



Paranoia and belief updating during the COVID-19 crisis

Praveen Suthaharan^{1,12}, Erin J. Reed^{2,3,12}, Pantelis Leptourgos¹, Joshua G. Kenney¹, Stefan Uddenberg⁴, Christoph D. Mathys^{5,6,7,8}, Leib Litman⁹, Jonathan Robinson⁹, Aaron J. Moss⁹, Jane R. Taylor^{1,10}, Stephanie M. Groman¹ and Philip R. Corlett^{1,10,11}✉

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 made their belief updating more erratic. A proactive lockdown made 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 and the QAnon conspiracy theories. These beliefs were associated with erratic task behaviour and changed priors. Taken together, we found that real-world uncertainty increases paranoia and influences laboratory task behaviour.

Crisis, 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 (Fig. 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 preregister our experiments. Our interests evolved as the pandemic did. We chose to continue gathering data on participants' belief updating and 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 (Fig. 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 signified 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 was 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 after positive feedback⁵. We compared those data (gathered via the Amazon Mechanical Turk Marketplace in the USA between December 2017 and August 2018; Supplementary Table 1) to a new task version with identical contingencies but framed socially (Fig. 1a). Instead of selecting between decks of cards ('non-social task'), participants (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 participants with low paranoia after receiving positive feedback in both; however, there were no substantial behavioural differences between tasks (Supplementary Fig. 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; Supplementary Fig. 2b). There were also no differences in points

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, Institute for Biomedical Engineering, University of Zurich, Zurich, Switzerland. ⁷Swiss Federal Institute of Technology Zurich, Zurich, Switzerland. ⁸Scuola Internazionale Superiore di Studi Avanzati, Trieste, Italy. ⁹CloudResearch, New York, NY, USA. ¹⁰Department of Psychology, Yale University, New Haven, CT, USA. ¹¹Wu Tsai Institute, Yale University, New Haven, CT, USA. ¹²These authors contributed equally: Praveen Suthaharan, Erin J. Reed. ✉e-mail: philip.corlett@yale.edu

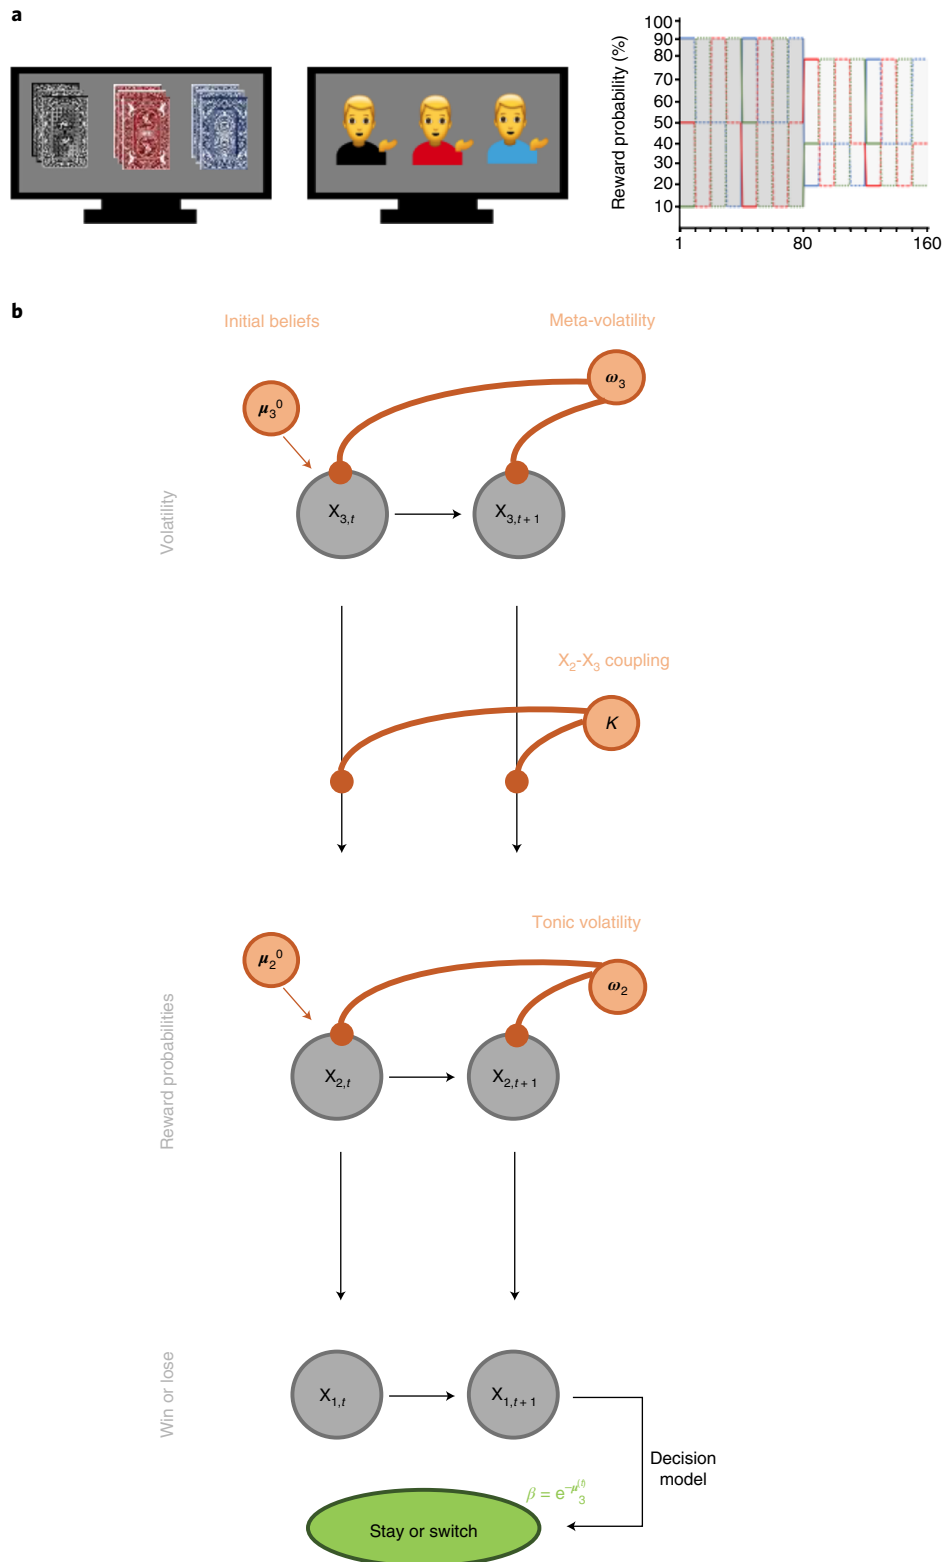


Fig. 1 | Probabilistic reversal learning and hierarchical Gaussian filter. Depictions of our behavioural tasks and computational model used to ascertain belief-updating behaviour. **a**, Non-social and social task stimuli and reward contingency schedule. **b**, Hierarchical model for capturing changes in beliefs under task environment volatility.

($BF_{10}=0.163$, strong evidence for the null hypothesis) or reversals achieved ($BF_{10}=0.210$, strong evidence for the null hypothesis) between 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 modelling

participants' choices, we could estimate latent processes⁹. We supposed that they continually updated a probabilistic representation of the task (a generative model), which guided their behaviour^{10,11}. To estimate their generative models, we identified: (1) a set of prior assumptions about how events are caused by the environment (the perceptual model); and (2) the behavioural consequences of their posterior beliefs about options and outcomes (the response model^{10,11}). Inverting the response model also entailed inverting the perceptual model and yielded a mapping from task cues to the beliefs that caused participants' responses^{10,11} (Fig. 1b).

The perceptual model (Fig. 1b) consists of three hierarchical layers of belief about the task, represented as probability distributions that 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 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 (that is, '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 and underwrites perceived change in the underlying statistics of the environment (that is, 'the world is changing'), which may call for more wholesale belief revision. The layers of beliefs are fed through a sigmoid response function (Fig. 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 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 (Fig. 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 , the 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 —the evolution rate of reward contingencies. Thus, we observed an impact of paranoia on behaviour 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 a higher learning rate for unexpected events

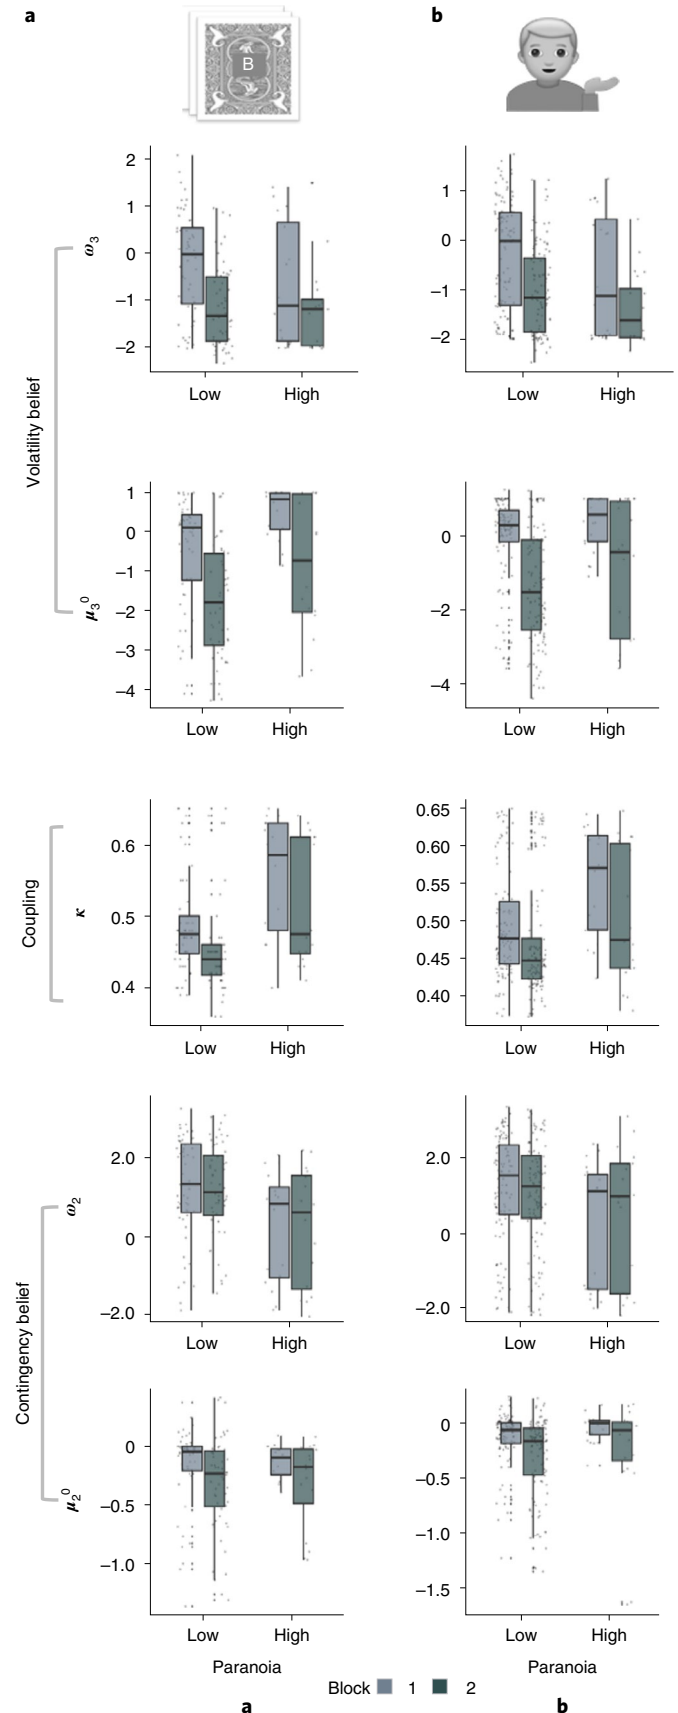


Fig. 2 | Prepandemic (n = 202) social and non-social reversal learning. a, Non-social task (n = 72), volatility beliefs, coupling and contingency beliefs. **b**, Social task (n = 130), volatility beliefs, coupling and contingency beliefs. In both tasks high paranoia subjects exhibit elevated priors for 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: $F_{(1, 198)} = 36.58$, $p = 0.042$), were slower to update those beliefs (ω_2 ; group: $F_{(1, 198)} = 12.986$, $p = 0.053$, 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: $F_{(1, 198)} = 21.45$, $p = 0.008$, block: $F_{(1, 198)} = 30.281$, $p = 0.001$, $\eta_p^2 = 0.031$) between volatility and contingency beliefs. The centre horizontal lines within the plots represent the median values, the boxes span from the 25th to the 75th percentile and the whiskers extend to 1.5x the interquartile range.

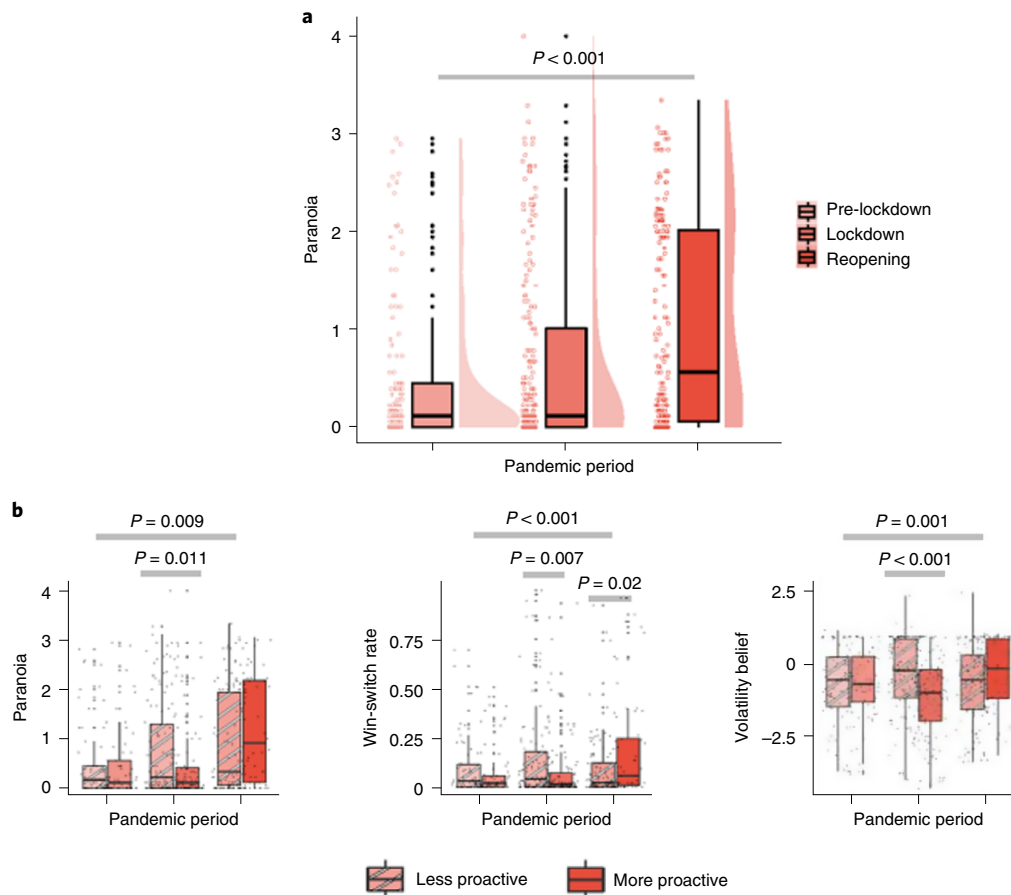


Fig. 3 | Paranoia, state proactivity, task behaviour and belief updating during a pandemic. Paranoia increased as the pandemic progressed ($F_{(2, 530)} = 14.7$, $P < 0.001$, $\eta_p^2 = 0.053$). **a**, Self-rated paranoia ($n = 533$) before the pandemic, during lockdown and after reopening. **b**, We observed a main effect of the 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. The centre horizontal lines within the plots represent the median values, the boxes span from the 25th to the 75th percentile and the whiskers extend to 1.5x the interquartile range.

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 March 2020–17 July 2020; Supplementary Tables 2 and 3). We did not preregister 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 (Fig. 3a; $F_{(2, 530)} = 14.7$, $P < 0.001$, $\eta_p^2 = 0.053$). Anxiety increased (Supplementary Fig. 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.

In the USA, states responded differently to the pandemic; some instituted lockdowns early and broadly (more proactive), whereas others closed later and reopened sooner (less proactive) (equation (1) and Supplementary Fig. 3). When they reopened, some states mandated mask-wearing (more proactive) while 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 the pandemic period (Fig. 3b; $F_{(2, 527)} = 4.948$, $P = 0.007$, $\eta_p^2 = 0.018$) and a state proactivity by period interaction (Fig. 3b; $F_{(2, 527)} = 4.785$, $P = 0.009$, $\eta_p^2 = 0.018$) for paranoia and win-switch behaviour (Fig. 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 (Fig. 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 (Fig. 3b; $t_{227} = 2.57$, $P = 0.011$, Cohen's $d = 0.334$, 95% confidence interval (CI) = 0.071–0.539), less erratic win-switching (Fig. 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 (Fig. 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 (Fig. 3b; $t_{67} = -2.39$, $P = 0.02$, Cohen's $d = 0.483$, 95% CI = -0.164 to -0.015). Furthermore, participants in mandate states had higher contamination fear (Supplementary Fig. 4; $t_{101} = -2.89$, $P = 0.005$, Cohen's $d = 0.471$, 95% CI = -0.655 to -0.121).

None of the other pandemic or policy effects on parameters (priors or learning rates) survived false discovery rate (FDR) correction for multiple comparisons. Therefore, we carried win-switch

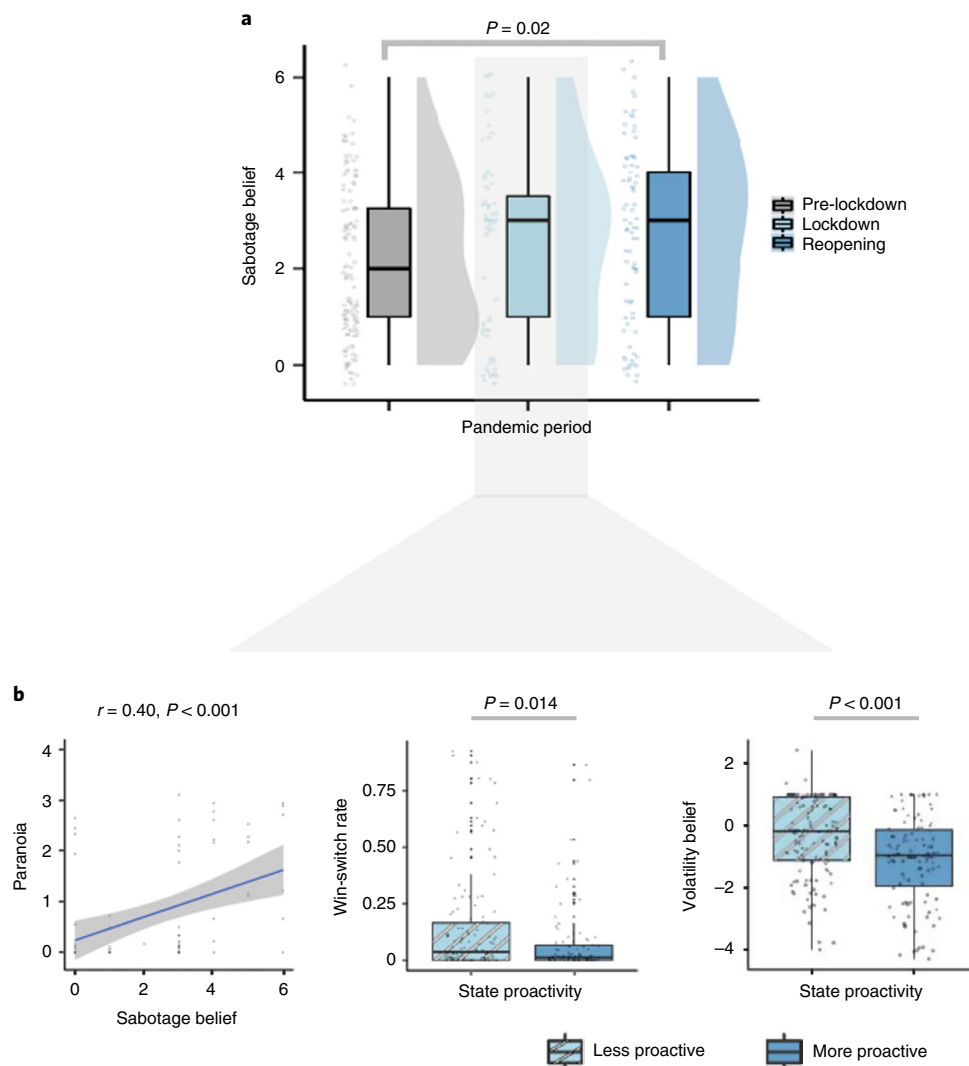


Fig. 4 | Sabotage belief and the effects of lockdown (social task; $n = 280$). **a**, Sabotage belief, the conviction that an avatar-partner deliberately caused a loss in points, increased from prelockdown to reopening ($t_{145} = -2.35$, $P = 0.02$, Cohen's $d = 0.349$, 95% CI = -1.114 to -0.096). **b**, Self-rated paranoia in the real world correlated with sabotage belief in the task (Fig. 3b; $r = 0.4$, $P < 0.001$). During lockdown, when proactive state responses were associated with decreased self-rated paranoia, the win-switch rate ($t_{216} = 2.73$, $P = 0.014$, Cohen's $d = 0.351$, 95% CI = 0.019 – 0.117) and μ_3^0 ($t_{223} = 4.2^0$, $P < 0.001$, Cohen's $d = 4.299$, 95% CI = 0 – 1.647) were significantly lower in participants from states with more vigorous lockdown. Analysis was performed on individuals who responded to the sabotage question. The centre horizontal lines within the plots represent the median values, the boxes span from the 25th to the 75th percentile and the whiskers extend to $1.5\times$ the interquartile range. The grey shaded region in **b** (leftmost graph) represents the 95% confidence interval for predictions from a linear model.

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 (Fig. 4a; $t_{145} = -2.35$, $P = 0.02$, Cohen's $d = 0.349$, 95% CI = -1.114 to -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 (Fig. 4b; $r = 0.4$, $P < 0.001$). During lockdown, when proactive state responses were associated with decreased self-rated paranoia, win-switch rate ($t_{216} = 2.73$, $P = 0.014$, Cohen's $d = 0.351$, 95% CI = 0.019 – 0.117) and μ_3^0 ($t_{223} = 4.20$, $P < 0.001$, Cohen's $d = 4.299$, 95% CI = 0 – 1.647) were significantly lower in participants from states with more vigorous lockdown (Fig. 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 (Supplementary Fig. 5). The data must be longitudinal but they need not follow the same participants¹⁴. It is essential to demonstrate that—before implementation—the areas adopting different policies are matched in terms of the trends in the variable being compared (parallel trends assumption).

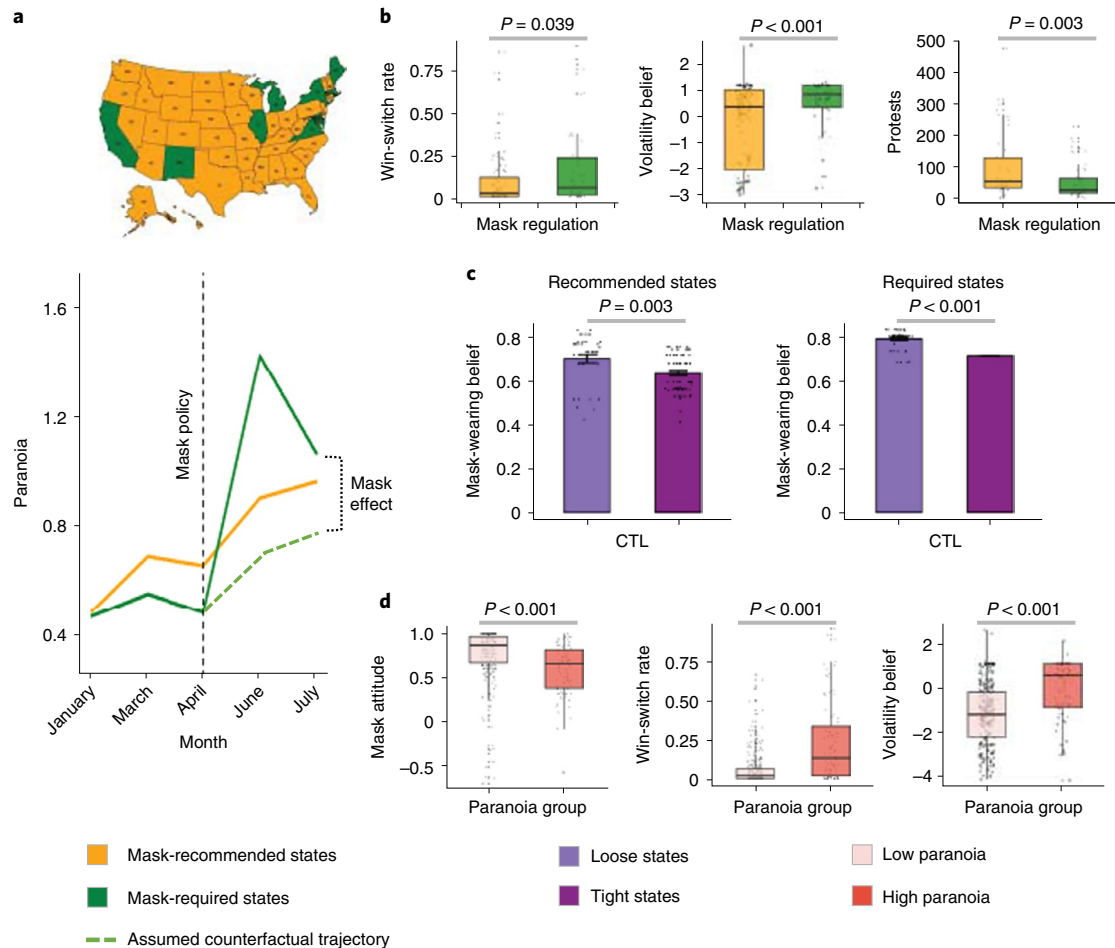


Fig. 5 | Effects of mask policy on paranoia and belief updating. We observed a significant increase in paranoia and perceived volatility, especially in states that issued a state-wide mask mandate. **a**, Map of the US states colour-coded to their respective mask policy ($n_{rec} = 40$, $n_{req} = 11$) and a DiD analysis (bottom: $n = 533$, $\delta_{DiD} = 0.396$, $P = 0.038$) of mask rules suggested a 40% increase in paranoia in states that mandated mask-wearing. **b**, Win-switch rate (left: $n = 172$, $n_{rec} = 120$, $n_{req} = 52$, $t_{67} = -2.4$, $P = 0.039$, Cohen's $d = 0.483$, 95% CI = -0.164 to -0.015) and volatility belief (middle: $t_{141} = -3.7$, $P < 0.001$, Cohen's $d = -3.739$, 95% CI = 0 to -1.585) were higher in mask-mandating states but more protests per day occurred in mask-recommended states (right: $n = 110$, $n_{rec} = 55$, $n_{req} = 55$, $t_{83} = 3.10$, $P = 0.0027$, Cohen's $d = 0.591$, 95% CI = 17.458 – 80.142). **c**, Effects of CTL in mask-recommended states (left; $n = 120$, $n_{loose} = 38$, $n_{tight} = 82$, $t_{57} = 3.06$, $P = 0.003$, Cohen's $d = 0.663$, 95% CI = 0.022 – 0.107) and mask-required states (right; $n = 52$, $n_{loose} = 48$, $n_{tight} = 4$ ($t_{47} = 12.84$, $P < 0.001$, Cohen's $d = 1.911$, 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 participants with high paranoia are less inclined to wear masks in public (left: $t_{158} = 4.59$, $P < 0.001$, Cohen's $d = 0.520$, 95% CI = 0.091 – 0.229), have more promiscuous switching behaviour (middle: $t_{138} = -6.40$, $P < 0.001$, Cohen's $d = 1.148$, 95% CI = -0.227 to -0.120) and elevated prior beliefs about volatility (right: $t_{138} = -6.04$, $P < 0.001$, Cohen's $d = -6.041$, 95% CI = 0 to -2.067). In **b** and **d**, the centre horizontal lines within the plots represent the median values, the boxes span from the 25th to the 75th percentile and the whiskers extend to $1.5\times$ the interquartile range.

Using the pretreatment outcomes, we cannot reject the null hypothesis that pretreatment 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 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 (before 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-participant design, so it is important to establish that there were no demographic differences (age, sex, ethnicity) 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; sex, $\chi^2 = 0.37$, d.f. = 1 , $P = 0.54$, $BF_{10} = 0.11$, anecdotal evidence for null hypothesis; ethnicity, 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.




Fig. 6 | Predicting paranoia from pandemic features. Regression model predictions ($n=172$) in states where masks were recommended (left) versus mandated (right). Paranoia predictions based on estimated state mask-wearing (x axis, low 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. The shaded regions represent the 95% confidence interval.

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

Variation in rule following relates to paranoia. To unpack the DiD result, we further explored whether any other features might illuminate the variation in paranoia by local mask policy¹⁹. There were state-level cultural differences – measured by the index of cultural tightness and looseness (CTL)¹⁹—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 states tend towards the Democrats¹⁹. Since 2020 was a politically tumultuous time and the pandemic was politicized, we thought it prudent to incorporate politics into our analyses.

We also tried to assess whether people were following the mask-wearing rules. We acquired independent survey data gathered in the USA from 250,000 respondents who, between 2 and 14 July, 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 (Fig. 5c).

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

Through backward linear regression with removal (equation (3)), we fitted 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 (Fig. 6 and 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 March 2020–17 July 2020; Supplementary 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 (Fig. 5d; $t_{158}=4.59$, $P<0.001$, Cohen's $d=0.520$, 95% CI = 0.091 – 0.229). Again, the win-switch rate was significantly higher in individuals with high paranoia (Fig. 5d; $t_{138}=-6.40$, $P<0.001$, Cohen's $d=1.148$, 95% CI = -0.227 to -0.120), as was their prior belief about volatility (Fig. 5d; $t_{138}=-6.04$, $P<0.001$, Cohen's $d=-6.041$, 95% CI = 0 to -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 the attitudes of individuals 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 (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 Fig. 6; $r=0.47$, $P<0.001$), which is 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 after the murder of George Floyd. These protests began on 24 May 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 24 May through to 31 July 2020 in mask-recommended versus mask-mandating states ($t_{83}=3.10$, $P=0.0027$, Cohen's $d=0.591$, 95% CI = 17.458 – 80.142). This suggests that 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 (Fig. 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-mandating versus mask-recommend states. We compared the change in unemployment from lockdown to reopening in mask-mandating versus mask-recommend states and found no significant differ-

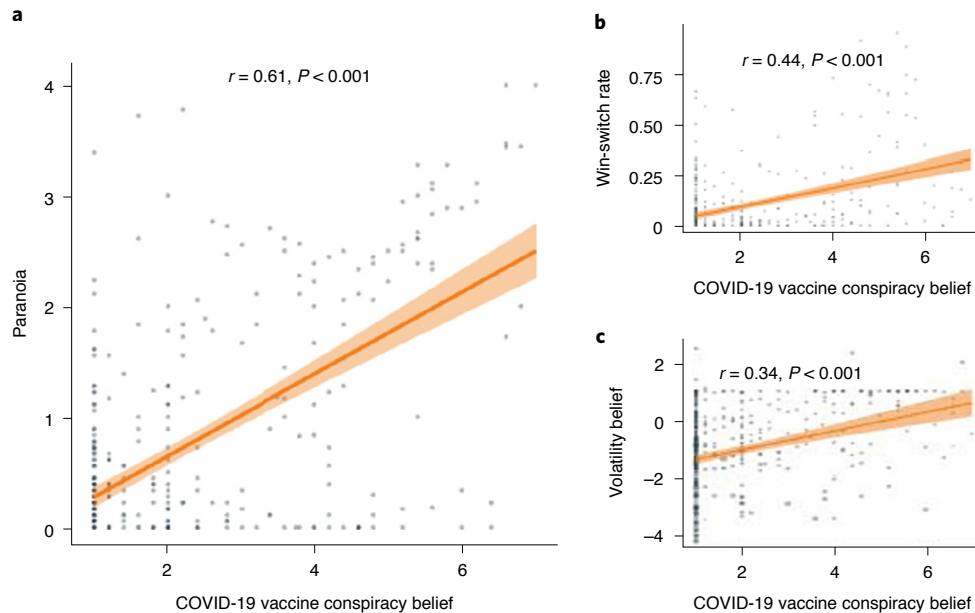


Fig. 7 | Relating vaccine conspiracy beliefs to paranoia and task behaviour. We assayed the COVID-19 vaccine conspiracy beliefs of individuals ($n = 403$) to investigate the underlying relationships to behaviour. **a**, Individuals with higher paranoia endorsed more vaccine conspiracies ($r = 0.61$, $P < 0.001$). **b**, COVID conspiracy beliefs were correlated with erratic task behaviour ($r = 0.44$, $P < 0.001$). **c**, Perturbed volatility priors ($r = 0.34$, $P < 0.001$). Analysis was performed on individuals who responded to COVID vaccine conspiracy questions. The shaded regions represent the 95% confidence interval.

ence ($t_{17} = -1.85$, $P = 0.08$, $BF_{10} = 1.04$, 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 described 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 found no differences in demographic variables across our study periods (prepandemic, lockdown, reopening, sex: $F_{(2,523)} = 0.341$, $P = 0.856$, $\eta_p^2 = 0.001$, $BF_{10} = 0.03$, 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; ethnicity: $F_{(2,520)} = 1.10$, $P = 0.856$, $\eta_p^2 = 0.004$, $BF_{10} = 0.06$, 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 Fig. 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 Fig. 8). Finally, 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 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: prepandemic, $r = 0.76$, $P = 2.2 \times 10^{-8}$; lockdown, $r = 0.78$, $P = 5.8 \times 10^{-9}$; reopening, $r = 0.81$, $P = 8.5 \times 10^{-10}$

(Supplementary Fig. 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; sex, $P = 0.77$, $BF_{10} = 0.13$, moderate evidence for the null hypothesis; ethnicity, $P = 0.34$, $BF_{10} = 0.20$, moderate evidence for the null hypothesis) versus postreopening (age, $P = 0.57$, $BF_{10} = 0.21$, moderate evidence for the null hypothesis; sex, $P = 0.77$, $BF_{10} = 0.19$, moderate evidence for the null hypothesis; ethnicity, $P = 0.07$, $BF_{10} = 0.55$, anecdotal evidence for the null hypothesis) for mask-mandating 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. While correlated, paranoia and conspiracy beliefs are not synonymous²⁴. Therefore, we also assessed conspiracy beliefs about a potential COVID vaccine in the follow-up study (Supplementary Table 5). We found that conspiracy beliefs about a vaccine correlated significantly with paranoia (Fig. 7a; $r = 0.61$, $P < 0.001$) and that such beliefs were associated with erratic task behaviour (Fig. 7b; win-switch rate: $r = 0.44$, $P < 0.001$) and perturbed volatility priors (Fig. 7c; $r = 0.34$, $P < 0.001$) in an identical manner to mask concerns and paranoia more broadly. In the UK, early in the pandemic, conspiracy theorizing was associated with higher paranoia and less adherence to public health countermeasures²⁵. We replicated and extended those findings to the USA and provided mechanistic suggestions centred 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 (Supplementary Table 5). We assessed participants' 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,

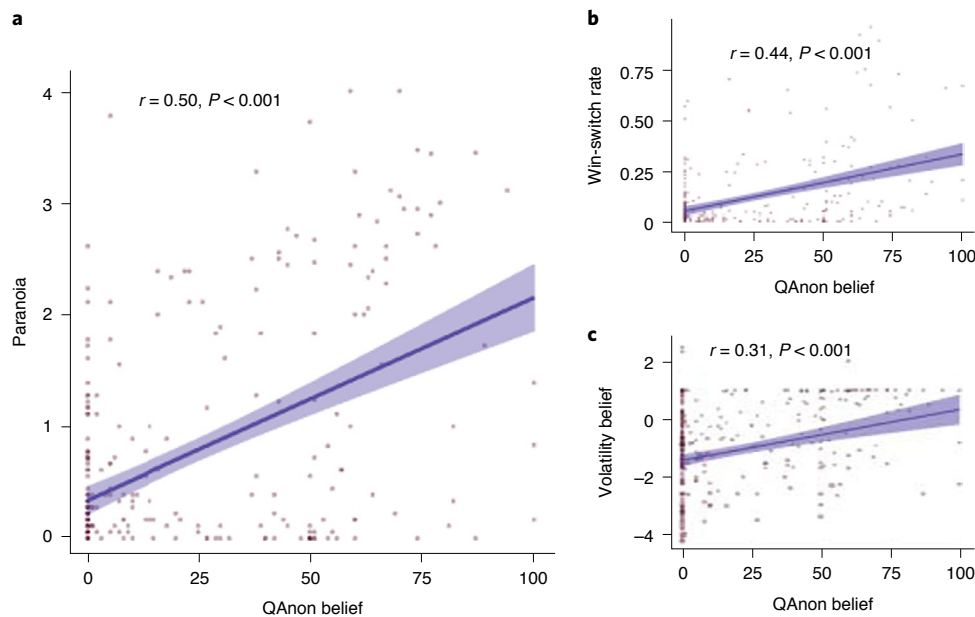


Fig. 8 | Relating QAnon beliefs to paranoia and task behaviour. **a**, Individuals ($n=307$) with higher paranoia endorsed more QAnon beliefs ($r=0.5$, $P<0.001$). **b,c**, Similarly, QAnon beliefs were strongly correlated with erratic task behaviour ($r=0.44$, $P<0.001$) (**b**) and perturbed volatility priors ($r=0.31$, $P<0.001$) (**c**). Analysis was performed on individuals who responded to the QAnon questions. The shaded regions represent the 95% confidence interval.

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 a stronger belief in QAnon. QAnon belief and paranoia more broadly were highly correlated (Fig. 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 (Fig. 8b; $r=0.44$, $P<0.001$) and volatility belief (Fig. 8c; $r=0.31$, $P<0.001$) just like paranoia. Supplementary Fig. 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, attested to more paranoia, evinced more win-switching and 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 has been 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 (CTL¹⁹) as a possible contributor to the increased paranoia that we observed. State tightness may originate in response to threats such as 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 priori-

tize 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 that have been implicated in paranoia in the laboratory^{4,27-29}.

Mask wearing is a collective action problem, wherein most people are conditional co-operators, 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. In this study, we present behavioural evidence for a connection between real-world volatility, conspiracy theorizing, paranoia and hesitant attitudes towards 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 antisemitic belief that Jewish people were poisoning wells and causing the pandemic was sadly extremely common¹⁷. The acquired immune deficiency syndrome (AIDS) epidemic was associated with a number of conspiracies related to public health measures but less directly. For

example, people believed that human immunodeficiency virus was created through the polio vaccination programme 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 seat belt laws³³. Like masks in a viral pandemic, seat belts 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 should be worn was associated with public outcry³³. People were concerned about the imposition on their freedom³³. They complained that seat belts 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 Fig. 6). We suggest this finding is consistent with previous reports that people with persecutory delusions tend to evince an intentional bias towards 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. Although 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 showed that our results are not driven by other real-world events or issues with our sampling—there are several important caveats to our conclusions. We did not preregister 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-subject or panel data¹⁴. Our DiD analysis leveraged some tentative causal claims despite being based on between-subject data¹⁴. Mask-recommended states were culturally tighter although of course cultural tightness did not change during the course of our study. Tightness interacted with mandate and adherence to mask-wearing policy (Fig. 6). The baseline difference in tightness would have worked against the effects we observed, not in their favour. Indeed, our multiple regression analysis found no evidence for an effect of tightness on paranoia in states without a mask mandate (Fig. 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, although 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 (for example, bull versus bear markets), then they may yield different results, not because of inadequacy of the theory or experiment but because participants' behaviour was being modulated by heretofore underappreciated 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 the 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 (no. 2000026290). Written 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 Mechanical Turk while providing additional security for easy recruitment²³. Sample sizes were determined based on our previous 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 (Behavioural tasks section). Participants were compensated with USD\$6 for completion and a bonus of USD\$2 if they scored in the top 10% of all respondents.

Pandemic study. A total of 605 participants were collected and divided into 202 prelockdown participants, 231 lockdown participants and 172 reopening participants. Of the 202, we included the 72 (16 with high paranoia) participants who completed the non-social task (described in a previous publication⁷). The paranoia of those participants was self-rated with the Structured Clinical Interview for DSM-IV Axis I Personality Disorders (SCID-I) paranoid trait questions, which are strongly overlapping and correlated with the Green et al. scale⁶. We recruited 130 (20 with high paranoia) participants who completed the social task. Similarly, of the 231, we recruited 119 (20 with high paranoia) and 112 (30 with high paranoia) participants who completed the non-social and social tasks, respectively. Lastly, of the 172, we recruited 93 (35 with high paranoia) and 79 (35 with 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 study⁷. We excluded a total of 163 submissions—18 from prelockdown (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 individuals with low paranoia and 91 were individuals with high paranoia. Similar exclusion and inclusion criteria were applied for recruitment; most notably, we leveraged CloudResearch's newly added Data Quality feature that only allows vetted high-quality participants—individuals who passed their screening measures—into our study. This systematically cleaned all poor participants from our sample pool.

Behavioural tasks. Participants completed a three-option probabilistic reversal learning task with a non-social (card deck) or social (partner) domain frame. For the non-social domain frame, 3 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 change⁴¹. For the social domain frame, 3 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 as possible towards their group project. Like the non-social domain frame, 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, unbeknown to the participants, the underlying contingencies transitioned 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 had changed.

Questionnaires. After task completion, questionnaires were administered via Qualtrics. We queried demographic information (age, sex, educational attainment, ethnicity) and mental health questions (past or present diagnosis, medication use), SCID-II⁴², Beck's Anxiety Inventory⁴³, Beck's Depression Inventory⁴⁴, Dimensional Obsessive-Compulsive Scale⁴⁵, and critically, the revised Green et al. Paranoid Thoughts Scale⁴⁶, which separates clinically from non-clinically paranoid individuals based on the receiver operating characteristic curve-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 (Supplementary Fig. 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 calculate mask attitude (values < 0 indicate attitude against mask-wearing and values > 0 indicate attitude in favour of mask-wearing) to identify group differences in paranoia. To compute an individual's coronavirus vaccine conspiracy belief, we aggregated 5 vaccine-related questions from the 48-item coronavirus conspiracy questionnaire: (1) the coronavirus vaccine will contain microchips to control 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 World Health Organization already has a vaccine and are withholding it. We adopted a 7-point scale: (1) strongly disagree; (2) disagree; (3) somewhat disagree; (4) neutral; (5) somewhat agree; (6) agree; and (7) strongly agree. 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 rates⁴⁹, confirmed COVID-19 cases⁵⁰ and mask-wearing²⁰ in the USA. For 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 to measure the difference in unemployment between the prepandemic and pandemic periods, respectively. For confirmed cases, the *New York Times* has published cumulative counts of coronavirus cases since January 2020.

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 the 2 and 14 July²⁰. 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: California, New Mexico, Michigan, Illinois, New York, Massachusetts, Rhode Island, Maryland, Virginia, Delaware and Maine 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 (<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 2020.

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 where the participant switched after positive feedback divided by the number of trials where they received positive feedback. Lose-stay rates were calculated as the number of trials where a participant persisted after negative feedback divided by the 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-19 (ref. 51). 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 (that is, introduced date-baseline date); *E*: number of days before the order was lifted since it was introduced (that is, expiration date-introduced date) where the baseline date is defined as the date at which the first SAH order was implemented (Supplementary Fig. 3). California was the first to enforce the order on 19 March 2020 (that is, 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 prelockdown 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 prelockdown data instead had no impact on our findings (Supplementary Table 6).

Causal inference. To measure the attribution of mask policy on paranoia, we adopted a 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 by the following equation:

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

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 preintervention difference in paranoia between both 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, CTL index and mask-wearing belief—and one categorical state variable, that is, mask policy. We fitted 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. For the full model:

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

For the reduced model:

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

Computational modelling. The Hierarchical Gaussian Filter (HGF) toolbox v5.3.1 is freely available for download in the Translational Algorithms for

Psychiatry-Advancing Science package at <https://translationalneuromodeling.github.io/tapas>^{10,11}. We installed and ran the package in MATLAB and Statistics Toolbox Release 2016a (MathWorks).

We estimated perceptual parameters individually for the first and second halves of the task (that is, blocks 1 and 2). Each participant's choices (that is, 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 three-level HGF multi-armed bandit configuration for our perceptual model and paired it with the softmax- μ_0 3 decision model.

Statistics. Statistical analyses and effect size calculations were performed with an alpha of 0.05 and two-tailed P values in RStudio v.1.3.959.

Bayes factors (BF_{10}) were reported for non-significant t -tests and analyses of variance (ANOVAs) to provide additional evidence of no effect (or no differences)⁵² We defined the null hypothesis (H_0) as there being no difference in the means of behaviour/demographics between groups ($H_0: \mu_1 - \mu_2 = 0$) and the alternative hypothesis (H_1) as a difference ($H_0: \mu_1 - \mu_2 \neq 0$). Interpretations of the BF_{10} were adopted from Lee and Wagenmakers⁵³.

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-squared exact tests (two groups) or Monte Carlo tests (more than two groups). Correlations were computed with Pearson's r .

HGF parameter estimates and behavioural patterns (win-switch and lose-stay rates) were analysed by repeated measures and split-plot ANOVAs (that is, block designated as within-participant factor; pandemic, paranoia group and social versus non-social condition as between-participants factors). Model parameters were corrected for multiple comparisons using the Benjamini-Hochberg⁵⁴ method with an FDR of 0.05 in ANOVAs across experiments.

Visualization. All data visualization were produced in RStudio. Some were adopted from the raincloud plot theme⁵⁵.

Reporting Summary. Further information on research design is available in the Nature Research Reporting Summary linked to this article.

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>.

Received: 12 January 2021; Accepted: 6 July 2021;

Published online: 27 July 2021

References

- van Rooijen, J. W. & Douglas, K. M. Conspiracy theories as part of history: the role of societal crisis situations. *Mem. Stud.* **10**, 323–333 (2017).
- Smallman, S. Whom do you trust? Doubt and conspiracy theories in the 2009 influenza pandemic. *J. Int. Glob. Stud.* **6**, 2 (2015).
- Raihani, N. J. & Bell, V. An evolutionary perspective on paranoia. *Nat. Hum. Behav.* **3**, 114–121 (2019).
- Feehey, E. J., Groman, S. M., Taylor, J. R. & Corlett, P. R. Explaining delusions: reducing uncertainty through basic and computational neuroscience. *Schizophr. Bull.* **43**, 263–272 (2017).
- Reed, E. J. et al. Paranoia as a deficit in non-social belief updating. *eLife* **9**, e56345 (2020).
- Freeman, D. et al. The revised Green et al., Paranoid Thoughts Scale (R-GPTS): psychometric properties, severity ranges, and clinical cut-offs. *Psychol. Med.* **51**, 244–253 (2021).
- Soltani, A. & Izquierdo, A. Adaptive learning under expected and unexpected uncertainty. *Nat. Rev. Neurosci.* **20**, 635–644 (2019).
- Ryder, A. G., Costa, P. T. & Bagby, R. M. 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).
- Corlett, P. R. & Fletcher, P. C. Computational psychiatry: a Rosetta Stone linking the brain to mental illness. *Lancet Psychiatry* **1**, 399–402 (2014).
- Mathys, C., Daunizeau, J., Friston, K. J. & Stephan, K. E. A Bayesian foundation for individual learning under uncertainty. *Front. Hum. Neurosci.* **5**, 39 (2011).
- Mathys, C. D. et al. Uncertainty in perception and the hierarchical Gaussian filter. *Front. Hum. Neurosci.* **8**, 825 (2014).
- 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. USA* **115**, E11446–E11454 (2018).
- Marinescu, I. E., Lawlor, P. N. & Kording, K. P. Quasi-experimental causality in neuroscience and behavioural research. *Nat. Hum. Behav.* **2**, 891–898 (2018).
- Angrist, J. D. & Pischke, J.-S. *Mostly Harmless Econometrics* (Princeton Univ. Press, 2008).
- Jaeger, D. A., Joyce, T. J. & Kaestner, R. A. A cautionary tale of evaluating identifying assumptions: did reality TV really cause a decline in teenage childbearing? *J. Bus. Econ. Stat.* **38**, 317–326 (2020).
- Goodman-Bacon, A. & Marcus, J. Using difference-in-differences to identify causal effects of COVID-19 policies. *Surv. Res. Methods* **14**, 153–158 (2020).
- Cohn, N. *The Pursuit of the Millennium* (Oxford Univ. Press, 1961).
- Cameron, A. C. & Miller, D. L. A practitioner's guide to cluster-robust inference. *J. Hum. Resour.* **50**, 317–372 (2015).
- Harrington, J. R. & Gelfand, M. J. Tightness-looseness across the 50 United States. *Proc. Natl Acad. Sci. USA* **111**, 7990–7995 (2014).
- Katz, J., Sanger-Katz, M. & Quealy, K. *Estimates from The New York Times, based on roughly 250,000 interviews conducted by Dynata from July 2 to July 14 (The New York Times and Dynata, 2020); https://github.com/nytimes/covid-19-data/tree/master/mask-use*
- Blakemore, S. J., 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).
- Moss, A. J., Rosenzweig, C., Robinson, J. & Litman, L. Demographic stability on Mechanical Turk despite COVID-19. *Trends Cogn. Sci.* **24**, 678–680 (2020).
- Litman, L., Robinson, J. & Abberbock, T. TurkPrime.com: a versatile crowdsourcing data acquisition platform for the behavioral sciences. *Behav. Res. Methods* **49**, 433–442 (2017).
- 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. *Eur. J. Soc. Psychol.* **48**, 909–926 (2018).
- Freeman, D. et al. Coronavirus conspiracy beliefs, mistrust, and compliance with government guidelines in England. *Psychol. Med.* <https://doi.org/10.1017/S0033291720001890> (2020).
- Colombo, M. Two neurocomputational building blocks of social norm compliance. *Biol. Philos.* **29**, 71–88 (2014).
- Corlett, P. R. et al. Disrupted prediction-error signal in psychosis: evidence for an associative account of delusions. *Brain* **130**, 2387–2400 (2007).
- Corlett, P. R., Taylor, J. R., Wang, X.-J., Fletcher, P. C. & Krystal, J. H. Toward a neurobiology of delusions. *Prog. Neurobiol.* **92**, 345–369 (2010).
- Romaniuk, L. et al. Midbrain activation during Pavlovian conditioning and delusional symptoms in schizophrenia. *Arch. Gen. Psychiatry* **67**, 1246–1254 (2010).
- Ostrom, E. Collective action and the evolution of social norms. *J. Econ. Perspect.* **14**, 137–158 (2000).
- Worobey, M. et al. Origin of AIDS: contaminated polio vaccine theory refuted. *Nature* **428**, 820 (2004).
- Gonsalves, G. & Staley, P. Panic, paranoia, and public health—the AIDS epidemic's lessons for Ebola. *N. Engl. J. Med.* **371**, 2348–2349 (2014).
- Giubilini, A. & Savulescu, J. Vaccination, risks, and freedom: the seat belt analogy. *Public Health Ethics* **12**, 237–249 (2019).
- Robertson, L. Road death trend in the United States: implied effects of prevention. *J. Public Health Pol.* **39**, 193–202 (2018).
- Heyes, C. & Pearce, J. M. Not-so-social learning strategies. *Proc. R. Soc. B* **282**, 20141709 (2015).
- Freeman, D. et al. Concomitants of paranoia in the general population. *Psychol. Med.* **41**, 923–936 (2011).
- 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).
- Henco, L. et al. Bayesian modelling captures inter-individual differences in social belief computations in the putamen and insula. *Cortex* **131**, 221–236 (2020).
- 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).
- Heyes, C. Précis of cognitive gadgets: the cultural evolution of thinking. *Behav. Brain Sci.* **42**, E169 (2019).
- DiGrazia, J. The social determinants of conspiratorial ideation. *Socius* **3**, 237802311668979 (2017).
- American Community Survey* (United States Census, 2017); <https://www.census.gov/acs/www/data/data-tables-and-tools/data-profiles/2017/>
- Gelfand, M. J. et al. Differences between tight and loose cultures: a 33-nation study. *Science* **332**, 1100–1104 (2011).

44. Beck, A. T., Epstein, N., Brown, G. & Steer, R. A. An inventory for measuring clinical anxiety: psychometric properties. *J. Consult. Clin. Psychol.* **56**, 893–897 (1988).
45. Beck, A. T., Ward, C. H., Mendelson, M., Mock, J. & Erbaugh, J. An inventory for measuring depression. *Arch. Gen. Psychiatry* **4**, 561–571 (1961).
46. Abramowitz, J. S. et al. Assessment of obsessive-compulsive symptom dimensions: development and evaluation of the Dimensional Obsessive-Compulsive Scale. *Psychol. Assess.* **22**, 180–198 (2010).
47. Knotek, E. 2nd et al. Consumers and COVID-19: survey results on mask-wearing behaviors and beliefs. *Economic Commentary* <https://doi.org/10.26509/frbc-ec-202020> (2020).
48. Enders, A. et al. Who supports QAnon? A case study in political extremism https://www.joescinski.com/uploads/7/1/9/5/71957435/qanon_2-4-21.pdf (2021)
49. Ettlinger, M. & Hensley, J. COVID-19 economic crisis: by state. *Carsey School of Public Policy* <https://carsey.unh.edu/COVID-19-Economic-Impact-By-State> (2021).
50. *An ongoing repository of data on coronavirus cases and deaths in the U.S.* (*The New York Times*, 2020); <https://github.com/nytimes/covid-19-data>
51. Status of lockdown and stay-at-home orders in response to the coronavirus (COVID-19) pandemic. *Ballotpedia* [https://ballotpedia.org/Status_of_lockdown_and_stay-at-home_orders_in_response_to_the_coronavirus_\(COVID-19\)_pandemic,_2020](https://ballotpedia.org/Status_of_lockdown_and_stay-at-home_orders_in_response_to_the_coronavirus_(COVID-19)_pandemic,_2020) (2020).
52. Gelman, A. & Stern, H. The difference between ‘significant’ and ‘not significant’ is not itself statistically significant. *Am. Stat.* **60**, 328–331 (2006).
53. Lee, M. D & Wagenmakers, E.-J. *Bayesian Cognitive Modeling: A Practical Course* (Cambridge Univ. Press, 2013).
54. Hochberg, Y. & Benjamini, Y. More powerful procedures for multiple significance testing. *Stat. Med.* **9**, 811–818 (1990).
55. Allen, M. et al. Raincloud plots: a multi-platform tool for robust data visualization. *Wellcome Open Res.* **4**, 63 (2021).

Acknowledgements

This work was supported by the Yale University Department of Psychiatry, the Connecticut Mental Health Center and Connecticut State Department of Mental Health and Addiction Services. It was funded by an International Mental Health Research

Organization/Janssen Rising Star Translational Research Award, an Interacting Minds Center (Aarhus) Pilot Project Award, National Institute of Mental Health (NIMH) grant no. R01MH12887 (P.R.C.), NIMH grant no. R21MH120799-01 (P.R.C. and S.M.G.) and an Aarhus Universitets Forskningsfond Starting Grant (C.D.M.). E.J.R. was supported by the National Institutes of Health (NIH) Medical Scientist Training Program training grant no. GM007205, National Institute of Neurological Disorders and Stroke Neurobiology of Cortical Systems grant no. T32 NS007224 and a Gustavus and Louise Pfeiffer Research Foundation Fellowship. S.U. received funding from an NIH T32 fellowship (no. MH065214). S.M.G. and J.R.T. were supported by a National Institute on Drug Abuse grant no. 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.

Author contributions

P.S., P.R.C., E.J.R., J.R.T. and S.M.G. conceived the study. E.J.R., J.G.K., C.D.M., P.L. and S.U. contributed the task materials and analysis code. A.J.M., J.R. and L.L. contributed the data. P.S. acquired the data. P.S. and P.R.C. analysed the data. All authors wrote and edited the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41562-021-01176-8>.

Correspondence and requests for materials should be addressed to P.R.C.

Peer review information *Nature Human Behaviour* thanks Ryan Balzan, Michael Moutoussis and the other, anonymous, reviewer(s) for their contribution to the peer review of this work. Peer reviewer reports are available.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

© The Author(s), under exclusive licence to Springer Nature Limited 2021

Reporting Summary

Nature Research wishes to improve the reproducibility of the work that we publish. This form provides structure for consistency and transparency in reporting. For further information on Nature Research policies, see our [Editorial Policies](#) and the [Editorial Policy Checklist](#).

Statistics

For all statistical analyses, confirm that the following items are present in the figure legend, table legend, main text, or Methods section.

n/a Confirmed

- The exact sample size (n) for each experimental group/condition, given as a discrete number and unit of measurement
- A statement on whether measurements were taken from distinct samples or whether the same sample was measured repeatedly
- The statistical test(s) used AND whether they are one- or two-sided
Only common tests should be described solely by name; describe more complex techniques in the Methods section.
- A description of all covariates tested
- A description of any assumptions or corrections, such as tests of normality and adjustment for multiple comparisons
- A full description of the statistical parameters including central tendency (e.g. means) or other basic estimates (e.g. regression coefficient) AND variation (e.g. standard deviation) or associated estimates of uncertainty (e.g. confidence intervals)
- For null hypothesis testing, the test statistic (e.g. F , t , r) with confidence intervals, effect sizes, degrees of freedom and P value noted
Give P values as exact values whenever suitable.
- For Bayesian analysis, information on the choice of priors and Markov chain Monte Carlo settings
- For hierarchical and complex designs, identification of the appropriate level for tests and full reporting of outcomes
- Estimates of effect sizes (e.g. Cohen's d , Pearson's r), indicating how they were calculated

Our web collection on [statistics for biologists](#) contains articles on many of the points above.

Software and code

Policy information about [availability of computer code](#)

Data collection Questionnaires were administered via the Qualtrics® survey platform (Qualtrics Labs, Inc., Provo, UT). Behavioral data were collected through custom task code. Both were linked to a Amazon Mechanical Turk Human Intelligence Task (HIT; <https://requester.mturk.com>).

Data analysis We performed computational modeling using the HGF toolbox v5.3.1 (<https://translationalneuromodeling.github.io/tapas/>) in MATLAB and Statistics Toolbox Release 2016a (MathWorks®, Natick, MA). Statistical analyses were completed in IBM SPSS Statistics, Version 25 (IBM Corp., Armonk, NY). Analyses and figures were generated with RStudio.

For manuscripts utilizing custom algorithms or software that are central to the research but not yet described in published literature, software must be made available to editors and reviewers. We strongly encourage code deposition in a community repository (e.g. GitHub). See the Nature Research [guidelines for submitting code & software](#) for further information.

Data

Policy information about [availability of data](#)

All manuscripts must include a [data availability statement](#). This statement should provide the following information, where applicable:

- Accession codes, unique identifiers, or web links for publicly available datasets
- A list of figures that have associated raw data
- A description of any restrictions on data availability

Data and code are available:

<https://github.com/psuthaharan/covid19paranoia>. Figures 2-8 have associated raw data. All links to publicly available data that were used in the analyses are in the manuscript.

Field-specific reporting

Please select the one below that is the best fit for your research. If you are not sure, read the appropriate sections before making your selection.

Life sciences Behavioural & social sciences Ecological, evolutionary & environmental sciences

For a reference copy of the document with all sections, see [nature.com/documents/nr-reporting-summary-flat.pdf](https://www.nature.com/documents/nr-reporting-summary-flat.pdf)

Behavioural & social sciences study design

All studies must disclose on these points even when the disclosure is negative.

Study description	We collected quantitative behavioral data (i.e., participant choices on a structured task) and questionnaire responses regarding mental health and demographics. We collected this data online replication through Amazon Mechanical Turk (MTurk) before and during the evolving coronavirus pandemic in 2020. We also leveraged publicly available data to motivate and unpack our findings.
Research sample	Demographic summary for the N=1,010 individuals in our sample: <ul style="list-style-type: none"> a. Age <ul style="list-style-type: none"> i. min age = 19 ii. max age = 73 iii. median age = 35 b. Gender <ul style="list-style-type: none"> i. 40.40% identified as female ii. 58.51% identified as male c. Hispanic <ul style="list-style-type: none"> i. 10% identified as Hispanic ii. 89.50% identified as non-Hispanic d. Race (White: 78.22% vs. Non-White: 21.68%) <ul style="list-style-type: none"> i. 78.22% identified as White ii. 19.6% identified as Black iii. 4.85% identified as Asian iv. 3.27% identified as Multi-racial v. 0.99% identified as American-Indian vi. 1.98% identified as Other e. Education (HS: 12.28% vs. College: 87.62%) <ul style="list-style-type: none"> i. 12.28% identified as highschool or equivalent ii. 18.71% identified as some college iii. 11.68% identified as associate iv. 42.57% identified as bachelors v. 11.78% identified as masters vi. 1.58% identified as doctorate vii. 1.28% identified as postdoc
Sampling strategy	This was a stratified random sample. Based on effects present in our previous work (e.g., Powers, Mathys, and Corlett, Science, 2017, Reed et al, 2020, eLife), we recruited participants until we reached >20 participants with high paranoia per task version. Task effects were replicated internally across pandemic periods, as were the associations with paranoia, mask beliefs, and vaccine conspiracy beliefs (total N=1010). In order to assuage concerns about sampling we leveraged data from CloudResearch from 7,293 experiments comprising 2.5 million participants. We find that our data were representative of the global sample (in terms of age, demographic features, and geographical location), and that the global sample (from which we drew our participants) did not change significantly throughout the evolving pandemic.
Data collection	Task and questionnaire were administered remotely via web browser on desktop computers or laptops through the Amazon Mechanical Turk (MTurk) workplace. The task was not enabled for use on cell phones, tablets, or other devices. Researchers were not present. The presence of other individuals is unknown.
Timing	Pre-pandemic data were acquired between August 2018 and end of January 2020 Lock-down Data were gathered between March 16th and June 31st 2020 Reopening data were gathered in July 2020 Replication data were gathered in August, September, October and November 2020
Data exclusions	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
Non-participation	No participants dropped out
Randomization	Participants were not allocated into experimental groups

Reporting for specific materials, systems and methods

We require information from authors about some types of materials, experimental systems and methods used in many studies. Here, indicate whether each material, system or method listed is relevant to your study. If you are not sure if a list item applies to your research, read the appropriate section before selecting a response.

Materials & experimental systems

n/a	Involvement	Material/System
<input checked="" type="checkbox"/>	<input type="checkbox"/>	Antibodies
<input checked="" type="checkbox"/>	<input type="checkbox"/>	Eukaryotic cell lines
<input checked="" type="checkbox"/>	<input type="checkbox"/>	Palaeontology and archaeology
<input checked="" type="checkbox"/>	<input type="checkbox"/>	Animals and other organisms
<input type="checkbox"/>	<input checked="" type="checkbox"/>	Human research participants
<input checked="" type="checkbox"/>	<input type="checkbox"/>	Clinical data
<input checked="" type="checkbox"/>	<input type="checkbox"/>	Dual use research of concern

Methods

n/a	Involvement	Method
<input checked="" type="checkbox"/>	<input type="checkbox"/>	ChIP-seq
<input checked="" type="checkbox"/>	<input type="checkbox"/>	Flow cytometry
<input checked="" type="checkbox"/>	<input type="checkbox"/>	MRI-based neuroimaging

Human research participants

Policy information about [studies involving human research participants](#)

Population characteristics

See above

Recruitment

Participants were recruited through a "Human Intelligence Task" (HIT) advertisement in the Amazon Mechanical Turk online marketplace.

Ethics oversight

Yale University Institutional Review Board (IRB), Human Investigations Committee (HIC)

Note that full information on the approval of the study protocol must also be provided in the manuscript.