A normative account of self-deception, overconfidence, and paranoia

Rossi-Goldthorpe, R.A.^{1,2}, Leong, Y.C.³, Leptourgos, P.¹, & Corlett, P.R.^{1*}

1. Department of Psychiatry, Yale University, New Haven, CT, USA

2. Interdepartmental Neuroscience Program, Yale University, New Haven, CT, USA

3. Helen Wills Neuroscience Institute, University of California, Berkeley, CA, USA

* Correspondence: Philip R. Corlett philip.corlett@yale.edu

Abstract: Self-deception, paranoia, and overconfidence involve misbeliefs about self, others, and world. They are often considered mistaken. Here we explore whether they might be adaptive, and further, whether they might be explicable in normative Bayesian terms. We administered a difficult perceptual judgment task with and without social influence (suggestions from a cooperating or competing partner). Crucially, the social influence was uninformative. We found that participants heeded the suggestions most under the most uncertain conditions and that they did so with high confidence, particularly if they were more paranoid. Model fitting to participant behavior revealed that their prior beliefs changed depending on whether the partner was a collaborator or competitor, however, those beliefs did not differ as a function of paranoia. Instead, paranoia, self-deception, and overconfidence were associated with participants' perceived instability of their own performance. These data are consistent with the idea that self-deception, paranoia, and overconfidence flourish under uncertainty, and have their roots in low self-esteem, rather than excessive social concern. The normative model suggests that spurious beliefs can have value – self-deception is irrational yet can facilitate optimal behavior. This occurs even at the expense of monetary rewards, perhaps explaining why self-deception and paranoia contribute to costly decisions which can spark financial crashes and costly wars.

Introduction

People lie to others, but they also lie to themselves. We might better deceive others by more convincingly deceiving ourselves¹. More fundamentally, self-deception may serve to protect self-esteem². We deceive ourselves into believing that we are kinder, fairer, and more proficient than average¹. Such overconfidence can be adaptive both intra- and interpersonally – a modicum of overconfidence can increase performance² and persuasiveness³. However, there can be substantial personal and social costs to self-deception as too much self-deception can cause an individual to make poor decisions^{4, 5}. These costs are particularly high when the individual is in a position of power or leadership – self-deception and the accompanying overconfidence have plagued many leaders and colored their decisions, contributing to decisions that fuel destructive wars and economic collapse.

Similarly, paranoia – the belief that others have malicious intentions towards us – shares many of the hallmarks of self-deception⁶. Paranoia may protect self-esteem by explaining away failings in terms of others' malintent^{7, 8}. By polarizing the social world into malevolent and benevolent agents, paranoia may facilitate in-group versus out-group distinctions and solidify group identity⁶. Both effects of paranoia might be accomplished through self-deception, via its direct inflationary effects on self-image and indirectly, via the impact of overconfidence. Confident people are believed more, and their advice is more likely to be followed^{9, 10}. By convincing others, overconfident people further reinforce their own misbeliefs^{11, 12}.

Both paranoia and self-deception appear to flourish under uncertainty. Ambiguity creates room for the selective reframing of evidence in order to self-serve¹³. Indeed, self-deception about one's own abilities may decrease after prolonged exposure to an objective performance assessment¹⁴. Furthermore, perceptual and response biases are maximized in the context of ambiguous stimuli¹⁵. Likewise, times of great uncertainty, from terrorist attacks¹⁶ to viral pandemics, are fertile grounds for paranoia and conspiracy thinking (blaming misfortune on powerful outgroups)¹⁷. In laboratory tasks, paranoid individuals expect more volatility but also fail to adapt to and learn appropriately from volatility¹⁸. Hence, the conditions that foment self-deception and paranoia are similar. It is as yet unclear whether they share underlying psychological mechanisms, and whether they are similarly sensitive to uncertainty or social affiliative processes. If the processes underlying both paranoia and self-deception are similar, solutions to either could be utilized to reduce the other as well. Additionally, a shared mechanism might suggest that paranoia could amplify self-deceptive behaviors, thus bolstering misbeliefs and causing more distress.

Here we explore the specific relationships between paranoia and self-deception, in the context of uncertain perceptual decision-making under social influence. Using computational modeling that explicitly quantifies the contributions of social and non-social information to decisions, we sought to delineate empirically whether and how self-deception and over-confidence are related to paranoia, with implications for their psychological

Normative self-deception

functions and underlying mechanisms. There is debate as to whether phenomena like self-deception and overconfidence violate Bayesian principles of belief updating^{19, 20}, and thence challenge accounts of paranoia and delusions that appeal to those principles^{19, 20}. By fitting normative computational models – based on Bayesian mechanisms^{21, 22} – to our data, we sought to determine whether and how a Bayesian account could explain self-deception, overconfidence, and paranoia. In particular, we investigated how self-deceptive choices and the confidence of those self-deceptive choices differed on the basis of paranoia. This behavioral contrast can then be teased apart using Bayesian belief updating, and a more precise measure of what aspect of belief updating is altered in paranoia leading to overconfident self-deception can be understood.

Behavioral data

We employed an image categorization task first described by Leong et al. (2019)¹⁵. Participants classified merged images of faces and scenes, as either containing more face or scene, and they subsequently expressed their confidence in their choice. These “chimeric” images ranged from 100% face and 0% scene to 100% scene and 0% face over 80 trials (*C1 Phase*). After the conclusion of the C1 phase, participants were informed they were either working with a partner who was either a collaborator ($n=329$), or a competitor ($n=334$), who would be placing bets on whether the next image would be mostly face or mostly house (Figure 1A). The payoff matrix (Figure 1B) illustrates the economic outcomes for each decision – the optimal strategy in both conditions is to classify the images objectively as the monetary amount is contingent upon the accuracy of the classification. In the cooperation condition, the participant would receive a monetary bonus if their partner’s bet was correct, in addition to the earnings from correctly classifying the image. In the competition condition, the participants would lose money if their partner’s bet was correct. Crucially, regardless of the accuracy of the bet, the reward maximizing strategy is to classify the images correctly. As a result, there is no benefit to the participant for biased responses based on the bet. The participants saw their partner’s bet before seeing the image, before providing their classification and confidence again (*C2 Phase*). They classified the same images they saw in *C1*. In this initial study, 50% of the bets were accurate, thus there was no informational value in the bet, and thus no information in favor of trusting the other’s bets.

We recruited participants ($N=663$) from the Amazon Mechanical Turk Marketplace (MTurk). To determine they were engaged with that task rather than choosing randomly, we fit generalized linear models (GLMs) to their choices. We replicated two behavioral effects observed in Leong et al., (2019). First, as the percentage of scene in the chimera increased, the probability of responding scene followed an s-shaped psychometric curve, indicating that in general, participants were able to categorized the chimera’s accurately. Second, there was a motivational bias on perception: the bets influenced the participants’ choices differently based on experimental condition (cooperation vs. competition). Generalized linear mixed-effects models (GLMEs) revealed a significant *bet x group*

Normative self-deception

interaction (GLME: $z=-8.663$, $p<2e-16$). Participants in the cooperation condition were more likely to agree with the bet while participants in the competition condition were more likely to disagree with the bet (Figure 2B), indicating that participants were motivated to respond with whatever classification aligned with the role of the other individual. This bias suggests an underlying self-deceptive response where individuals lie to themselves in order to respond with the motivationally-driven response that might conflict with their original objective classification.

Self-Deception

We defined self-deception after Mijović-Prelec and Prelec (2010)²³. If a classification changed between sessions (C1 and C2) to either agree with the bet (cooperation condition) or to disagree with the bet (competition condition) the response was self-deceptive. The raw self-deception score was simply the number of such classification changes. We were also interested in confidence in these responses – to explore whether participants were merely guessing when they changed their minds to conform to or defect from the bets. As self-deception involves lying to oneself to protect one's self-image, self-deception overlaps significantly with overconfidence. We created a metric where a high measure of self-deception would reflect not only a high percentage of self-deceptive trials, but also where those responses were given with high confidence, as this would represent the most self-deceptive response. A self-deceptive response with lower confidence reflects some internal recognition that the individual is not fully believing their own lies. We normalized each participant's confidence on self-deceptive trials by their confidence during initial classification (C1). To compute self-deception we multiplied their raw score (number of deception trials) by their normalized confidence on those trials. Hence, high self-deceivers engaged in more self-deceptive responses with elevated confidence relative to their own confidence baseline.

Paranoia and Self-Deception

We defined high paranoia as scoring above a clinically meaningful cutoff on the revised Green et al. Paranoid Thoughts Scale (R-GPTS), a self-rating scale²⁴. Analysis of variance in self-deceptive choice behavior revealed a main effect of paranoia (high or low), a main effect of group (competition or collaboration) but no paranoia by group interaction.

High paranoia participants also made more self-deceptive choices ($F(1, 659) = 13.65$, $p_{\text{bonf}}=0.0007155$), and were also more confident on those trials ($F(1,620) = 81.691$, $p_{\text{bonf}}<2e-16$). The difference between groups remained significant when we examined confidence-normalized self-deception score ($F(1, 620) = 58.0612$, $p_{\text{bonf}}= 2.8659e-13$, Figure 3A, B). We also found that the cooperation group had increased confidence-weighted self-deception ($F(1, 620) = 15.0085$, $p_{\text{bonf}}=3.442e-4$)— people were more likely to confidently self-deceive to conform to their partners' bet in the cooperation group relative to defecting from the bet in the competition

Normative self-deception

group. There was no group by paranoia interaction, suggesting that centering in vs out-group membership was not differentially impacted by paranoia (Figure 3C-D).

Which trials engender self-deception?

We explored how the objective proportions of face and scene in a chimera impacted self-deception. We analyzed the fraction of self-deceptive responses for each image category. Across paranoia groups, most self-deceptive responses occurred for the most ambiguous images (50/50 scene-face, Figure 4). A GLMM model showed a significant interaction between objective image ambiguity and paranoia group (GLMM: $z = -5.908$, $p = 3.47e-09$) – the high paranoia group evinced self-deception to the slightly less ambiguous stimuli .

Computational modeling

This task is structured so that participants should ideally evaluate each image independently and ignore the bet. However, as mentioned above, there is indeed a bias where the individual either aligns or disagrees with the bet. This existence of a motivational bias implies that the individual must be attributing some sort of value to the bet, and as a result, they might update this value as they gain more information through these trials. Analogously, participants might erroneously infer that image classification might depend on previous images, and update their beliefs about the images accordingly. There is no aspect of the task that suggests participants should learn about underlying contingencies, and as a result, a belief-updating model would reflect that participants are erroneously learning about the bet and images. In particular, these inaccurate associations might contribute to self-deceptive behavior. Self-deception could be facilitated by the extrapolation of non-existent relationships – it's easier to lie to yourself when you create a false narrative. Here, that might be that the bet and images should be viewed as changing, connected events – as a result, changing one's answer is in response to these underlying connections that must be considered, rather than self-deception.

Since we choose to examine how incorrect learning might be occurring with self-deception, we decided to employ a meta-Bayesian approach where experimenters make inferences about the inferences of participants – observing the observer^{21, 22}. To do this, we utilized a 2-layer Hierarchical Gaussian Filter (HGF)^{21, 22} model to the behavioral data to explore belief updating, paranoia, and self-deception. The HGF uses a combination of a perceptual and response model to which connects sensory cues to beliefs – inversion of the models allows analysis of how participant's binary responses instantiate their beliefs about the task. The perceptual model contains hierarchical layers representing beliefs, where lower belief layers update higher belief layers through precision-weighted prediction errors

Normative self-deception

Behavioral results indicate that bet and image impact behavior. Therefore, we choose to employ an HGF with two streams of information processing²⁵: social (the bet) and non-social (the image)²⁶. The two streams were integrated and then passed to a softmax decision model. We convolved choice behavior with the model, and inverted it to estimate individual model parameters.

The first perceptual level for the social stream captures participants' perceived accuracy of the bet ($x_{1,s}$) while the first level for the non-social stream represented the participants' belief that the image was mostly scene or mostly face ($x_{1,ns}$). The second level in the social stream ($x_{2,s}$) represented the tendency of bet accuracy over trials, while the second non-social stream ($x_{2,ns}$) represented the tendency for the image to be more face or more scene over trials. Both second level beliefs were modulated by their respective variances, ω_s and ω_{ns} . In the social stream, ω_s directly impacted the $x_{1,s}$ social belief by controlling the steepness of the logistic sigmoid function computing $x_{1,s} - \omega_s$ is a multiplicative term on the exponent, and as a result, a larger exponent causes sigmoid function to become steeper. Therefore, a higher ω_s caused the probability of perceiving the bet as accurate or inaccurate to approach more extreme values (0 or 1) in the first level due to the increased steepness, as well as increased the variability of the 2nd level belief. In other words, when ω_s was higher, the social information was more likely to dominate choice relative to non-social information (Figure 6A). In the non-social stream, ω_{ns} reflected the variance at the second level (tendency), while first level estimations $\hat{\mu}_{1,ns}$ about the image classification were directly modulated by the scene percentage of the previous trial, implementing a recency bias. This recency bias captured the ambiguity of the previous image – when the previous image was more ambiguous (50/50 scene face), the term incorporating the previous scene percentage was zero, down weighting the impact of the prior trial, causing $x_{2,ns}$ to have a greater influence on the estimation, $\hat{\mu}_{1,ns}$. In this way, ambiguous images received a negligible recency bias so that individual differences in the perception of these images did not impact the prediction. Only non-ambiguous images should impact the prediction, $\hat{\mu}_{1,ns}$. When the previous image was unambiguously face or scene, the recency bias term dominated and 2nd level information was ignored in favor of the recent sensory evidence. When ω_{ns} increased, the estimate for $x_{2,ns}$ contributed more to the computation of $\hat{\mu}_{1,ns}$ – down weighting the impact of objective sensory evidence on choice (Figure 6B). Both streams of processing were also controlled by the priors at the second level, $\mu_{2,s}^0$ and $\mu_{2,ns}^0$, representing the initial beliefs about the tendency of the bets to be accurate and the tendency of the image to be more face or scene, respectively. An increased $\mu_{2,s}^0$ would represent an increased initial perception of the tendency of the bets to be accurate.

Modeling results

We found a significant difference in the initial beliefs (priors) at $x_{2,s}$ between the cooperation and competition groups. The cooperation group had a significantly elevated $\mu_{2,s}^0$ compared to the competition group ($F(1, 654) =$

Normative self-deception

16.7405, $p_{\text{bonf}}=0.000145$, Figure 7A). The elevation in $\mu_{2,s}^0$ in the cooperation group represents a stronger initial belief that the bet would be more accurate, aligning with, and perhaps underwriting the observed motivational bias effect. There was a significantly increased ω_{ns} in high paranoia participants ($F(1, 654) = 18.6837$, $p_{\text{bonf}} = 5.349e-5$) (Figure 7B). When ω_{ns} is greater, the recency bias term on the 1st level (the influence of the sensory inputs) carries less weight on the prediction of the image classification. As a result, increased variability of the image tendencies reflects an overweighting of this perceived overall tendency (which optimally, should not be utilized as these are independent events), while the sensory evidence is underweighting due to these quantities being inversely related in their contributions to the prediction. It could represent a lack of trust in one's abilities or experiences, independent of others' advice. The high paranoia group also evinced an elevated ω_s ($F(1,654) = 9.425$, $p_{\text{bonf}}=0.006687$), representing a more unstable belief about the tendency of the bet to be accurate (Figure 7C). In both groups we found significant correlations between ω_{ns} and confidence-weighted self-deception, illustrated in Figure 8. While the low paranoia group did display a significant correlation (Pearson's $r = 0.282$, $p = 1.535e-9$), the high paranoia group showed a stronger correlation (Pearson's $r = 0.462$, $p = 6.163e-11$), this difference was significantly different (Fisher's z-transformed r , $p=0.0228$), thus, self-deception independent of paranoia level is driven by ω_{ns} (perceived unreliability of ones' own choices), but the drive is stronger in high paranoia participants. This higher perceived unreliability in one's own choices (and consequentially, discounting current sensory information) combined with an unstable belief about the tendencies of the bet accuracy enhances self-deception in high paranoia participants.

Bet Manipulation

In order to better characterize the impact of social influences on perceptual decisions, we manipulated the accuracy of the bets in a follow up study ($n=324$). In the original work and our first experiment, bets were 50% accurate, reflecting a partner who was performing at chance. We increased bet accuracy to 75%, reflecting a partner whose bets were better than chance.

This manipulation significantly impacted self-deception and confidence. The raw number of self-deceptive trials and normalized confidence in self-deception decreased in the high paranoia group relative to experiment 1 (Independent samples t-test; raw self-deception: $t(192.12) = 3.0756$, $p_{\text{bonf}} = 0.004814$; confidence: $t(96.42) = 2.5655$, $p_{\text{bonf}} = 0.02368$), while remaining unchanged in low paranoia participants (Independent samples t-test; raw self-deception: $t(679.42) = 2.1347$, $p_{\text{bonf}} = 0.06628$; confidence: $t(461.33) = -0.934$, $p_{\text{bonf}} = 0.7016$). This indicates that the high paranoia participants were sensitive to their partners' abilities (Figure 9).

The decrease in self-deception was manifest in model parameter estimates as well. We fit the same 2-layer HGF to the new dataset. The difference in ω_{ns} in high paranoia we found in the original experiment was preserved in

Normative self-deception

the follow-up ($F(1, 319) = 6.4532$, $p_{\text{bonf}}=0.014$), and there was significant no interaction between paranoia group and bet accuracy ($F(1, 977) = 0.451$, $p_{\text{bonf}} = 1$, Figure 10C). High paranoia participants evinced elevated variability in the tendency to perceive the image as face or scene, manifest as an overweighting of stimulus tendency rather than current sensory evidence. In contrast, there was no difference in ω_s ($F(1,319) = 0.0225$, $p_{\text{bonf}} = 1$, Figure 10B) based on paranoia group. This suggests that increasing the partners' accuracy caused high paranoia participants to perceive less social volatility and behave less self-deceptively, centering the impact of such volatility on the belief formation, updating and behavior or paranoid participants, whose challenges are exacerbated by more ambiguous and uncertain situations. We found no difference between the two experiments in terms of the number of trials on which participants could have self-deceived (instances in which bet differed from C1 classifications, Independent samples t-test: $t(532.55) = 0.21728$, $p = 0.8281$). Furthermore, as in experiment 1, stimulus ambiguity drove self-deception – the 50/50 scene/face stimuli were most likely to engender self-deception. However, in experiment 2, the self-deception to less ambiguous cues was less pronounced, obviating the difference between high and low paranoia groups. In experiment 2, we replicated the effect of group manipulation (cooperation versus competition) on initial social beliefs. We found a main effect of experimental group, however, it did not survive Bonferroni correction for multiple comparisons ($F(1, 319) = 4.6709$, $p_{\text{uncorrected}} = 0.03112$, $p_{\text{bonf}} = 0.09336$, Figure 9A). Again, these prior beliefs were no different between the high and low paranoia participants. This replicated lack of interaction is hard to reconcile with models of paranoia that rely on coalitional cognition, since we found no effect of paranoia on coalition or competition⁶.

Model selection and validation

We selected the most appropriate model via Bayesian Model Selection (BMS). For the model space shown in Table 3, we compared a variety of perceptual and response models. With the exception of the Rescorla-Wagner model and a model for random-responding, the model space included a number of 2-layer HGFs with parallel processing of social and non-social stimuli. Determining the bias terms in the model cannot be determined without comparing the various versions and necessitates a quantitative method to determine which model provided a better fit to the data. Due to the high number of models, we used family-wise comparison to narrow down a winning perceptual and winning response model. Family BMS for the perceptual model space yielded a winning model of P1 (HGF with bias term of a scaled- ω_s), with a protected exceedance probability of 1 (Table 4). We found a winning response model of R1 (softmax with decision-noise only) with a protected exceedance probability of 0.9786 (Table 5). Correspondingly, our winning model was M1, which used a P1 perceptual model and R1 response model. Note that as well as including variations of the 2-layer HGF for a perceptual model, we also utilized a Rescorla-Wagner model and a model for random responding, both of which were outperformed by the 2-layer HGF models.

Normative self-deception

Model Simulations

We utilized each individual parameter set found in Experiment 1 to simulate responses for each participant. After simulating, we used the simulated responses to invert the original model in order to validate our findings regarding group differences in parameters. We averaged simulated parameter sets for each participant, and compared the values to parameters found based on their actual responses. Each parameter of interest was significantly correlated with its simulated companion (Figure 11). The group differences based on paranoia group membership for both ω_{ns} ($F(1, 656) = 6.072, p = 0.014$) and ω_s ($F(1, 656) = 15.95, p = 0.724e-5$) were preserved in the simulated parameter sets, as well as the main effect of experimental group on the social priors ($F(1, 651) = 12.51, 0.000432$). Successful parameter recovery and recapitulation of the observed group effects reassures us that we have the appropriate model.

Normative versus non-normative models

The Rescorla and Wagner (1972) rule centers prediction error in learning²¹. Cues have associations with valued outcomes and those associations are updated by mismatches between the associative predictions and the experienced outcomes (prediction errors) weighted by fixed associability parameters that correspond to the salience of the cues and outcomes²⁷. Despite its success, the model is non-normative and heuristic²¹. It does not conform to the principles of probability theory and often performs poorly in real-world situations where outcomes and states must be inferred under uncertainty²¹. The Rescorla-Wagner rule had poor fit to our data, compared to the HGF (see Tables 3 and 4). A normative Bayesian model provided a better account of self-deception and overconfidence and their association with paranoia.

Self-esteem, Paranoia & Overconfidence

The initial classification phase (C1) measures each participant's objective classification ability. We ranked participants on this metric. Next, we ranked participants on their perceived choice reliability during C2 (ω_{ns}). Computing the difference in these ranks gives a metric of participants' insight into their performance. Having a large difference in these ranks (rank of ω_{ns} >> rank of C1-score) corresponds to an overly pessimistic view of oneself. We find a significant correlation between the rank difference and confidence-weighted self-deception (High paranoia: Pearson's $r=0.554, p = 6.163e-16$; Low paranoia: Pearson's $r=0.343, p = 1.618e-13$) which suggests that low self-confidence and diminished ability increase the incidence and confidence of self-deceptive responses (Figure 12). These correlations were significantly different (Fisher's z-transformed $r, p=0.0027$), suggesting paranoid participants take the opportunity to bolster their view of themselves. Overconfidence and self-deception protect against negative self-image, but at an economic cost.

Demographics & Confounds

Normative self-deception

Paranoia often correlates with demographic features and other affective states that might be driving our findings. We found significant correlations between depression and paranoia (Pearson's $r = 0.529$, $p = <2e-16$), and between anxiety and paranoia (Pearson's $r = 0.612$, $p = <2e-16$). In order to dissect the impact of anxiety and depression on the key model parameter (ω_{ns}), we performed a multiple regression with self-ratings of paranoia, anxiety and depression as predictors. Both GPTS (paranoia) and BAI (anxiety) scores were significant predictors of ω_{ns} , while BDI (depression) score were not significant predictors. In order to dissociate the effects of anxiety and paranoia on these parameters, we split participants into four groups: Those with high paranoia and high anxiety (1), high paranoia and low anxiety (2), low paranoia and high anxiety (3), and low paranoia and low anxiety (4). We found that the high paranoia/high anxiety group (1) were no different from the high paranoia/low anxiety group (2) in ω_{ns} ($F(3, 659) = 6.457$, $p = 0.00026$, Post-hoc Tukey test, $p = 0.345$) while the low paranoia/high anxiety group (3) had a significantly lower ω_{ns} compared to the high paranoia/high anxiety group (1) (Post-hoc Tukey test, $p = 0.0034$, Figure 13. Though anxiety and paranoia are highly correlated, paranoia appears more responsible for the group differences in self-deception and the associated model parameters.

Consistent with previous reports, our high paranoia group had higher proportions of black/African-American as well as Hispanic/Latino participants. The high paranoia group had a higher proportion of participants with bachelors and masters degrees, as well as a higher percentage of participants in medium-high income brackets. We also found that a lower percentage of the high paranoia group reported no diagnosis of a psychiatric disorder, while having a higher percentage of people who refused to share this information. The high paranoia group were younger than the low paranoia group (Table 1).

In order to determine how these demographic differences might influence model parameters, we performed ANCOVAs using demographics (race, ethnicity, age, gender), psychiatric diagnosis and medication usage, and socioeconomic factors (income, education) as covariates (Table 2). All effects of paranoia group on ω_{ns} were robust to the inclusion of all the covariates, as was the effect of experimental group on initial beliefs about the bets.

Discussion

People with high paranoia made more high-confidence self-deceptive responses during challenging perceptual decisions under the social influence of a collaborating or competing partner. They overrode their previous choices to agree with collaborators and defect from competitors. This effect was attenuated by making the partner's bets more accurate. We fit a computational model which captured how participants estimated and weighted the influence of current and historical sensory data as well as current and historical social inputs. In this framework self-deception in paranoia was not driven by changes in initial prior weighting of social information (though such

Normative self-deception

priors did distinguish the group working with a collaborator from the group working against a competitor). Rather, the increased self-deception in high paranoia participants was driven by two processes: (1) an underweighting of current sensory inputs relative to the prevailing tendencies from recent trials and (2) an overweighting of the partners' current bet relative to the history of bet accuracy. Taken together, these data are consistent with self-deception flourishing in high paranoia as a result of a lack of confidence in one's own perceptual inferences, coupled with an excessive influence of social suggestions (regardless of affiliation). We observed less self-deception when the partners' bets were more accurate, suggesting that self-deception is particularly likely in paranoid participants when self (non-social) and others (social) are experienced as unreliable sources of information.

Some have argued that motivated reasoning and self-deception contradict Bayesian accounts of belief updating, suggesting instead that biased beliefs are really preferences—things that people desire to be true, and that they are driven by identity (what defines people and their important groups like political parties)¹⁹. Others, have pushed back, suggesting instead that these biases might be understood in terms of differences in (i) perceived reliability of evidence or evidence sources²⁰, (ii) prior beliefs²⁰, or (iii) deriving utility from beliefs and their consistency per se²⁸

The hierarchical Gaussian Filter approach is inherently Bayesian^{21, 22}, since it rests on sequential updating of beliefs according to Bayes' theorem, where beliefs represent inferences about hidden states of the environment (self, others, and external stimuli) in the form of posterior probability distributions, incorporating estimates of estimation uncertainty and environmental uncertainty^{21, 22}. Taking this approach, we found that over-confident self-deception and paranoia can indeed be explained in Bayesian terms: as changes in learning rates and relative weightings of social information, in response to pessimistic estimates about one's own proficiency in perceptual judgments, particularly under high stimulus ambiguity. This model outperformed a simpler non-normative heuristic model²⁷ which neither fit nor simulated our observations. Thus, we find evidence in favor of suggestion (i) and indirectly, in favor of suggestion (iii). Group identity drove changes in prior weightings, however, contrary to suggestion (ii) and coalitional accounts of paranoia, we did not see those prior beliefs contributing significantly to self-deception and paranoia in our data.

Taken together, our data and modeling results suggest that self-deception, a form of biased belief updating, can indeed be explained in normative Bayesian terms. Furthermore, paranoia, albeit on a continuum in the general population (rather than persecutory delusion per se) yields similarly to a normative explanation. This addresses detractors of predictive coding models of delusions¹⁹ and provides empirical support for the proposed solutions to the challenge of biased beliefs²⁰. More broadly, our work joins the growing literature wherein normative Bayesian

Normative self-deception

models can explain apparently irrational phenomena – like political polarization²⁹. Much like perceptual illusions can be cast in Bayesian terms³⁰, decision-making phenomena that appear to depart from normative Bayesian principles, and demand heuristic explanations – evoking the plesitocene evolutionary environment³¹ - may likewise yield to Bayesian explanation, albeit one that incorporates the utility of self-esteem and implements illusory self-regard³¹. Our data suggest the value of this ego defense was greater than the monetary incentives we offered. Future work will explore this trade-off more directly.

In experiment 2, we found that decreasing the ambiguity of the social information (increasing the fidelity of the partner bets) was also impactful. Under social comparison theory, individuals are compelled to improve their performance and minimize discrepancies between their own and others' performance, generating competitive behaviour³². As we describe presently, uncertainty can prompt social comparison^{32, 33}. However, comparison concerns decrease dramatically when uncertainty about one's ranking relative to others is removed³⁴. We contend that increasing the accuracy of partners' bets in experiment 2, neutralized high-confidence self-deception because it made the discrepancy between participant and partner performance clearer and rendered self-deception less necessary, warranted, or appropriate.

In prior work we have related paranoia to deficits in non-social belief updating under uncertainty across species¹⁸. In a between group study we found no difference in the relationships between social and non-social belief updating and paranoia³⁵. However, it would be absurd to suggest that paranoia – a deeply social concern – has no social component. We argued instead that paranoia is underwritten by deficits in domain-general belief-updating which are accommodated by overarching social explanations, such as the ministrations of powerful others. Others have examined how volatile streams of social and non-social information are combined in the brain²⁶. Perturbations in those processes have been related to paranoia, learning rates and beliefs about volatility^{36, 37}. We replicate and extend that work and our own here, delineating how deficits in non-social inference can lead to overconfident self-deception under social influence. Notably however, we did not establish whether there were changes in paranoia as a result of experiencing the task and evincing self-deception. Future work with the task might establish – using momentary assessment for example – whether experiencing spurious social suggestion during a difficult perceptual task exacerbates paranoia.

Our modeling work was consistent with self-deception impacting self-esteem and thence over-confidence in high paranoia participants. However, our task did not have a conduit for that over-confidence – in terms of convincing others of one's insights or abilities³. A task with reciprocal exchange between participants would be enlightening. Differing self-deception when confidence is communicated between partners would be consistent with a role for self-deception in deceiving others as well as self¹. In an advice-giving task, patients with schizophrenia were

Normative self-deception

overconfident in their own advice, particularly those with delusions²⁹. Our data suggest this effect might be driven by self-deception secondary to an experience of one's own perceptual unreliability.

The overconfidence that self-deception engenders can be adaptive, increasing ambition, morale, resolve, and generating a self-fulfilling prophecy that actually increases the probability of success³⁸. However, it also leads to faulty assessments, unrealistic expectations and hazardous decisions³⁸. Too much self-deception in key opinion-leaders engenders hubris, market bubbles, financial collapses, policy failures, disasters and costly wars³⁸. With this task and our quantitative individual-level model of self-deception, we can predict which individuals might be susceptible, and we can optimize or militate against it. Our data suggest numerous mitigation opportunities. Boosting self-esteem – by conditioning positive self-associations – appears to mollify paranoia³⁹, it ought to similarly diffuse self-deception. Providing feedback on classification performance might give a more realistic assessment of the reliability of one's own judgments and emphasizing successful classifications may similarly prevent self-deception and possibly paranoia. Clarifying partner abilities relative to self also mitigated self-deception in high paranoia participants. This suggests that leaders and advisers who are more reliably proficient and perhaps less overconfident themselves, are less likely to augur self-deception and paranoia in their followers. The mechanisms we identify presently might underwrite the rise of bizarre conspiracy theories, expressed with extreme confidence in the face of overwhelming contradictory evidence.

It remains unclear whether self-deceptive responses have their locus in sensory processes or in post-sensory judgement and decision-making⁴⁰. Gathering neuroimaging data during task performance will answer that question¹⁵. Using neural decoding we can, with confidence, discern whether someone is perceiving a face or a house⁴¹. If the partner bet induces self-deception, we can use decoding methodology to conclude whether or not neural responses in cortices that represent face and house are actually altered, or whether alternatively, self-deception has its locus in down-stream post-perceptual judgment mechanisms, and we can align those data with model parameter estimates to establish how these mechanisms are instantiated in the brain and how they contribute to perception versus judgement.

Given the debate about self-deception and delusions⁴², it will be important to establish whether the same effects are present in people with confirmed delusional beliefs. Recent work on advice giving by people with schizophrenia suggests that patients with delusions are over confident in their advice²⁹. We suggest that our data are consistent with the possibility that delusion (albeit on the extreme end of a continuum of paranoia) might entail self-deception. At the same time – in light of our data – delusion and self-deception may not violate epistemic rationality⁴³ and might harbor adaptive function⁴⁴.

Normative self-deception

In silico, evolutionary simulations suggest that overconfidence increases under uncertainty². Our data provide empirical support for this prediction. Furthermore, we suggest that such over-confidence flourishes in the face of difficult decisions when one's foregoing performance is perceived as unreliable, even when the overconfidence is not communicated to the partner (our participants are not bluffing their partners). Novel threats (like pandemics), complex technologies, and new and untested leaders, all entail uncertainty. When costs are low, some judicious self-deception may be adaptive. However, owing to their volatility, their capacity for unpredicted and unpredictable change, these are also the most dangerous situations – in which costs can spike and overconfidence can become fatal. We have tracked how real-world volatility increases paranoia³⁵, and we have observed its impact on overconfidence; in poor institutional and individual responses to the pandemic. Removing the opportunity for self-deception (with less ambiguous decisions and more reliable advice) and otherwise assuaging self-esteem should mollify overconfidence and paranoia.

Methods

Participants (N = 719) were recruited for experiment 1 online via CloudResearch – an online research platform that provides access to Amazon MTurk's worker population with enhanced control of recruitment (cite TurkPrime paper). The study was available to workers with a 90% of higher HIT approval rate, only in the United States. Only workers who did not attempt experiment 1 were allowed to attempt experiment 2. Responses were examined after to verify human participation – participants who declined more than 30% of the survey responses were automatically rejected. In addition, free response answers containing nonsensical answers were also flagged for review and were rejected. We also excluded responses on the basis of random responding. Our task includes pictures that are clearly scene or face (no ambiguity), and provide a control to see if participants were completing the task in good faith or randomly pressing keys. We identified 48 individuals who failed over 50% of these control trials (a majority of which also answered each trial with the same response) and excluded them from the analysis, as we could not say that they truly attempting to correctly categorize the images. 23 of the submissions containing nonsensical answers to free response questions were included in the 48 submissions discarded from random responding. In addition, 8 participants were discarded due to missing identifier keys, as task data could not be correctly matched to survey data. Only participants who passed these checks had their data included for analysis of self-deception, as random responding might artificially inflate self-deception. For experiment 1, our total sample (N = 663) of complete submissions included 334 participants for the competition condition and 329 participants for the cooperation condition.

For experiment 2, we applied the same requirements for complete submissions as in experiment 1. Participants (N = 327) were recruited through CloudResearch, this time using the new Data Quality feature. This allowed only

Normative self-deception

high quality participants into our study. As a result, only 3 submissions were excluded, 2 for missing identifier information and 1 for random responding patterns following our criteria above.

Behavioral task

Participants (N = 663) completed a perceptual categorization task with 2 phases. In the first classification phase (C1), participants were asked to categorization each image as a face or a scene, and provide a self-report measure of confidence after each classification. There were 80 trials in C1, with varying images of face and scene. Images were provided by Leong and colleagues (2019). The distribution of trials for each scene percentage was as follows: 4 trials (0% scene), 6 trials (35% scene), 8 trials (40% scene), 14 trials (45% scene), 16 trials (50% scene), 14 trials (55% scene), 8 trials (60% scene), 6 trials (65 % scene), 4 trials (100% scene). Stimuli were randomized and each participant was randomly assigned to a different stimulus set (9 total) to control for the effects of different faces. Confidence was reported using a 7 item Likert-Scale.

After completing C1, participants were informed that they would be working with a collaborator, or competing against a competitor. Participants were told the other would be making a bet before they were shown an image, and were told the various payments for answering the same or different to the bet. Participants completed only one version of the experiment. Participants were not told any information about how accurate the bets would be – for experiment 1, exactly half of the bets aligned with the correct answer. Participants were then presented with a bet from their collaborator or competitor, after which an image from C1 was shown. After each image, confidence was then assessed. There were 80 trials in C2, since every image shown in C1 was again shown to the participants. The underlying bet accuracy was manipulated to be 75% correct in experiment 2, however, participants were not given any extra information than participants in experiment 1.

The task was administered via CloudResearch, from a web browser link. There was no minimum time to respond, as poor internet connections and other demands can impede the workers ability to answer within a short time frame. Workers joined the HIT without knowing which version they were doing.

Questionnaires

After the task was completed, workers were redirected to a Qualtrics survey. The survey included demographic information (age, gender, income, educational level, ethnicity, and race) as well as mental health questions (diagnosis, medication use). Inventories were also included such as the Revised Green et al. Paranoia Thoughts Scale (R-GPTS)²⁴, Beck's Anxiety Inventory (BAI)⁴⁵, Beck's Depression Inventory (BDI)⁴⁶, and the Pathological Narcissism Inventory (PNI)⁴⁷. We also included free response questions asking about their experience and what they thought the experiment was measuring in order to look with nonsense linguistic patterns indicating a non-

Normative self-deception

human response. Participants who scored 11 or higher on the R-GPTS persecution scale were classified as high paranoia as this is the ROC-recommended clinical cutoff.

Behavioral Analysis

Motivational bias was assessed through fitting participant choice data to a general linear mixed effects model (GLMM) using the lme4 package in R. GLMMs were also fit to choice data using only scene percentage as a variable in order to confirm that classifications were related to the objective scene percentage (rather than random responding).

Self-deceptive responses were classified using the patterns of response determined in Mijović-Prelec and Prelec²³. We chose to use their metric capturing the total amount of self-deceptive responses rather than incorporating another one of their response types (such as “inconsistent” trials) as these trial types were not yoked in the design – a pattern of C1 response and bet did not result in the participant making a decision that was either self-deceptive or inconsistent. Our confidence weighted self-deception metric was computed as follows:

$$CWSD = raw\ self - deception\ score * \left(\frac{\text{mean confidence on self-deceptive trials}}{\text{mean confidence on C1 trials}} \right) \quad (1)$$

As a result, each score accounted for individual differences in self-reported confidence, and captured differences in confidence on self-deceptive trials.

Computational Modeling

We adapted open-access code for an HGF with 2 streams of processing²⁵ in MATLAB 2018a (MathWorks R, Natick, MA). The HGF requires the combination of the generative model of the agent’s inferences (perceptual model) with a response model incorporating the agent’s actions. The model is inverted and parameter values are computed from the participants choices. Our perceptual model was a 2-layered HGF, while we used a softmax model for binary choices for the response model.

The first level of the generative model ($x_{1,s}$ and $x_{1,ns}$) represent the beliefs about the accuracy of the bet (1=correct, 0=incorrect) and the image category (1=scene, 0=face), respectively. Then, the second level describes the perception of the tendency of the first level, so the tendency for the bet to be correct ($x_{2,s}$) and the tendency of the image category ($x_{2,ns}$). The 2nd level has a Markov-like dependence where the estimate of $x_{2,s}$ and $x_{2,ns}$ are updated from their respective values on the previous time step. In contrast, the first level beliefs are computed directly from 2nd level at time t , through a logistic sigmoid where:

Normative self-deception

$$\hat{\mu}_{1,s}(t) = \frac{1}{1 + e^{-\eta \omega_s \mu_{2,s}(t-1)}} \quad (2)$$

$$\hat{\mu}_{1,ns}(t) = \frac{1}{1 + e^{-(\mu_{2,ns}(t-1) + \text{recency bias})}} \quad (3)$$

Since the streams of processing used different response spaces (bet accuracy vs. image category), the binary choices used to invert the model were coded by answering with the bet (1) or answering against the bet (0). This response coding could then be transformed into the separate spaces to invert the model in both processing streams.

The specific formulations of equations 2 and 3 were deduced from model comparison. First, since a prediction of an unknown image logically might be influenced by the previous image, we incorporated a recency bias that biased the non-social prediction towards the previous image depending upon the ambiguity of the image. On the social side, there was no obvious way to incorporate how the bet might bias perceptual inferences. We decided to explore how adding a bias term on the logistic sigmoid connecting $x_{2,s}$ and $x_{1,s}$ might help explain the behavioral effect of motivated perception. Since there was no established way to recapitulate this effect, we tried to incorporate the bias in a few logical ways: an additive term on the exponent (shifting the inflection point of the psychometric curve), a multiplicative term on the exponent (shifting the steepness of the psychometric curve), or a combination of those terms. We determined that a multiplicative term provided better fit, but upon further exploration we found that this term had a high correlation with ω_s (Pearson's $r = 0.998$, $p < 2.2 \cdot 10^{-16}$). As a result, we tried to replace this multiplicative term with ω_s (perceptual model P3). However, the high correlation does not account for the fact that the relationship between these terms might not be one-to-one, and as a result, the best-fitting model was found to have a bias term that was a linear scaling of ω_s as a multiplicative term on the exponent. A diagram detailing the impact of these added parameters is shown in Figure 13.

Although mapping the 2nd to first level was different between the two streams, the computations by which the beliefs evolved on the 2nd level were the same for the 2 processing streams. The belief at the second level (μ_2), is updated by the precision-weighted prediction error from the first level:

$$\Delta \mu_2(t) \propto \frac{\hat{\pi}_2(t)}{\pi_2(t)} \delta_1(t) \quad (4)$$

where δ_1 is the prediction error at the first level and $\pi_2(t)$ is the precision of the posterior second level belief.

Normative self-deception

$$\delta_1(t) = \mu_1(t) - \hat{\mu}_1(t) \quad (5)$$

The first level predicted belief ($\hat{\mu}_1$) is determined by the logistic sigmoid above (equation 2, 3), and the prediction error generated incorporates the model inputs (bet accuracy and scene percentage) for the respective processing streams for the current trial.

In order to combine the two information streams, the combined belief, $b(t)$, was computed as a linear combination of the predictions of the first level beliefs, weighted by their precisions.

$$b(t) = \frac{\hat{\pi}_{1,s}(t)\hat{\mu}_{1,s}(t) + \hat{\pi}_{1,ns}(t)\hat{\mu}_{1,ns}(t)}{\hat{\pi}_{1,s}(t) + \hat{\pi}_{1,ns}(t)} \quad (6)$$

This combined belief was then fed into a softmax function to compute the probability of agreeing with the bet:

$$P(y(t) = 1) = \frac{b^\beta}{b^\beta + (1 - b)^\beta} \quad (7)$$

As an alternate response model, we also examined the effect of adding a bias term to induce unequal weighting of the two streams as seen in Diaconescu et al. (2014)⁴⁸ as well as Henco et al. (2020)²⁵, which ultimately did not produce a good fit to the behavioral data (model comparison values displayed in Table 5 for response models).

Model Selection

Bayesian model selection was conducted in MATLAB using the VBA_groupBMC function in the VBM toolbox⁴⁹. Posterior probabilities, exceedance probabilities, and protected exceedance probabilities for the various perceptual and response models are displayed in Tables 4 and 5. The winning model was a combination of the winning perceptual model (P1) and the winning response model (R1).

Simulations

For each participant's estimated parameter set, we simulated 30 sets of response vectors. Then, we inverted the model using each simulated response vector as the choice data, and generated 30 simulated parameter sets for each participant. We were able to compare the simulated parameter sets to the original parameter set for each participant. All group differences were maintained using the simulated parameters, demonstrating that the group differences in parameters could be recovered and were fairly resilient to changes in exact response patterns. All parameters of interest were significantly correlated with their simulated companions. Of note, a large majority of the images were very ambiguous, and as a result, an objective classification would follow a Bernoulli distribution,

Normative self-deception

so we might expect a number of responses to have high variability in response. Despite these trials having high randomness, group differences are still preserved, even when the responses are different.

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 1.2.5033.

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. Correlations were computed with Pearson's rho.

Model parameters and self-deception scores were analyzed using ANOVAs, with Bonferroni correction for multiple-comparison (as needed). We performed ANCOVAs for model parameters using three sets of covariates: (1) demographics (age, gender, ethnicity, and race); (2) mental health factors (medication usage, diagnostic category); (3) and metrics and correlates of global cognitive function (educational attainment, income).

Normative self-deception

1. von Hippel, W. & Trivers, R. The evolution and psychology of self-deception. *Behav Brain Sci* **34**, 1-16; discussion 16-56 (2011).
2. Johnson, D.D. & Fowler, J.H. The evolution of overconfidence. *Nature* **477**, 317-320 (2011).
3. Schwardmann, P. & van der Weele, J. Deception and self-deception. *Nat Hum Behav* **3**, 1055-1061 (2019).
4. Dunning, D., Heath, C. & Suls, J.M. Flawed Self-Assessment: Implications for Health, Education, and the Workplace. *Psychol Sci Public Interest* **5**, 69-106 (2004).
5. Barber, B.M., Odean, T. The courage of misguided convictions *Financ. Analysts J.* **55**, 41-55 (1999).
6. Raihani, N.J. & Bell, V. An evolutionary perspective on paranoia. *Nat Hum Behav* **3**, 114-121 (2019).
7. Kinderman, P. & Bentall, R.P. Self-discrepancies and persecutory delusions: evidence for a model of paranoid ideation. *J Abnorm Psychol* **105**, 106-113 (1996).
8. Kaney, S. & Bentall, R.P. Persecutory delusions and the self-serving bias. Evidence from a contingency judgment task. *J Nerv Ment Dis* **180**, 773-780 (1992).
9. Penrod, S., Cutler, B. Witness confidence and witness accuracy: Assessing their forensic relation. *Psychology, Public Policy, and the Law* **1**, 817-845 (1995).
10. Zarnoth, P. & Sniezek, J.A. The Social Influence of Confidence in Group Decision Making. *J Exp Soc Psychol* **33**, 345-366 (1997).
11. Krebs, D.L., Laird, . Judging Yourself as You Judge Others: Moral Development and Exculpation. *Journal of Adult Development* **5** (1998).
12. Denton, K., Zabatany, L. Age differences in support processes in conversations between friends. *Child Development* **67**, 1360-1373 (1996).
13. Sloman, S.A., Fernbach, P.M. & Hagmayer, Y. Self-deception requires vagueness. *Cognition* **115**, 268-281 (2010).
14. Chance, Z., Gino, F., Norton, M.I. & Ariely, D. The slow decay and quick revival of self-deception. *Front Psychol* **6**, 1075 (2015).
15. Leong, Y.C., Hughes, B.L., Wang, Y. & Zaki, J. Neurocomputational mechanisms underlying motivated seeing. *Nat Hum Behav* **3**, 962-973 (2019).
16. van Prooijen, J.W. & Douglas, K.M. Conspiracy theories as part of history: The role of societal crisis situations. *Mem Stud* **10**, 323-333 (2017).
17. Smallman, S. Whom do You Trust? Doubt and Conspiracy Theories in the 2009 Influenza Pandemic. *Journal of International and Global Studies* **6**, 1-24 (2015).
18. Reed, E.J., et al. Paranoia as a deficit in non-social belief updating. *Elife* **9** (2020).
19. Williams, D. Hierarchical Bayesian models of delusion. *Conscious Cogn* **61**, 129-147 (2018).
20. Tappin, B.M. & Gadsby, S. Biased belief in the Bayesian brain: A deeper look at the evidence. *Conscious Cogn* **68**, 107-114 (2019).
21. Mathys, C., Daunizeau, J., Friston, K.J. & Stephan, K.E. A bayesian foundation for individual learning under uncertainty. *Frontiers in human neuroscience* **5**, 39 (2011).
22. Mathys, C.D., et al. Uncertainty in perception and the Hierarchical Gaussian Filter. *Frontiers in human neuroscience* **8**, 825 (2014).
23. Mijovic-Prelec, D. & Prelec, D. Self-deception as self-signalling: a model and experimental evidence. *Philos Trans R Soc Lond B Biol Sci* **365**, 227-240 (2010).
24. Green, C.E., et al. Measuring ideas of persecution and social reference: the Green et al. Paranoid Thought Scales (GPTS). *Psychol Med* **38**, 101-111 (2008).
25. 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).

Normative self-deception

26. Behrens, T.E., Hunt, L.T., Woolrich, M.W. & Rushworth, M.F. Associative learning of social value. *Nature* **456**, 245-249 (2008).
27. Rescorla, R.A., Wagner, A.R. A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and non-reinforcement. in *Classical conditioning II: Current research and theory* (ed. A.H. Black, Prokasy, W.F.) (Appleton-Century-Crofts, New York, 1972).
28. Mobius, M., Niederle, M., Niehaus, P., Rosenblat, T. Managing Self-Confidence: Theory and Experimental Evidence. *National Bureau of Economic Research Working Paper Series II*, 11-14 (2012).
29. Hertz, U., Bell, V., Barnby, J., Quillin, A., Bahrami, B. The Communication of Metacognition for Social Strategy in Psychosis: An Exploratory Study. *Schizophrenia Bulletin Open* **1** (2020).
30. Geisler, W.S. & Kersten, D. Illusions, perception and Bayes. *Nat Neurosci* **5**, 508-510 (2002).
31. Haselton, M.G. & Nettle, D. The paranoid optimist: an integrative evolutionary model of cognitive biases. *Pers Soc Psychol Rev* **10**, 47-66 (2006).
32. Festinger, L. A theory of social comparison processes. *Human Relations* **7**, 117-140 (1954).
33. Garcia, S.M., Tor, A. & Schiff, T.M. The Psychology of Competition: A Social Comparison Perspective. *Perspect Psychol Sci* **8**, 634-650 (2013).
34. Garcia, S.M., Tor, A. Rankings, standards, and competition: Task vs. scale comparisons. *Organizational Behavior and Human Decision Processes* **102**, 95-108 (2007).
35. Suthaharan, P.R., E. J.,....Corlett, P.R. Paranoia and Belief Updating During a Crisis. *Psycharxiv Preprint* (2020).
36. Diaconescu, A.O., Wellstein, K.V., Kasper, L., Mathys, C. & Stephan, K.E. Hierarchical Bayesian models of social inference for probing persecutory delusional ideation. *J Abnorm Psychol* **129**, 556-569 (2020).
37. Wellstein, K.V., et al. Inflexible social inference in individuals with subclinical persecutory delusional tendencies. *Schizophr Res* **215**, 344-351 (2020).
38. McKay, R.T. & Dennett, D.C. The evolution of misbelief. *Behav Brain Sci* **32**, 493-510; discussion 510-461 (2009).
39. Espinosa, R., Valiente, C., Varese, F. & Bentall, R.P. Can We Ameliorate Psychotic Symptoms by Improving Implicit Self-Esteem? A Proof-of-Concept Experience Sampling Study of an Evaluative Classical Conditioning Intervention. *J Nerv Ment Dis* **206**, 699-704 (2018).
40. Firestone, C. & Scholl, B.J. Cognition does not affect perception: Evaluating the evidence for "top-down" effects. *Behav Brain Sci* **39**, e229 (2016).
41. Kay, K.N., Naselaris, T., Prenger, R.J. & Gallant, J.L. Identifying natural images from human brain activity. *Nature* **452**, 352-355 (2008).
42. McKay, R., Langdon, R. & Coltheart, M. "Sleights of mind": delusions, defences, and self-deception. *Cogn Neuropsychiatry* **10**, 305-326 (2005).
43. Bortolotti, L. & Mameli, M. Self-Deception, Delusion and the Boundaries of Folk Psychology. *Humanamente* **20**, 203-221 (2012).
44. Fineberg, S.K. & Corlett, P.R. The doxastic shear pin: delusions as errors of learning and memory. *Cogn Neuropsychiatry* **21**, 73-89 (2016).
45. 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).
46. Beck, A.T., Ward, C.H., Mendelson, M., Mock, J. & Erbaugh, J. An inventory for measuring depression. *Archives of general psychiatry* **4**, 561-571 (1961).
47. Pincus, A.L., et al. Initial construction and validation of the Pathological Narcissism Inventory. *Psychol Assess* **21**, 365-379 (2009).
48. Diaconescu, A.O., et al. Inferring on the intentions of others by hierarchical Bayesian learning. *PLoS Comput Biol* **10**, e1003810 (2014).

Normative self-deception

49. Rigoux, L., Stephan, K.E., Friston, K.J. & Daunizeau, J. Bayesian model selection for group studies - revisited. *Neuroimage* **84**, 971-985 (2014).

Normative self-deception

Figure 1

Task structure for C2 phase. **A**, sequence of task for the 2 conditions. **B**, payoff matrices for both conditions. Participants should ideally classify the image objectively (as they did in the initial classification phase) without using the bet to inform their decision.

Normative self-deception

813
814
815
816
817

818
819
820
821

Figure 2

Participant’s choices displayed a motivational bias. The bet x group interaction shows that participants in the cooperation group tended to align with the bet (higher probability of answering scene when the bet was scene), while the competition group tended to disagree with the bet (higher probability of responding with scene when the bet was face).

826
827
828
829
830
831
832
833
834

Normative self-deception

Figure 3

Self-deception and confidence is different between groups. **A**, the high paranoia group had elevated raw self-deception scores (percentage of self-deceptive responses). **B**, mean confidence on those self-deceptive trials was elevated in high paranoia participants. **C**, the confidence-weighted self-deception, which controls for individual variation in baseline confidence, is higher in the high paranoia group. **D**, confidence-weighted self-deception is also elevated in the cooperation group relative to the competition group. * $P \leq 0.05$, ** $P \leq 0.01$, *** $P \leq 0.001$.

Normative self-deception

Figure 4

Self-deceptive responses occurred more with ambiguous images. The high paranoia group self-decieved more on slightly less ambiguous images than the low paranoia group.

Figure 5

The 2-level HGF with parallel processing streams for social and non-social stimuli. The choice data is fed into the model, which is inverted to obtain parameter estimations for an individual. The perceptual model includes the both the social and non-social information, which is then used to compute the combined belief, b . This combined belief is the input to the response model. More details are given in the methods.

Normative self-deception

Figure 6

Modulation of ω_s and ω_{ns} impact image categorization. **A**, increasing ω_s causes the prediction about the accuracy of the bet ($\hat{\mu}_{1,s}$) to become closer to an extreme (1 or 0). This tilts the combined belief towards this prediction. **B**, increasing ω_{ns} causes the recency bias to have less of an effect on the prediction about the image categorization ($\hat{\mu}_{1,ns}$) while perceived tendency (second level belief) dominates the prediction. The effect is stronger when the recent image is very ambiguous.

Normative self-deception

Figure 7

Estimated parameters show differences based on paranoia group and experimental group. **A**, the cooperation group has an elevated prior for the social information ($\mu_{2,s}^0$) compared to the competition group. **B**, the variance of the perceived tendency of image categorization (ω_{ns}) is increased in high paranoia group as well as **C**, the variance of the perceived tendency of bet accuracy (ω_s). * $P \leq 0.05$, ** $P \leq 0.01$, *** $P \leq 0.001$.

Normative self-deception

Figure 8

ω_{ns} (variance of $x_{2,ns}$) is correlated with confidence weighted self-deception. The correlation is statistically stronger in the high paranoia group. $*P \leq 0.05$, $**P \leq 0.01$, $***P \leq 0.001$.

Normative self-deception

Figure 9

Self-deception and confidence-weighted self-deception is different between the two experiments only in the high paranoia group. **A**, raw self-deception scores are lower in the high paranoia group with more accurate bets than the high paranoia group with less accurate bets. There is no difference in the low paranoia group. **B**, mean confidence on self-deceptive trials decreases in the high paranoia group with higher bet accuracy. * $P \leq 0.05$, ** $P \leq 0.01$, *** $P \leq 0.001$.

939

940

Figure 10

Group parameter differences in the two experiments differ. **A**, both experiments show a difference in social priors ($\mu_{2,s}^0$), with the cooperation group having an increased prior compared to the competition group. **B**, the difference in the variance of the 2nd level belief on the social side (ω_s) between paranoia groups disappears when the bet accuracy is higher in experiment 2. **C**, the difference between paranoia groups in the variance of the 2nd level belief about the image (ω_{ns}) is maintained in both experiments. * $P \leq 0.05$, ** $P \leq 0.01$, *** $P \leq 0.001$.

957

958

959

960

961

962

963

Figure 11

Parameters of interest were correlated with their simulated values obtained from simulation of responses and inversion of the model.

Figure 12

Difference between the relative perceived choice reliability (rank of ω_{ns}) and relative objective classification performance (rank of C1 accuracy) is correlated with confidence-weighted self-deception. A higher ω_{ns} represents a lower perceived choice reliability (analogously, a higher perceived choice unreliability), so individuals scoring high on this rank difference have high perceived unreliability and low ability. This relationship is significantly stronger in the high paranoia group compared to the low paranoia group.

Normative self-deception

Figure 13

High paranoia is responsible for elevated ω_{ns} independent of anxiety. The high paranoia high anxiety group showed similar values of ω_{ns} to the high paranoia low anxiety group, while the low paranoia high anxiety group had a significantly decreased ω_{ns} .

Table 1

Demographics table for experiment 1

	Low paranoia (n=469)	High paranoia (n=194)	Statistic	P-value
Age	38.6 (10.96)	35.92 (9.28)	-3.1899(396.16)	0.001537
Gender			5.49(3)	0.1393
% Female	42.86%	47.42%		
% Male	54.16%	51.03%		
% Other/not specified	2.98%	1.52%		
Education			53.39(7)	3.1e-9
% High school degree or equivalent	9.38%	7.22%		
% Associate degree	12.16%	6.7%		
% Bachelor's degree	43.28%	44.85%		
% Master's degree	7.68%	26.2%		
% Doctorate or professional degree	1.7%	1.03%		
% Completed some college	22.43%	13.85 %		
% Other/not specified	3.41%	0%		
Ethnicity			21.965(2)	1.7e-5
% Hispanic or Latino	8.1%	20.6%		
% Not Hispanic or Latino	89.13%	78.35%		
% Not specified	2.77%	1.03%		
Race			36.19(5)	8.69e-7
% White	76.76%	65.98%		
% Black or African American	9.38%	22.16%		
% Asian	4.26%	5.67%		
% American Indian or Alaska Native	0.21%	2.57%		
% Multiracial	6.18%	3.61%		
% Other/not specified	3.2%	0%		
Income			17.15(6)	0.0087
% Less than \$20,000	18.55%	12.88%		
% \$20,000 to \$34,999	17.91%	17.01%		
% \$35,000 to \$49,999	16.41%	22.16%		
% \$50,000 to \$74,999	24.73%	30.93%		
% \$75,000 to \$99,999	11.94%	13.92%		
% Over \$100,000	6.18%	2.58%		
% Decline to answer	4.3%	0.52%		
Psychiatric Diagnosis			53.39(7)	3.1e-9
% No psychiatric diagnosis	61.19%	36.08%		
% Schizophrenia spectrum	0.43%	0.515%		

Normative self-deception

% Anxiety/Depression only	8.32%	13.92%		
% Bipolar disorder	1.49%	0%		
% Multiple/not specified	10%	9.28%		
% Personality disorder	0.64%	3.6%		
% Decline to state	17.9%	36.6%		
% Medicated	6.39%	5.15%	0.186(1)	0.6659
BAI	5.77(7.83)	21.68(15.49)	13.6(234.84)	<2.2e-16
BDI	8.08(10.49)	23.38(15.39)	12.68(270.18)	<2.2e-16
GPTS	0.22(0.4)	2.26(0.69)	38.55(252.11)	<2.2e-16

Table 2

ANCOVAs for model parameters

Effect	Df	ω_{ns}		ω_s		$\mu_{2,s}^0$	
		F	p-value	F	p-value	F	p-value
<u>Demographics</u>							
Age	1, 628	0.003	0.4856	0.188	0.665	1.929	0.165
Gender	3, 628	2.517	0.5723	1.065	0.363	1.695	0.167
Race	5, 628	2.302	0.0434	0.813	0.541	0.472	0.797
Ethnicity	2, 628	19.166	8.33e-9	11.193	1.67e-5	0.013	.987
Paranoia Group	1, 624	13.871	0.000214	7.149	0.0077	1.586	0.208
Version	1, 627	6.705	0.00984	1.792	0.181	29.079	1.613e-7
Paranoia Group*Version	1, 622	0.071	0.790	2.303	0.12967	1.670	0.197
<u>Mental Health Factors</u>							
Psychotropic Medication	1, 655	2.160	0.142	0.090	0.76388	0.538	0.463
Diagnosis	6, 655	3.518	0.00196	3.423	0.00246	0.121	0.994
Paranoia Group	1, 649	16.217	6.32e-5	4.593	0.03247	3.483	0.0624
Version	1, 654	7.503	0.00633	0.844	0.35856	27.101	2.59e-7
Paranoia Group*Version	1, 647	0.316	0.57424	3.505	0.06164	2.883	0.09
<u>Cognitive Ability</u>							
Income	6, 649	1.622	0.13832	1.353	0.2313	1.191	0.309
Education	7, 649	2.494	0.0156	2.320	0.0242	0.788	0.598
Paranoia Group	1, 643	14.161	0.000183	5.673	0.0175	2.259	0.133
Version	1, 648	6.888	0.008881	0.872	0.3507	26.831	2.97e-7
Paranoia Group*Version	1, 641	0.365	0.5457	3.642	0.0568	2.274	0.132

Table 3

Model Space

Response	Decision-noise only	Decision-noise and bias term for unequal weighting
Perceptual		
HGF with scaled ω_s bias term on 1 st layer (P1)	M1	M2
HGF with $\omega_s + \zeta_s$ bias term on 1 st layer (P2)	M3	M4
HGF with non-scaled ω_s bias term on 1 st layer (P3)	M5	M6
HGF with bias of constant ζ_s on 1 st layer (P4)	M7	M8
HGF with no bias terms on social layers (P5)	M9	M10
HGF with scaled ω_s bias term on 1 st layer with constant ω_s and ω_{ns} (P6)	M11	M12
Rescorla-Wagner (P7)	M13	M14
Random responding (P8)	M15	M16

Table 4

Family-wise Bayesian model selection for perceptual models

Perceptual Model Family	P1	P2	P3	P4	P5	P6	P7	P8
Posterior probability	0.5277	0.2236	0.0002	0.0002	0.2447	0.0002	0.0032	0.0002
Exceedance probability	1	0	0	0	0	0	0	0
Protected exceedance probability	1	0	0	0	0	0	0	0

Table 5

Family-wise Bayesian model selection for response models

Response Model Family	R1	R2
Posterior probability	0.5390	0.4610
Exceedance Probability	0.9786	0.0214
Protected exceedance probability	0.9786	0.0214