



Published in final edited form as:

Nat Hum Behav. 2021 September ; 5(9): 1190–1202. doi:10.1038/s41562-021-01176-8.

Paranoia and belief updating during the COVID-19 crisis

Praveen Suthaharan^{#1}, Erin J. Reed^{#2,3}, Pantelis Leptourgos¹, Joshua G. Kenney¹, Stefan Uddenberg⁴, Christoph D. Mathys^{5,6,7}, Leib Litman⁸, Jonathan Robinson⁸, Aaron J. Moss⁸, Jane R. Taylor^{1,9}, Stephanie M. Groman¹, Philip R. Corlett^{1,9,10,*}

¹Department of Psychiatry, Connecticut Mental Health Center, Yale University, New Haven, CT, USA

²Interdepartmental Neuroscience Program, Yale School of Medicine, New Haven, CT, USA.

³Yale MD-PhD Program, Yale School of Medicine, New Haven, CT, USA.

⁴Booth School of Business, University of Chicago, Chicago, IL, USA.

⁵Interacting Minds Center, Aarhus University, Aarhus, Denmark

⁶Translational Neuromodeling Unit (TNU), Institute for Biomedical Engineering, University of Zurich and ETH Zurich, Zurich, Switzerland

⁷Scuola Internazionale Superiore di Studi Avanzati (SISSA), Trieste, Italy

⁸CloudResearch, 65-30 Kissena Blvd, Queens College Tech Incubator, CEP Hall 2, Flushing, NY 11367

⁹Department of Psychology, Yale University, New Haven, CT, USA

¹⁰Wu Tsai Institute, Yale University, New Haven, CT, USA

These authors contributed equally to this work.

Abstract

The COVID-19 pandemic has made the world seem less predictable. Such crises can lead people to feel that others are a threat. Here, we show that the initial phase of the pandemic in 2020 increased individuals' paranoia and rendered their belief updating more erratic. A proactive lockdown rendered people's belief updating less capricious. However, state-mandated mask wearing increased paranoia and induced more erratic behaviour. This was most evident in states where adherence to mask wearing rules was poor but where rule following is typically more common. Computational analyses of participant behaviour suggested that people with higher paranoia expected the task to be more unstable. People who were more paranoid endorsed conspiracies about mask-wearing and potential vaccines, as well as the QAnon conspiracy theory.

Users may view, print, copy, and download text and data-mine the content in such documents, for the purposes of academic research, subject always to the full Conditions of use: <https://www.springernature.com/gp/open-research/policies/accepted-manuscript-terms>

*Corresponding Author philip.corlett@yale.edu.

Author Contributions

PS, PC, ER, JT, and SG, conceived of the study. ER, JK, CM, PL and SU contributed task materials and analysis code. AM, JR, and LL contributed data. PS and PC analyzed the data. All authors wrote and edited the manuscript.

Competing interests

The authors declare no competing interests.

These beliefs were associated with erratic task behaviour and changed priors. Taken together, we find that real-world uncertainty increases paranoia and influences laboratory task behaviour.

Introduction

Crises, from terrorist attacks¹ to viral pandemics, are fertile grounds for paranoia², the belief that others bear malicious intent towards us. Paranoia may be driven by altered social inferences³, or by domain-general mechanisms for processing uncertainty^{4, 5}. The COVID-19 pandemic increased real-world uncertainty and provided an unprecedented opportunity to track the impact of an unfolding crisis on human beliefs.

We examined self-rated paranoia⁶ alongside social and non-social belief updating in computer-based tasks (Figure 1a), spanning three time periods: before the pandemic lockdown, during lockdown, and into reopening. We further explored the impact of state-level pandemic responses on beliefs and behaviour. We hypothesized that paranoia would increase during the pandemic, perhaps driven by the need to explain and understand real-world volatility¹. Furthermore, we expected that real-world volatility would change individuals' sensitivity to task-based volatility, causing them to update their beliefs in a computerized task accordingly⁵. Finally, since different states responded more or less vigorously to the pandemic, and the residents of those states complied with those policies differently, we expected that efforts to quell the pandemic would change perceived real-world volatility, and thus paranoid ideation and task-based belief updating. We did not pre-register our experiments. Our interests evolved as the pandemic did. We chose to continue gathering data on participants' belief updating and to leverage publicly available data in an effort to explore and explain the differences we observed.

Results

Relating paranoia to task-derived belief updating

We administered a probabilistic reversal learning task. Participants chose between options with different reward probabilities to learn the best option (Figure 1b)⁷. The best option changed and, part way through the task, the underlying probabilities became more difficult to distinguish, increasing unexpected uncertainty and blurring the distinction between probabilistic errors and errors that signify a shift in the underlying contingencies. Participants were forewarned that the best option may change, but not when or how often⁷. Hence, the task assayed belief formation and updating under uncertainty⁷. The challenge is to harbour beliefs that are robust to noise but sensitive to real contingency changes⁷. Before the pandemic, people who were more paranoid (scoring in the clinical range on standard scales^{6, 8}) were more likely to switch their choices between options, even following positive feedback⁵. We compared those data (gathered via the *Amazon Mechanical Turk Marketplace* in the U.S.A. between December 2017 and August 2018; see Supplementary Table 1) to a new task version with identical contingencies, but framed socially (Figure 1a). Instead of selecting between decks of cards ('non-social task'), participants (see Supplementary Table 1) chose between three potential collaborators who might increase or decrease their score. These data were gathered during January 2020, before the World Health Organization

declared a global pandemic. Participants with higher paranoia switched more frequently than low paranoia participants after receiving positive feedback in both; however, there were no significant behavioral differences between tasks (Supplementary figure 2a; win-switch rate: $F_{(1, 198)}=0.918$, $p=0.339$, $\eta_p^2=0.0009$, $BF_{10}=1.07$, anecdotal evidence for null hypothesis of no difference between tasks lose-stay rate: $F_{(1, 198)}=3.121$, $p=0.08$, $\eta_p^2=0.002$, $BF_{10}=3.24$, moderate evidence for the alternative hypothesis, a difference between tasks). See also Supplementary figure 2b; there were also no differences in points ($BF_{10}=0.163$, strong evidence for the null hypothesis) or reversals achieved ($BF_{10}=0.210$, strong evidence for the null hypothesis) between the social and non-social tasks).

Computational modelling

Probabilistic reversal learning involves decision making under uncertainty. The reasons for decisions may not be manifest in simple counts of choices or errors. By modeling participants' choices, we can estimate latent processes⁹. We suppose they continually update a probabilistic representation of the task (a generative model) which guides their behavior^{10, 11}. To estimate their generative models, we identify: (1) a set of prior assumptions about how events are caused by the environment (the perceptual model), and (2) the behavioral consequences of their posterior beliefs about options and outcomes (the response model^{10, 11}). Inverting the response model also entails inverting the perceptual model, and yields a mapping from task cues to the beliefs that cause participants' responses^{10, 11} (Figure 1b).

The perceptual model (Figure 1b) is comprised of three hierarchical layers of belief about the task, represented as probability distributions which encode belief content and uncertainty: (1) reward belief (what was the outcome?), (2) contingency beliefs (what are the current values of the options [decks/collaborators]?), and (3) volatility beliefs (how do option values change over time?). Each layer updates the layer above it in light of evolving experiences, which engender prediction errors and drive learning proportionally to current variance. Each belief layer has an initial mean μ^0 , which for simplicity we will refer to as the prior belief, although strictly speaking the prior belief is the Gaussian distribution with mean μ^0 and variance σ^0 . ω_2 and ω_3 encode the evolution rate of the environment at the corresponding level (contingencies and volatility). Higher values imply a more rapid tonic level of change. The higher the expected uncertainty (i.e., 'I expect variable outcomes'), the less surprising an atypical outcome may be, and the less it drives belief updates ('this variation is normal'). κ captures sensitivity to perceived phasic or unexpected changes in the task. κ underwrites perceived change in the underlying statistics of the environment (i.e. 'the world is changing'), which may call for more wholesale belief revision. The layers of beliefs are fed through a sigmoid response function (Figure 1b). We made the response model temperature inversely proportional to participants' volatility belief - rendering decisions more stochastic under higher perceived volatility. Using this model we have previously demonstrated identical belief updating deficits in paranoid humans and rats administered methamphetamine⁵, and that this model better captures participants' responses compared to standard reinforcement-learning models⁵, including models that weight positive and negative prediction errors differently¹².

For ω_3 (evolution rate of volatility) we observed a main effect of group (Figure 2; $F_{(1, 198)}=4.447$, $p=0.036$, $\eta_p^2=0.014$) and block ($F_{(1, 198)}=38.89$, $p < 0.001$, $\eta_p^2=0.064$), but no effect of task or three-way interaction. Likewise we found group and block effects, for μ_3^0 - the volatility prior - (group: $F_{(1, 198)}=8.566$, $p=0.004$, $\eta_p^2=0.035$; block: $F_{(1, 198)}=161.845$, $p < 0.001$, $\eta_p^2=0.11$) and κ , expected uncertainty learning rate, (group: $F_{(1, 198)}=21.45$, $p < 0.001$, $\eta_p^2=0.08$; block: $F_{(1, 198)}=30.281$, $p < 0.001$, $\eta_p^2=0.031$) but no effect of task or three-way interactions. We found a group effect ($F_{(1, 198)}=12.986$, $p < 0.001$, $\eta_p^2=0.053$) but no task, block or interaction effects on ω_2 - evolution rate of reward contingencies. Thus, we observed an impact of paranoia on behavior and model parameters that did not differ by the social or non-social framing of the task. People with higher paranoia expected more volatility and reward, initially, had higher learning rate for unexpected events, but slower learning from expected uncertainty, and reward, regardless of whether they were learning about cards or people.

How the evolving pandemic impacted paranoia and belief updating

After the pandemic was declared we continued to acquire data on both tasks (19/03/20 - 17/07/20; see Supplementary Table 2 and Supplementary Table 3). We did not pre-register our experiments. We examined the impact of real-world uncertainty on belief updating in a computerized task.

The onset of the pandemic was associated with increased self-reported paranoia from January 2020 through the lockdown, peaking during reopening (Figure 3a; $F_{(2, 530)}=14.7$, $p < 0.001$, $\eta_p^2=0.053$). Anxiety increased (Supplementary Figure 1; $F_{(2, 529)}=4.51$, $p=0.011$, $\eta_p^2=0.017$) but, the change was less pronounced than paranoia, suggesting a particular impact of the pandemic on beliefs about others.

Within the U.S.A., states responded differently to the pandemic; some instituted lockdowns early and broadly (more proactive), whereas others closed later and reopened sooner (less proactive). See Equation 1 and Supplementary Figure 3. When they reopened, some states mandated mask wearing (more proactive). Others did not (less proactive). We conducted exploratory analyses to discern the impact of lockdown and reopening policies on task performance and belief updating.

We observed a main effect of pandemic period (Figure 3b; $F_{(2, 527)}=4.948$, $p=0.007$, $\eta_p^2=0.018$) and a state proactivity by period interaction (Figure 3b; $F_{(2, 527)}=4.785$, $p=0.009$, $\eta_p^2=0.018$) for paranoia and win-switch behaviour (Figure 3b; main effect: $F_{(2, 527)}=3.270$, $p=0.039$, $\eta_p^2=0.012$; interaction: $F_{(2, 527)}=8.747$, $p < 0.001$, $\eta_p^2=0.032$) and volatility priors (Figure 3b; $F_{(2, 527)}=8.623$, $p=0.001$, $\eta_p^2=0.032$).

Early in the pandemic, vigorous lockdown policies (closing early, extensively, and remaining closed) were associated with less paranoia (Figure 3b; $t_{227}=2.57$, $p=0.011$, $d=0.334$, 95% CI=(0.071, 0.539)), less erratic win-switching (Figure 3b; $t_{216}=2.73$, $p=0.007$, Cohen's $d=0.351$, 95% CI=(0.019, 0.117)), and weaker initial beliefs about task volatility (Figure 3b; $t_{217}=4.22$, $p < 0.001$, Cohen's $d=0.561$, 95% CI=(0.401, 1.10)) compared to participants in states that imposed a less vigorous lockdown.

At reopening, paranoia was highest, and participants' task behaviour was most erratic in states that mandated mask wearing (Figure 3b; $t_{67} = -2.39$, $p = 0.02$, $d = 0.483$, 95% CI = $(-0.164, -0.015)$). Furthermore, participants in mandate states had higher contamination fear (Supplementary Figure 4; $t_{101} = -2.89$, $p = 0.005$, $d = 0.471$, 95% CI = $(-0.655, -0.121)$).

None of the other pandemic or policy effects on parameters (priors or learning rates) survived false discovery rate correction for multiple comparisons. We therefore carried win-switch rates and initial beliefs (or priors) about volatility into subsequent analyses.

We asked participants in the social task to rate whether or not they believed that the avatars had deliberately sabotaged them. Reopening was associated with an increase in self-reported sabotage beliefs (Figure 4a; $t_{145} = -2.35$, $p = 0.02$, $d = 0.349$, 95% CI = $(-1.114, -0.096)$). There were no significant main effects or interactions. Given the effects of the pandemic and policies on paranoia and task behaviour, we explored the impact of lockdown policy on behaviour in the social task, specifically. Self-rated paranoia in the real world correlated with sabotage belief in the task (Figure 4b; $r = 0.4$, $p < 0.001$). During lockdown, when proactive state responses were associated with decreased self-rated paranoia, win-switch rate ($t_{216} = 2.73$, $p = 0.014$, $d = 0.351$, 95% CI = $(0.019, 0.117)$) and μ_3^0 ($t_{223} = 4.20$, $p < 0.001$, $d = 4.299$, 95% CI = $(0, 1.647)$) were significantly lower in participants from states with more vigorous lockdown (Figure 4b). As paranoia increased with the pandemic, so did task derived sabotage beliefs about the avatars. Participants in states that locked down more vigorously engaged in less erratic task behaviour, and had weaker initial volatility beliefs.

Paranoia is induced by mask-wearing policies

Following a quasi-experimental approach to causal inferences (developed in econometrics and recently extended to behavioural and cognitive neuroscience¹³), we pursued an exploratory difference-in-differences (DiD) analysis (following Equation 2) to discern the effects of state mask-wearing policy on paranoia. A DiD design compares changes in outcomes before and after a given policy takes effect in one area, to changes in the same outcomes in another area that did not introduce the policy¹⁴ (see Supplementary Figure 5). The data must be longitudinal, but they needn't follow the same participants¹⁴. It is essential to demonstrate that – prior to implementation - the areas adopting different policies are matched in terms of the trends in the variable being compared (parallel trends assumption). Using pre-treatment outcomes, we cannot reject the null hypothesis that pre-treatment trends of the treated and control states developed in parallel ($\lambda = -0.1$, $p = 0.334$). This increases our confidence that the parallel trend assumption also holds in the treatment period. However, such analyses are not robust to baseline demographic differences between the treatment groups¹⁵. Before pursuing such an analysis, it is important to establish parity between the two comparator locations¹⁶, so that any differences can be more clearly ascribed to the policy that was implemented. We believe such parity applies in our case. First, there were no significant differences at baseline in the number of cases or deaths in states that went on to mandate versus recommend mask wearing (cases, $t_{10} = -1.22$, $p = 0.25$, $BF_{10} = 2.3$, anecdotal evidence for null hypothesis; deaths, $t_{10} = -1.14$, $p = 0.28$, $BF_{10} = 2.02$, anecdotal evidence for null hypothesis). Furthermore, paranoia is held to flourish during periods of economic inequality¹⁷. There were no baseline differences in unemployment rates in April

(prior to the mask policy onset) between states that mandated masks versus states that recommended mask wearing ($t_{16}=-0.81$, $p=0.43$, $BF_{10}=0.42$, anecdotal evidence for null hypothesis). We employed a between participants design, so it is important to establish that there were no demographic differences (age, gender, race) in participants from states that mandated versus participants from states that recommended mask-wearing (age, $t=-1.46$, $d.f. = 42.5$, $p=0.15$, $BF_{10}=0.105$, anecdotal evidence for null hypothesis; gender, $\chi^2=0.37$, $d.f.=1$, $p=0.54$, $BF_{10}=0.11$, anecdotal evidence for null hypothesis; race, Fisher's exact test for count data, $p=0.21$, $BF_{10}=0.105$, anecdotal evidence for null hypothesis). On these bases, we chose to proceed with the DiD analysis.

We implemented a non-parametric cluster bootstrap procedure, which is theoretically robust to heteroscedasticity and arbitrary patterns of error correlation within clusters, and to variation in error processes across clusters¹⁸. The procedure reassigns entire states to either treatment or control and recalculates the treatment effect in each reassigned sample, generating a randomization distribution.

Mandated mask wearing was associated with an estimated 40% increase in paranoia ($\delta_{DID} = 0.396$, $p=0.038$), relative to states in which mask wearing was recommended but not required (Figure 5a, Supplementary Figure 5). This increase in paranoia was mirrored as significantly higher win-switch rates in participant task performance (Figure 5b; $t_{67} = -2.4$, $p = 0.039$, $d = 0.483$, 95% CI = (-0.164, -0.015)) as well as stronger volatility priors (Figure 5b; $t_{141} = -3.7$, $p < 0.001$, $d = -3.739$, 95% CI = (0, -1.585)). The imposition of a mask mandate appears to have increased paranoia.

Variation in rule following relates to paranoia

In order to unpack the DiD result, we further explored whether any other features might illuminate the variation in paranoia by local mask policy¹⁹. There are state-level cultural differences – measured by the Cultural Tightness and Looseness (CTL) index¹⁹ – with regards to rule following and tolerance for deviance. Tighter states have more rules and tolerate less deviance, whereas looser states have few strongly enforced rules and greater tolerance for deviance¹⁹. This index also provides a proxy for state politics. Tighter states tend to vote Republican, looser tend towards the Democrats¹⁹. Since 2020 was a politically tumultuous time, and that the pandemic became politicized, we thought it prudent to incorporate politics into our analyses.

We also tried to assess whether people were following the mask rules. We acquired independent survey data gathered in the U.S.A. from 250,000 respondents who, between July 2 and July 14, were asked: *How often do you wear a mask in public when you expect to be within six feet of another person?*²⁰ These data were used to compute an estimated frequency of mask wearing in each state during the reopening period (Figure 5c).

We found that in culturally tighter states where mask wearing was mandated, mask wearing was lowest ($t_{47}=12.84$, $p < 0.001$, $d = 1.911$, 95% CI = (0.064,0.088)). Furthermore, even in states where mask wearing was recommended, mask wearing was lowest in culturally tighter states ($t_{57}=3.06$, $p=0.003$, $d = 0.663$, 95% CI = (0.022,0.107)).

Through backward linear regression with removal (see Equation 3), we fit a series of models attempting to predict individuals' self-rated paranoia from the features of their environment, including whether they were subject to a mask mandate or not, the cultural tightness of their state, state-level mask-wearing, and Coronavirus cases in their state. In the best fitting model ($F_{(11,160)}=1.91, p=0.04$) there was a significant three way interaction between mandate, state tightness and perceived mask wearing ($t_{24}=-2.4, p=0.018$) – paranoia was highest in mandate state participants living in areas that were culturally tighter, where fewer people were wearing masks (Figure 6, Supplementary Table 4). Taken together, our DiD and regression analyses imply that mask-wearing mandates and their violation, particularly in places that value rule following, may have increased paranoia and erratic task behaviour. Alternatively, the mandate may have increased paranoia in culturally conservative states, culminating in less mask wearing.

Paranoia relates to beliefs about mask-wearing

In a follow-up study, we attempted a conceptual replication, recruiting a further 405 participants (19/03/20 - 17/07/20; see Extended Data Table 4), polling their paranoia, their attitudes toward mask-wearing, and capturing their belief updating under uncertainty with the probabilistic reversal learning task. Individuals with high paranoia were more reluctant to wear masks and reported wearing them significantly less (Figure 5d; $t_{158} = 4.59, p < 0.001, d = 0.520, 95\% \text{ CI} = (0.091, 0.229)$). Again, win-switch rate was significantly higher in high paranoia individuals (Figure 5d; $t_{138} = -6.40, p < 0.001, d = 1.148, 95\% \text{ CI} = (-0.227, -0.120)$), as was their prior belief about volatility (Figure 5d; $t_{138} = -6.04, p < 0.001, d = -6.041, 95\% \text{ CI} = (0, -2.067)$), confirming the links between paranoia, mask hesitancy, erratic task behaviour and expected volatility that our DiD analysis suggested. Our data across the initial study and replication imply that paranoia flourishes when individuals' attitudes conflict with what they are being instructed to do, particularly in areas where rule following is more common – paranoia may be driven by a fear of social reprisals for one's anti-mask attitudes.

Sabotage beliefs in the non-social task

Our domain-general account of paranoia⁵ suggests that performance on the non-social task should be related to paranoia, which we observed previously⁵ and presently. In the same follow-up study (see Supplementary Table 5) we asked participants to complete the non-social probabilistic reversal learning task and, at completion, to rate their belief that the inanimate non-social card decks were sabotaging them. Participants' self-rated paranoia correlated with their belief that the cards were sabotaging them (Supplementary Figure 6; $r = 0.47, p < 0.001$), consistent with reports that people with paranoid delusions imbue moving polygons with nefarious intentions²¹.

Other changes coincident with the onset of mask policies

In addition to the pandemic, other events have increased unrest and uncertainty, notably the protests following the murder of George Floyd. These protests began on May 24th 2020 and continue, occurring in every US state. To explore the possibility that these events were contributing to our results, we compared the number of protest events in mandate and recommended states in the months before and after reopening. There were significantly

more protests per day from May 24th through July 31st 2020 in mask-recommended states versus mask-mandated states ($t_{83}=3.10$, $p=0.0027$, $d = 0.591$, 95% CI = (17.458, 80.142)). This suggests the effect of mask mandates we observed was not driven by the coincidence of protests and reopening. Protests were less frequent in states whose participants had higher paranoia (Figure 5b).

Furthermore, there were no significant differences in cases ($t_{12}=-1.45$, $p=0.17$, $BF_{10}=1.63$, anecdotal evidence for null hypothesis) or deaths ($t_{11}=-1.64$, $p=0.13$, $BF_{10}=6.21$, moderate evidence for alternative hypothesis) at reopening in mask-mandate vs mask-recommend states. We compared the change in unemployment from lockdown to reopening in mask-mandate vs mask-recommend states and found no significant difference ($t_{17}=-1.85$, $p=0.08$, $BF_{10}=1.04$, anecdotal evidence for null hypothesis).

Changes in the participant pool did not drive the effects

Given that the pandemic has altered our behaviour and beliefs, it is critical to establish that the effects we describe above are not driven by changes in sampling. For example, with lockdown and unemployment, more people may have been available to participate in online studies. We find no differences in demographic variables across our study periods (pre-pandemic, lockdown, reopening, gender, $F_{(2,523)}= 0.341$, $p = 0.856$, $\eta_p^2 = 0.001$, $BF_{10}=0.03$, strong evidence for null hypothesis; age, $F_{(2,522)}= 2.301$, $p = 0.404$, $\eta_p^2 = 0.009$, $BF_{10}=0.19$, moderate evidence for null hypothesis; race, $F_{(2,520)}= 1.10$, $p = 0.856$, $\eta_p^2 = 0.004$, $BF_{10}=0.06$, strong evidence for null hypothesis; education, $F_{(2,530)}= 0.611$, $p = 0.856$, $\eta_p^2 = 0.002$, $BF_{10}=0.04$, strong evidence for null hypothesis; employment, $F_{(2,529)}= 0.156$, $p = 0.856$, $\eta_p^2 = 0.0006$, $BF_{10}=0.03$, strong evidence for null hypothesis; income, $F_{(2,523)}= 1.31$, $p = 0.856$, $\eta_p^2 = 0.005$, $BF_{10}=0.08$, strong evidence for null hypothesis; medication, $F_{(2,408)}= 0.266$, $p = 0.856$, $\eta_p^2 = 0.001$, $BF_{10}=0.04$, strong evidence for null hypothesis; mental and neurological health, $F_{(2,418)}= 3.36$, $p = 0.288$, $\eta_p^2 = 0.016$, $BF_{10}=0.620$, anecdotal evidence for null hypothesis; Supplementary Figure 7). Given that the effects we describe depend on geographical location, we confirm that the proportions of participants recruited from each state did not differ across our study periods ($\chi^2=6.63$, d.f. = 6, $p=0.34$, $BF_{10}=0.16$, moderate evidence for null hypothesis, Supplementary Figure 8). Finally, in order to assuage concerns that the participant pool changed as the result of the pandemic, published analyses confirm that it did not²². Furthermore, in collaboration with CloudResearch²³, we ascertained location data spanning our study periods from 7,293 experiments comprising 2.5 million participants. The distributions of participants across states match those that we recruited, and the mean proportion of participants in a state across all studies in the pool for each period correlates significantly with the proportion of participants in each state in the data we acquired for each period: pre-pandemic, $r = 0.76$, $p = 2.2E-8$; lockdown, $r = 0.78$, $p = 5.8E-9$; reopening, $r = 0.81$, $p = 8.5E-10$ (Supplementary Figure 7). Thus, we did not, by chance, recruit more participants from mask-mandating states or tighter states, for example. Furthermore, focusing on the data that went into the DiD, there were no demographic differences pre- (age, $p=0.65$, $BF_{10}=0.14$, moderate evidence for the null hypothesis; gender, $p=0.77$, $BF_{10}=0.13$, moderate evidence for the null hypothesis; race, $p=0.34$, $BF_{10}=0.20$, moderate evidence for the null hypothesis) versus post-reopening (age, $p=0.57$, $BF_{10}=0.21$, moderate evidence for the null hypothesis; gender,

$p=0.77$, $BF_{10}=0.19$, moderate evidence for the null hypothesis; race, $p=0.07$, $BF_{10}=0.55$, anecdotal evidence for the null hypothesis) for mask-mandate versus mask-recommended states. Taken together with our task and self-report results, these control analyses increase our confidence that during reopening, people were most paranoid in the presence of rules and perceived rule breaking, particularly in states where people usually tend to follow the rules.

Paranoia versus conspiracy theorizing

Whilst correlated, paranoia and conspiracy beliefs are not synonymous²⁴. Therefore, we also assessed conspiracy beliefs about a potential COVID vaccine in the follow-up study (see Supplementary Table 5). We found that conspiracy beliefs about a vaccine correlated significantly with paranoia (Figure 7a; $r=0.61$, $p < 0.001$), and that such beliefs were associated with erratic task behaviour (Figure 7b; win-switch rate: $r=0.44$, $p < 0.001$) and perturbed volatility priors (Figure 7c; $r=0.34$, $p < 0.001$) in an identical manner to mask concerns and paranoia more broadly. In the U.K., early in the pandemic, conspiracy theorizing was associated with higher paranoia and less adherence to public health countermeasures²⁵. We replicate and extend those findings to the U.S.A., and we provide mechanistic suggestions centered on domain-general belief updating mechanisms: priors on volatility and learning rates.

To further address how politics might have contributed to our results, we gathered more data in September 2020 (see Supplementary Table 5). We assessed participant's performance on the probabilistic reversal learning task, and we also asked them to rate their belief in the QAnon conspiracy theory. QAnon is a right-wing conspiracy theory, concerned with the ministrations of the deep-state, prominent left-wing politicians, and Hollywood entertainers. Its adherents believe that those individuals and organizations are engaged in child trafficking and murder, for the purposes of extracting and consuming the adrenochrome from the children's brains. They believe Donald Trump is part of a plan with the army to arrest and indict politicians and entertainers. We found that people who identify as Republican had stronger belief in QAnon. QAnon belief and paranoia more broadly were highly correlated (Figure 8a; $r=0.5$, $p < 0.001$). Furthermore, QAnon belief correlated with COVID conspiracy theorizing ($r=0.5$, $p < 0.001$). Finally, QAnon endorsement correlated with win-switch behaviour (Figure 8b; $r=0.44$, $p < 0.001$) and volatility belief (Figure 8c; $r=0.31$, $p < 0.001$) just like paranoia. Supplementary Figure 9 depicts the effect of political party affiliation on QAnon belief, paranoia, win-switch behaviour, and volatility belief. People who identified as Republican were more likely to endorse the QAnon conspiracy, they attested to more paranoia, they evinced more win-switching, and they had stronger initial beliefs about task volatility. Taken together, our data suggest that personal politics, local policies, and local political climate all contributed to paranoia and aberrant belief updating.

Discussion

The COVID-19 pandemic was associated with increased paranoia. The increase was less pronounced in states that enforced a more proactive lockdown, and more pronounced at reopening in states that mandated mask-wearing. Win-switch behaviour and volatility

priors tracked these changes in paranoia with policy. We explored cultural variations in rule following (cultural tightness or looseness¹⁹) as a possible contributor to the increased paranoia that we observed. State tightness may originate in response to threats like natural disasters, disease, territorial, and ideological conflict¹⁹. Tighter states typically evince more coordinated threat responses¹⁹. They have also experienced greater mortality from pneumonia and influenza throughout their history¹⁹. However, paranoia was highest in tight states with a mandate, with lower mask adherence during reopening. It may be that societies that adhere rigidly to rules are less able to adapt to unpredictable change. Alternatively, these societies may prioritize protection from ideological and economic threats over a public health crisis, or perhaps view the disease burden as less threatening.

Our exploratory analyses suggest that mandating mask-wearing may have caused paranoia to increase, altering participants' expected volatility in the tasks (μ_3^0). Follow-up exploratory analyses suggested that in culturally tighter states with a mask mandate, those rules were being followed less (fewer people were wearing masks), which was associated with greater paranoia. Violations of social norms engender prediction errors²⁶ which have been implicated in paranoia in the laboratory^{4, 27-29}.

Mask-wearing is a collective action problem, wherein most people are *conditional cooperators*; generally willing to act in the collective interest as long as they perceive sufficient reciprocation by others³⁰. Perceiving others refusing to follow the rules and failing to proffer reciprocal protection appears to have contributed to the increase in paranoia we observed. Indeed, paranoia, a belief in others' nefarious intentions, also correlated with reluctance to wear a mask, and with endorsement of vaccine conspiracy theories. Finally, people who do not want to abide by the mask-wearing rules might be paranoid about being caught violating those rules.

The 2020 election in the USA politicized pandemic countermeasures. In follow-up studies conducted in September 2020 we found that paranoia correlated with endorsement of the far-right QAnon conspiracy theory, as did task related prior beliefs about volatility. We suggest that the rise of this conspiracy theory was driven by the volatility that people experienced in their everyday lives during the pandemic. This has long been theorized historically. Here we present behavioural evidence for a connection between real-world volatility, conspiracy theorizing, paranoia, and hesitant attitudes toward pandemic countermeasures. Evidence relating real-world uncertainty to paranoia and conspiracy theorizing has, thus far, been somewhat anecdotal and largely historical. For example, during the Black Death, the conspiratorial anti-semitic belief that Jewish people were poisoning wells and causing the pandemic was sadly extremely common¹⁷. The AIDS epidemic was associated with a number of conspiracies related to public health measures – but less directly. For example, people believed that HIV was created through the polio vaccination program in Africa³¹. More broadly, the early phases of the AIDS epidemic were associated with heightened paranoia concerning homosexuals and intravenous drug users³². Perhaps the closest relative to our mask mandate result involves seatbelt laws³³. Like masks in a viral pandemic, seatbelts are (and continue to be) extremely effective at preventing serious injury and death in road traffic accidents³⁴. However, the introduction of State Laws prescribing that they are worn was associated with public outcry³³. People were concerned about the

imposition on their freedom³³. They complained that seatbelts were particularly dangerous when cars accidentally entered bodies of water. The evidence shows seatbelt wearing, like mask wearing, is not associated with excess fatality.

Paranoia is, by definition, a social concern. It must be undergirded by inferences about social features. Our data suggest that paranoia increases greatly when social rules are broken, particularly in cultures where rule-following is valued. However, we do not believe this is license to conclude that domain-specific coalitional mechanisms underwrite paranoia as some have argued³. Rather, our data show that both social and non-social inferences under uncertainty (particularly prior beliefs about volatility) are similarly related to paranoia. Further, they are similarly altered by real-world volatility, rules and rule breaking. We suggest that some social inferences are instantiated by domain-general mechanisms^{5, 35}. Our follow-up study demonstrating that people imputed nefarious intentions to the decidedly inanimate card decks tends to support this conclusion (Supplementary Figure 6). We suggest this finding is consistent with previous reports that people with persecutory delusions tend to evince an intentional bias toward animated polygons²¹. More broadly, paranoia often relates to domain-general belief updating biases³⁶, and thence to domain-specific social effects³⁷. Indeed, when tasks have both social and non-social components, there are often no differences in the weightings of these components between patients with schizophrenia and controls^{38, 39}. However, we cannot make definitive claims about the domain-general nature of paranoia. Though our social task was not preferentially related to paranoia, it may be that it was not social enough. There are clearly domain-specific social mechanisms⁴⁰. We should examine the relationships between paranoia and these more definitively social tasks, and will do so in future.

While we independently (and multiply) replicated the associations between concerns about interventions that might mitigate the pandemic, paranoia and task behaviour, and we show that our results are not driven by other real-world events, or issues with our sampling, there remain a number of important caveats to our conclusions. We did not pre-register our experiments, predictions, or analyses. Nor did we run a within-subject study through the pandemic periods. Our DiD analysis should be considered exploratory. DiD analyses require longitudinal but not necessarily within-subjects or panel data¹⁴. Our DiD analysis does leverage some tentative causal claims, despite being based on between-subjects data¹⁴. Mask Recommended states were culturally tighter though of course both cultural tightness did not change during the course of our study. Tightness did interact with mandate and adherence to mask wearing policy (Figure 6). The baseline difference in tightness would have worked against the effects we observed, not in their favor. Indeed, our multiple regression analysis found no evidence for an effect of tightness on paranoia in states without a mask-mandate (Figure 6). Critically, we do not know if any participant, or anyone close to them, was infected by COVID-19, so our work cannot speak to the more direct effects of infection. There are of course other factors that changed as a result of the pandemic. Unemployment increased dramatically, though not significantly more in mandate states. Historically, conspiracies peak not only during uncertainty, but also during periods of marked economic inequality¹⁷. Internet searches for conspiracy topics increase with unemployment⁴¹. The patterns of behaviour we observed may have also been driven by economic uncertainty, although our data militate against this interpretation somewhat, since

Gini coefficients⁴² (a metric of income inequality) did not differ between mandate and recommend states ($t_{19}=-1.60$, $p=0.13$). Finally, our work is based entirely in the USA. In future work we will expand our scope internationally. Cultural features⁴³ and pandemic responses vary across nations. This variance should be fertile grounds in which to replicate and extend our findings.

We highlight the impact that societal volatility and local cultural and policy differences have on individual cognition. This may have contributed to past failures to replicate in psychological research. If replication attempts were conducted under different economic, political or social conditions (bull versus bear markets, for example), then they may yield different results, not because of inadequacy of the theory or experiment but because the participants' behaviour was being modulated by heretofore under-appreciated stable and volatile local cultural features.

Per predictive processing theories⁴, paranoia increased with increases in real-world volatility, as did task-based volatility priors. Those effects were moderated by government responses. On one hand, proactive leadership mollified paranoia during lockdown, by tempering expectations of volatility. On the other hand, mask mandates enhanced paranoia during reopening by imposing a rule that was often violated. These findings may help guide responses to future crises.

Methods

All experiments were conducted at the Connecticut Mental Health Center in strict accordance with Yale University's Human Investigation Committee who provided ethical review and exemption approval (HIC# 2000026290). Informed consent was provided by all research participants.

Experiment.

A total of 1,010 participants were recruited online via CloudResearch – an online research platform that integrates with MTurk while providing additional security for easy recruitment²³. Sample sizes were determined based on our prior work with this task, platform, and computational modelling approach. Two studies were conducted to investigate paranoia and belief updating: a pandemic study and a replication study. Participants were randomized to one of two task versions (see Behavioural tasks). Participants were compensated with \$6 for completion and a bonus of \$2 if they scored in the top 10% of all respondents.

Pandemic study.—A total of 605 participants were collected, divided into 202 pre-lockdown participants, 231 lockdown participants, and 172 reopening participants. Of the 202, we included the 72 (16 high paranoia) participants who completed the non-social task (described in a prior publication⁵). Those participants paranoia was self-rated with the SCID-II paranoid trait questions, which are strongly overlapping and correlated with the Green et al scale⁶. We recruited 130 (20 high paranoia) participants who completed the social task. Similarly, of the 231, we recruited 119 (20 high paranoia) and 112 (30 high paranoia) participants who completed the non-social and social tasks, respectively. Lastly,

of the 172, we recruited 93 (35 high paranoia) and 79 (35 high paranoia) participants who completed the non-social and social tasks, respectively. In addition to CloudResearch's safeguard from bot submissions, we implemented the same study advertisement, submission review, approval and bonusing as described in our previous study⁵. We excluded a total of 163 submissions – 18 from pre-lockdown (social only), 34 from lockdown (non-social and social), and 111 from reopening (non-social and social). Of the 18, 17 were excluded based on incomplete/nonsensical free-response submissions and 1 for insufficient questionnaire completion. Of the 34, 29 were excluded based on incomplete/nonsensical free-response submissions and 5 for insufficient questionnaire completion. Of the 111, all were excluded based on incomplete/nonsensical free-response submissions. Submissions with grossly incorrect completion codes were rejected without further review.

Replication study.—We collected a total of 405 participants of which 314 were low paranoid individuals and 91 were high paranoid individuals. Similar exclusion and inclusion criteria were applied for recruitment; most notably, we leveraged CloudResearch's newly added *Data Quality* feature which only allows vetted high-quality participants – individuals who have passed their screening measures – into our study. This systematically cleaned all poor participants from our sample pool.

Behavioural tasks.

Participants completed a 3-option probabilistic reversal-learning task with a non-social (card deck) or social (partner) domain frame. **Non-social:** Three decks of cards were presented for 160 trials, divided evenly into 4 blocks. Each deck contained different amounts of winning (+100) and losing (−50) cards. Participants were instructed to find the best deck and earn as many points as possible. It was also noted that the best deck could change¹¹. **Social:** Three avatars were presented for 160 trials, divided evenly into 4 blocks. Participants were advised to imagine themselves as students at a university working with classmates to complete a group project, where some classmates were known to be unreliable – showing up late, failing to complete their work, getting distracted for personal reasons – or deliberately sabotage their work. Each avatar either represented a helpful (+100) or hurtful (−50) partner. We instructed participants to select an avatar (or partner) to work with to gain as many points towards their group project. Like the non-social, they were instructed that the best partner could change. For both tasks, the contingencies began as 90% reward, 50% reward, and 10% reward with the allocation across deck/partner switching after 9 out of 10 consecutive rewards. At the end of the second block, unbeknownst to the participants, the underlying contingencies transition to 80% reward, 40% reward, and 20% reward – making it more difficult to discern whether a loss of points was due to normal variations (probabilistic noise) or whether the best option has changed.

Questionnaires.

Following task completion, questionnaires were administered via Qualtrics, we queried demographic information (age, gender, educational attainment, ethnicity, and race) and mental health questions (past or present diagnosis, medication use), *Structured Clinical Interview for DSM-IV Axis II Personality Disorders* (SCID-II)⁸, Beck's Anxiety Inventory (BAI)⁴⁴, Beck's Depression Inventory (BDI)⁴⁵, the Dimensional Obsessive-Compulsive

Scale (DOCS)⁴⁶, and critically, the revised Green et al., Paranoid Thoughts Scale (R-GPTS)⁶ – which separates clinically from non-clinically paranoid individuals based on the ROC-recommended cut-off score of 11. We also polled participants' beliefs about the social task ('Did any of the partners deliberately sabotage you?') – on a Likert scale from 'Definitely not' to 'Definitely yes'. We later added the same item for the non-social task ('Did you feel as though the decks were tricking you?') to investigate sabotage belief differences between tasks (see Supplementary Figure 6).

In a follow-up study, we adopted a survey⁴⁷ that investigated individual US consumers' mask attitude and behaviour and a survey²⁵ of COVID-19 conspiracies. The 9-item mask questionnaire was used for our study to compute mask attitude (values < 0 indicate attitude against mask-wearing and values > 0 indicate attitude in favor of mask-wearing) for identifying group differences in paranoia. To compute an individual's coronavirus vaccine conspiracy belief, we aggregated five vaccine-related questions from the 48-item coronavirus conspiracy questionnaire:

1. The coronavirus vaccine will contain microchips to control the people.
2. Coronavirus was created to force everyone to get vaccinated.
3. The vaccine will be used to carry out mass sterilization.
4. The coronavirus is bait to scare the whole globe into accepting a vaccine that will introduce the 'real' deadly virus.
5. The WHO already has a vaccine and are withholding it.

We adopted a 7-point scale: strongly disagree (1), disagree (2), somewhat disagree (3), neutral (4), somewhat agree (5), agree (6) and strongly agree (7). A higher score indicates greater endorsement of a question.

QAnon.

To measure beliefs about the QAnon conspiracy, we used a questionnaire that polled respondents political attitudes⁴⁸, in particular towards QAnon.

Additional features.

Along with the task and questionnaire data, we examined state-level unemployment rate⁴⁹, confirmed COVID-19 cases⁵⁰, and mask wearing²⁰ in the USA. **Unemployment.** The *Carsey School of Public Policy* reported unemployment rates for the months of February, April, May and June in 2020. We utilized the rates in April and June as our markers for measuring the difference in unemployment between the pre-pandemic period and pandemic period, respectively. **Confirmed cases.** The *New York Times* published cumulative counts of coronavirus cases since January.

Mask wearing.

Similarly, at the request of the *New York Times*, *Dynata* – a research firm – conducted interviews on mask use across the USA and obtained a sample of 250,000 survey respondents between July 2 and July 14²⁰. Each participant was asked: *How often do you*

wear a mask in public when you expect to be within six feet of another person? The answer choices to the question included *Never*, *Rarely*, *Sometimes*, *Frequently*, and *Always*.

Mask Policies.

According to the Philadelphia Inquirer: <https://fusion.inquirer.com/health/coronavirus/covid-19-coronavirus-face-masks-infection-rates-20200624.html>, 11 states mandated mask-wearing in public: CA, NM, MI, IL, NY, MA, RI, MD, VA, DE, and ME at the time of our reopening data collection. The other states from which we recruited participants recommended mask wearing in public.

Protests.

We accessed the publicly available data from the armed conflict location and event data project (ACLED, <https://acleddata.com/special-projects/us-crisis-monitor/>), which has been recording the location, participation, and motivation of protests in the US since the week of George Floyd's murder in May.

Behavioural analysis.

We analysed tendencies to choose alternative decks after positive feedback (win-switch) and select the same deck after negative feedback (lose-stay). Win-switch rates were calculated as the number of trials in which the participant switched after positive feedback divided by the number of trials in which they received positive feedback. Lose-stay rates were calculated as number of trials in which a participant persisted after negative feedback divided by total negative feedback trials.

Lockdown proactivity metric.

We also defined a proactivity metric (or score) to measure how inadequately or adequately a state reacted to COVID-19⁵¹. This score was calculated based on when a state's stay-at-home (SAH) order was introduced (I) and when it expired (E):

I : number of days from baseline to when the order was introduced (i.e., introduced date-baseline date).

E : number of days before the order was lifted since it was introduced (i.e., expiration date – introduced date).

where baseline date is defined as the date at which the first SAH order was implemented (See Supplementary Figure 3). California was the first to enforce the order on March 19th, 2020 (i.e., baseline date = 1).

We calculated proactivity as follows:

$$\rho = \begin{cases} \frac{1}{1 + \frac{I}{E}} & \text{if } E \geq I > 0 \\ 0 & \text{if } E = 0 \text{ and } I = 0 \end{cases} \quad (1)$$

This function gives states with early lockdown ($I \rightarrow 1$) and sustained lockdown ($E \rightarrow \infty$) a higher proactivity score ($\rho \rightarrow 1$) while giving states that did not issue state-wide SAH orders ($E = 0; I = 0$) a score of 0.

Therefore, our proactivity (ρ) metric – either 0 (never lockdown, less proactive) or ranging from 0.5 (started lockdown, less proactive) to 1 (started lockdown, more proactive) – offers a reasonable approach for measuring proactive state interventions in response to the pandemic.

In our analyses, for lockdown, we separated less proactive and more proactive states at the median. For reopening, states that mandated mask-wearing were designated more proactive and states that recommended mask-wearing were designated less proactive.

We set the proactivity of the pre-lockdown data to be the proactivity of the lockdown response that would be enacted once the pandemic was declared. Using the reopening proactivity designation for the pre-lockdown data instead had no impact on our findings (see Supplementary Table 6).

Causal inference.

To measure attribution of mask policy on paranoia, we adopt a difference-in-differences (DiD) approach. The DiD model we used to assess the causal effect of mask policy on paranoia in states that either recommended or required masks to be worn in public is represented below by the following equation:

$$P_{it} = \alpha + \beta t_i + \lambda M_i + \delta(t_i * M_i) + \epsilon_{it} \quad (2)$$

where P_{it} is the paranoia level for individual i and time t , α is the baseline average of paranoia, β is the time trend of paranoia in the control group, λ is the pre-intervention difference in paranoia between both the control and treatment groups, and δ is the mask effect. The control and treatment groups, in our case, represent states that recommend and require mask-wearing, respectively. The interaction term between time and mask policy represents our DiD estimate.

Multiple regression analysis.

We conducted a multiple linear regression analysis, attempting to predict paranoia based on three continuous state variables – number of COVID-19 cases, cultural tightness and looseness (CTL) index, and mask-wearing belief – and one categorical state variable – mask policy. We fit a 15-predictor paranoia model and performed backward stepwise regression to find the model that best explains our data. Below we illustrate the full 15-predictor model and the resulting reduced 11-predictor model:

Full model:

$$\hat{y} = \beta_0 + \beta_1 * X_{CASES} + \beta_2 * X_{POLICY} + \beta_3 * X_{CTL} + \beta_4 * X_{MASK} + \beta_5 * X_{CASES*POLICY} + \beta_6 * X_{CASES*CTL} + \beta_7 * X_{POLICY*CTL} + \beta_8 * X_{CASES*MASK} + \beta_9 * X_{CTL*MASK} + \beta_{10} * X_{CTL*MASK} + \beta_{11} * X_{CASES * POLICY * CTL} + \beta_{12} * X_{CASES * POLICY * MASK} + \beta_{13} * X_{CASES * CTL * MASK} + \beta_{14} * X_{POLICY * CTL * MASK} + \beta_{15} * X_{CASES * POLICY * CTL * MASK}$$

Reduced model:

$$\hat{y} = \beta_0 + \beta_1 * X_{CASES} + \beta_2 * X_{POLICY} + \beta_3 * X_{CTL} + \beta_4 * X_{MASK} + \beta_5 * X_{CASES*POLICY} + \beta_6 * X_{CASES*CTL} + \beta_7 * X_{POLICY*CTL} + \beta_8 * X_{POLICY*MASK} + \beta_9 * X_{CTL*MASK} + \beta_{10} * X_{CASES * POLICY * CTL} + \beta_{11} * X_{POLICY * CTL * MASK} \quad (3)$$

Computational modeling.

The Hierarchical Gaussian Filter (HGF) toolbox v5.3.1 is freely available for download in the TAPAS package at <https://translationalneuromodeling.github.io/tapas>^{10, 11}. We installed and ran the package in MATLAB and Statistics Toolbox Release 2016a (MathWorks®, Natick, MA).

We estimated perceptual parameters individually for the first and second halves of the task (i.e., blocks 1 and 2). Each participant's choices (i.e., deck 1, 2, or 3) and outcomes (win or loss) were entered as separate column vectors with rows corresponding to trials. Wins were encoded as '1', losses as '0', and choices as '1', '2', or '3'. We selected the autoregressive 3-level HGF multi-arm bandit configuration for our perceptual model and paired it with the softmax-mu03 decision model.

Statistics.

Statistical analyses and effect size calculations were performed with an alpha of 0.05 and two-tailed p-values in RStudio: Integrated Development Environment for R, Version 1.3.959.

Bayes Factors (BF₁₀) were reported for nonsignificant t-tests and ANOVAs to provide additional evidence of no effect (or no differences)⁵². We define the null hypothesis (H₀) as there being no difference in the means of behavior/demographics between groups (H₀: μ₁ - μ₂ = 0), and the alternative hypothesis (H₁) as a difference (H₀: μ₁ - μ₂ ≠ 0). Interpretations of the BF₁₀ were adopted from Lee and Wagenmakers, 2013⁵³.

Independent samples t-tests were conducted to compare questionnaire item responses between high and low paranoia groups. Distributions of demographic and mental health characteristics across paranoia groups were evaluated by Chi-Square Exact tests (two groups) or Monte Carlo tests (more than 2 groups). Correlations were computed with Pearson's rho.

HGF parameter estimates and behavioural patterns (win-switch and lose-stay rates) were analyzed by repeated measures and split-plot ANOVAs (i.e., block designated as within-subject factor; pandemic, paranoia group, and social versus non-social condition as between-subject factors). Model parameters were corrected for multiple comparisons using the Benjamini Hochberg⁵⁴ method with a false discovery rate of 0.05 in analyses of variance across experiments.

Data availability

The data that support this paper are available at <https://github.com/psuthaharan/covid19paranoia>

Code availability

The code used to analyse the data and generate the figures is available at <https://github.com/psuthaharan/covid19paranoia>

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgements

This work was supported by the Yale University Department of Psychiatry, the Connecticut Mental Health Center (CMHC) and Connecticut State Department of Mental Health and Addiction Services (DMHAS). It was funded by an IMHRO / Janssen Rising Star Translational Research Award, an Interacting Minds Center (Aarhus) Pilot Project Award, NIMH R01MH12887 (P.R.C.), NIMH R21MH120799-01 (P.R.C. & S.M.G.), and an Aarhus Universitets Forskningsfond Starting Grant (C.D.M.). E.J.R. was supported by the NIH Medical Scientist Training Program Training Grant, GM007205; NINDS Neurobiology of Cortical Systems Grant, T32 NS007224; and a Gustavus and Louise Pfeiffer Research Foundation Fellowship. S.U. received funding from an NIH T32 fellowship (MH065214). S.M.G. and J.R.T. were supported by NIDA DA DA041480. The funders had no role in study design, data collection and analysis, decision to publish or preparation of the manuscript. L.L., J.R., and A.J.M. are employees of CloudResearch. We dedicate this work to the late Bob Malison, whose enthusiasm and encouragement galvanized us during uncertain times.

References

1. van Prooijen JW & Douglas KM Conspiracy theories as part of history: The role of societal crisis situations. *Mem Stud* 10, 323–333 (2017). [PubMed: 29081831]
2. Smallman SWhom do You Trust? Doubt and Conspiracy Theories in the 2009 Influenza Pandemic. *Journal of International and Global Studies* 6, 1–24 (2015).
3. Raihani NJ & Bell V An evolutionary perspective on paranoia. *Nat Hum Behav* 3, 114–121 (2019). [PubMed: 30886903]
4. Feeney EJ, Groman SM, Taylor JR & Corlett PR Explaining Delusions: Reducing Uncertainty Through Basic and Computational Neuroscience. *Schizophr Bull* 43, 263–272 (2017). [PubMed: 28177090]
5. Reed EJ, et al. Paranoia as a deficit in non-social belief updating. *Elife* 9 (2020).
6. Freeman D, et al. The revised Green et al., Paranoid Thoughts Scale (R-GPTS): psychometric properties, severity ranges, and clinical cut-offs. *Psychol Med*, 1–10 (2019).
7. Soltani A & Izquierdo A Adaptive learning under expected and unexpected uncertainty. *Nat Rev Neurosci* (2019).

8. Ryder AG, Costa PT & Bagby RM Evaluation of the SCID-II personality disorder traits for DSM-IV: coherence, discrimination, relations with general personality traits, and functional impairment. *J Pers Disord* 21, 626–637 (2007). [PubMed: 18072864]
9. Corlett PR, Fletcher PC Computational Psychiatry: A Rosetta Stone linking the brain to mental illness. *Lancet Psychiatry* (2014).
10. Mathys C, Daunizeau J, Friston KJ & Stephan KE A bayesian foundation for individual learning under uncertainty. *Frontiers in human neuroscience* 5, 39 (2011). [PubMed: 21629826]
11. Mathys CD, et al. Uncertainty in perception and the Hierarchical Gaussian Filter. *Frontiers in human neuroscience* 8, 825 (2014). [PubMed: 25477800]
12. Lefebvre G, Nioche A, Bourgeois-Gironde S & Palminteri S Contrasting temporal difference and opportunity cost reinforcement learning in an empirical money-emergence paradigm. *Proc Natl Acad Sci U S A* 115, E11446–E11454 (2018). [PubMed: 30442672]
13. Marinescu IE, Lawlor PN & Kording KP Quasi-experimental causality in neuroscience and behavioural research. *Nat Hum Behav* 2, 891–898 (2018). [PubMed: 30988445]
14. Angrist JA, Pischke J-S. *Mostly Harmless Econometrics* (Princeton University Press, Princeton, 2008).
15. Jaeger DA, Joyce TJ, Kaestner RA Cautionary Tale of Evaluating Identifying Assumptions: Did Reality TV Really Cause a Decline in Teenage Childbearing? *Journal of Business & Economic Statistics* 38 (2020).
16. Goodman-Bacon A, Marcus J Using Difference-in-Differences to Identify Causal Effects of COVID-19 Policies. *Survey Research Methods* 14, 153–158 (2020).
17. Cohn N *The Pursuit of the Millennium* (Oxford University Press, Oxford, 1961).
18. Cameron AC, Miller DLA *Practitioner’s Guide to Cluster-Robust Inference*. *The Journal of Human Resources* 50, 317–372 (2015).
19. Harrington JR & Gelfand MJ Tightness-looseness across the 50 united states. *Proc Natl Acad Sci U S A* 111, 7990–7995 (2014). [PubMed: 24843116]
20. Dynata TNY Estimates from The New York Times, based on roughly 250,000 interviews conducted by Dynata from July 2 to July 14. (2020).
21. Blakemore SJ, Sarfati Y, Bazin N & Decety J The detection of intentional contingencies in simple animations in patients with delusions of persecution. *Psychol Med* 33, 1433–1441 (2003). [PubMed: 14672252]
22. Moss AJ, Rosenzweig C, Robinson J, Litman L Demographic Stability on Mechanical Turk Despite COVID-19. *Trends Cogn Sci* 24 (2020).
23. Litman L, Robinson J, & Abberbock T TurkPrime.com: A versatile crowdsourcing data acquisition platform for the behavioral sciences. *Behavior research methods* 49, 433–442 (2017). [PubMed: 27071389]
24. Imhoff R, Lamberty P How paranoid are conspiracy believers? Toward a more fine-grained understanding of the connect and disconnect between paranoia and belief in conspiracy theories. *European Journal of Social Psychology* 48, 909–926 (2018).
25. Freeman D, et al. Coronavirus conspiracy beliefs, mistrust, and compliance with government guidelines in England. *Psychol Med*, 1–13 (2020).
26. Colombo M Two neurocomputational building blocks of social norm compliance. *Biological Philosophy* 29, 71–88 (2014).
27. Corlett PR, et al. Disrupted prediction-error signal in psychosis: evidence for an associative account of delusions. *Brain : a journal of neurology* 130, 2387–2400 (2007). [PubMed: 17690132]
28. Corlett PR, Taylor JR, Wang XJ, Fletcher PC & Krystal JH Toward a neurobiology of delusions. *Progress in neurobiology* 92, 345–369 (2010). [PubMed: 20558235]
29. Romaniuk L, et al. Midbrain activation during Pavlovian conditioning and delusional symptoms in schizophrenia. *Archives of general psychiatry* 67, 1246–1254 (2010). [PubMed: 21135324]
30. Ostrom E Collective Action and the Evolution of Social Norms. *Journal of Economic Perspectives* 14, 137–158 (2000).
31. Worobey M, et al. Origin of AIDS: contaminated polio vaccine theory refuted. *Nature* 428, 820 (2004). [PubMed: 15103367]

32. Gonsalves G & Staley P Panic, paranoia, and public health--the AIDS epidemic's lessons for Ebola. *N Engl J Med* 371, 2348–2349 (2014). [PubMed: 25372947]
33. Giubilini A & Savulescu J Vaccination, Risks, and Freedom: The Seat Belt Analogy. *Public Health Ethics* 12, 237–249 (2019). [PubMed: 32082418]
34. Robertson LRoad death trend in the United States: implied effects of prevention. *J Public Health Policy* 39, 193–202 (2018). [PubMed: 29531304]
35. Heyes C & Pearce JM Not-so-social learning strategies. *Proceedings. Biological sciences / The Royal Society* 282 (2015).
36. Freeman D, et al. Concomitants of paranoia in the general population. *Psychol Med* 41, 923–936 (2011). [PubMed: 20735884]
37. Pot-Kolder R, Veling W, Counotte J & van der Gaag M Self-reported Cognitive Biases Moderate the Associations Between Social Stress and Paranoid Ideation in a Virtual Reality Experimental Study. *Schizophr Bull* 44, 749–756 (2018). [PubMed: 29040776]
38. Henco L, et al. Bayesian modelling captures inter-individual differences in social belief computations in the putamen and insula. *Cortex* 131, 221–236 (2020). [PubMed: 32571519]
39. Henco L, et al. Aberrant computational mechanisms of social learning and decision-making in schizophrenia and borderline personality disorder. *PLoS Comput Biol* 16, e1008162 (2020). [PubMed: 32997653]
40. Heyes C Précis of Cognitive Gadgets: The Cultural Evolution of Thinking. *Behav Brain Sci* 42 (2018).
41. DiGrazia J The Social Determinants of Conspiratorial Ideation. *Socius: Sociological Research for a Dynamic World* (2017).
42. Bureau U American Community Survey. (2017).
43. Gelfand MJ, et al. Differences between tight and loose cultures: a 33-nation study. *Science* 332, 1100–1104 (2011). [PubMed: 21617077]
44. Beck AT, Epstein N, Brown G & Steer RA An inventory for measuring clinical anxiety: psychometric properties. *J Consult Clin Psychol* 56, 893–897 (1988). [PubMed: 3204199]
45. Beck AT, Ward CH, Mendelson M, Mock J & Erbaugh J An inventory for measuring depression. *Archives of general psychiatry* 4, 561–571 (1961). [PubMed: 13688369]
46. Abramowitz JS, et al. Assessment of obsessive-compulsive symptom dimensions: development and evaluation of the Dimensional Obsessive-Compulsive Scale. *Psychol Assess* 22, 180–198 (2010). [PubMed: 20230164]
47. Knotek II E, Schoenle R, Dietrich A, Müller G, Myrseth KOR, & Weber M . Consumers and COVID-19: Survey Results on Mask-Wearing Behaviors and Beliefs. *Economic Commentary* (2020).
48. Enders A, Uscinski JE et al. . Who Supports QAnon? A Case Study in Political Extremism. available at joescinski.com (2021).
49. Policy, T.C.S.o.P. Unemployment Rate by State. (2020).
50. Times NY An ongoing repository of data on coronavirus cases and deaths in the U.S . (2020).
51. Ballotopedia. Status of lockdown and stay-at-home orders in response to the coronavirus (COVID-19) pandemic. (2020).
52. Gelman A, Stern H The Difference Between “Significant” and “Not Significant” is not Itself Statistically Significant. *American Statistician* 60 (2006).
53. Lee MD, Wagenmakers EJ Bayesian cognitive modeling: A practical course (Cambridge University Press, 2013).
54. Hochberg Y & Benjamini Y More powerful procedures for multiple significance testing. *Stat Med* 9, 811–818 (1990). [PubMed: 2218183]

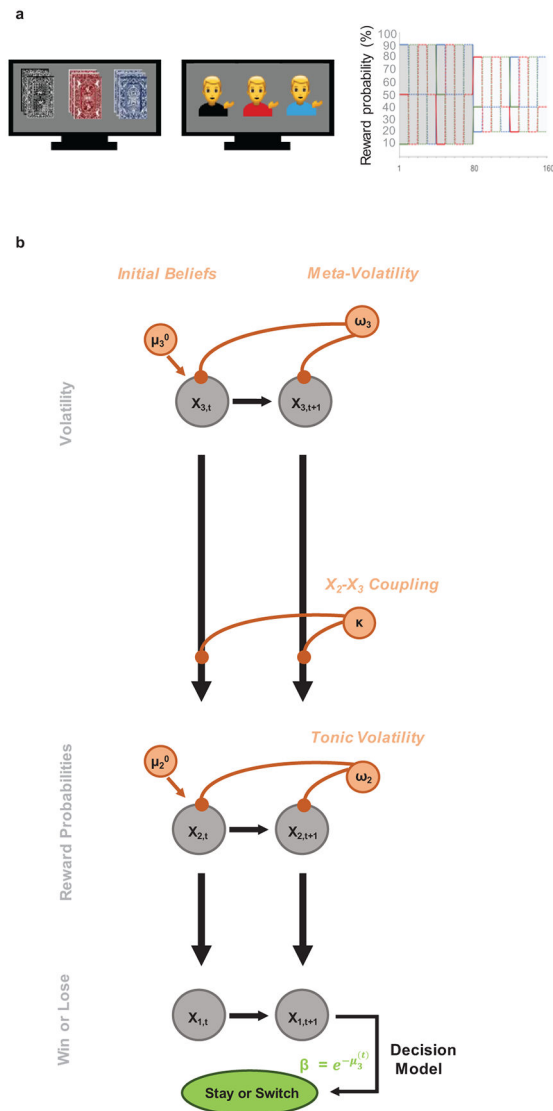


Figure 1. Probabilistic Reversal Learning and Hierarchical Gaussian Filter.

Depictions of our behavioural tasks and computational model used to ascertain belief-updating behaviour. **a**, Non-social and social task stimuli and reward contingency schedule.

b, Hierarchical model for capturing changes in beliefs under task environment volatility.

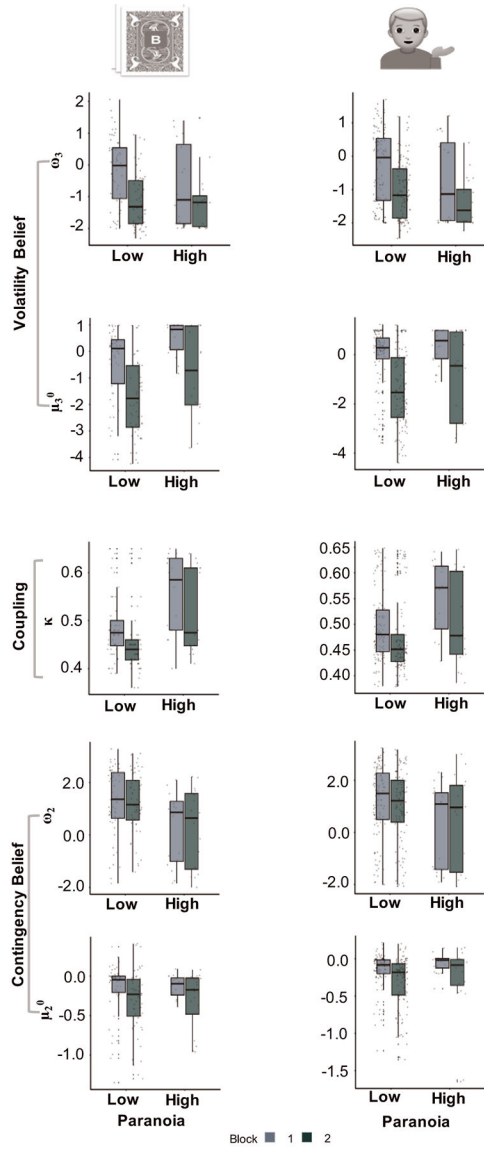


Figure 2. Pre-pandemic (N=202) social and non-social reversal learning.

a, non-social and social task stimuli and reward contingency schedule. **b**, in both non-social ($n=72$) and social ($n=130$) tasks, high paranoia subjects exhibit elevated priors for volatility (μ_3^0 ; group: $F_{(1, 198)}=8.566$, $p=0.004$, $\eta_p^2=0.035$; block: $F_{(1, 198)}=161.845$, $p < 0.001$, $\eta_p^2=0.11$) and contingency (μ_2^0 ; block: $F_{(1, 198)}=36.58$, $p < 0.001$, $\eta_p^2=0.042$), were slower to update those beliefs (ω_2 ; group: $F_{(1, 198)}=12.986$, $p < 0.001$, $\eta_p^2=0.053$, ω_3 ; group: $F_{(1, 198)}=4.447$, $p=0.036$, $\eta_p^2=0.014$, block: $F_{(1, 198)}=38.89$, $p < 0.001$, $\eta_p^2=0.064$), and had higher coupling (κ group: $F_{(1, 198)}=21.45$, $p < 0.001$, $\eta_p^2=0.08$, block: $F_{(1, 198)}=30.281$, $p < 0.001$, $\eta_p^2=0.031$) between volatility and contingency beliefs. **c**, 3-level HGF perceptual model (gray) with a softmax decision model (green). Level 1 (x_1): trial-by-trial perception of win or loss feedback. Level 2 (x_2): stimulus-outcome associations (i.e., deck values). Level 3 (x_3): perception of the overall reward contingency context. The impact of phasic volatility upon x_2 is captured by κ (i.e., coupling). Tonic volatility modulates x_3 and x_2 via ω_3 and ω_2 , respectively. μ_3^0 is the initial value of the third level volatility belief. **Box-plots:** Centre

lines show the medians; box limits indicate the 25th and 75th percentiles; whiskers extend 1.5 times the interquartile range from the 25th and 75th percentiles, outliers are represented by dots; crosses represent sample means; data points are plotted as open circles.

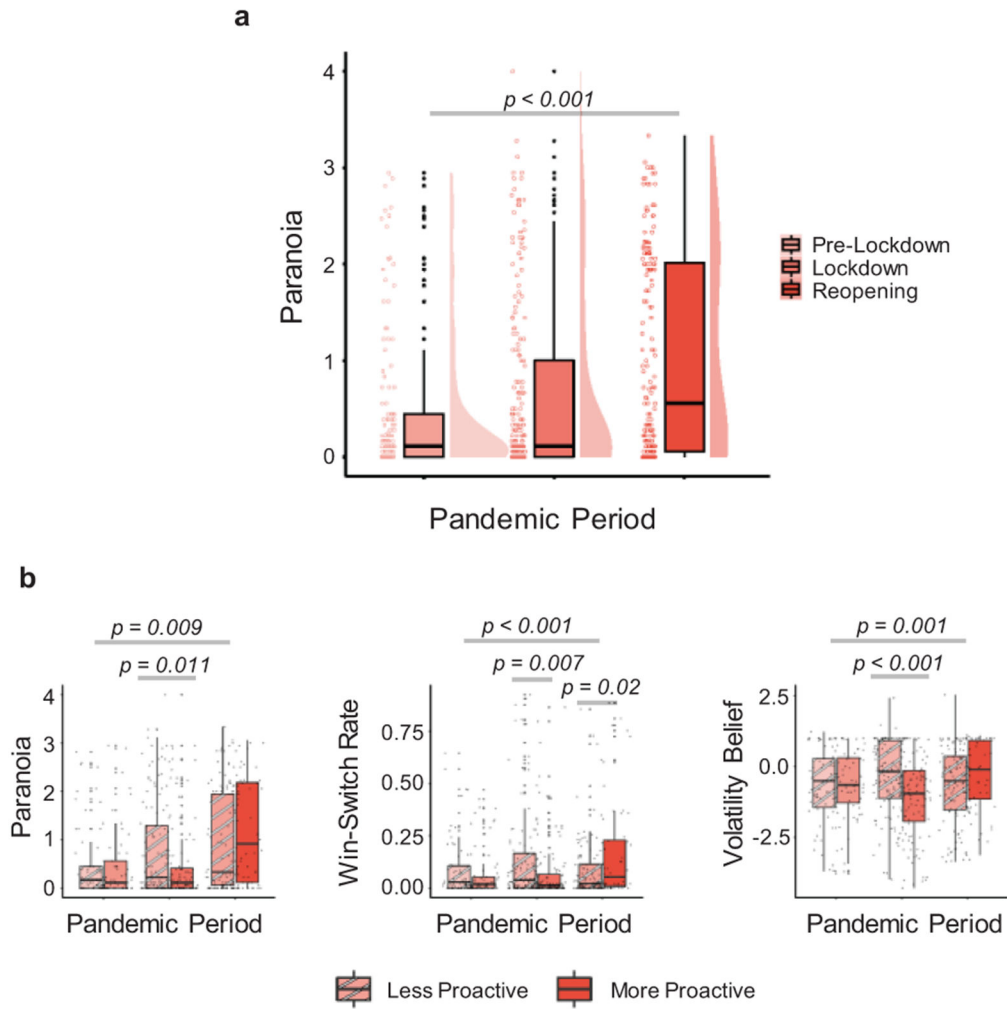


Figure 3. Paranoia, state proactivity, task behaviour, and belief updating during a pandemic. Paranoia increased as the pandemic progressed ($F_{(2, 530)}=14.7, p < 0.001, \eta_p^2=0.053$). **a**, self-rated paranoia ($N=533$), prior to the pandemic, during lockdown and following reopening. **b**, We observed a main effect of pandemic period ($F_{(2, 527)}=4.948, p = 0.007, \eta_p^2=0.018$) and a state proactivity by period interaction ($F_{(2, 527)}=4.785, p = 0.009, \eta_p^2=0.018$) for paranoia and win-switch behaviour (main effect: $F_{(2, 527)}=3.270, p = 0.039, \eta_p^2=0.012$; interaction: $F_{(2, 527)}=8.747, p < 0.001, \eta_p^2=0.032$) and volatility priors ($F_{(2, 527)}=8.623, p = 0.001, \eta_p^2=0.032$) we observed significant interactions between pandemic period and the proactivity of policies.

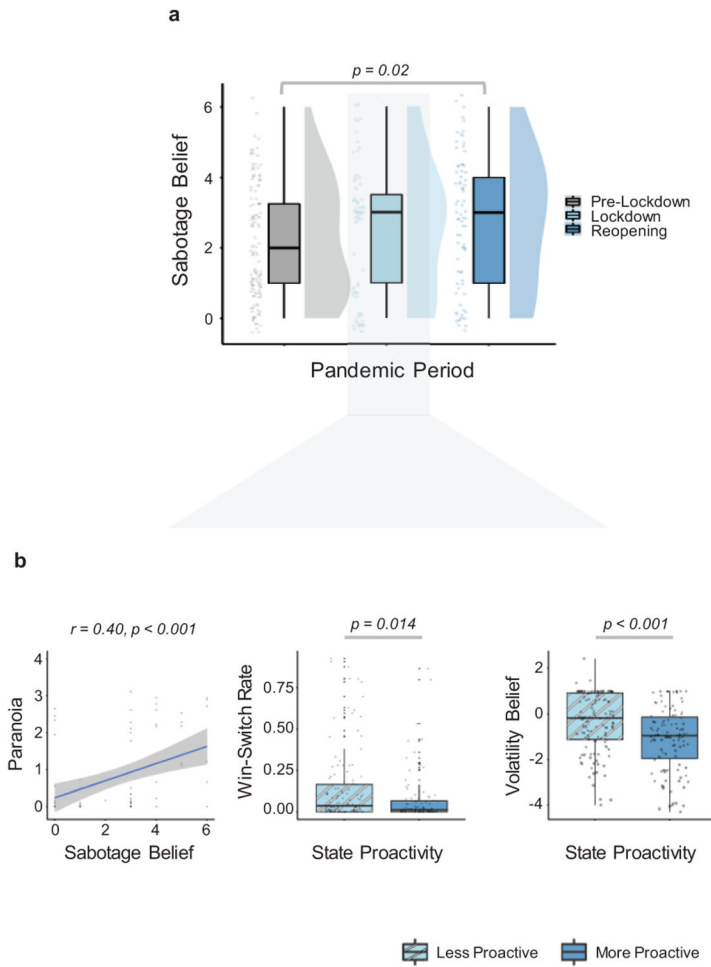


Figure 4. Sabotage belief and the effects of lockdown (social task; N=280).

a, sabotage belief, the conviction that an avatar-partner deliberately caused a loss in points, increased from pre-lockdown to reopening ($t_{145} = -2.35$, $p = 0.02$, $d = 0.349$, 95% CI = $(-1.114, -0.096)$). **b**, Self-rated paranoia in the real world correlated with sabotage belief in the task (Figure 3b; $r = 0.40$, $p < 0.001$). During lockdown, when proactive state responses were associated with decreased self-rated paranoia, win-switch rate ($t_{216} = 2.73$, $p = 0.014$, $d = 0.351$, 95% CI = $(0.019, 0.117)$) and μ_3^0 ($t_{223} = 4.20$, $p < 0.001$, $d = 4.299$, 95% CI = $(0, 1.647)$) were significantly lower in participants from states with more vigorous lockdown. Analysis performed on individuals who responded to sabotage question.

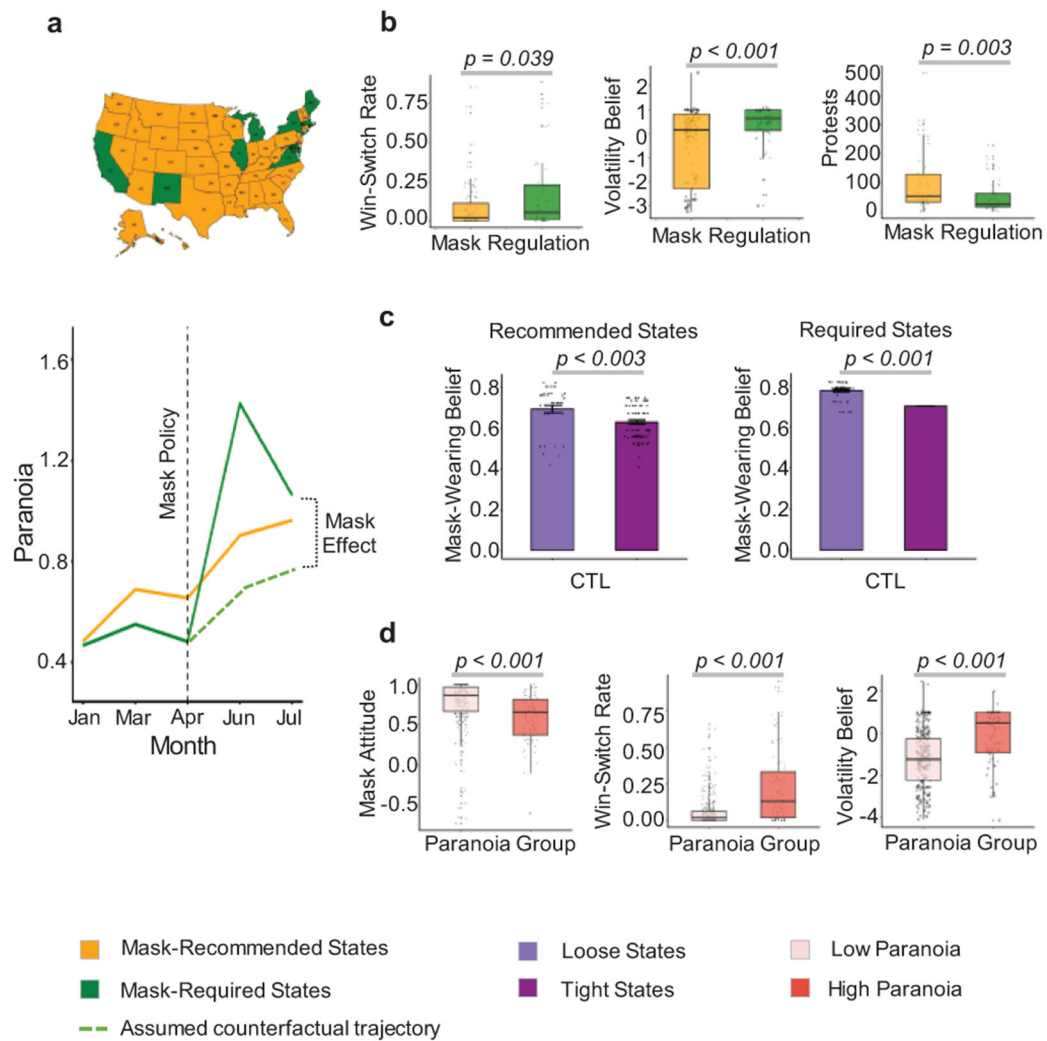


Figure 5. Effects of mask policy on paranoia and belief-updating.

We observe a significant increase in paranoia and perceived volatility, especially in states that have issued a state-wide mask mandate. **a**, Map of the US states color-coded to their respective mask policy ($n_{rec}=40$, $n_{req}=11$) and a Differences-in-Differences analysis (bottom; $N=533$, $\delta_{DID}=0.396$, $p=0.038$) of mask rules suggests a 40% increase in paranoia in states that mandate mask-wearing. **b**, Win-switch rate (left; $N=172$, $n_{rec}=120$, $n_{req}=52$, $t_{67}=-2.4$, $p=0.039$, $d=0.483$, 95% CI = $(-0.164, -0.015)$) and volatility belief (middle; $t_{141}=-3.7$, $p<0.001$, $d=-3.739$, 95% CI = $(0, -1.585)$) are higher in mask-mandate states, but more protests per day occurred in mask-recommended states (right; $N=110$, $n_{rec}=55$, $n_{req}=55$, $t_{83}=3.10$, $p=0.0027$, $d=0.591$, 95% CI = $(17.458, 80.142)$). **c**, Effects of Cultural Tightness and Looseness (CTL) in mask-recommended states (left; $N=120$, $n_{loose}=38$, $n_{tight}=82$, $t_{57}=3.06$, $p=0.003$, $d=0.663$, 95% CI = $(0.022, 0.107)$) and mask-required states (right; $N=52$, $n_{loose}=48$, $n_{tight}=4$, $t_{47}=12.84$, $p<0.001$, $d=1.911$, 95% CI = $(0.064, 0.088)$) implicating violation of social norms in the genesis of paranoia. **d**, Follow-up study ($N=405$, $n_{low}=314$, $n_{high}=91$) illustrating that high paranoia participants are less inclined to wear masks in public (left; $t_{158}=4.59$, $p<0.001$, $d=0.520$, 95% CI = $(0.091, 0.229)$), have more

promiscuous switching behaviour (middle; $t_{138} = -6.40$, $p < 0.001$, $d = 1.148$, 95% CI = $(-0.227, -0.120)$) and elevated prior beliefs about volatility (right; $t_{138} = -6.04$, $p < 0.001$, $d = -6.041$, 95% CI = $(0, -2.067)$).

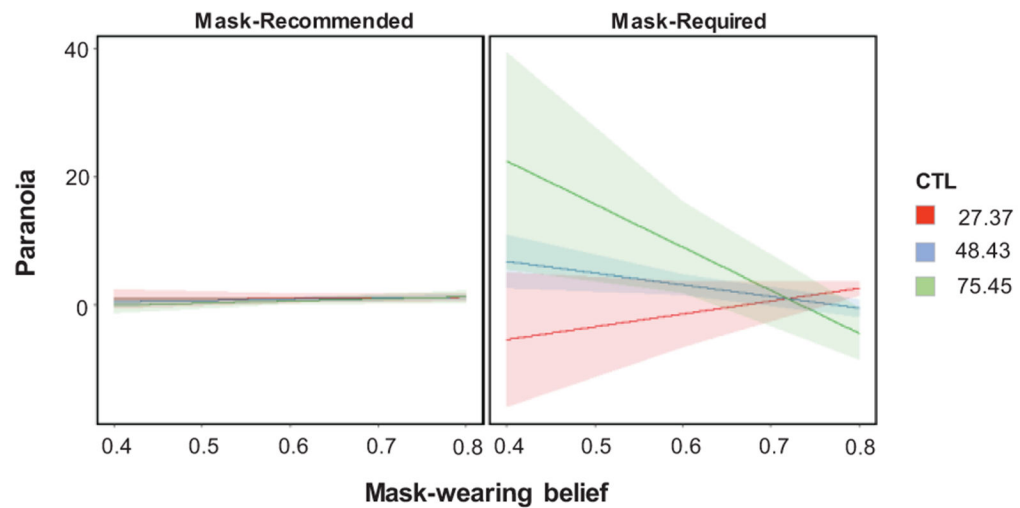


Figure 6. Predicting paranoia from pandemic features.

Regression model predictions (N=172) in states where masks were recommended (Left Panel) versus mandated (right panel). Paranoia predictions based on estimated state mask-wearing (x-axis, low mask-wearing to high mask-wearing) and cultural tightness. **Red** – Loose states, that do not prize conformity, **Blue** - states with median tightness, **Green** – tight states that are conservative and rule-following. Paranoia is highest when mask wearing is low, in culturally tight states with a mask-wearing mandate ($F_{(11,160)}=1.91, p=0.04$). Values represent high, median, and low estimated state tightness.

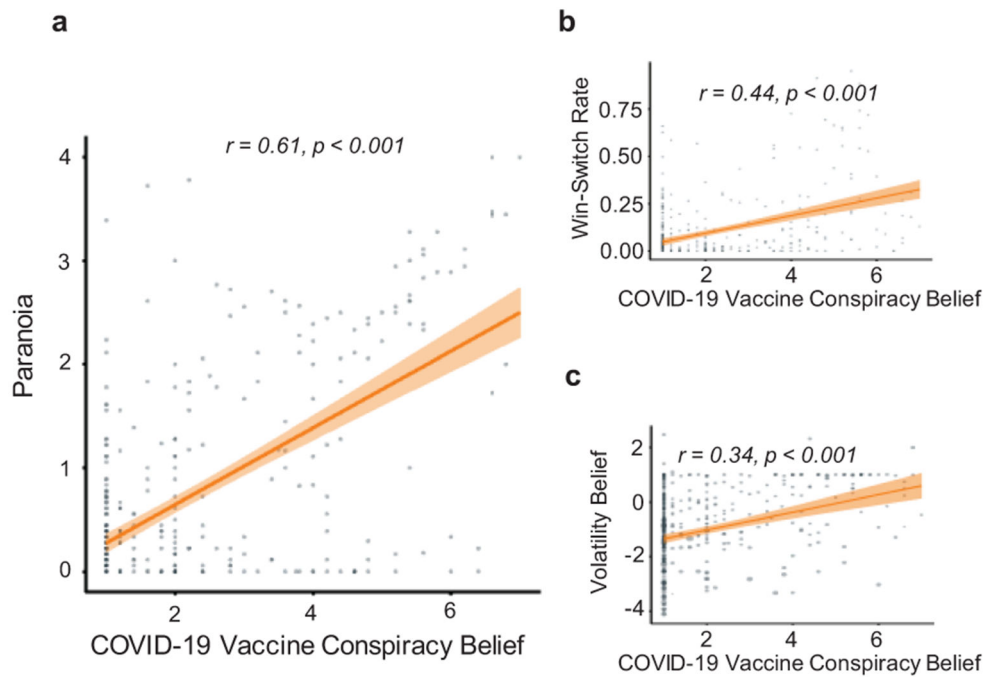


Figure 7. Relating vaccine conspiracy beliefs to paranoia and task behaviour.

We assayed individual's (N=403) COVID-19 vaccine conspiracy beliefs to investigate underlying relationships to behaviour. **a**, We find individuals with higher paranoia endorse more vaccine conspiracies ($r=0.61, p<0.001$). **b**, COVID conspiracy beliefs were correlated with erratic task behaviour ($r=0.44, p<0.001$), and **c**, perturbed volatility priors ($r=0.34, p<0.001$). Analysis performed on individuals who responded to covid vaccine conspiracy questions.

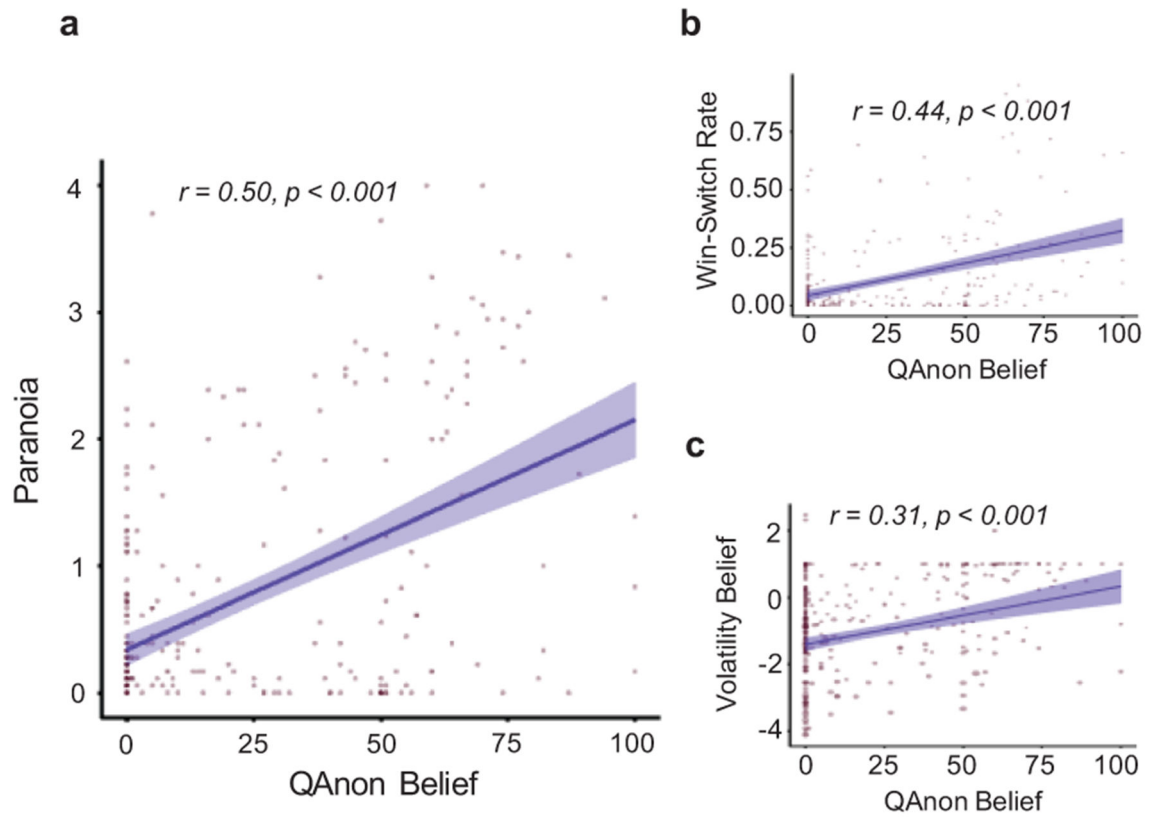


Figure 8. Relating QAnon beliefs to paranoia and task behaviour.

(a) We find individuals ($N=307$) with higher paranoia endorsed more QAnon beliefs ($r=0.5$, $p<0.001$). Similarly, (b) QAnon beliefs were strongly correlated with erratic task behaviour ($r=0.44$, $p<0.001$) and (c) perturbed volatility priors ($r=0.31$, $p<0.001$). Analysis performed on individuals who responded to QAnon questions.