption, $\bar{c} > \max_\theta g_n(\theta) > g_n(\theta)(1-\alpha_n^*)$ for all θ, n. Hence, F having full support on $(0, \bar{c})$ implies that $F(g_n(\theta)(1-\alpha_n^*)) < F(\bar{c}) = 1$. □

Proof of Proposition 2. The first part of the proposition directly follows from the fact that the RHS of eq. (7) is increasing in $g_n(\theta)$ for all $\alpha_n^* \in [0, 1]$. Hence, since the RHS is strictly decreasing in α_n, an increase in θ would require α_n^* to go up until the LHS equals the RHS.

To prove the second part, we characterize $\sigma_n(\theta)$ as a fixed point. By eq. (7) and Proposition 1, $\sigma_n(\theta)$ is the unique solution $\sigma \in (0, 1)$ to equation

$$\sigma = F(g_n(\theta)(1 - \gamma F(g_n(\theta)) - (1 - \gamma)\sigma)). \tag{D.1}$$

Note that the RHS of (D.1) is decreasing in σ. Hence, to prove that $\sigma_n(\theta)$ is strictly increasing if $e_F(x) < 1$ it suffices to pin down the necessary and sufficient condition for the RHS to be strictly increasing in θ.

Denote by $h(\theta)$ expression $f(g_n(\theta)(1-\gamma F(g_n(\theta))-(1-\gamma)\sigma))$. Differentiating the RHS of (D.1) w.r.t. θ we obtain

$$\begin{aligned} &h(\theta)\left[g_n'(\theta)(1 - \gamma F(g_n(\theta)) - (1 - \gamma)\sigma) - \gamma g_n(\theta) f(g_n(\theta)) g_n'(\theta)\right] \\ &\quad = h(\theta) g_n'(\theta)\left[1 - \gamma F(g_n(\theta)) - (1 - \gamma)\sigma - \gamma g_n(\theta) f(g_n(\theta))\right]. \end{aligned}$$

Since $h(\theta)$ and $g_n'(\theta)$ are strictly positive, this expression is positive if and only if

$$1 - \gamma F(g_n(\theta)) - (1 - \gamma)\sigma - \gamma g_n(\theta) f(g_n(\theta)) > 0. \tag{D.2}$$

We can rewrite this condition as

$$1 - F(g_n(\theta)) + (1 - \gamma)(F(g_n(\theta)) - \sigma) - \gamma g_n(\theta) f(g_n(\theta)) > 0.$$

Dividing by $1 - F(g_n(\theta))$ and rearranging we obtain

$$1 + (1-\gamma)\frac{F(g_n(\theta)) - \sigma}{1 - F(g_n(\theta))} > \gamma \frac{g_n(\theta) f(g_n(\theta))}{1 - F(g_n(\theta))}.$$

Equation (D.1) implies that $\sigma < F(g_n(\theta))$ for all θ so the LHS of this expression is strictly greater than one. Accordingly, we can find $K > 1/\gamma$ given by

$$K = \frac{1}{\gamma}\left(1 + (1-\gamma)\frac{F(g_n(\theta)) - \sigma_n(\theta)}{1 - F(g_n(\theta))}\right)$$

such that condition (D.2) is satisfied if and only if $e_F(g_n(\theta)) < K$. □

Proof of Lemma 1. First, notice that g being increasing implies that there exists a unique $\hat{\theta}$ such that $g(\hat{\theta}) = c_i$ for all $c_i \in [g(0), g(1)]$. In addition, since $p_n(0|a,\cdot)$ is increasing for all a then there exist $\underline{p}_i, \overline{p}_i$ such that $p_n(0|1,\hat{\theta}) = \underline{p}_i$ and $p_n(0|0,\hat{\theta}) = \overline{p}_i$. But then, since a rational agent chooses $a = 1$ if $c_i < g(\theta)$, we must have that

$$c_i = g(\hat{\theta}) = p_n(0|0,\hat{\theta}) - p_n(0|1,\hat{\theta}) = \overline{p}_i - \underline{p}_i.$$

Finally, the continuity of $p_n(0|a,\cdot)$ yields the expressions in eq. (8). □

E Survey Questions

Infection Status:

Q1: Have you been tested for the coronavirus since [DATE OF PREVIOUS SURVEY] (when you last took our coronavirus survey? If so, what was the result?

1. I have been tested and I tested positive (I had coronavirus)
2. I have been tested and I tested negative (I did not have coronavirus)
3. I have been tested and I do not know the result
4. I have not been tested

Q2: Whether or not you have had a coronavirus test, has a doctor or another healthcare professional diagnosed you as having or probably having the coronavirus since [DATE OF PREVIOUS SURVEY]?

1. Yes
2. No
3. Unsure

Q3 [if $Q1 \neq 1$ & $Q2 \neq 1$]: Do you think you've been infected with the coronavirus since [DATE OF PREVIOUS SURVEY]?

1. Yes
2. No

Perceived Infection Risk:

On a scale of 0 to 100 percent, what is the chance that you will get the coronavirus in the next three months? If you're not sure, please give your best guess.
[0%-100% Visual Linear Scale]

Beliefs about Effectiveness of Preventive Behaviors:

How **effective** are the following actions for keeping you safe from coronavirus

[Random order of the items on the list]

	Extremely Ineffective	Somewhat Ineffective	Somewhat Effective	Extremely Effective	Unsure
Wearing a face mask such as the one shown here 					
Praying					
Washing your hands with soap or using hand sanitizer frequently					
Seeing a doctor if you feel sick					
Seeing a doctor if you feel healthy but worry that you were exposed					
Avoiding public spaces, gatherings and crowds					
Avoiding contact with people who could be high-risk					
Avoiding hospitals and clinics					
Avoiding restaurants					
Avoiding travel					

Adoption of Behaviors:

Which of the following have you done in the **last seven days** to keep yourself safe from coronavirus Only consider actions that you took or decisions that you made personally.

- Worn a mask or other face covering

  [Yes/No]

- Prayed

  [Yes/No]

- Washed your hands with soap or used hand sanitizer several times per day

  [Yes/No]

- Visited a doctor

  [Yes/No]

- Avoided public spaces, gatherings, or crowds

  [Yes/No]

- Avoided contact with people who could be high-risk

  [Yes/No]

- Avoided eating at restaurants

  [Yes/No]

Sources of Information:

Which of the following information sources have you used to learn about the coronavirus in the **past seven days**?

[Random order of the items on the list]

- Local public health officials such as officials from your county health department

  [Yes/No]

- The US Department of Health and Human Services (HHS)

  [Yes/No]

- The Centers for Disease Control and Prevention (CDC)

  [Yes/No]

- The World Health Organization (WHO)

  [Yes/No]
- Your contacts on social media (Facebook, Twitter, etc.)

  [Yes/No]
- Your close friends and members of your family

  [Yes/No]
- Your coworkers, classmates, or other acquaintances

  [Yes/No]
- Your physician

  [Yes/No]
- Public television and radio

  [Yes/No]
- Fox News

  [Yes/No]
- CNN

  [Yes/No]
- MSNBC

  [Yes/No]
- NBC News

  [Yes/No]
- ABC News

  [Yes/No]
- Your local newspaper

  [Yes/No]
- National newspapers such as the New York Times, The Washington Post, and USA Today [Yes/No]
- Your local TV news

  [Yes/No]

- President Trump

  [Yes/No]

- Vice President Pence

  [Yes/No]

Individual Characteristics:

A comprehensive set of demographic variables is provided with each UAS survey. The complete list can be found at this link.
Political affiliation was obtained from the UAS election surveys, accessible at this link.