Interactions between attributions and beliefs at high temporal resolution : evidence from a novel task

Elena Zamfir, Peter Dayan
MPI for Biological Cybernetics, Tuebingen

Abstract

Inferring causes of the good and bad events that we experience is part of the process of building models of our own capabilities and of the world around us. Making such inferences can be difficult because of complex reciprocal relationships between attributions of the causes of particular events, and beliefs about the capabilities and skills that influence our role in bringing them about. Abnormal causal attributions have long been studied in connection with psychiatric disorders, notably depression and paranoia; however, the mechanisms behind attributional inferences and the way they can go awry are not fully understood. We administered a novel, challenging, game of skill to a substantial population of online participants, and collected trial-by-trial time series of both their beliefs about skill and attributions about the causes of the success and failure of real experienced outcomes. We found reciprocal relationships that provide empirical confirmation of the attribution-self representation cycle theory. This highlights the dynamic nature of the processes involved in attribution, and validates a framework for developing and testing computational accounts of attribution-belief interactions.

1 Introduction

When people succeed or fail at achieving their goals, they typically use beliefs about their *abilities* or skills to *attribute* the result to their own internal efforts or to external factors. Conversely, people typically use appraisals of past successes and failures to inform their beliefs about themselves (Bentall, 2003; Bentall et al., 2001).

A simple concrete example can be used to illustrate these two processes at work: consider a student who finds out they got a low mark in a math assignment. If the student has a strong belief in their mathematical ability or general intelligence they might assign this low grade to the assignment being excessively difficult or confusing – an *external* cause. Alternatively, they might consider themselves to have been tired and unable to properly focus when working on their homework – an *internal*, but *transient* cause. In contrast, a student lacking confidence in their math abilities might interpret the low grade as a result of their stupidity, and attribute the failure to this *internal* and *stable* factor – if one is stupid, one is likely to remain so in the future. Thus beliefs can influence causal attributions. Conversely, attributions can influence beliefs: assigning the low grade to external factors or to tiredness will most likely not affect the student's estimation of their ability in math; however attributing the low grade to stupidity is likely to strengthen their belief in their inability, and impact their motivation for studying this topic in the future.

Crucially, such appraisals do not occur in isolation: students will get a large number of various grades on their assignments or exams throughout their studies, creating the potential for interactions between successive such appraisals. Most students will successfully manage the emotional impact of such micro triumphs or disasters, maintaining their well-being while learning about and striving to improve their knowledge and skill base. However more extreme or even tragic path dependencies (Dayan et al., 2020) are also possible: if a student who is slightly prone to low self-esteem suffers from a string of such failures, each could amplify the effect of the next, by strengthening the student's belief in their inadequacies, making them more likely to interpret feedback in a negative way and therefore further erode their self esteem. Getting trapped in such vicious cycles could drain the students' energy and motivation, making it increasingly difficult to maintain a healthy and productive approach to their studies, and affecting their mental health.

This example illustrates how pervasive such inferences are in everyday situations, but also shows the potential for reciprocal connections between causal attributions and beliefs at the level of individual events. If such interactions at high temporal resolution are indeed present, the accumulation of their effects could produce rich dynamics, ranging from the maintenance of a homeostatic regime to vicious cycles spiralling out of control (Bentall, 2003; Bentall et al., 2001).

Abnormal beliefs and attributional patterns have been topics of particular interest in connection with psychiatric disorders, and a wealth of theoretical proposals has been put forth about the psychological mechanisms involved, and how they malfunction in various disorders. Early theories explaining the role of attributions in the generation and maintenance of depression postulated the existence of stable, trait-like "attributional styles" (the reformulated learned helplessness theory of depression; Abramson et al., 1978) or "cognitive styles" – propensities for making specific attributions for events and inferences about their consequences and the self (the hopelessness theory of depression; Abramson et al., 1989) and a "feed-forward" causal chain model linking negative events to symptoms of depression.

Much of the research inspired by the theories of depression associated with learned helplessness (Maier and Seligman, 1976; Seligman, 1972; Seligman and Maier, 1967) measures attributions through questionnaires involving hypothetical meaningful life events, a paradigmatic example being the Attributional Style Questionnaire (ASQ (Peterson et al., 1982)). In cases where real outcomes were used, the typical experimental design involved either groups differing based on Beck Depression Inventory (BDI) scores or other measures of depression, or groups subjected to mood inducing manipulations (Allan et al. (2007); Alloy and Abramson (1982); Alloy et al. (1981)), who were asked to perform a task and then report their attribution for the outcome. In these paradigms, attributions were often measured as contingency judgements (Allan et al., 2007; Alloy and Abramson, 1979; Benassi and Mahler, 1985; Carson et al., 2010; Kaney and Bentall, 1992; Martin et al., 1984; Presson and Benassi, 2003; Vázquez, 1987; Venkatesh et al., 2018). That is, participants estimated the degree of control they exercised over the appearance of stimuli in simple lab-based tasks involving button presses and contingent or non-contingent presentation of visual stimuli. Less often, the effect of attributions was measured as participants' evaluations of their own performance, or expectations of future performance, in tasks framed as either skill or chance tasks (Campbell and Sedikides, 1999; Moore and Fresco, 2012), in an attempt to manipulate participants' attributions of the outcomes they experienced. Research involving relationships between attributional styles and beliefs about the self has found that measures of self-esteem can account for variation in attributional styles in both the general population and in psychiatric patients (Romney, 1994; Tennen and Herzberger, 1987; Tennen et al., 1987). Notably, in accordance with the underlying assumptions about trait-like attributional styles, attributions were manipulated and/or measured at the *condition level*, producing aggregate rather than trial-by-trial measures.

More recently, the dynamic nature of the processes involved was brought into focus in Bentall (2003); Bentall et al. (2001)'s attribution-self-representation cycle theory (ASRC), a model of paranoid ideation relying on interactions between attributions and beliefs about the self. Bentall et al. (2001) postulate that rather than having a trait-like attributional style, individuals use current beliefs about the self, or readily available stored knowledge about the self, along with attributional signposts in situational information (Kelley, 1967) when making causal attributions. According to this theory, the process of attribution formation involves a search for explanations that starts with explanations involving the self and terminates when a suitable explanation is found; on the other hand, once an attribution is formed, it primes representations of the self that are consistent with it. Thus, along with effects of attributions on beliefs about the self, Bentall et al. recognise the possibility of effects in the opposite direction, leading to a system with fluctuating components and potentially complex effects of interactions between them - the attribution-self-representation cycle (Bentall, 2003). The system can be influenced by relatively stable factors, such as individual differences in stored knowledge about the self, motivational biases, tendency to attend to specific types of information or ability to understand others, as well as fluctuating factors determining the relative availability of information in different circumstances.

Within this framework, the self-serving or self-enhancing attributional biases that healthy people display (Alloy and Abramson (1979); Bentall and Kaney (2005); Lyons et al. (2020); Martin et al. (1984); Tillman and Carver (1980); Vázquez (1987); for reviews, see Blaine and Crocker (1993); Campbell and Sedikides (1999); Mezulis et al. (2004)) function as homeostatic mechanisms for the maintenance of beliefs about the self within healthy parameters; the absence or disruption of these homeostatic mechanisms leaves people vulnerable to aberrant protective mechanisms or vicious cycles (where negative internal attributions lead to a worsening of self-beliefs, leading to further negative attributions), producing and or maintaining depressive or paranoid symptoms. The dynamic nature of both attributions and self-beliefs, and the fact that they exert reciprocal influences on each other can be expected to produce complex patterns of relationships between them (Adams et al., 2021). In turn, the inherent difficulty of predicting such relationships cuuld account for some of the inconclusive or contradictory findings of studies aimed at testing previous theories, whose predictive power was duly low (for reviews, see Liu et al. (2015); Robins and Hayes (1995)).

Research investigating the effect of manipulating experience on subsequent attributions produced evidence that this phenomenon is indeed more dynamic than previously considered (Bentall and Kaney, 2005; Dunning et al., 1995; Forgas et al., 1990). In a series of experiments in which participants' mood was manipulated (either through false feedback in an experimental task or through exposure to emotionally charged short films), Forgas et al. (1990) measured attributions made by healthy participants either for hypothetical situations or for their own real exam results and found that participants given a positive mood made more internal and stable attributions for positive outcomes, and less internal and stable attributions for negative outcomes, than participants given negative mood. Similarly, in a study involving clinical populations, Bentall and Kaney (2005) administered the expanded ASQ questionnaire to depressed, manic and normal participants before and after exposing them to a mild failure experience (an anagram solving task which included unsolvable items). Negative internality scores increased after the failure experience for both groups of patients (although not for the healthy controls, who might be less vulnerable to the effects of experiencing such mild failures). Conversely, situations involving potential threats to self-esteem promoted self-serving attributions (Dunning et al., 1995).

These studies showed that attribution-making tendencies can vary, *at the task timescale*, within the same population, and that factors such as mood or achievement can exert an effect on them. These results highlight the need for more precise investigations, specifically designed to uncover

the dynamics of attributions and beliefs about the self at high temporal resolution. This would allow for a quantitative test of theoretical accounts based on time-varying interactions between these variables, such as the ASRC (Bentall, 2003; Bentall et al., 2001; Bentall and Kaney, 2005). To this end, we designed and administered a novel engaging task of skill that participants learned to perform whilst also providing us with time series of both attributions and beliefs about skill. This procedure differs from its predecessors in three key ways. First, participants experienced a series of real successes and failures as they learned. Second, we collected time series of *participants' own* attributions for outcomes and beliefs about skill at a fine, trial-level, temporal granularity. Third: we did not manipulate the content of participants' attributions or skill estimates. Importantly, although participants' responses were relatively unconstrained, their experience in the task was influenced by measurable and controllable external parameters, and their performance could be be objectively quantified. Finally, we aimed to make the task engaging enough for participants to care about the outcomes and their progress, and thus report meaningful attributions and beliefs about the self.

We found evidence of trial-level reciprocal effects of attributions and beliefs which could support the dynamics postulated by the ASRC: participants updated their skill beliefs more after outcomes attributed internally than after ones attributed externally, and internal attributions for wins and losses varied with participants' previous skill estimates - they took more credit for wins and less blame for losses with increasing skill. Moreover, individual differences in these effects correlated with traditional questionnaire-based measures of self esteem, locus of control and attributional style.

The rest of this paper is organised as follows: in section 2.1 we present simulations which introduce the formalism we will use throughout and illustrate the richness of dynamics that can be supported by trial-level reciprocal attribution-beliefs interactions. We then present the task and experimental design in section 2.2. We report the results of our analyses of the real data in sections 2.3 -focused on skill reports, 2.4-focused on attributions, and 2.5- analyses of questionnaire responses. Finally, in section 3 we reflect on the task and discuss our results, directions for future work and conclusions.

2 Results

2.1 Simulations of an artificial agent

In this section we present simulations of simple artificial agents endowed with evolving beliefs about skill and a mechanism for making causal attributions. We introduce our formalisation of these variables and their interactions, and illustrate some of the behavioural patterns supported by reciprocal trial-by-trial effects between them. The point is not to provide a detailed model of the empirical investigation that follows in section 2.2 and beyond, but rather to provide a formalized underpinning for the interpretation of that experiment, and to illustrate some of the relationships that we will examine in data from human participants.

Consider an agent endowed with beliefs (s^t) about its skill in performing a task (beliefs that evolve over time t) together with the ability to make attributions (a^t) for the outcomes (o^t) it experiences. Further suppose that it aims to learn how skilled it is. The agent repeatedly performs the task, gaining probabilistic binary feedback about its outcome on every trial ($o^t \in \{0, 1\}$). We further assume the agent uses its belief about its skill to attribute each outcome

internally ($a^t = 1$) or externally ($a^t = 0$), and that it uses information from the outcome and its attribution to update its belief about its skill. This expresses the reciprocity at the heart of the ASRC.

Generation of attributions We assume that the more skilled the agent believes itself to be, the more likely it is to attribute wins internally and the less likely it is to attribute losses internally. We formalise this by letting the likelihood of internal attributions be controlled by a sigmoid function of skill for wins, and an inverse sigmoid functions for losses. Biases for internal vs external attributions are controlled by the sigmoid indifference points x_{0w}, x_{0l}, while sensitivity to skill is controlled by slope parameters β_w, β_l, for wins and losses respectively:

$$\begin{gathered} p(a^t = 1|o^t, s^t) = \begin{cases} \sigma(\beta_w * (s^t - x_{0w})) & \text{if } o^t = 1 \\ 1 - \sigma(\beta_l * (s^t - x_{0l})) & \text{if } o^t = 0 \end{cases} \quad \text{where} \\ \sigma(x) = \frac{1}{1 + e^{-x}} \text{ is the sigmoid function.} \end{gathered} \tag{1}$$

Skill belief updates Conversely, the agent uses attributions to inform the updating of its belief about skill. On every trial, the agent computes the "prediction error", the difference between its expectation of winning given its current belief about skill and the actual outcome experienced, and corrects its internal skill estimate so as to reduce this error. The magnitude of the correction is determined by a learning rate, which we assume depends on the outcome being used for updating and on its attribution, such that both the outcome valence and the attribution modulate learning:

$$\begin{gathered} s^{t+1} = s^t + \alpha_{a^t o^t} \delta_s, \text{ where} \\ \delta_s = o^t - \sigma(s^t) \end{gathered} \tag{2}$$

and $\alpha_{a^t o^t}$ is one of four parameters, corresponding to the previous outcome x attribution combination.

As noted, this simplified setting includes the potential for reciprocal effects between beliefs about skill and attributions at the trial-by-trial level, and illustrates aspects of the models we will use to fit the real data, as discussed below. To illustrate the effects these interactions can have on dynamics, we performed simulations of the agent described above, varying parameters and outcomes experienced.

Figure 1a shows the skill belief evolution of 100 simulation runs of an agent with the same parameters ($x_{0w} = 0, x_{0l} = 0, \beta_w = 2, \beta_l = 2, \alpha_{11} = 0.1, \alpha_{01} = 0.05, \alpha_{10} = 0.1, \alpha_{00} = 0.05$) and the same initial skill ($s^0 = 0$), experiencing the same fixed sequence of outcomes (drawn randomly with equal probability between wins and losses on every trial), such that random variations are only introduced via the probabilistic sampling of attribution in equation 1. The resulting variability evidently increases over trials.

The data from our human participants include just the time series of skill beliefs, trial outcomes, and attributions. We might hope to unearth evidence about the underlying relationships by examining their interactions. The trial-level effect of skill on attribution is duly illustrated in Figure 1b, which shows the proportion of internal attributions for wins and losses as a function of current skill belief quintile across all trials for all the agents shown in Figure 1a. The bias to attribute wins to the self (i.e., internally) given the belief that the skill is higher, and vice-versa for losses, is apparent. Conversely, Figure 1c illustrates the trial-level effect of attribution on skill as the average skill belief update for each outcome x attribution combination, across the same runs. As is true of the causal process, internal attributions have a greater effect on skill than external ones.

(a) Skill estimates evolution

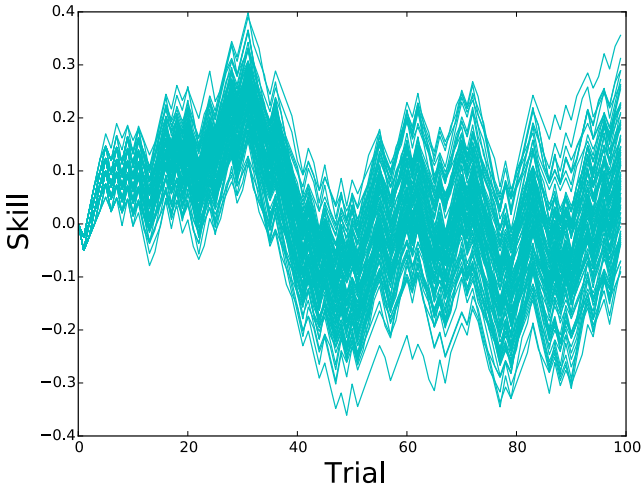

(b) Trial by trial effect skill → attribution

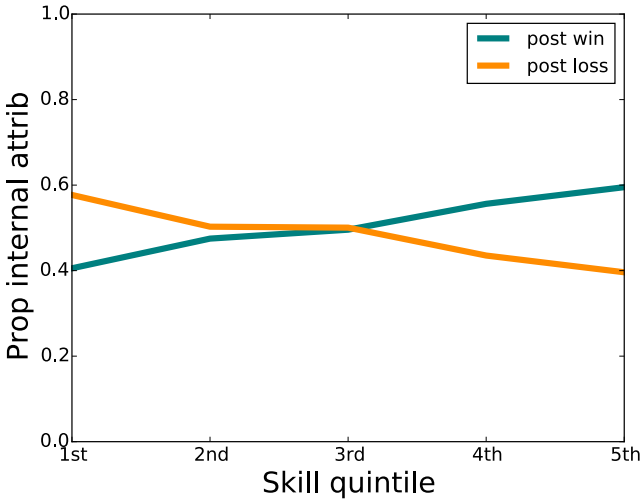

(c) Trial by trial effect attribution → skill

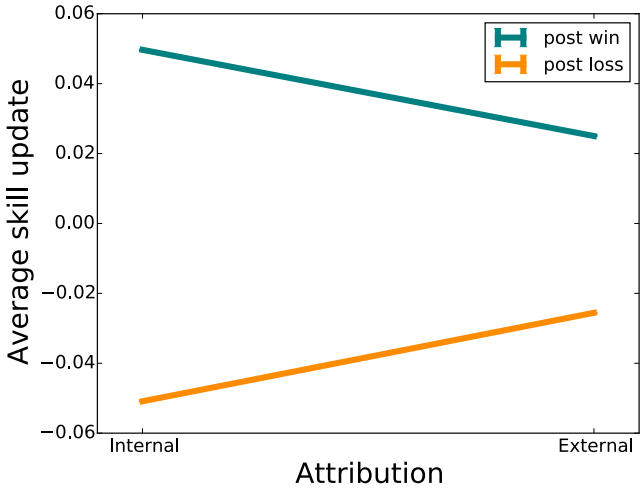

(d) Attribution → skill effect decay

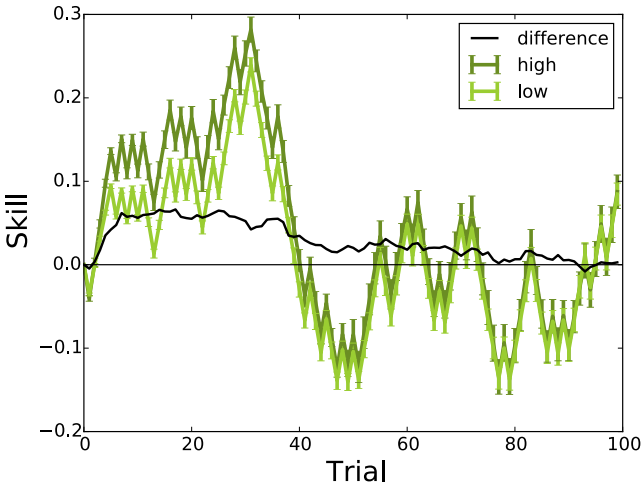

(e) Skill → attribution effect decay

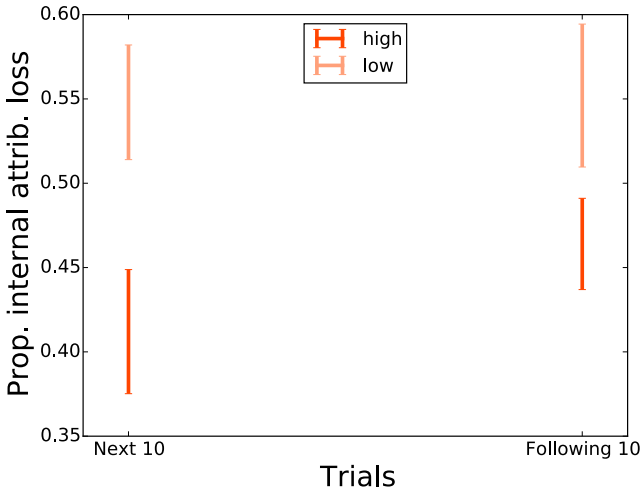

Figure 1: Reciprocal influences between skill and attributions, simulation 1. a) Skill evolution, 100 runs of agent with parameters $x_{0w} = 0, x_{0l} = 0, \beta_w = 2, \beta_l = 2, \alpha_{11} = 0.1, \alpha_{01} = 0.05, \alpha_{10} = 0.1, \alpha_{00} = 0.05$, see text for details. b) Effect of skill on attribution at the individual trial level, across all runs plotted in a). c) Effect of attribution on skill update at the individual trial level, across all runs plotted in a). d) Time decay of attribution effect. See text for details. e) Time decay of skill effect. See text for details.

Note that variability in skill beliefs is maintained and amplified despite the fact that, at least for these parameter values, differences produced by variability in early attributions decay in time, as do differences in attribution tendencies between runs with high vs low initial skill beliefs. Figure 1d shows the mean ± s.e.m of skill evolution for the simulation runs plotted in Figure 1a, separated (median split) according to the proportion of internal attributions for *the first 5 wins*. The effect decays steadily after trial 40, and is gone by trial 90; thus the divergence between runs is not produced by sensitivity to small random differences in initial attributions. Nor is it a result of high sensitivity to differences in skill beliefs: the effect of skill belief on subsequent attributions also decays in time, as illustrated in Figure 1e. Here simulation runs have been separated according to the skill value at trial 50, at which point there are already substantial differences between runs with high (dark red-top 25 % of the runs) and low skill (light red - bottom 25% of the runs). To illustrate the decaying effect of skill differences on attributions we computed, for these two sets of runs, the proportion of internal attributions for losses at two later time points (left: proportion of internal attributions for the first batch of 10 losses after trial 50, right: proportion of internal attributions for the second batch of 10 losses after trial 50) - plotted in figure 1e. Thus the sustained and even increasing divergence between the different runs is not produced in this case merely by sensitivity to initial conditions, or by a positive feedback loop. Nor is it due to randomness in attributions alone: divergence in skill beliefs is reduced in the absence of a reciprocal effect of skill on attributions: this is illustrated in figure 2a and b, which shows (yellow) the evolution of skill beliefs for 100 agents in which attribution is matched for average propensity with the full agents (cyan), but is independent of skill (all other parameters being equal). Instead, it is the sustained mutual interactions between attributions and skill beliefs that gradually amplify the divergence between runs introduced by randomness in attributions.

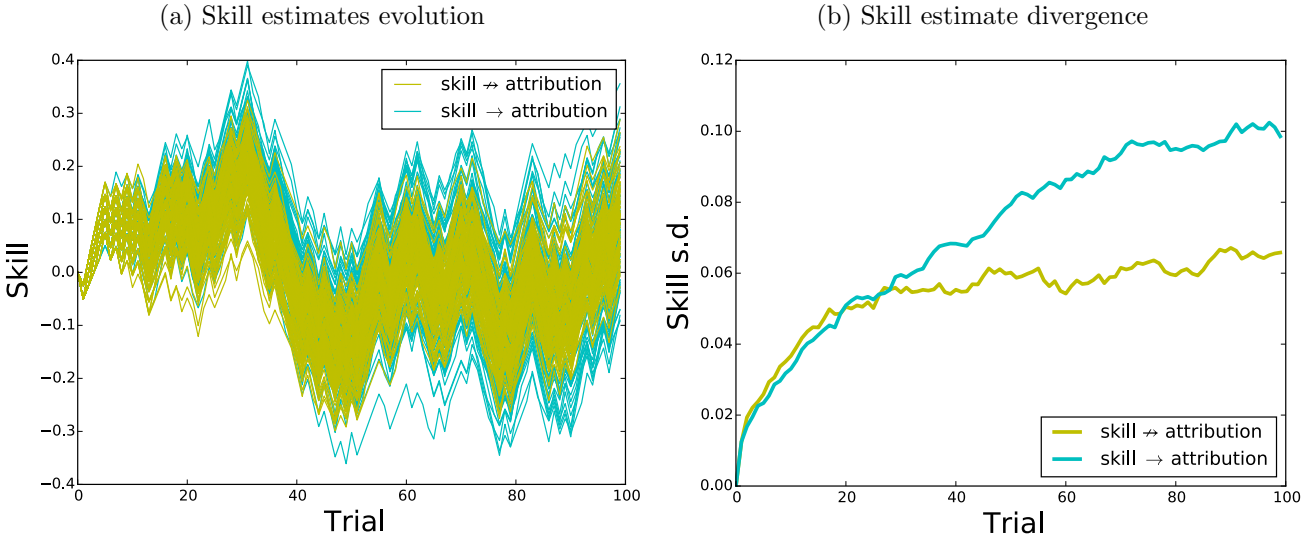

Figure 2: Removing reciprocal connections in simulation 1. a) Skill evolution, runs with (cyan, same as in Figure 1a) and without (yellow) skill effect on attribution, see text for details. b) S. d. of runs in a.

For other parameter settings, reciprocal dependencies between skill beliefs and attributions can produce positive feedback loops that dramatically amplify small differences between simulation runs. This is illustrated in Figure 3, which shows 300 simulation runs of an agent starting from the same initial skill ($s^0 = 0$) and experiencing the same fixed sequence of outcomes as the runs in Figure 1, but with a different set of parameters ($x_{0w} = -0.03, x_{0l} = 0, \beta_w = 100, \beta_l = 1, \alpha_{11} = 0.15, \alpha_{01} = 0, \alpha_{10} = 0.15, \alpha_{00} = 0$). Due to a combination of high sensitivity to skill

(a) Skill estimates evolution

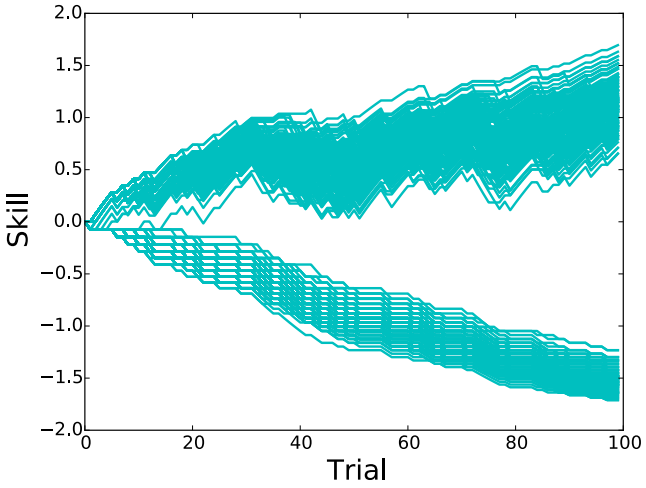

(b) Trial by trial effect skill → attribution

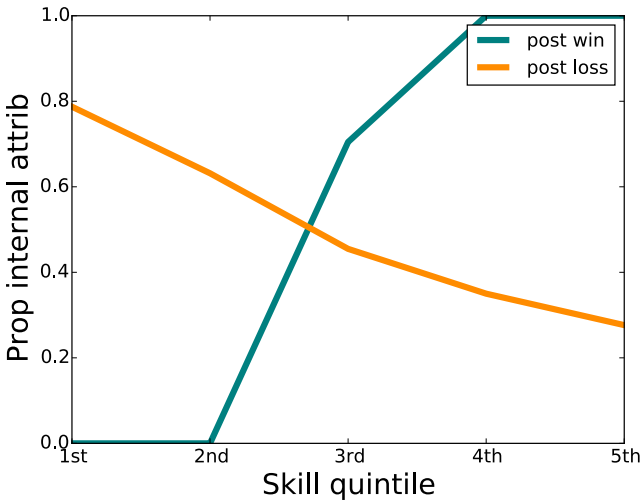

(c) Trial by trial effect attribution → skill

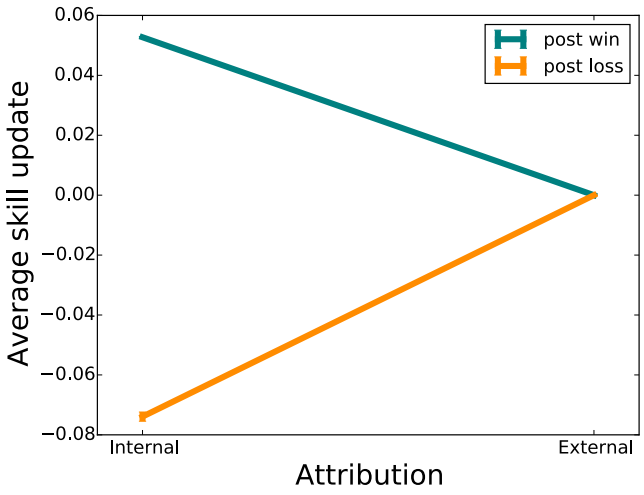

(d) Attribution → skill effect in time

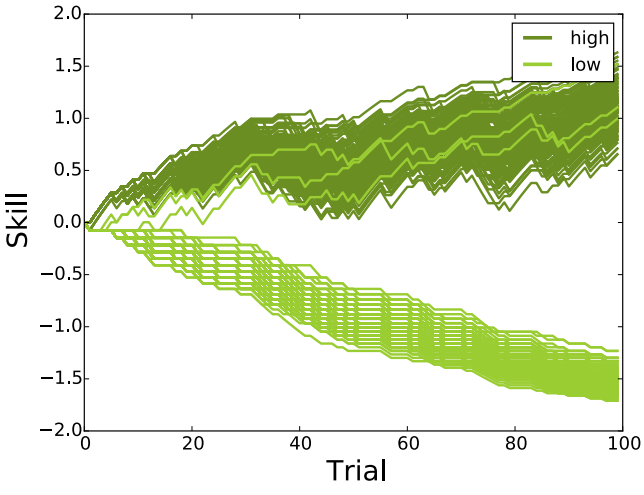

(e) Skill → attribution effect in time

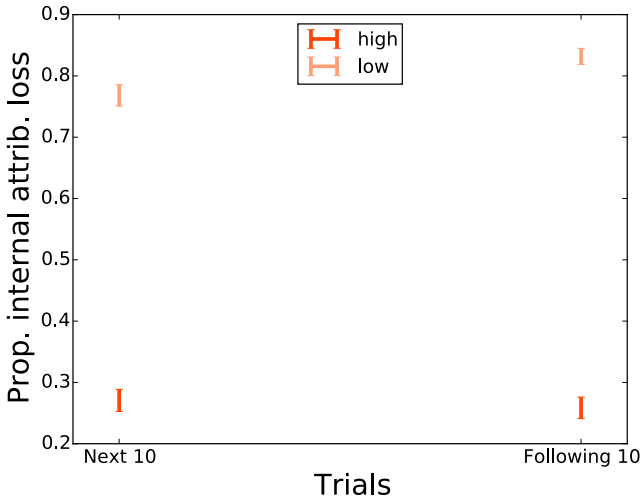

Figure 3: Reciprocal influences between skill and attributions, simulation 2. a) Skill evolution, 300 agents with parameters $x_{0w} = -0.03, x_{0l} = 0, \beta_w = 100, \beta_l = 1, \alpha_{11} = 0.15, \alpha_{01} = 0, \alpha_{10} = 0.15, \alpha_{00} = 0$, see text for details. b) Effect of skill on attribution at the individual trial level, across all runs plotted in a). c) Effect of attribution on skill update at the individual trial level, across all runs plotted in a). d) Amplification of attribution effect. See text for details. e)Time evolution of skill effect. See text for details.

in attribution-making (Figure 3b) and large differences between learning from internally and externally attributed outcomes (Figure 3c) differences produced by variability in early attributions are amplified, rather than dampened, in time: Figure 3d shows the skill evolution for the simulation runs in Figure 3a, separated (median split) according to the proportion of internal attributions for *the first 5 wins*. The effect of skill on attribution is also propagated in time, rather than decaying as in the previous example: Figure 3e reproduces Figure 1e, showing the proportion of internal attributions for two successive batches of 10 losses after trial 50, for the runs with the top 25% (dark red) and bottom 25% (light red) values of skill belief at trial 50.

(a) Skill estimates evolution

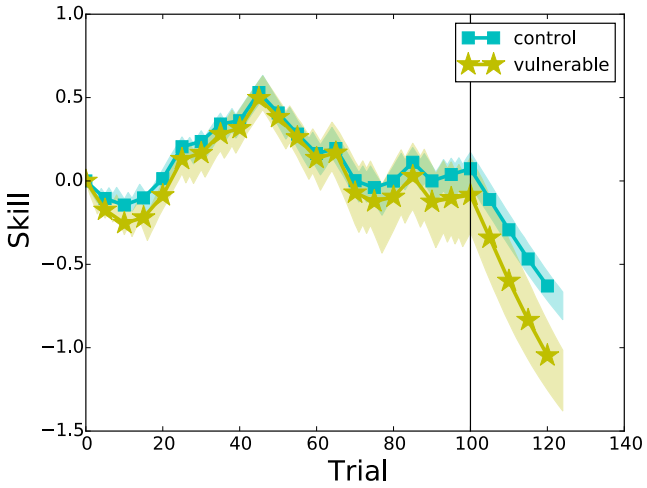

(b) Trial by trial effect skill → attribution

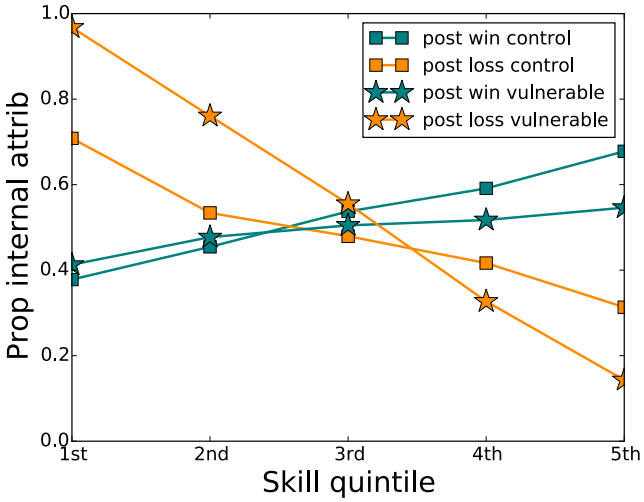

(c) Trial by trial effect attribution → skill

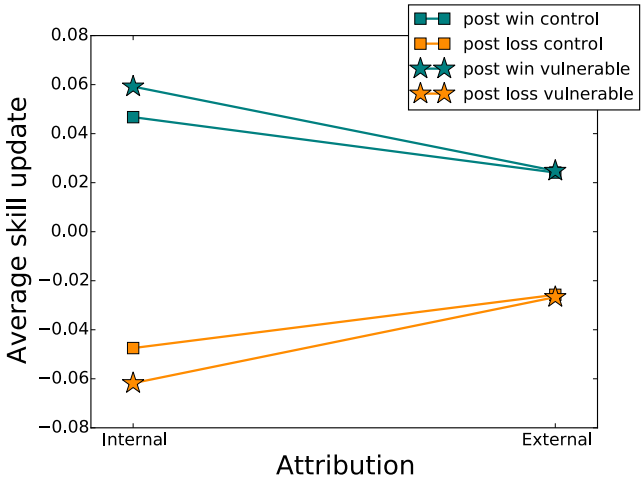

Figure 4: Latent vulnerability, simulation 3. a) Skill evolution for control (cyan, squares, $x_{0w} = 0, x_{0l} = 0, \beta_w = 2, \beta_l = 2, \alpha_{11} = 0.1, \alpha_{01} = 0.05, \alpha_{10} = 0.1, \alpha_{00} = 0.05$) and vulnerable (yellow, stars, $\beta_w = 0.5, \beta_l = 5, \alpha_{11} = 0.12, \alpha_{10} = 0.15$, all other parameters equal to control) agents, see text for details. b) Effect of skill on attribution at the individual trial level. c) Effect of attribution on skill update at the individual trial level.

Trial-level reciprocal effects also enable settings of the parameters which endow agents with "latent" vulnerability to negative outcomes - vulnerabilities that are not apparent in relatively benign conditions, when agents experience balanced positive and negative outcomes, but become apparent when they hit a streak of negative outcomes (Bentall, 2003; Bentall et al., 2001; Bentall and Kaney, 2005). This is illustrated in Figure 4, which shows 100 runs of "control" agents ($x_{0w} = 0, x_{0l} = 0, \beta_w = 2, \beta_l = 2, \alpha_{11} = 0.1, \alpha_{01} = 0.05, \alpha_{10} = 0.1, \alpha_{00} = 0.05$) and 100 runs

of "vulnerable" agents ($\beta_w = 0.5, \beta_l = 5, \alpha_{11} = 0.12, \alpha_{10} = 0.15$, all other parameters equal to "control") starting from the same initial skill belief value ($s^0 = 0$) and experiencing the same fixed sequence of outcomes as in Figure 1, with an additional "negative streak" of 25 losses at the end. In this particular case, compared to "control" agents, vulnerable agents are less sensitive to skill when making attributions for wins and more sensitive to skill when making attributions for losses (Figure 4b), as well as more sensitive to attribution when updating their skill beliefs (Figure 4c). This leads to vulnerable agents' beliefs growing increasingly negative with respect to their healthy counterparts during the negative streak period. This cartoon model of depression could be expanded to include agents' willingness to engage in tasks which could yield additional (positive and negative) information about their abilities (Dayan et al., 2020). If this additional component was related to the agents' current beliefs - a plausibly realistic scenario- it could further amplify differences between "healthy" and "vulnerable" agents, by impairing the latter's ability to recover from negative events even more: agents with low skill beliefs would have reduced motivation to engage in potentially rewarding activities (activities yielding positive information about the self), and could therefore get stuck in low belief states.

These simulations illustrate some of the complexity in belief dynamics which can be produced by reciprocal connections between belief updating and outcome attributions at the individual trial level. Disentangling the contributions of various parameters and or of the history of previous experience in such systems is challenging, particularly in real data, where both functional relationships between variables and parameter values need to be inferred, and history is only partially known. However such simulations can be used to provide a roadmap for dissecting real life systems and placing their study on a firm footing. In our case, the first step is establishing whether reciprocal effects between attributions and beliefs at the individual trial level can at all be detected in real life data. This is the focus of the rest of this paper.

2.2 Experimental design

In order to investigate whether reciprocal trial-level effects between attributions and beliefs can be detected in human participants' behaviour as postulated by the ASRC, we exposed participants in our experiment to a situation similar to the one encountered by the simulated agents presented above: they had to repeatedly perform a task, evaluate their outcomes and learn about their skill.

The task we used is a remunerated game of skill inspired by the "Penguin Pursuit" game on the Lumosity "brain training" platform `https://www.lumosity.com/en/`. Participants are presented, on every trial, with a maze (Figure 5a) through which they need to move a token red square from the starting position to the finish position, marked by a trophy; they do so by using the arrow keys to control the token. If the token reaches the goal in the limited time allocated, a win is indicated by the appearance of a smiley face; otherwise a frowning one appears to indicate loss. During each trial, repeatedly and unpredictably, the maze rotates and the correspondence between arrow keys and the direction of movement on the screen changes, according to the following rules: the maze is equipped with a "North" direction, marked on the screen by a compass needle; the arrow keys always move the token toward the corresponding cardinal directions of the maze (Up towards "North", down towards "South", right towards "East", left towards "West"). Initially, the maze's "North" corresponds to the top of the screen; however when the maze rotates, the compass needle rotates with it, such that "North" can point to any of the four directions on the screen, and accordingly pressing the up key no longer moves the token up on the screen, but towards the maze's "North", wherever that is during each rotation. Participants

(a) Task

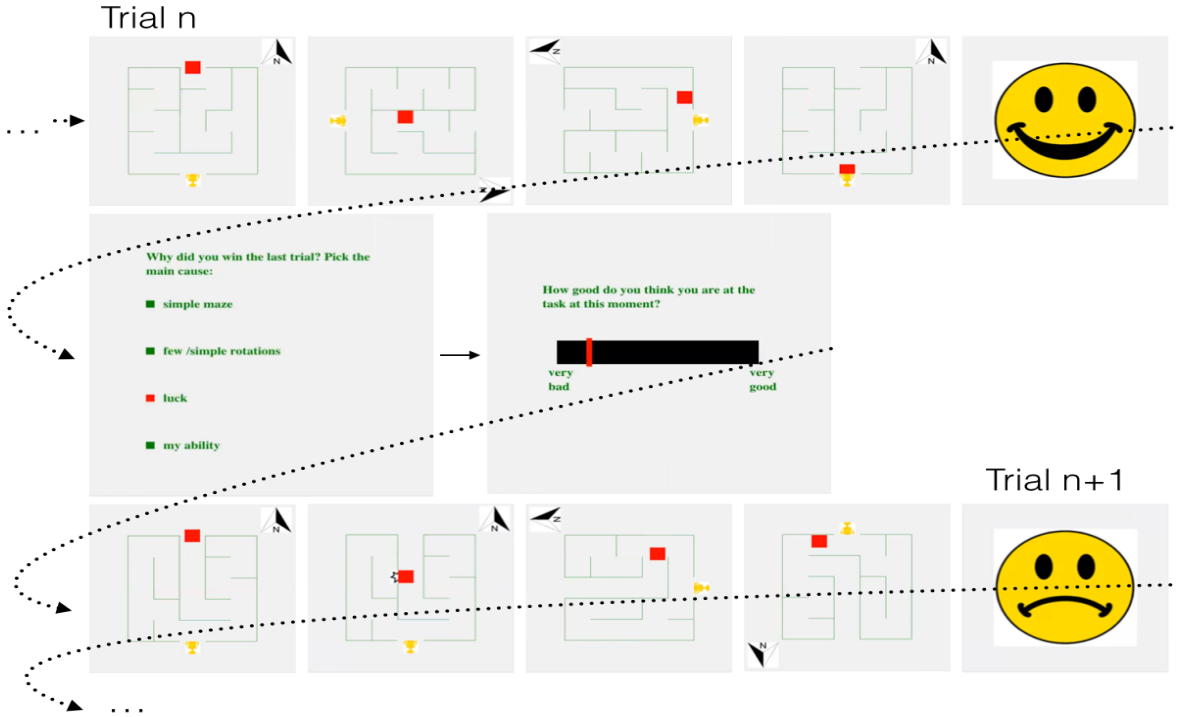

(b) Performance overview

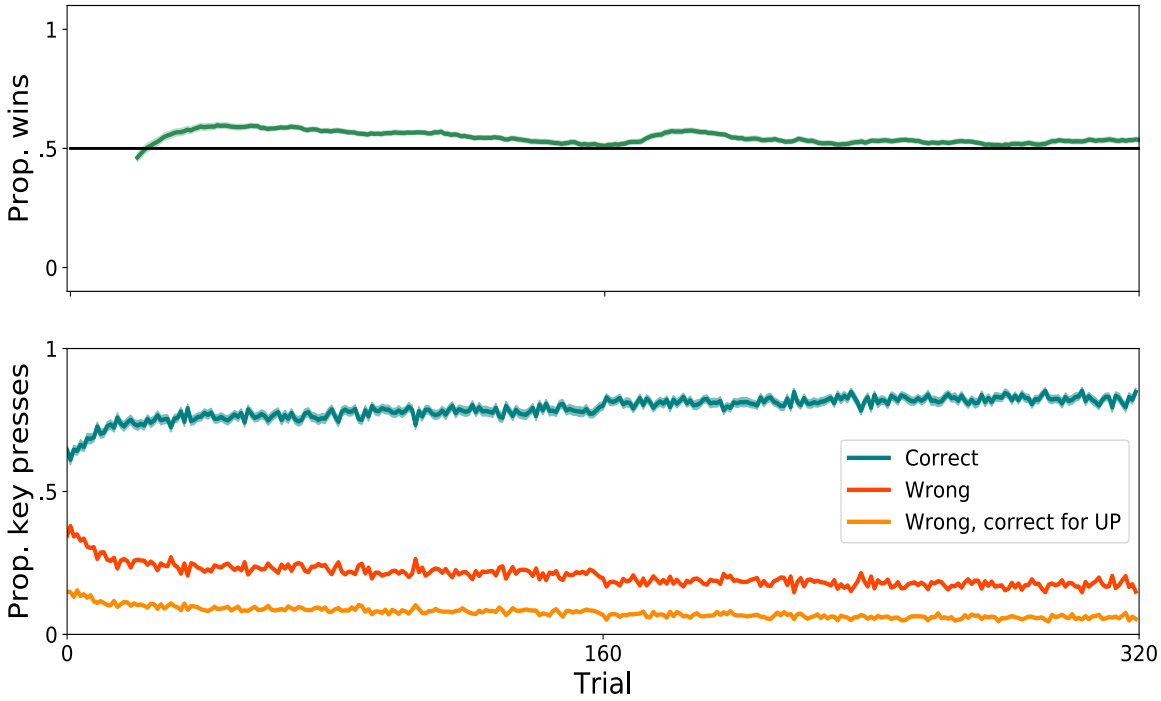

Figure 5: a)Task structure and example trials: four frames sampled for illustration purposes from two trials are displayed in order on the top and bottom rows. After every two trials, participants are asked to attribute the latest outcome to one of 4 given causes, and then to report how good they believe themselves to be at the task. Dotted arrow indicate the flow of time (see text for details). b) Evolution of performance across trials, mean ± s.e.m. across participants; top: running average of the proportion of wins, sliding window 20 trials, bottom: per trial proportion of correct key presses, wrong key presses and wrong, but correct for the normal UP orientation.

therefore have to learn to adapt quickly to the change in the correspondence between key presses and resulting movements on the screen.

Trial difficulty is determined by the size of the maze, the frequency of rotations per trial and the time available. These parameters are adapted to maintain performance at roughly 50% (Figure 5b), using a double staircase procedure in order to prevent participants from gaming this adaptation. See Supplementary Materials 1 for a detailed description.

(a) Skill estimates evolution: summary

(b) Skill estimate evolution: example participants

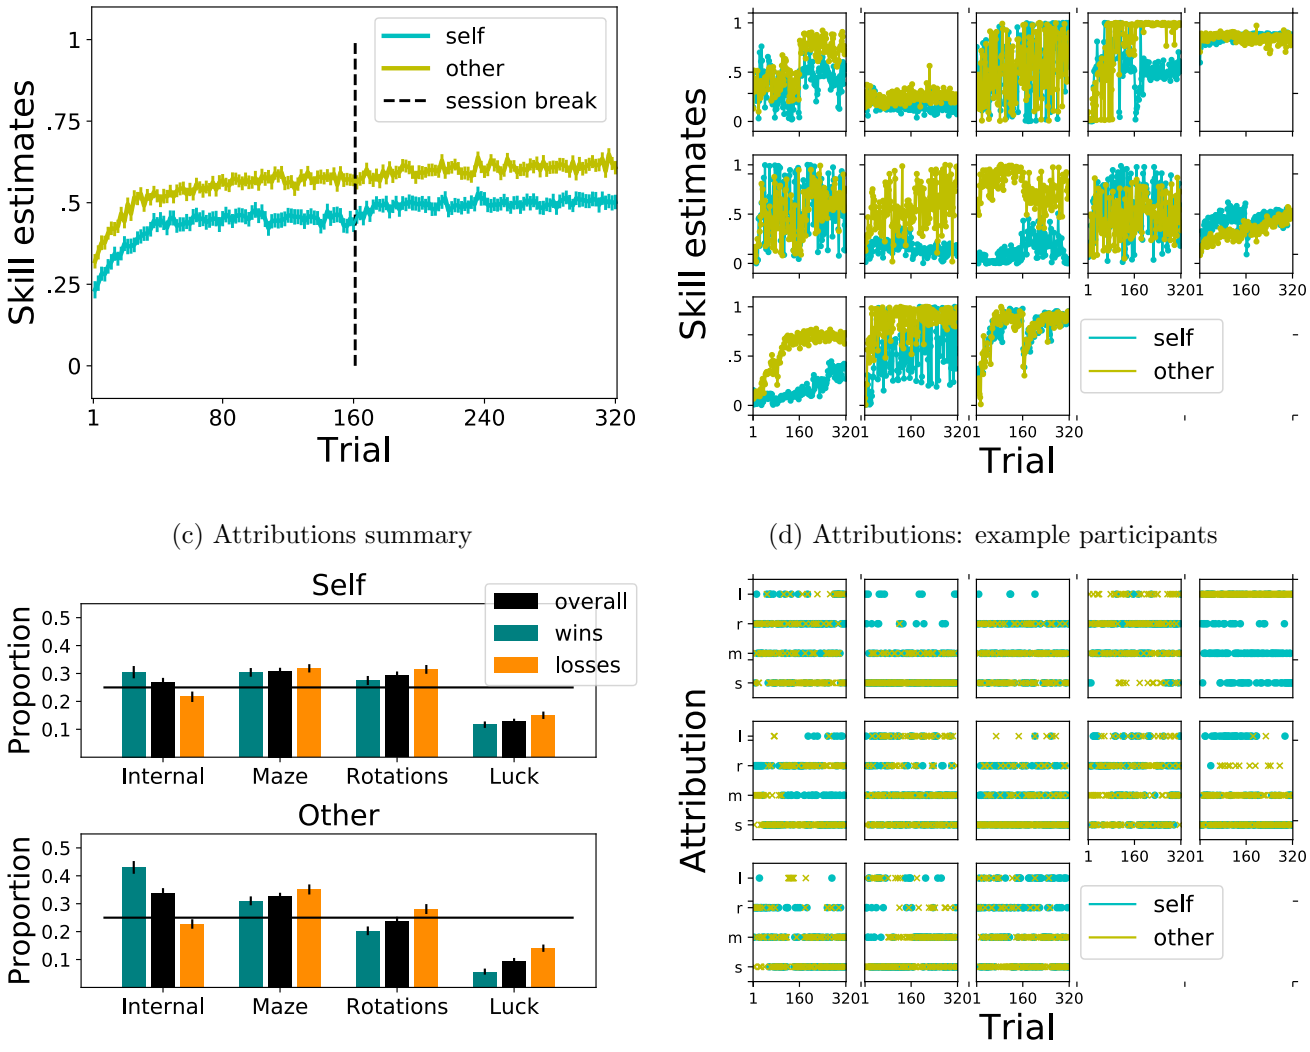

(c) Attributions summary

(d) Attributions: example participants

Figure 6: Skill estimates and attributions overview. a) Evolution of skill estimates across trials, mean ± s. e. m. across participants. b) Evolution of skill estimates for individual participants, chosen to illustrate variability. c) Attribution proportions, mean ± s.e.m across participants, overall and conditioned on outcomes. d) Time series of attributions for individual participants, chosen to illustrate variability, not the same as in b.

Every two trials, immediately after seeing the trial outcome, participants are asked to provide a causal attribution for the outcome and then to estimate their current skill in playing the game. Participants are informed that a total time of 20s is available to answer each question, and that reward for the following trial would be withheld if they fail to provide an answer in this time limit. Attributions are elicited with multiple-choice questions with "ability", "maze", "rotations", "luck" as the options (see Methods section 4.3 for a detailed description). These

options reflect the internal vs external aspect of attribution, with the latter being separated into two different quantifiable parameters of the task (maze complexity vs rotations), and the option of blaming or crediting luck. Skill estimates are reported with a slider on a continuous scale from "very bad" to "very good" (Figure 5a).

The experiment involves two conditions: first participants play the game themselves and answer questions related to their performance; we refer to this as the self condition. Participants are then told that they will watch trials recorded from another participant having previously participated in the experiment and they are asked to make attributions, evaluate skill and bet on this "other"'s performance; this is the other condition. In fact, each participant is shown their own recorded trials from the self condition, to control for the precise sequence of difficulties and outcomes encountered. In order to reduce the likelihood of participants recognising their own performance, mazes are left-right mirrored when played back. Due to the length of the task, each condition was split into two sessions, performed on two successive days; there was a 7 day delay between conditions.

Three questionnaires were administered immediately prior to the first session: the Attributional Style questionnaire (ASQ) (Peterson et al., 1982), the Levenson Locus of Control questionnaire (Levenson, 1974) and the Rosenberg Self esteem scale (Rosenberg, 1965).

Figure 6 shows a summary of participants' skill estimates and attributions (a, c), as well as full time series of responses for a selection of individual participants (b, d), illustrating individual differences.

2.3 Analyses of skill estimates

We performed both model-agnostic and model-dependent analyses in order to test whether participants' attributions had detectable effects on their reported beliefs about skill; we hypothesised that outcomes attributed internally would have a larger impact on beliefs than outcomes attributed externally.

Model agnostic analyses As expected, we found an effect of outcome on skill updates, with significantly larger updates after wins than after losses for both self (paired $t = 14.91$, two-sided $p < 1/5000, d = 2.65$, see Methods section 4.4 for details) and other (paired $t = 15.24$, two-sided $p < 1/5000, d = 3.1$), Figure 7. We therefore tested for an effect of attribution (internal vs external) on skill updates conditioned on outcome and found significant effects of attributions on both wins (self: paired $t = 4.14$, two-sided $p < 1/5000, d = 0.52$; other: paired $t = 7.09$, two-sided $p < 1/5000, d = 1.01$) and losses (self: paired $t = -3.47$, two-sided $p = 0.0008, d = -0.43$; other: paired $t = -4.06$, two-sided $p < 1/5000, d = -0.62$) in the expected directions: participants' reported skill estimates increased more after wins attributed internally than after wins attributed externally, and decreased more after losses attributed internally than after losses attributed externally (Figure 7a).

(a) Outcome and attribution effect on skill updates

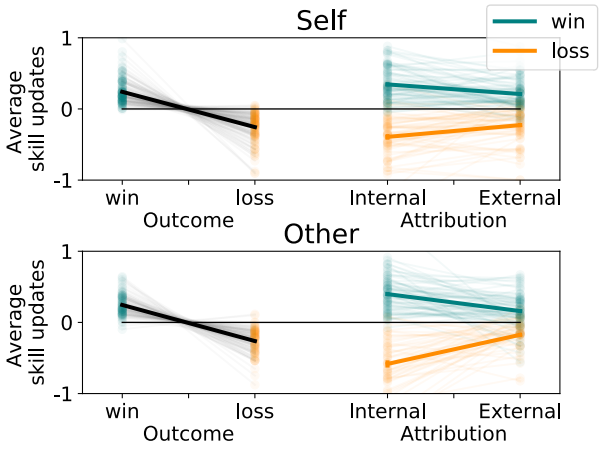

(b) Winning skill model parameters self

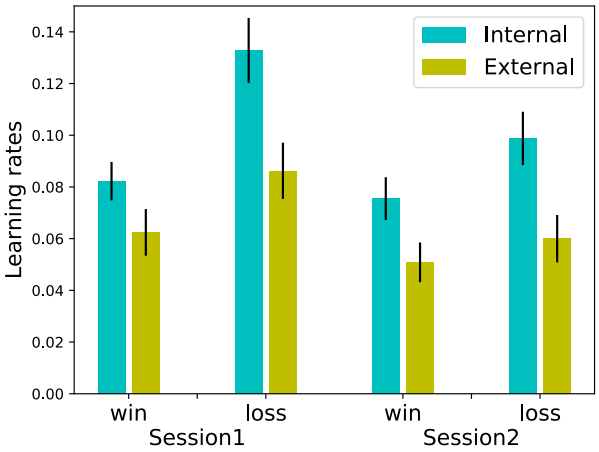

(c) Skill estimates model comparison

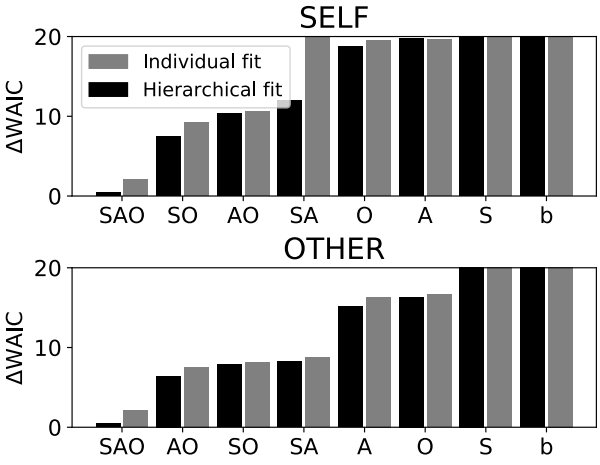

(d) Winning skill model parameters other

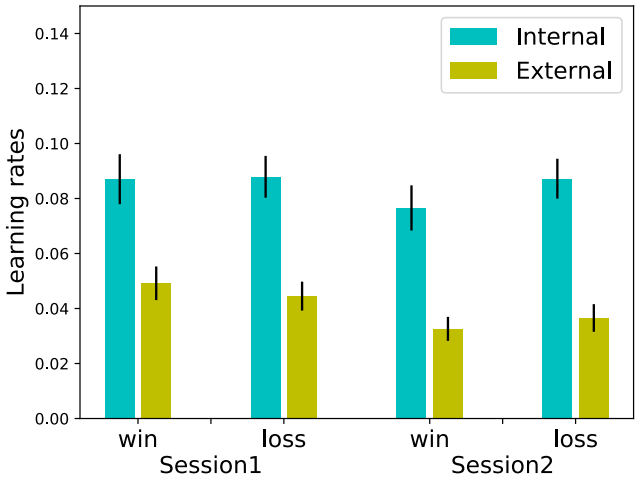

Figure 7: Effect of outcome and attribution on skill updates. a) Model agnostic analyses: faded lines and dots represent individual participants, bold lines represent mean± s.e.m across participants. b), d) Learning rates, winning model for skill. b): self, d): other. Cyan: learning rates for internal attribution, yellow: learning rates for external attributions. c) Skill estimates model comparison. Top: self, bottom: other. Difference in WAIC scores from each model to the preferred one. Smaller WAIC scores indicate better models.

Model comparison We next turned to trial-by-trial models of participant's skill estimates. We compared variants of Rescorla-Wagner models, built on the assumption that participants maintain an internal estimate of skill, which they adjust as a result of the outcomes they experience (in the self condition)/they observe (in the other condition). Thus on every trial the prediction error – the difference between the outcome and the current estimate of skill – is used to update the underlying skill estimate. The factor weighing the prediction error's contribution – the learning rate – is what differs between models (see complete model description in Methods section 4.5). Due to our main interest in the effect of attributions on skill beliefs, one change to the baseline model was to allow the learning rate to be different for different attributions (model A). Because participants' responses were not constrained, attribution and outcome could not be entirely orthogonalised in the data, and participants showed a preference for internal attributions for wins and external attributions for losses (see 2.4); therefore to control for any effect of outcome disguising as an effect of attribution in model A, we also allowed learning rates to be different for wins vs losses (model O). Finally, because each condition was divided between

two successive days (see 2.2), we introduced an additional variation to the baseline model by allowing learning rates to differ between the two sessions (S). We therefore compared a family of models in which learning rates were allowed to vary orthogonally along three directions- outcome, attribution, session- leading to a total of 8 models. See Methods section 4.5 for full details on the model fitting and model comparison procedures.

For both self and other, model comparison favoured a hierarchical version of the full model (self: WAIC = -92.15, ΔWAIC = -6.99 to the next best model, other: WAIC = -104.86, ΔWAIC = -5.95 to the next best model, Figure 7c), indicating that allowing learning rates to differ for internal vs external attributions improved the model evidence, despite a higher penalty for the increased number of parameters. The best model was well able to fit the data (self: $r^2 = 0.44 \pm 0.24$ across participants, worst $r^2 = 0.01$, best $r^2 = 0.93$; other: $r^2 = 0.42 \pm 0.24$, worst $r^2 = 0.01$, best $r^2 = 0.95$; see Supplementary Figure 3 for example best and worst participant fits).

Model parameters Results of model-agnostic analyses indicated that outcomes attributed internally had a stronger impact on participants' subsequent skill estimates than outcomes attributed externally. Analyses of parameters from the best skill model confirmed this result (Figure 7b;d): for both self and other we found higher learning rates for internal than for external attributions (self: paired $t = 8.63$ one-sided $p < 1/5000$, other: paired $t = 12.08$, one-sided $p < 1/5000$, p-values computed from permutation tests). In fact, the difference between learning rates for internal and external attributions was significant for each outcome x session combination (see Supplementary Materials 4 for full results). This result was also confirmed by analyses of fits with shuffled attribution responses (see Supplementary Figure 4 for details and results).

In addition, learning rates were lower for wins than for losses for self (paired $t = -6.34$, two-sided $p < 1/5000$ from permutation test), but not for other (paired $t = 0.82$, two-sided $p = 0.41$); this effect of outcome in the self condition was stronger for internal than for external attributions (paired $t = -2.65$, two-sided $p = 0.008$).

Thus both model-agnostic and model dependent analyses identified significant effects of attribution on skill estimates, outcomes attributed internally having a stronger effect on beliefs than those attributed externally. In addition, model-based analyses indicated stronger learning from losses than from wins in the self condition, particularly in the case of internally attributed outcomes.

2.4 Analyses of attributions

We next performed both model-agnostic and model-dependent analyses of attributions. Our main question was whether participants' reported beliefs about skill had detectable effects on their attributions; our hypothesis was that participants would make more internal attributions for wins and less internal attributions for losses when believing they possess high rather than low skill. We also investigated the effects of objective task features and performance, which we expected to be reflected in attributions to the relevant response options provided; specifically, we expected to find effects of path length on attributions to maze, and effects of the frequency of rotations on attributions to rotations. All statistical tests were performed using permutations to check for significance, as detailed in the methods section.

(a) Effects of skill estimates, task features and objective performance on attributions for self

(b) Attributions model comparison

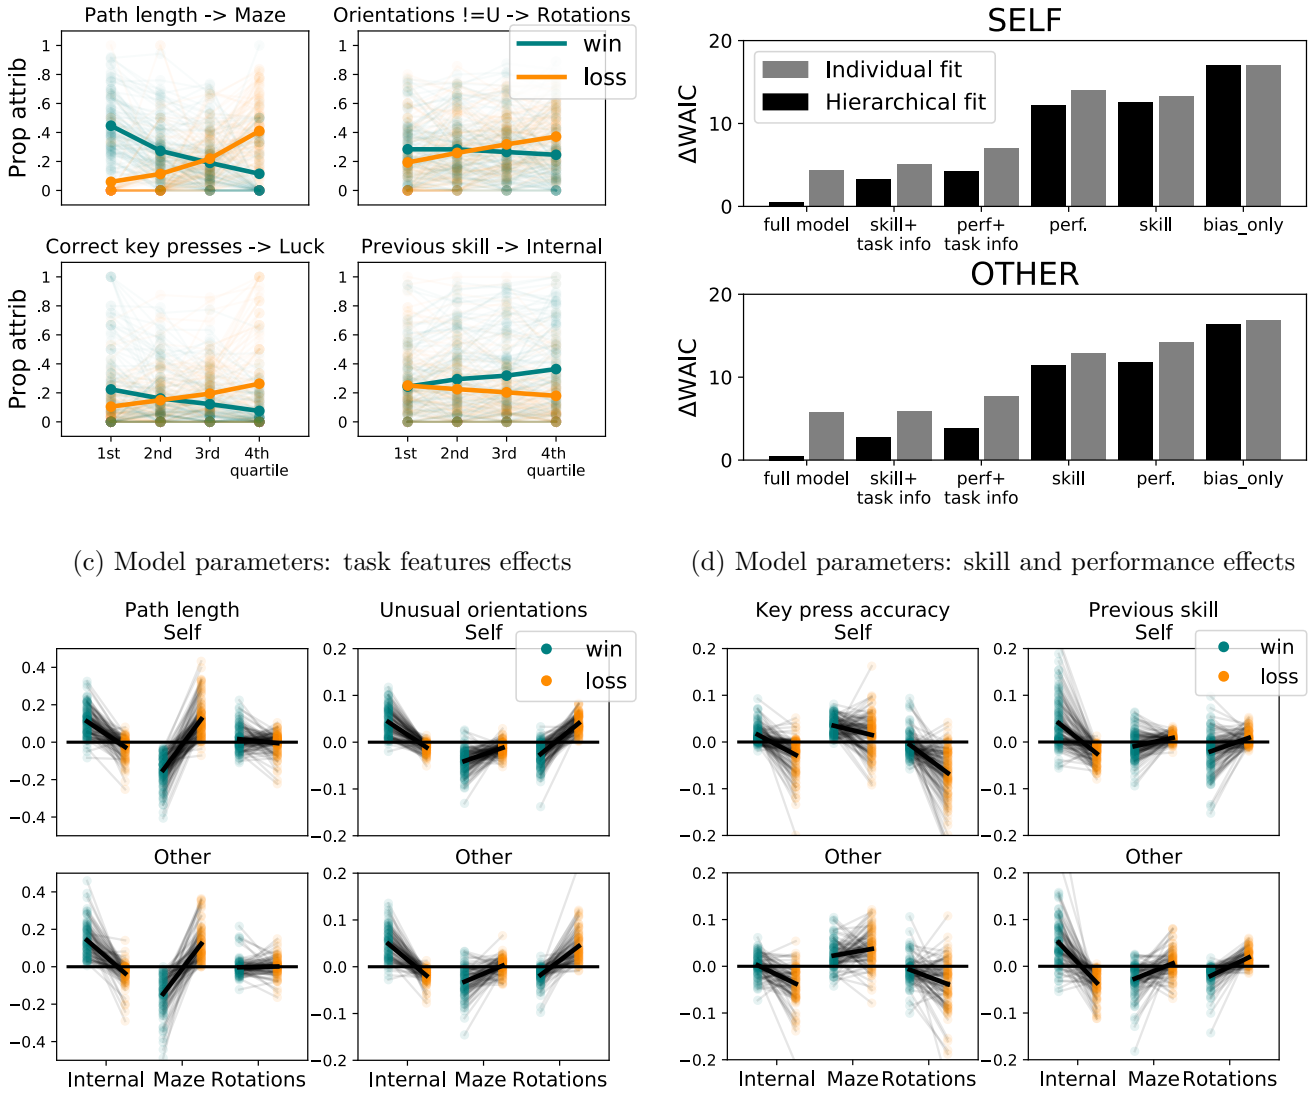

(c) Model parameters: task features effects

(d) Model parameters: skill and performance effects

Figure 8: Attribution analyses: a) Features of interest and attributions summary self: faded lines represents individual participants, bold lines represent mean ± s.e.m across participants. Orange: losses, teal: wins. b) Attribution model comparison. Top: self, bottom: other. Difference in WAIC scores from each model to the preferred one. Smaller WAIC scores indicate better models. c) Winning attribution model parameters. Effects of path length (left) and proportion of unusual orientations (right). Top: self, bottom: other. d)Effects of key press accuracy (left) and reported skill (right). Top: self, bottom: other.

Model-agnostic analyses Consistent with previous results (Campbell and Sedikides, 1999; Mezulis et al., 2004) we found that overall participants made more internal attributions for wins than for losses in both conditions (self: paired $t = 3.1$, two-sided $p = 0.0036, d = 0.54$; other: paired $t = 6.69$, two-sided $p < 10^{-5}, d = 1.07$, p-value computed from permutation tests, see Methods section 4.4 for details),with a significantly larger effect for other than for self (paired $t = 3.77$, two-sided $p = 0.0004, d = 0.42$). Analyses of the effects of objective task and performance measures on participants' attributions to the relevant attribution options revealed that participants were sensitive to the task manipulations, as well as to their own performance. Participants made more attributions to the maze for losses and fewer for wins as the length of the correct path through the maze increased (significant interactions of path length with

outcomes on attributions to maze: self: $F = 190.03, p < 1/5000$, other: $F = 127.12, p < 1/5000$, Figure 8 a for self and Supplementary Figure 5 for other; F statistics computed as for two-way repeated measures ANOVA (Howell, 2012), with p-values estimated from permutation tests, see Methods section 4.4). Similarly, the more the maze was displayed in an unusual orientation during the trial, the more participants blamed rotations for losses and the less they credited rotations for wins (significant interactions of prop non up orientations with outcomes on attributions to rotations: self: $F = 37.71, p < 1/5000$, other: $F = 22.02, p < 1/5000$, Figure 8a for self and Supplementary Figure 5 for other). Participants' attributions were also sensitive to objective performance: participants increased their assignment of losses to bad luck and decreased crediting wins to good luck with increasing accuracy in their key presses (significant interactions of prop correct key presses with outcomes on attributions to luck: self: $F = 53.23, p < 1/5000$, other: $F = 9.12, p = 0.0006$, Figure 8a for self and Supplementary Figure 5 for other). These results confirm the effectiveness of task manipulations, and show that participants understood the attributions options provided, using them reasonably.

Analyses of the effect of skill on internal attributions confirmed our hypotheses, identifying significant interactions between outcome and previous skill estimate (self: $F = 33.92, p < 1/5000$, other: $F = 29.51, p < 1/5000$, Figure 8a for self and Supplementary Figure 5 for other), with more internal attributions for wins and less internal attributions for losses with increasing skill estimates.

Model comparison In model-agnostic analyses we considered the effect of several variables on attributions *separately*. We next turned to trial-by-trial models of participants' attributions in order to investigate the effects we observed in model-agnostic analyses while controlling for the other factors. Here as well, we were mainly interested in the effect of reported skill estimates on attributions.

We used linear classification models, in which the contributions of the features of interest are first linearly combined into a score for each response option $s_o = \boldsymbol{w}_o \cdot \boldsymbol{f}$, with option-specific weights $\boldsymbol{w}_o$ as parameters, and then the resulting scores are passed through a softmax function to obtain response probabilities for each option $p(o) = \frac{\exp(s_o)}{\Sigma_{o \in O} \exp(s_o)}$ (see Methods section 4.6 for full model specifications). Because at least some of the features are expected to have different, potentially opposite effects on attributions for wins vs losses, as indeed it emerged from model-agnostic analyses of the data, all models were equipped with separate parameters for wins and losses. Because participants have to choose one of the 4 available response options, models were equipped with independent parameters for *three* of the options, and scores for the 4 options were constrained to sum to 0; thus parameters and preference for the "Luck" option were entirely derived from parameters and preferences for the other options, and are not analysed nor reported separately below. The models we compared varied in that they included different potential features: no features – baseline propensity for each attribution option; models including bias and one of the following set of features: reported skill, performance features, performance features and task features, reported skill and task features; and the full model: reported skill, performance features, task features. See Methods section 4.6 for a detailed account of the models, model fitting and model comparison procedures.

For both self and other, model comparison favoured the full model in its hierarchical version (self: WAIC = 155.21, ΔWAIC = -2.81 to the next best model, other: WAIC = 145.38, ΔWAIC = -2.29 to the next best model, see figure 8b), indicating that including reported skill as a feature over and above objective performance and task features improved the model evidence, despite a higher penalty for the increased number of parameters. The winning model was able to predict participants' responses above chance (chance level = 0.25, median probability of true answer

self: 0.48 ± 0.14, ranging from: 0.27 to 0.93; other 0.53 ± 1.64, ranging from 0.28 to 0.99).

Model parameters We then performed analyses of model parameters from the winning attribution model. These fall into two categories: baseline preferences for different attribution options, as captured by bias parameters, and the effects of various features, as captured by the corresponding feature weights in the model (in both cases, direct comparisons between raw parameters are not meaningful, due to the presence of different sigmoid functions for wins and losses; all effects reported were obtained by transformations of the raw parameters to correct for this, as detailed in Methods section 4.6). These confirmed the results of model-agnostic analyses, but also revealed additional insights, as detailed below.

Consistent with model agnostic results, baseline preference for internal attributions was significantly higher for wins than for losses in both conditions (self: paired $t = 4.77$, two-sided $p < 1/5000, d = 0.61$, other: paired $t = 10.34$, two-sided $p < 1/5000, d = 1.36$), with a significantly stronger effect for other than for self (other vs self paired $t = 4.65$, two-sided $p < 1/5000, d = 0.58$).

Recapitulating the results of model-agnostic analyses we found that for both self and other increasing path length significantly reduced the likelihood of attributing wins to the maze and significantly increased the likelihood of attributing losses to the maze, with the opposite pattern for internal attributions (Figure 8 c, results in Supplementary Table 2). Similarly, increasing proportion of unusual orientations of the maze during trial significantly reduced the likelihood of attributing wins to rotations and increased the likelihood of attributing losses to rotations,with the opposite pattern for internal attributions (Figure 8c, results in Supplementary Table 3).

The effects of reported skill on internal attributions were also in the expected directions, increasing the likelihood of attributing wins internally (self: $m = 0.04$, two-sided $p < 1/5000$ from sign permutation test; other $m = 0.05$, two-sided $p < 1/5000$) and decreasing the likelihood of attributing losses internally (self: $m = -0.03$, two-sided $p < 1/5000$; other $m = -0.04$, two-sided $p < 1/5000$; Figure 8d). Reported skill had qualitatively similar effects on the two external options: increasing skill increased the likelihood of blaming maze and rotations for losses and decreased the likelihood of crediting them for wins (Figure 8d, results in Supplementary Table 4). This is different from the effects of momentary performance: increasing trial-level key press accuracy increased the likelihood of attributions to maze and decreased the likelihood of attributions to rotations irrespective of outcome (Figure 8d, results in Supplementary Table 5).

Thus both model-agnostic and model dependent analyses identified skill reports as having significant effects on participants' attributions. In addition, model-dependent analyses revealed differences between reported skill and key press accuracy in the pattern of their results on external attribution options; these are consistent with the distinction between a momentary and specific performance measure on the one hand, and a more stable and general measure of ability on the other.

2.5 Relationships with questionnaire scores

In order to test whether our task taps into the attribution and belief dimensions typically investigated with well-established questionnaire measures we compared participants' questionnaire scores with summary statistics of their responses in our task and individual parameters from winning models.

(a) SE and average skill estimate

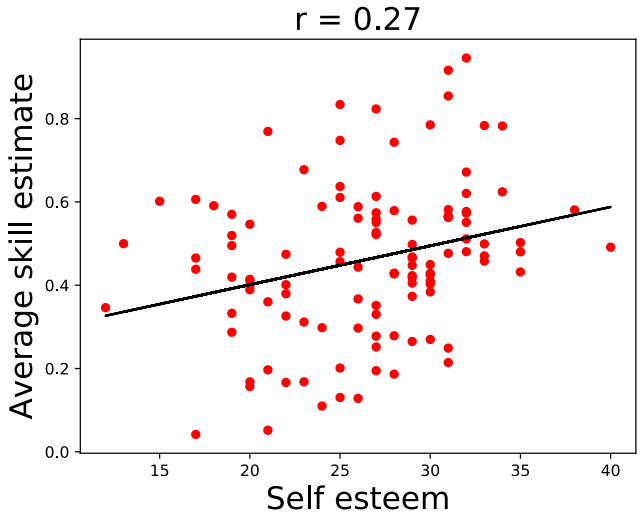

(b) SE and internal attributions post wins vs losses

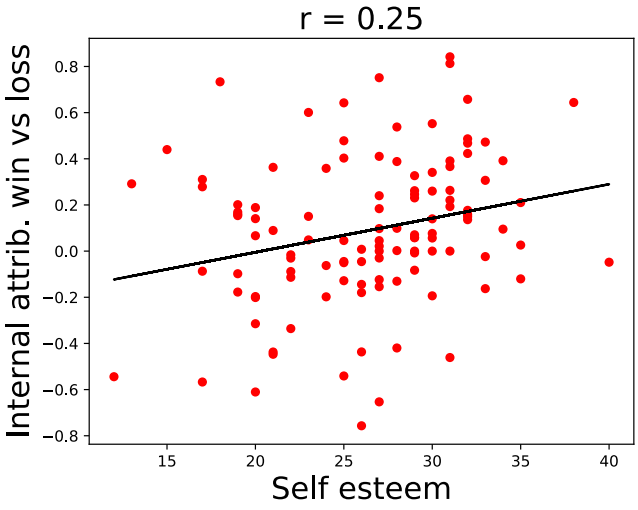

(c) SE and average skill update post wins vs losses

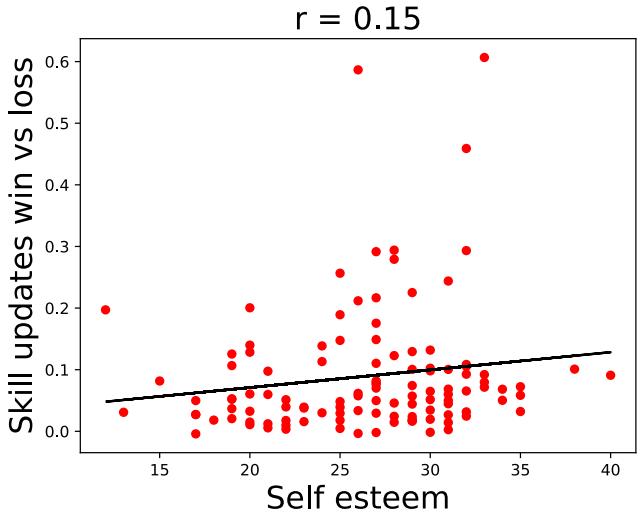

(d) LCi vs internal attributions

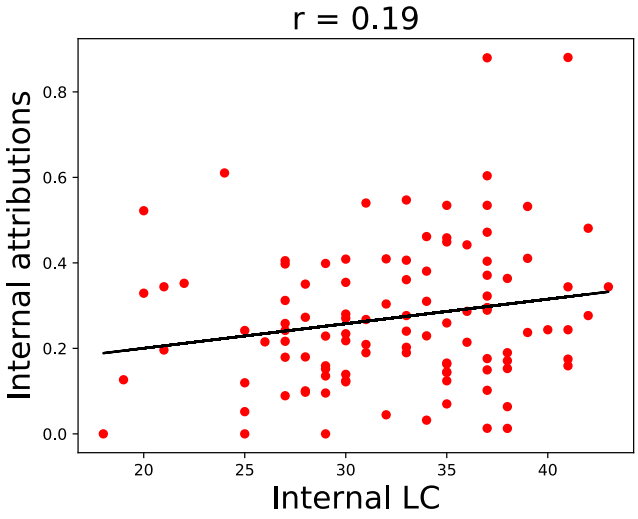

(e) ASQ I- vs internal attributions for losses

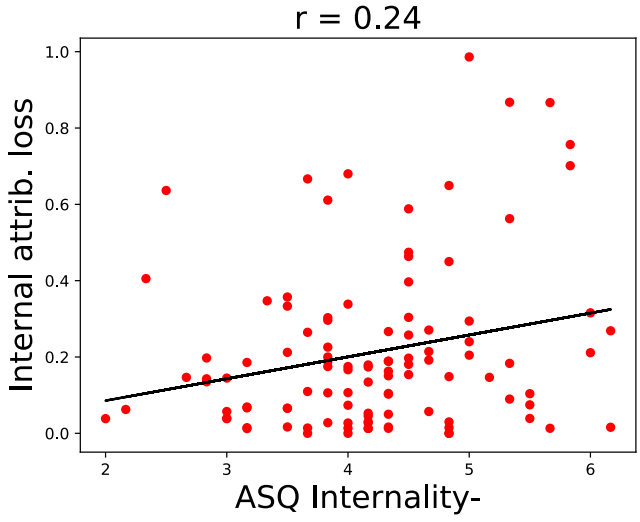

(f) ASQ I+ vs internal attributions for wins

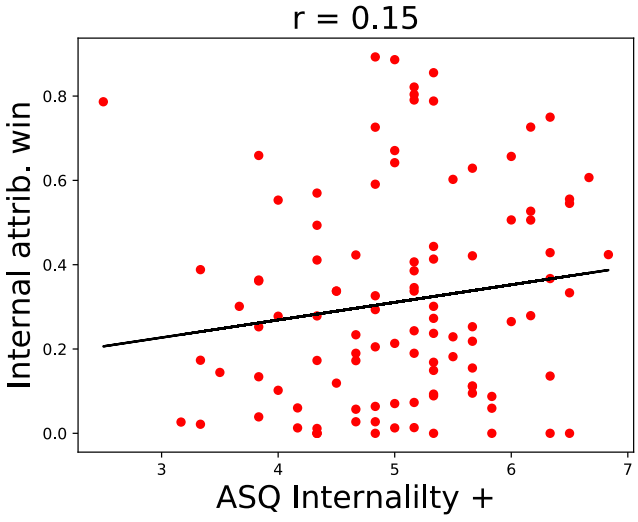

Figure 9: Behavioural measures vs questionnaire scores.

Questionnaires and behaviour in the task We started from some *a priori* expectations about connections between questionnaire measures and aspects of behaviour. Specifically, we expected participants with higher self-esteem scores to judge themselves as more skilled, to show larger differences in skill updates after wins vs after losses, and to show larger differences between the proportions of internal attributions for wins vs losses, compared to participants with lower self-esteem. We also expected to find a correlation between scores for internal locus of control and the proportion of internal attributions, as well as between scores on the internality positive and internality negative subscales of the ASQ and the corresponding proportions of internal attributions in our task.

We computed correlations between the relevant behavioural variable and relevant questionnaire score for or each of these hypothesised relationships and tested significance by performing permutation tests. For self esteem (SE), we found a significant positive correlation with the average skill estimate ($r = 0.27$, two-sided $p = 0.0026$, Figure 9a), and a significant positive correlation with the difference between the proportions of internal attributions for wins and losses ($r = 0.25, p = 0.0074$, Figure 9b), both in the expected directions; there was no significant relationship with the difference in skill updates after wins vs after losses ($r = 0.14, p = 0.1086$, Figure 9c). There was a positive correlation between the internal locus of control subscale and the proportion of internal attributions in the task ($r = 0.19, p = 0.0464$, Figure 9d). Finally, the proportion of internal attributions post losses in the task was positively correlated with the Internality for negative events subscale of the ASQ ($r = 0.24, p = 0.016$, Figure 9e); the relationship between the proportion of internal attributions for wins and the Internality for positive events subscale of the ASQ was not signficant ($r = 0.15, p = 0.1408$, Figure 9f).

Questionnaires and model parameters These relationships were confirmed by analyses of the related model parameters. Thus we found a positive correlations between SE and the initial skill parameter ($r = 0.23, p = 0.0284$), the difference between learning rates for wins and learning rates for losses ($r = 0.22, p = 0.0328$; see Methods section 4.8 for details of parameter computation) and the difference between learning rates for internal attributions for wins vs losses ($r = 0.27, p = 0.0058$) in the winning skill model; SE was also positively correlated with the difference between the baseline preference for internal attributions post wins vs losses extracted from the winning attribution model ($r = 0.26, p = 0.0098$). Scores on the internal subscale of the locus of control questionnaire (LCi) were positively correlated with the baseline preference for internal attributions ($r = 0.21, p = 0.0446$), and scores on the internality for negative events subscale of the ASQ were positively correlated with the baseline preference for internal attributions for losses ($r = 0.28, p = 0.0052$). The analogous relationship for positive events was not significant ($r = 0.15, p = 0.1276$).

3 Discussion

3.1 Summary

We developed a novel task to test for the presence of trial-level reciprocal effects in the generation of causal attributions and beliefs about self, as such high temporal resolution effects are central to attribution-self-representation cycle theory (Bentall, 2003; Bentall et al., 2001) and its account of the generation and maintenance of psychiatric disorders. In the context of learning a new competence, participants repeatedly experienced real outcomes and reported both causal attributions and beliefs about skill. In accordance with the theory, and at a trial-by-trial level,

participants changed their skill estimates more after outcomes attributed internally than after ones attributed externally. Conversely, with increasing reported skill, they took more credit for wins and less blame for losses, showing that attributions also varied with beliefs. Correlations between behaviour in the task and questionnaire-based measures of self-esteem, control, and attributional style suggest that the task indeed taps into some of the mechanisms investigated by these tools.

Model-dependent analyses confirmed these results: model comparison favoured a skill model with higher learning rates for internally attributed outcomes than externally attributed ones, and a model for attributions which included skill as a feature over and above objective task and performance measures. Analyses of model parameters recapitulated model-agnostic results, and provided additional insights into the mechanisms at work, notably an effect of valence on learning rates for self, but not for other, distinctions between the effect of trial-by-trial performance measure vs estimated skill on attributions, and relationships between individual differences in parameters and questionnaire scores.

Our results provide the first evidence at a high temporal resolution of the dynamical nature of, and interactions between, attribution and beliefs about skill, and lay the foundations for computational investigations of the behavioural regimes that such reciprocal interactions can sustain.

3.2 Reflections on the task

The novel task that we introduce in this work differs from previous work in several important ways. Participants' own attributions for outcomes and beliefs about skill are collected at fine temporal granularity, in the absence of any manipulation targeting their content. These attributions and beliefs about skill refer to real outcomes, which depend on participants' performance as well as on objective task manipulations, and which are experienced in a gameified context that promotes engagement and enhances motivation. While understanding the rules and playing the game are straightforward for participants, disentangling the extent to which various task aspects and the player's own performance contributed to the outcome is not. Participants are therefore faced with a realistically challenging problem in interpreting their experience to learn about how good they are at the task.

These novel aspects of the task enabled the investigation of trial-level effects of causal attributions on beliefs and the reverse, in a relatively realistic scenario involving real outcomes, beliefs and attributions, in a large population of online participants. There were, however, a number design choices that we could not readily base on evidence from previous work, and which could be improved in the future. Most importantly, we did not have access to any previously validated objective measures of difficulty and skill; instead, in order to maintain the balance between positive and negative outcomes, as well as to maintain subjects' engagement in the task, we adapted trial difficulty to subjects' performance level with a staircase procedure, involving the various aspects of the task that were, intuitively, the most relevant; we then attempted to extract a post hoc objective measure of difficulty from the data (see Supplementary material 2). The staircase procedure did achieve the practical desiderata of satisfying time constraints and providing balanced numbers of wins and losses. However, the steps of difficulty for some of the components were rather coarse, implying undesirably high trial-to-trial variations in difficulty (see Supplementary material 2). Given the evident viability of the framework, it would be possible in future work to explore at a finer granularity the components of the task that determine difficulty and to calibrate them by measuring the frequency of wins and losses for each setting

of these components in a large population of subjects.

Using an externally validated measure of difficulty together with a precisely calibrated staircase would allow participants' real underlying skill level to be closely tracked and objectively measured, enabling investigations into relationships between the accuracy of participants' beliefs about skill and their patterns of attributing outcomes, or questionnaire-based psychological measures. A number of questions that we could not address in this work could then be asked, such as whether participants display self-serving biases in their skill estimates, and, if so, whether such biases are associated with self-serving attribution patterns or with higher self-esteem scores. Finally, better empirical knowledge of task psychophysics could conceivably enable artificial simulation of agents with specified properties, which would remove the need to deceive participants about the "other" condition and allow the ordering between conditions to be balanced, as well as enable additional questions to be tackled. Introduction of multiple "other" conditions could be used to control for the effect of watching vs playing; evaluation of close "others" could be compared with evaluation of indifferent, hostile or artificial "others"; other social aspect of the processes involved could also be investigated, by e.g. introducing competition into the task, or by providing agents with a learning environment populated by real or artificial peers.

3.3 Self-serving and self-defeating mechanisms

Attributions in our data showed evidence for both self- and other-serving biases: in both conditions, participants made more internal attributions for wins than for losses. This pattern, in which people take more credit for positive outcomes than they take for negative ones, has often been reported before (Alloy and Abramson, 1979; Bentall and Kaney, 2005; Martin et al., 1984; Tillman and Carver, 1980; Vázquez, 1987) (see Campbell and Sedikides (1999); Mezulis et al. (2004) for reviews), and is thought to contribute to the maintenance of well being, promoting persistence and exploration in the face of negative feedback or failure.

Selective attendance to, or overweighing of, positive information, presumably in an attempt to maintain a positive view of oneself, has also been found in aspects of learning (Cazé and Van Der Meer, 2013; Cox et al., 2015; Frank et al., 2007; Lefebvre et al., 2017; Lyons et al., 2020; Sharot et al., 2011; Wrase et al., 2007; Yacubian et al., 2006) (see Blaine and Crocker (1993); Campbell and Sedikides (1999); Mezulis et al. (2004) for reviews). However this is not what we found in our data, where we observed instead heightened learning from negative feedback for self (but not for other): learning rates for skill associated with losses were higher than those associated with wins, with a larger effect for internally attributed outcomes than for externally attributed ones.

Heightened attention to negative feedback can be essential for survival in harsh environments, and might be particularly relevant in learning contexts where it can be used to improve performance (Maier and Seligman, 2016; Müller-Pinzler et al., 2019), There is also, accordingly, ample evidence in animals as well as humans of the privileged processing of negative feedback: negative, more than positive, feedback produces rapid and strong bodily responses, mobilising the organism for reaction; negative emotions produce more arousal than positive ones; negative events and information focus attention (see (Taylor, 1991) for a review); in humans, concepts for negative actions and consequences form earlier than their positive counterparts (Fincham, 1985), negative events are surveyed more for potential causal information (Bohner et al., 1988; Wong and Weiner, 1981), and they elicit more spontaneous causal attributional activity than positive ones (Peeters and Czapinski, 1990). Our result replicates that of (Müller-Pinzler et al., 2019) who, in a learning context somewhat similar to ours, found higher learning rates for negative

than for positive feedback, for self but not for other. The fact that in our data the effect was larger for internally attributed outcomes is consistent with an emphasis on negative feedback to improve performance while learning a novel task, since learning from internally attributed failures could be particularly informative for improving future performance.

The apparently contradictory pattern of biases in attribution and learning might be involved in maintaining balance between learning from negative feedback and maintaining a positive view of the self. In this scenario, the bias toward internal attributions for positive outcomes counteracts the threatening effects of heightened sensitivity to negative feedback, while this sensitivity promotes efficient learning from undesirable outcomes even in the presence of opposing motivational and emotional biases.

3.4 Loopy dynamics

As the name clearly states, the core concept of the attribution self-representations cycle theory is the cycle linking the two variables. In this work we focused on establishing the presence of each of the mutual effects, and have therefore analysed each of the two sides of the postulated cycle separately. This is a necessary first step, laying the ground for further research aimed at understanding the cycle dynamics and their involvement in mental health and disorders.

Indeed, models including "loopy", rather than linear connections between components have recently been the topic of several investigations. These illustrate the richness of the observable behaviours in such systems, and the way subtle modulations of their dynamics can produce important large scale effects (Adams et al., 2021; Dayan et al., 2020; Eldar and Niv, 2015). (Eldar and Niv, 2015) investigated interactions between emotional states (happy vs sad states induced by winning vs losing in a wheel of fortune draw) and learning about the rewarding properties of various slot machines in real human data and in simulations. Their results suggest that a positive-feedback effect of emotional state on the perception of outcomes may play a significant role in the emergence of mood instability. In a model of delusions, (Adams et al., 2021) show how small changes in "mood"- modelled as preference for one particular internal state- can, in conjunction with other factors (the quality of sensory information), dramatically alter behaviour, switching it from a relatively accurate pattern of inferences to one dominated by a "delusional" pattern of false and incorrigible inferences. In (Adams et al., 2021)'s case these effects were produced by reciprocal connections between mood, prior beliefs and actions: the agents' mood influenced their beliefs, which determined action choice; in the absence of poor quality external information action choice was, in turn, used as a source of information, which resulted in the strengthening of initial beliefs and mood.

Such effects could also occur from the connection between attributions and beliefs. The simulations we presented in section 2.1 illustrate how reciprocal effects can enrich the space of dynamics in even a simple model of two interacting variables, amplifying randomly generated differences between agents, or producing latent vulnerabilities. These simulation only scratch the surface of a much larger space, as such systems can display high sensitivity to changes in inputs and parameters, as well as path dependencies, i.e., sensitivity to the ordering of inputs(Adams et al., 2021; Dayan et al., 2020). There are a number of directions that need further exploration.

First, it is important to examine the parameter space systematically, i.e., to quantify the effects of exploring the ranges of individual parameters, as well as the effect of changes along multiple axes in the parameter space. Parameters of particular interest include attribution sensitivity to skill level (β parameters in our simulations), skill sensitivity to attributions (differences between

skill learning rates for internal vs external attributions in our simulations) and asymmetries between positive and negative outcomes (outcome-based learning rates in our simulation).

Going beyond examining parameters, one other particularly important direction is exploration of the space of models. This includes investigation of various linking functions between variables - e.g. in our simulations we chose *a priori* a sigmoid functional form of the dependencies of attributions on skill - but also higher level dimensions of variability, such as time varying models, or higher order dependencies between variables. Such extensions would allow formulation and testing of hypotheses involving multiple time-scales of interaction or threshold-like mechanisms that might be responsible for catastrophic dynamics. As an example, one can imagine testing, in this context, whether an association between particularly low levels of beliefs about self and particularly strong effects of negative attributions on beliefs could model the onset of depression (Bentall, 2003).

Finally, a crucial direction from a translational point of view is the investigation of perturbations and or environment manipulations that can produce qualitative shifts in the system dynamics, e.g. switches between "healthy" stable regimes and vicious cycles, or falls into of "delusional" (Adams et al., 2021) or other pathological attractor states (Bentall, 2003; Dayan et al., 2020). Harmful factors pushing the system away from equilibrium, and protective ones that prevent or dampen such dynamics (Haeffel and Vargas, 2011; Johnson et al., 2017; Liu et al., 2015; Robins and Hayes, 1995) would be particularly relevant for understanding the emergence of disorders and for developing preventive interventions or treatments. For instance, in connection with depression, simulations could be used to investigate differences between the effects of repeated small losses vs isolated exceptionally large losses, the effects of losses concentrated vs distributed in time (and their analogues for wins), and disproportionate effect of attributions for particularly significant events. Better understanding of such path dependencies associated with negative outcomes and attributions might yield insights into targets for therapeutic interventions or mechanisms for maximising their effectiveness.

3.5 Conclusions

To conclude, we propose a new framework for investigations into relationships between attributions and beliefs about self, and provide evidence of the dynamical nature of these variables and of reciprocal effects between them *at high temporal resolution*. We thereby validate one central tenet of the attribution self-representation cycle theory in healthy controls, helping bolster the important implications of this theory for the understanding of the emergence and maintenance of psychiatric disorders, as well as potentially for therapeutic approaches.

4 Methods

4.1 Participants

Participants (self: $N = 122$, 78 female, 41 male, 3 missing gender data, aged $18 - 35$ years, $M = 24.71, SD = 5.16$; other: $N = 92$, 58 female, 32 male, 2 missing gender data, aged $18 - 35$ years, $M = 24.93, SD = 5.18$) were recruited on Prolific. The experiment was conducted under ethics approval granted by Ethik-Kommission an der Medizinischen Fakultät der Eberhard-Karls-Universität und am Universitätsklinikum Tübingen.

4.2 Experiment timeline

Participants were directed to the experiment hosted on Gorilla®. After providing informed consent, participants completed the three questionnaires. They were then presented with the detailed instructions for the task, and answered verification questions to check understanding of the task instructions. Then then proceeded to the first session of the task. The second session in the self condition started with experimental trials directly and took part the following day. One week later, participants were provided with instructions for the other condition, which they they performed. Finally they performed the last session the following day. After the end of the last session participants were payed and provided with a feedback questions form, then with a separate question asking them to compare the two conditions. They were then fully debriefed.

4.3 Attribution and skill questions

Attributions were elicited with multiple-choice questions. After wins in the self condition, participants were asked "Why did you win the last trial? Pick the main cause:"; the response options were phrased as "simple maze", "few/ simple rotations", "luck", "my ability"; the corresponding version for the ability option in the other condition was "their ability". The question was adapted for losses by replacing "win" with "lose" and for the other condition by replacing "you" with "they". Response options for losses were phrased as "complex maze", "many/difficult rotations", "bad luck", "my lack of skill", respectively; and "their lack of skill" for the other condition.

Participants were then asked "How good do you think you are at the task at this moment?" ("you" replaced by "they" in the other condition) and responded using a slider on a continuous scale with extremes labeled "very bad" and "very good".

4.4 Model agnostic analyses

Model agnostic analyses of skill estimates were performed on skill updates, computed as follows: skill estimates were z-scored within participant, and differences between two successive z-scored skill estimates were entered into analyses as skill updates. For both model-agnostic and model-dependent analyses, attributions were relabelled as internal (self) vs external (maze complexity, rotations, luck). All tests were performed as permutation tests with 5000 samples, using the paired t-test statistic as the measure of interest and permuting labels within participant. E.g. to test for an effect of attribution conditioned on outcome =win, attribution labels (internal vs external) for all wins were permuted within participant; average update post internally attributed wins and average update post externally attributed wins were then computed for each participant; paired (int vs ext) t statistic was then computed across participants for each permutation. Two-sided p-values were computed across permutations.

Reported effect sizes are Hedge's corrected d

$$d = \left(1 - \frac{3}{4(l_1 + l_2) - 9}\right) \frac{\mu_1 - \mu_2}{\sqrt{\frac{(l_1-1)*\sigma_1^2 + (l_2-1)*\sigma_2^2}{l_1+l_2-2}}},$$

where μ_1 and μ_2 are the sample means, σ_1, σ_2 are the sample standard deviations, and l_1, l_2 are the sample sizes.

Model agnostic analyses of attributions were performed on attribution proportions, computed as follows: for a given level of a factor of interest and a given attribution option, the summary statistic we used was the proportion of attributions to the option, out of all attributions provided for the factor level: e.g. for the effect of outcome on internal attributions in the self condition we compared the proportions of attributions to self out of all attributions provided for wins, vs the proportion of attributions to self out of all attributions provided for losses. We used quartile discretisation for factors of interest other than outcome, which were continuous. For each participant and each factor of interest, discretisation was performed on the z-scored factor values. Reported p-values are estimated with permutation tests, performed by permuting labels within participant.

All tests reported as significant survived Benjamini Hochberg (Benjamini and Hochberg (1995)) correction for multiple comparisons, unless otherwise stated (82 tests, $\alpha = 0.05$, highest p-value under threshold = 0.0464).

4.5 Skill models

All skill models were variations of the basic Rescorla-Wagner model:

$$
\begin{aligned}
\delta_t &= o_t - s_{t-1} \\
s_t &= \begin{cases} s_{t-1} + \alpha * \delta_t & \text{if } t \neq t_0^{II} \\ s_{t-1} + \beta + \alpha * \delta_t & \text{otherwise, where} \end{cases} \\
t_0^{II} &= \text{index of first trial of the second session.} \\
s_t &= \text{underlying skill estimate at trial t} \\
o_t &= \text{outcome of trial t} \\
\alpha &= \text{learning rate} \\
\beta &= \text{effect of break between sessions}
\end{aligned}
$$

The reported skill estimate was assumed to be a noisy reading of the underlying skill estimate, drawn from a fixed Gaussian distribution $\mathcal{N}(0, 0.1)$. We refer to this as the baseline model. The rest of the models were obtained by allowing α to vary as follows:

- S: different learning rates for the first and second sessions (2 learing rates);
- O: different learning rates for wins and losses (2 learning rates);
- A: different learning rates for internal and external attributions plus a learning rate for outcomes with missing attributions (3 learning rates);
- SA: different learning rates for different attributions, separate for the two sessions (6 learning rates);
- SO: different learning rates for wins and losses, separate for the two sessions (4 learning rates);
- AO: different learning rates for different attributions, separately for wins and losses (6 learning rates);

- SAO: different learning rates for each combination of attribution and outcome, separately for the two sessions (12 learning rates).

Fitting and model comparison were performed separately for self and other. All models were coded in the pystan interface (https://pystan.readthedocs.io/en/latest/) to the STAN probabilistic programming language (https://mc-stan.org/), which we used to obtain samples from the posterior distribution over parameters and estimate model evidence (1000 iterations, 4 chains per model). All models were fitted in two versions: independently for each participant and as a hierarchical model over the entire population. For the hierarchical version, individual parameters were drawn from independent Beta distributions over each parameter at the population level:

$$\begin{aligned}
\theta_p^i &\sim Beta(\alpha_i, \beta_i), \text{ where} \\
\alpha_i, \beta_i &= \text{population level parameters} \\
p &\text{ indexes participant} \\
i &\text{ indexes latent parameter.}
\end{aligned}$$

4.6 Attribution models

All attribution models were built on the same underlying structure:

$$\begin{aligned}
s_{t,o} &= \boldsymbol{w}_o \cdot \boldsymbol{f}_t \forall o \in O \\
p_t(o) &= \frac{\exp(s_{t,o})}{\Sigma_{o \in O} \exp(s_{t,o})}, \text{ where} \\
s_{t,o} &= \text{score of response option o on trial t} \\
\boldsymbol{w}_o &= \text{feature weights for option o,} \\
&\quad \text{corresponding to the outcome on trial t-1} \\
\boldsymbol{f}_t &= \text{feature values on trial t} \\
O &= \text{set of available response options} \\
p_t(o) &= \text{probability of choosing option o on trial t}
\end{aligned} \tag{3}$$

All models had separate parameters for wins and losses. Because participants have to choose one of the 4 available response options, models were equipped with independent parameters for *three* of the options, and scores for the 4 options were constrained to sum to 0; thus parameters and preference for the Luck option were entirely derived from parameters and preferences for the other options and we do not report analyses of these parameters. Each model was defined by the features included:

- no features - baseline propensity for each attribution option;
- bias+ reported skill (1 feature, 12 parameters per participant);
- bias+ performance features: proportion of correct key presses out of all key presses in trial; proportion of pauses out of all frames in trial (2 features, 18 parameters per participant);
- bias+ performance features (as above) and task features: length of correct path through maze, proportion of frames with maze in unusual (not UP) orientation in trial (4 features, 30 parameters per participant);

- reported skill and task features (as above, 3 features, 24 parameters per participant);
- full model: reported skill, performance features, task features (5 features, 36 parameters per participant).

Each model was fitted in two versions: independently for each participant and in a hierarchical version, assuming a Normal distribution at the population level for each parameter. Fitting and model comparison were performed separately for self and other. All models were coded in the pystan interface (https://pystan.readthedocs.io/en/latest/) to the STAN probabilistic programming language (https://mc-stan.org/), which we used to obtain samples from the posterior distribution over parameters and estimate model evidence (1000 iterations, 4 chains per model).

Transforming feature weights into meaningful effects Because scores obtained from linear combinations of features are passed through a softmax to compute the probability of each attribution option, and because there are separate softmax transformations for wins and losses, direct comparison between feature weights for different options is not informative. As an example, the bias parameter for a given option does not directly and independently translate into the participant's preference for that option; it is the relationship between biases for the different options, entering into the softmax function, that determines preferences for the different options. In order to be able to meaningfully compare baseline preferences and the effects of various features we transformed these parameters as follows: for biases, we applied the softmax transformations to the bias parameters only, clamping all feature weights to 0, and compared the resulting probabilities. For a feature of interest x we computed xAo, its contribution to choosing attribution option A having encountered outcome o (e. g sI+ for the contribution of skill to making internal attributions for wins) as:

$$\begin{aligned}
xAo &= \frac{1}{T}\sum\nolimits_t \left.\frac{\partial p_t(A)}{\partial x}\right|_{x=0}, \text{where} \\
T &= \text{total number of trials} \\
p_t(A) &= \text{the probability of choosing attribution option A on trial t} \\
&= \frac{\exp(w_{Ao}\cdot f_t)}{\exp(w_{Io}\cdot f_t)+\exp(w_{Mo}\cdot f_t)+\exp(w_{Ro}\cdot f_t)+\exp(w_{Lo}\cdot f_t)}, \text{where} \\
f_t &= \text{feature values on trial t} \\
w_{Xo} &= \text{ weights for computing the score for attributions to X} \\
&\quad \text{after outcome o.}
\end{aligned} \tag{4}$$

Thus the derivative of $p_t(A)$, seen as a function of x, is evaluated at $x = 0$, and values are then averaged over all trials.

4.7 Model comparison

Model comparison was performed with the WAIC score (Watanabe, 2010), an approximation for the out-of-samples predictive log density, computed as

$$\begin{aligned}
\widehat{WAIC} = &-\frac{1}{P}\sum\nolimits_{p=1}^{P}\log\left(\frac{1}{S}\Sigma_{s=1}^{S}p(X_p|\boldsymbol{\theta}_p^s)\right) \\
&+\frac{1}{P}\sum\nolimits_{p=1}^{P}\left[\mathrm{Var}_{s=1}^{S}(\log p(X_p|\boldsymbol{\theta}_p^s))\right],
\end{aligned} \tag{5}$$

where $\left\{\boldsymbol{\theta}_p^1, \boldsymbol{\theta}_p^2 ... \boldsymbol{\theta}_p^S\right\}$ is the set of samples from the posterior distribution over the vector of individual parameters, $\boldsymbol{\theta}_p$ for participant p.

In all cases WAIC score comparisons preferred hierarchical versions of the models; we therefore only report analyses of these. Mean posterior parameters for individual participants were entered into further analyses.

4.8 Questionnaires

As detailed above, learning rates from the winning skill model varied along three axes: outcome, attribution and session. In analyses of relationships with questionnaire measures, parameters were averaged across the direction irrelevant for the analysis.

References

Abramson, L. Y., Metalsky, G. I., and Alloy, L. B. (1989). Hopelessness depression: A theory-based subtype of depression. *Psychological review*, 96(2):358.

Abramson, L. Y., Seligman, M. E., and Teasdale, J. D. (1978). Learned helplessness in humans: Critique and reformulation. *Journal of Abnormal Psychology*, 87(1):49–74.

Adams, R. A., Vincent, P., Benrimoh, D., Friston, K. J., and Parr, T. (2021). Everything is connected: Inference and attractors in delusions. *Schizophrenia Research*.

Allan, L. G., Siegel, S., and Hannah, S. (2007). The sad truth about depressive realism. *Quarterly Journal of Experimental Psychology*, 60(3):482–495.

Alloy, L. B. and Abramson, L. Y. (1979). Judgment of contingency in depressed and non-depressed students: Sadder but wiser? *Journal of Experimental Psychology: General*, 108(4):441–485.

Alloy, L. B. and Abramson, L. Y. (1982). Learned helplessness, depression, and the illusion of control. *Journal of Personality and Social Psychology*, 42(6):1114–1126.

Alloy, L. B., Abramson, L. Y., and Viscusi, D. (1981). Induced mood and the illusion of control. *Journal of Personality and Social Psychology*, 41(6):1129–1140.

Benassi, V. A. and Mahler, H. I. (1985). Contingency Judgments By Depressed College Students. Sadder But Not Always Wiser. *Journal of Personality and Social Psychology*, 49(5):1323–1329.

Benjamini, Y. and Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal statistical society: series B (Methodological)*, 57(1):289–300.

Bentall, R. P. (2003). *Madness explained: Psychosis and human nature*. Penguin UK.

Bentall, R. P., Corcoran, R., Howard, R., Blackwood, N., and Kinderman, P. (2001). Persecutory delusions: a review and theoretical integration. *Clinical psychology review*, 21(8):1143–1192.

Bentall, R. P. and Kaney, S. (2005). Attributional lability in depression and paranoia. *British Journal of Clinical Psychology*, 44(4):475–488.

Blaine, B. and Crocker, J. (1993). Self-esteem and self-serving biases in reactions to positive and negative events: An integrative review. In *Self-esteem*, pages 55–85. Springer.

Bohner, G., Bless, H., Schwarz, N., and Strack, F. (1988). What triggers causal attributions? The impact of valence and subjective probability. *European Journal of Social Psychology*, 18(4):335–345.

Campbell, W. K. and Sedikides, C. (1999). Self-threat magnifies the self-serving bias: A meta-analytic integration. *Review of general Psychology*, 3(1):23–43.

Carson, R. C., Hollon, S. D., and Shelton, R. C. (2010). Depressive realism and clinical depression. *Behaviour Research and Therapy*, 48(4):257–265.

Cazé, R. D. and Van Der Meer, M. A. (2013). Adaptive properties of differential learning rates for positive and negative outcomes. *Biological Cybernetics*, 107(6):711–719.

Cox, S. M., Frank, M. J., Larcher, K., Fellows, L. K., Clark, C. A., Leyton, M., and Dagher, A. (2015). Striatal D1 and D2 signaling differentially predict learning from positive and negative outcomes. *NeuroImage*, 109:95–101.

Dayan, P., Roiser, J. P., and Viding, E. (2020). The first steps on long marches: The costs of active observation.

Dunning, D., Leuenberger, A., and Sherman, D. A. (1995). A New Look at Motivated Inference: Are Self-Serving Theories of Success a Product of Motivational Forces? *Journal of Personality and Social Psychology*, 69(1):58–68.

Eldar, E. and Niv, Y. (2015). Interaction between emotional state and learning underlies mood instability. *Nature communications*, 6(1):1–10.

Fincham, F. D. (1985). Outcome Valence and Situational Constraints in the Responsibility Attributions of Children and Adults. *Social Cognition*, 3(2):218–233.

Forgas, J. P., Bower, G. H., and Moylan, S. J. (1990). Praise or blame? Affective influences on attributions for achievement. *Journal of Personality and Social Psychology*, 59(4):809–819.

Frank, M. J., Moustafa, A. A., Haughey, H. M., Curran, T., and Hutchison, K. E. (2007). Genetic triple dissociation reveals multiple roles for dopamine in reinforcement learning. *Proceedings of the National Academy of Sciences of the United States of America*, 104(41):16311–16316.

Haeffel, G. J. and Vargas, I. (2011). Resilience to depressive symptoms: The buffering effects of enhancing cognitive style and positive life events. *Journal of behavior therapy and experimental psychiatry*, 42(1):13–18.

Howell, D. C. (2012). *Statistical methods for psychology*. Cengage Learning.

Johnson, J., Panagioti, M., Bass, J., Ramsey, L., and Harrison, R. (2017). Resilience to emotional distress in response to failure, error or mistakes: A systematic review. *Clinical psychology review*, 52:19–42.

Kaney, S. and Bentall, R. P. (1992). Persecutory delusions and the self-serving bias: Evidence from a contingency judgment task. *Journal of Nervous and Mental Disease*.

Kelley, H. H. (1967). Attribution theory in social psychology. In *Nebraska symposium on motivation*. University of Nebraska Press.

Lefebvre, G., Lebreton, M., Meyniel, F., Bourgeois-Gironde, S., and Palminteri, S. (2017). Behavioural and neural characterization of optimistic reinforcement learning. *Nature Human Behaviour*, 1(4):1–9.

Levenson, H. (1974). Activism and Powerful Others: Distinctions Within the Concept of Internal-External Control. *Journal of Personality Assessment*, 38(4):377–383.

Liu, R. T., Kleiman, E. M., Nestor, B. A., and Cheek, S. M. (2015). The hopelessness theory of depression: A quarter-century in review. *Clinical Psychology: Science and Practice*, 22(4):345–365.

Lyons, B. A., McKay, A. M., and Reifler, J. (2020). High-status lobbyists are most likely to overrate their success. *Nature Human Behaviour*, 4(2):153–159.

Maier, S. F. and Seligman, M. E. (1976). Learned helplessness: Theory and evidence. *Journal of Experimental Psychology: General*, 105(1):3–46.

Maier, S. F. and Seligman, M. E. (2016). Learned helplessness at fifty: Insights from neuroscience. *Psychological review*, 123(4):349.

Martin, D. J., Abramson, L. Y., and Alloy, L. B. (1984). Illusion of control for self and others in depressed and nondepressed college students. *Journal of Personality and Social Psychology*, 46(1):125–136.

Mezulis, A. H., Abramson, L. Y., Hyde, J. S., and Hankin, B. L. (2004). Is there a universal positivity bias in attributions? A meta-analytic review of individual, developmental, and cultural differences in the self-serving attributional bias. *Psychological Bulletin*, 130(5):711–747.

Moore, M. T. and Fresco, D. M. (2012). Depressive realism: A meta-analytic review. *Clinical Psychology Review*, 32(6):496–509.

Müller-Pinzler, L., Czekalla, N., Mayer, A. V., Stolz, D. S., Gazzola, V., Keysers, C., Paulus, F. M., and Krach, S. (2019). Negativity-bias in forming beliefs about own abilities. *Scientific Reports*, 9(1):1–15.

Peeters, G. and Czapinski, J. (1990). Positive-negative asymmetry in evaluations: The distinction between affective and informational negativity effects. *European Review of Social Psychology*, 1(1):33–60.

Peterson, C., Semmel, A., Von Baeyer, C., Abramson, L. Y., Metalsky, G. I., and Seligman, M. E. (1982). The attributional style questionnaire. *Cognitive therapy and research*, 6(3):287–299.

Presson, P. K. and Benassi, V. A. (2003). Are depressive symptoms positively or negatively associated with the illusion of control? *Social Behavior and Personality*, 31(5):483–495.

Robins, C. J. and Hayes, A. M. (1995). The role of causal attributions in the prediction of depression. In Buchanan, G. M. and Seligman, M. E., editors, *Explanatory Style*, pages 71–97. Routledge.

Romney, D. M. (1994). Cross-validating a causal model relating attributional style, self-esteem, and depression: An heuristic study. *Psychological Reports*, 74(1):203–207.

Rosenberg, M. (1965). Society and the adolescent self-image [dissertation]. *NJ: Princeton Univ.*

Seligman, M. E. (1972). Learned helplessness. *Annual review of medicine*, 23(1):407–412.

Seligman, M. E. and Maier, S. F. (1967). Failure To Escape Traumatic Shock. *Journal of Experimental Psychology*, 74(1):1–9.

Sharot, T., Korn, C. W., and Dolan, R. J. (2011). How unrealistic optimism is maintained in the face of reality. *Nature Neuroscience*, 14(11):1475–1479.

Taylor, S. E. (1991). Asymmetrical effects of positive and negative events: the mobilization-minimization hypothesis. *Psychological bulletin*, 110(1):67.

Tennen, H. and Herzberger, S. (1987). Depression, Self-Esteem, and the Absence of Self-Protective Attributional Biases. *Journal of Personality and Social Psychology*, 52(1):72–80.

Tennen, H., Herzberger, S., and Nelson, H. F. (1987). Depressive Attributional Style: The Role of Self-Esteem. *Journal of Personality*, 55(4):631–660.

Tillman, W. S. and Carver, C. S. (1980). Actors' and observers' attributions for success and failure: A comparative test of predictions from Kelley's cube, self-serving bias, and positivity bias formulations. *Journal of Experimental Social Psychology*, 16(1):18–32.

Vázquez, C. (1987). Judgment of Contingency: Cognitive Biases in Depressed and Nondepressed Subjects. *Journal of Personality and Social Psychology*, 52(2):419–431.

Venkatesh, S., Moulds, M. L., and Mitchell, C. J. (2018). Testing for Depressive Realism in a Clinically Depressed Sample. *Behaviour Change*, 35(2):108–122.

Watanabe, S. (2010). Asymptotic equivalence of Bayes cross validation and widely applicable information criterion in singular learning theory. *Journal of Machine Learning Research*, 11:3571–3594.

Wong, P. T. and Weiner, B. (1981). When people ask "why" questions, and the heuristics of attributional search. *Journal of Personality and Social Psychology*, 40(4):650–663.

Wrase, J., Kahnt, T., Schlagenhauf, F., Beck, A., Cohen, M. X., Knutson, B., and Heinz, A. (2007). Different neural systems adjust motor behavior in response to reward and punishment. *NeuroImage*, 36(4):1253–1262.

Yacubian, J., Gläscher, J., Schroeder, K., Sommer, T., Braus, D. F., and Büchel, C. (2006). Dissociable systems for gain- and loss-related value predictions and errors of prediction in the human brain. *Journal of Neuroscience*, 26(37):9530–9537.

Interactions between attributions and beliefs at high temporal resolution : evidence from a novel task Supplementary materials

Elena Zamfir, Peter Dayan
MPI for Biological Cybernetics, Tuebingen

1 Staircase procedure

We used a double staircase procedure, with the stair used for each trial chosen randomly and updated based on the outcome as described below. Each stair was characterised by three variables:

- the overall size of the maze, n: if the maze is conceived as a square nXn matrix of "maze chambers", with each chamber having four possible walls, which can be present or absent, all the combinations of wall patterns that form a valid maze give the total set of available mazes for a given maze size n; there are four possible "levels" of maze size on our staircases, consisting in two values for n, one of which is randomly drawn, with equal probability, before generating the maze on each trial: level 0 has available n values $\{3, 5\}$, level 1 has available n values $\{5, 5\}$ level 2 has available n values $\{5, 7\}$ and level 3 has available n values $\{7, 7\}$.
- the average frequency of maze rotations during a trial $\overline{\nu}$: all trials start in the normal upright position and the first rotation, resulting in a randomly chosen orientation at an angle of 90, 180 or 270 degrees with respect to the upright one, happens 30 frames (1.5 seconds) later; a random number is then uniformly drawn from the interval $(\overline{\nu}-10, \overline{\nu}+10)$, representing the number of frames until the next rotation; the angle of the rotation is drawn randomly with equal probability from the three available options (90, 180, 270 degrees) every time a rotation happens. The available values for $\overline{\nu}$ were between 20 and 140 frames, and a staircase step was 10 frames.
- the available time for the trial, t: the available values for t are 10, 12 and 14 seconds, with the staircase step being 2 seconds.

Both staircases started with the same value for the available time, the maximum one of 14 seconds. One of the staircases started on level 1 of the maze dimensions available and on $\overline{\nu} = 30$ frames, the other started on level 2 of the maze dimensions available and on $\overline{\nu} = 80$ frames. These values were updated as follows: if still possible, $\overline{\nu}$ was increased (decreased) by 10 for wins and losses respectively; when the upper (lower) limit was reached, the maze dimension level was increased(decreased), if still possible; when the upper (lower) limit was reached for this as well, the available time t was decreased (increased) if possible.

2 Empirical difficulty and skill

We attempted to extract objective measures of both difficulty and skill from our data by first establishing a measure of difficulty, and then defining skill with respect to it. Intuitively, a difficult trial is one that a randomly picked subject in the population is likely to lose, while an easy trial is one that subjects are likely to win; subjects' skill is measured with respect to this difficulty - the more likely subjects are to win difficult trials, the more skilled they are considered to be. The staircase controls several objective dimensions along which trials vary, and which could contribute to their objective difficulty. We did not know *a priori* how these individual factors might contribute to an integrated difficulty score, instead, we attempted to infer the identities and weights of the different factors from data, assuming that difficulty predicts trial outcome. Specifically, given a set of objectively measurable task aspects that might act as factors in determining difficulty, $f_1, f_2, ..f_k$, we assumed that there is a stable set of weights, $w_1, w_2, .., w_k$, representing the contribution that they each have towards determining the outcome, such that on any trial t

$$
\begin{aligned}
p(o(t) = 1; \mathbf{w}) &= \sigma(\mathbf{w}_d^T \mathbf{f}_d(t)), \text{where} \\
\mathbf{w}_d^T &= (w_0, w_1, ...w_k) \\
\mathbf{f}_d^T(t) &= (1, f_1(t), ...f_k(t)) \\
f_i(t) &= \text{the measured value of factor } f_i \text{ at trial t} \\
o(t) &= \text{outcome at trial t, and} \\
\sigma(x) &= \frac{1}{1 + e^{-x}}, \text{ the sigmoid function.}
\end{aligned}
$$

Given $\mathbf{w}_d$, difficulty for trial t can be computed as $d(t) = -\mathbf{w}_d^T \mathbf{f}_d(t)$. In order to obtain, for each subject, an objective, external difficulty measure, not influenced by their own skill, we inferred $\mathbf{w}_d$ separately for each subject, by predicting outcomes for all remaining subjects[1]. We fitted all outcomes from all subjects with logistic regression models with several combinations of regressors and their interactions and, based on the cross-validation score, pickled the model which best explained outcomes across subjects. This included length of correct path to maze exit, proportion of non UP orientations, path length×orientation interaction, time available and necessary minimum speed as features. Figure 1, shows a summary and sanity checks of the resulting empirical difficulty measure.

Analysis of the accuracy of difficulty as a predictor of outcome suggests that the staircase might have failed to track subjects' skill levels: outcome prediction based only on difficulty values, according to the simple model $p(o(t) = 1) = \sigma(-d(t))$, is highly accurate, ranging from 0.69 to 0.91, with an average of 0.84 and s.d. of 0.04 across subjects(see figure 2). Tracking the subject's skill level would explore difficulty ranges where difficulty alone is insufficient as a predictor. The fact that difficulty alone is such a good predictor of outcomes implies that little room is left for skill; this is indeed what we found in our analyses aimed at defining skill, to which we turn next. We modelled objective skill as the evolving factor that intermediates between the objective difficulty of a trial and success, and attempted to extract an objective measure of skill from data by extending the previous difficulty-only model for outcomes:

[1]Here we present data from a pilot population of participants

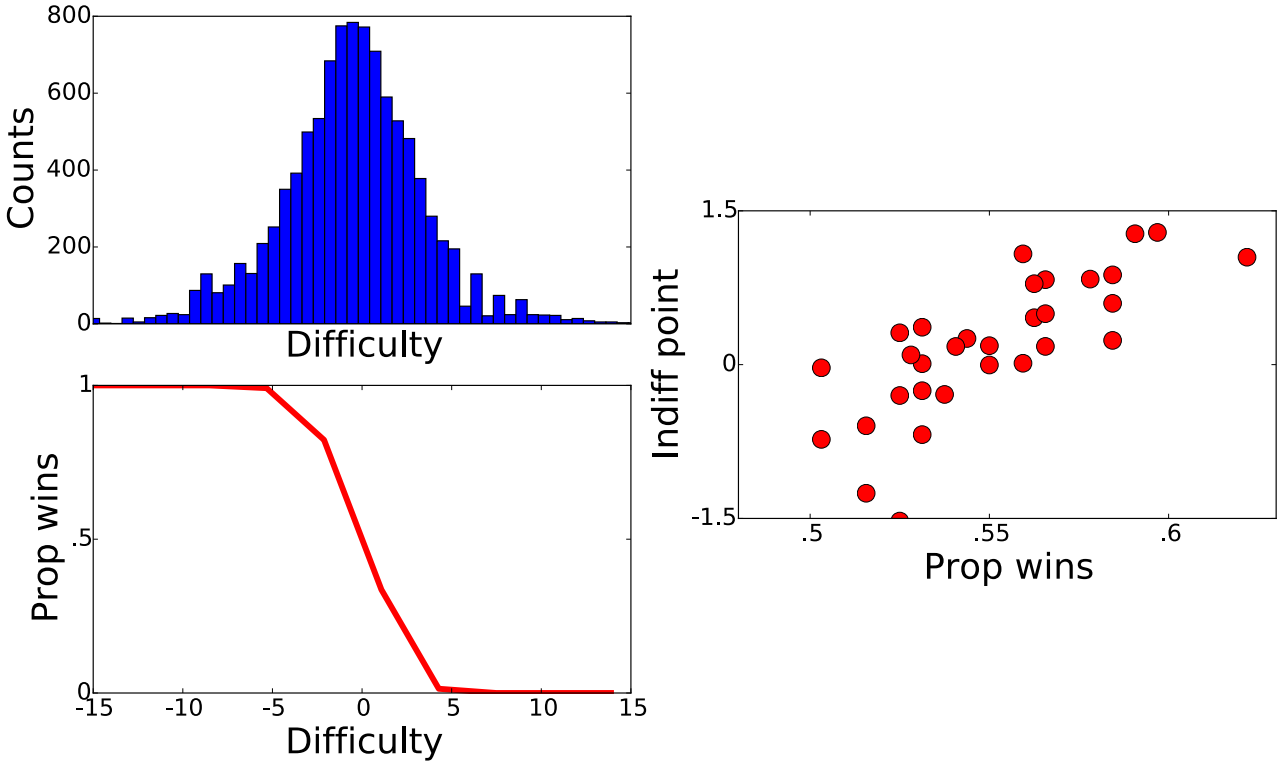

Figure 1: Difficulty measure: summary and sanity check. Left: data pooled from all subjects; top: distribution of difficulty values; bottom: relationship between difficulty and the proportion of wins. Right: relationship between difficulty indifference point - difficulty value for which subject is equally likely to win or lose the trial - and the proportion of trials won out of all trials; each dot represents a subject; $r^2 = 0.6$, p-value $= 3 * 10^{-7}$.

$$
\begin{aligned}
p(o(t) = 1; \mathbf{w}_p) &= \sigma(\mathbf{w}_p^T \mathbf{f}_p(t) + \mathbf{w}_d(t)), \text{ where} \\
o(t) &= \text{ outcome of trial t} \\
\mathbf{f}_p(t) &= \text{vector of performance regressors at trial t} \\
d(t) &= \text{ difficulty at trial t} \\
\mathbf{w}_p &= \text{performance weights, parameters} \\
\mathbf{w}_d &= \text{difficulty weight, fixed}
\end{aligned}
$$

For a given value of $\mathbf{w}_d$, fitting the above model to trials from one subject produces performance weights $\mathbf{w}_p$; these can then be used to obtain the trial by trial skill measure, computed as the performance contribution to the outcome prediction $s(t) = \mathbf{w}_p \mathbf{f}_p(t)$. We used three performance features, computed on a trial-by-trial basis, namely the proportion of pauses, the proportion of correct key presses, and the proportion of wrong key presses that would have been correct in the normal UP orientation. We compared the model's prediction accuracy for a range of negative values for $\mathbf{w}_d$, as well as for $\mathbf{w}_d = 0$, which is equivalent to using only performance features to predict outcomes (see figure 2); we also compared these accuracies with that obtained when using difficulty as the only predictor of outcomes. Note that the accuracy of the models including performance features is a training accuracy (and therefore likely an overestimation), as these models were fitted on the same individual subject data on which accuracy was computed; this is not the case for the difficulty-only model.

As illustrated in figure 2, this comparison showed that difficulty alone is overall more predictive of

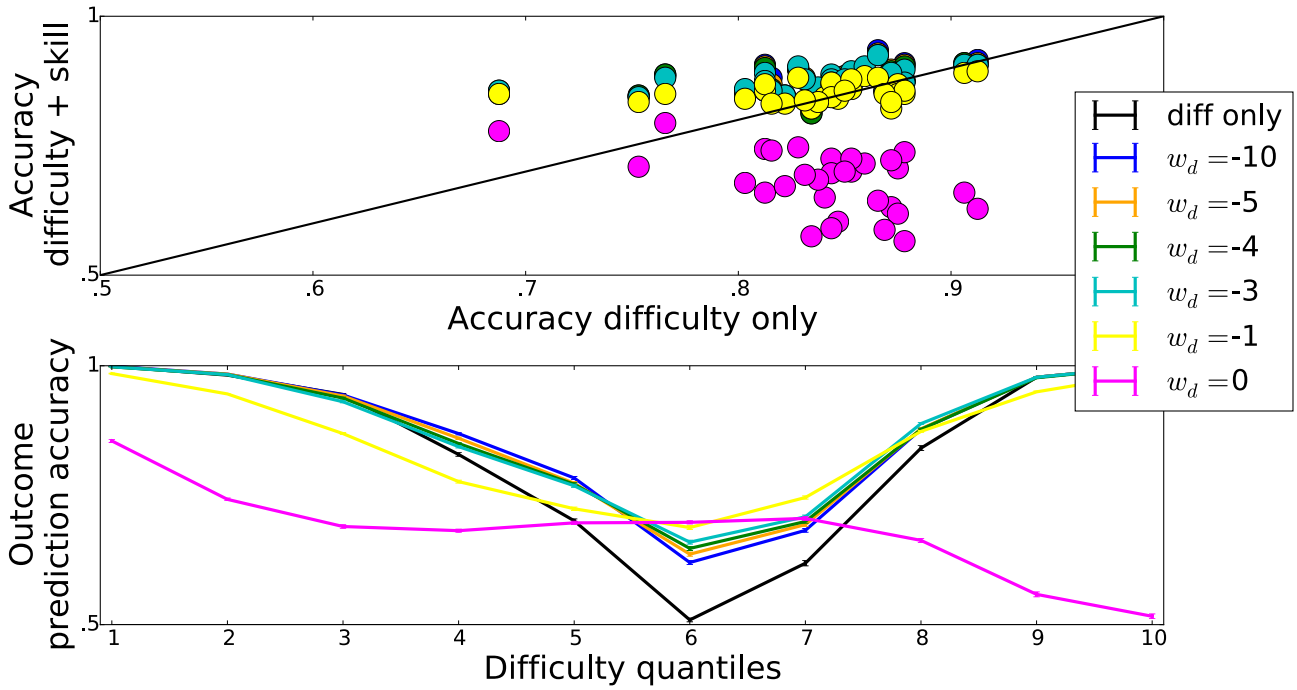

Figure 2: Accuracy for outcome prediction: difficulty only vs difficulty and skill models. Colours correspond to different $\mathbf{w}_d$ values in the skill and difficulty models; note that $\mathbf{w}_d = 0$ (purple) is equivalent to a model with skill only; black is used for the model with difficulty only. Top: overall accuracy; each dot represents one subject. Bottom: accuracy per difficulty level; mean $\pm$ s.e.m across subjects; difficulty was z-scored for each subject and discretised in 10-quantiles.

outcome than performance features alone, and that adding performance features to the difficulty-only model only marginally improves overall accuracy. Computing outcome prediction accuracy as a function of difficulty level provides a more detailed account of the models' performance, showing that there is only a narrow range of difficulty values for which difficulty alone fails to predict outcome, and where including performance features significantly improves prediction accuracy.

3 SAO skill model, quality of fit

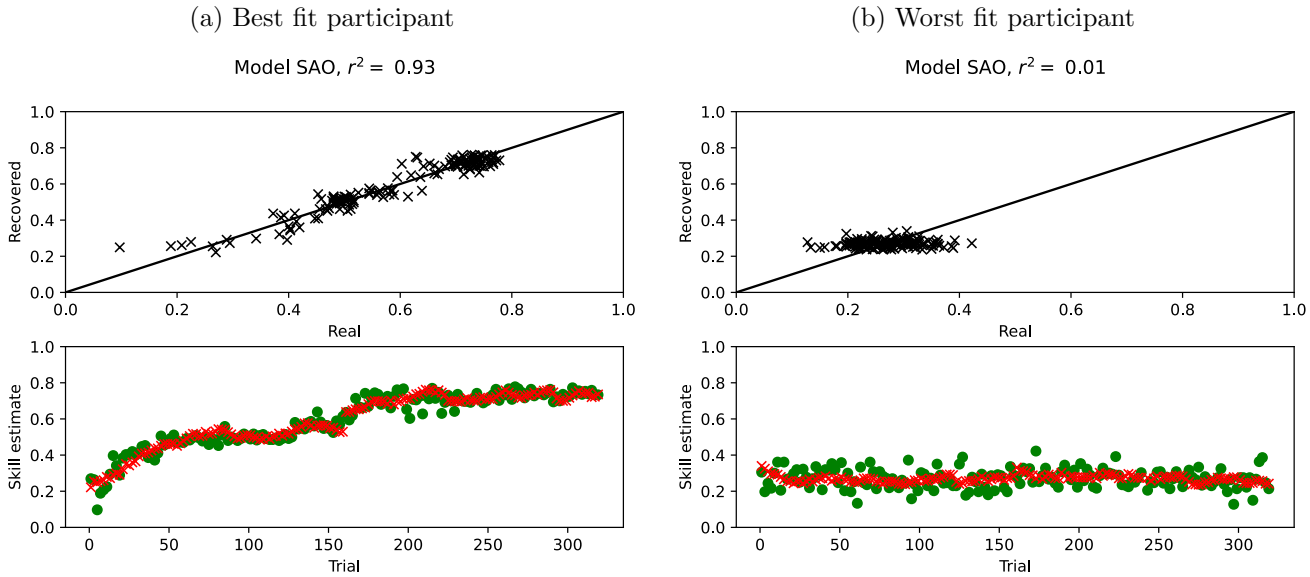

Figure 3: Example best (a) and worst (b) fit participants, SAO model of skill estimates, self condition. Top: participant responses vs underlying skill recovered with mean posterior parameters. Bottom: time series of participant responses and underlying skill recovered with mean posterior parameters

4 SAO skill model parameters

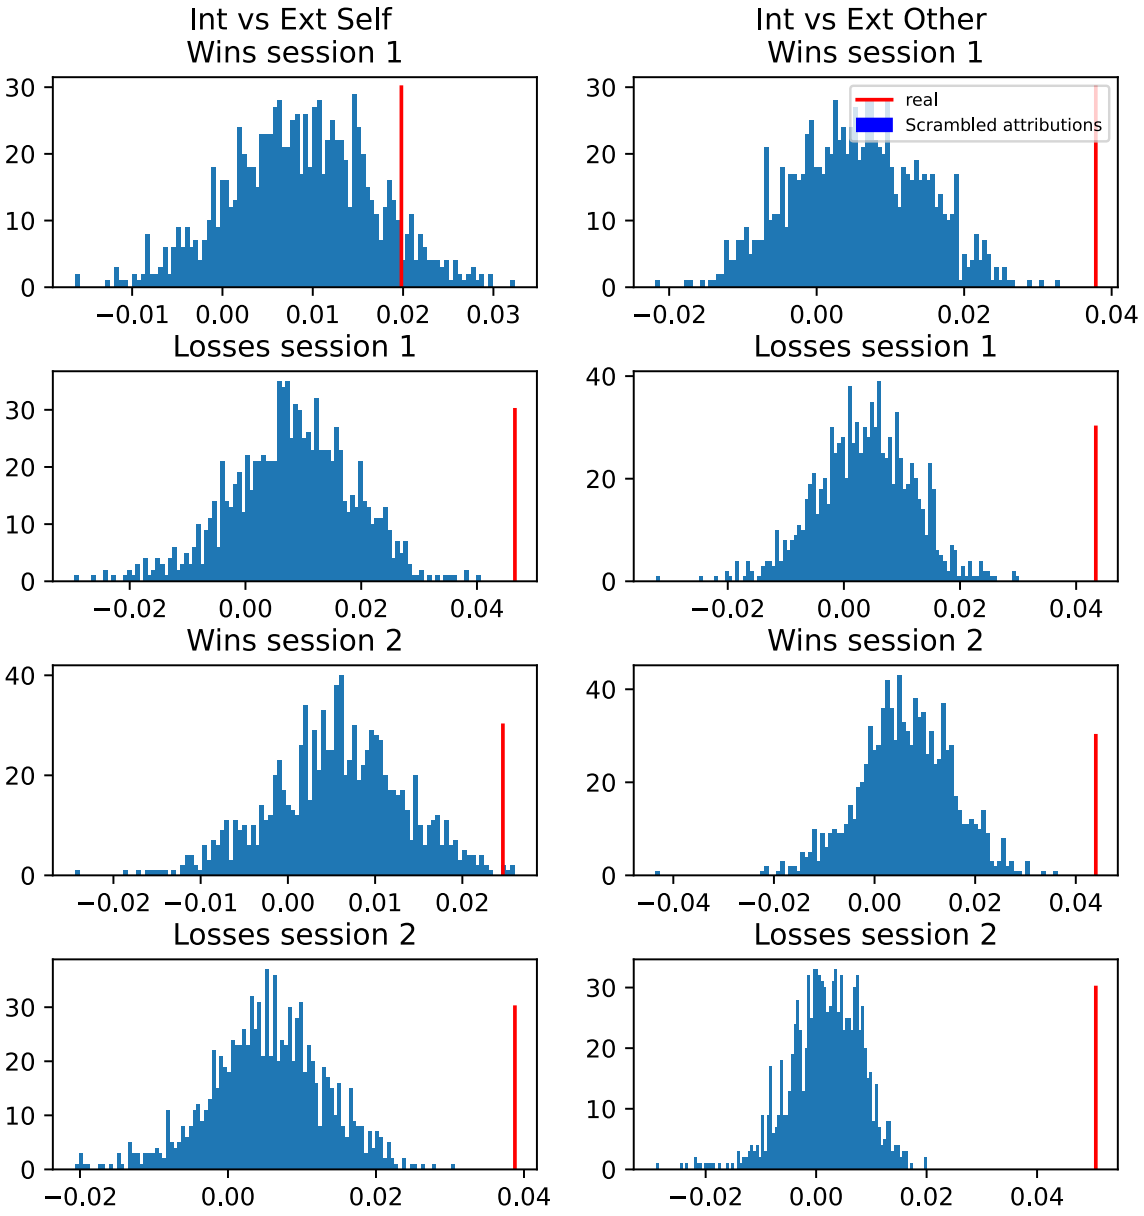

Figure 4: Attribution response shuffling: effect on difference between learning rates of the SAO model. Analyses were performed as follows: for each of 1000 permutations, attribution responses of each individual participant were shuffled and the SAO model was refitted. For each combination of outcome and session, the difference between the corresponding internal and external α parameters was averaged across participants. The resulting shuffle distribution is compared with the average difference obtained from fitting the real data. Left: self, right: other.

To test whether the attribution effects observed in model agnostic analyses were detectable in model parameters we performed permutation tests on mean posterior parameters from the winning model (Table 1). Furthermore, we refitted the model to data with scrambled attributions and compared the observed differences in learning rates to the ones obtained in the real data (Figure 4).

	Self		Other	
	Session 1	Session 2	Session 1	Session 2
Win	t = 2.41 p = 0.0056 d= 0.22	t=3.67 p=0.0002 d=0.28	t=4.45 p<1/5000 d=0.5	t=6.64 p<1/5000 d=0.69
Loss	t=5.88 p<1/5000 d=0.36	t=5.57 p<1/5000 d=0.36	t=5.98 p<1/5000 d=0.69	t=7.64 p<1/5000 d=0.84

Table 1: Permutation test results and effect sizes, comparisons between learning rates for internal vs external attributions, conditioned on outcome and session, SAO model. See Figure **??** in main text.

5 Model agnostic analyses of attributions: other

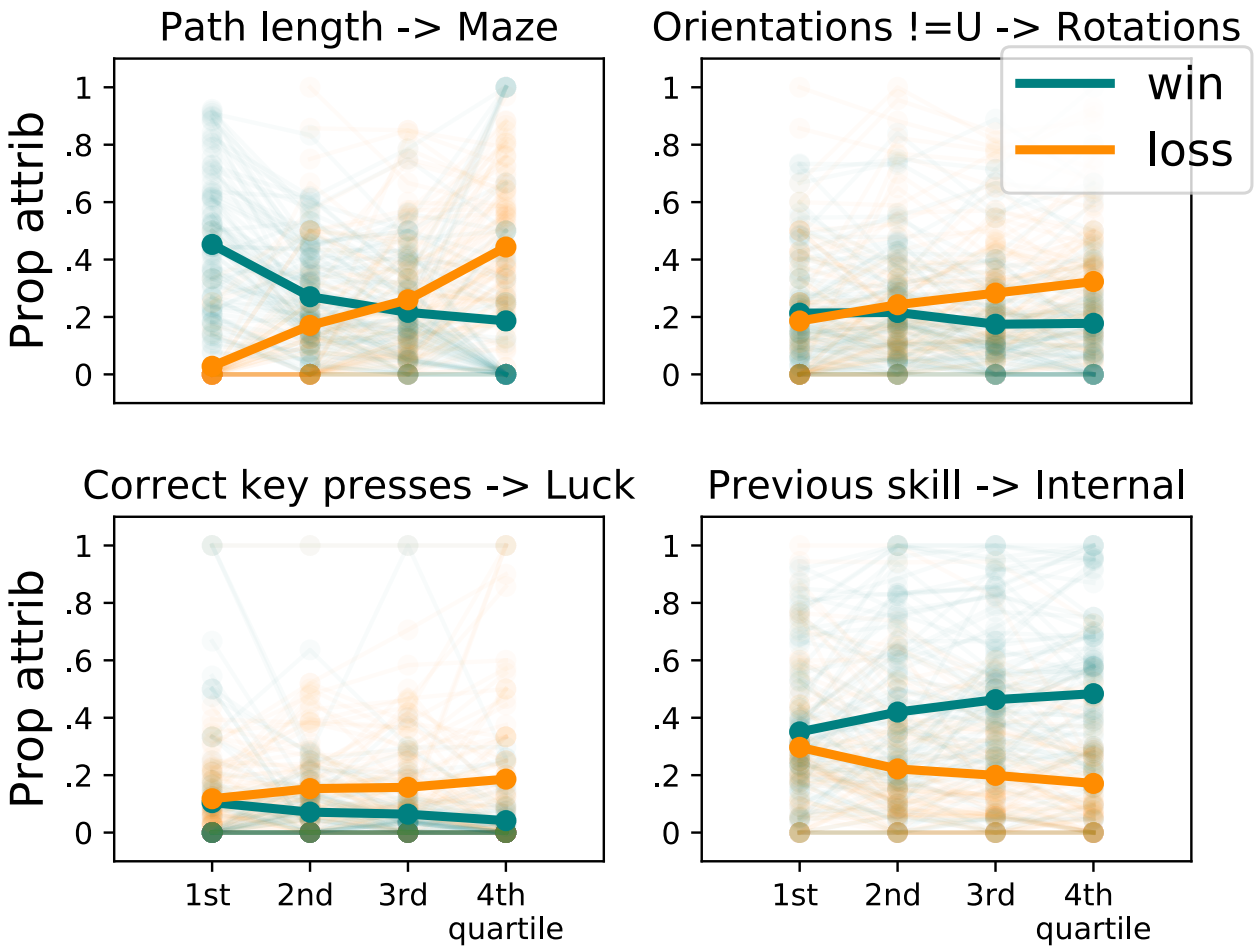

Figure 5: Features of interest and attributions summary other. Faded lines represents individual participants, bold lines represent mean ± s.e.m across participants. Orange: losses, teal: wins.

6 Model parameters, winning attribution model

To test wether feature effects extracted from the parameters of the winning models are significantly different from 0 we performed sign permutation tests on the mean across participants. Full results are in the tables below.

	Path length			
	Self		Other	
	Win	Loss	Win	Loss
Internal	m=0.11	m=-0.03	m=0.14	m=-0.04
	p<1/5000	p<1/5000	p<1/5000	p<1/5000
Maze	m=-0.15	m=0.12	m=-0.14	m=0.12
	p<1/5000	p<1/5000	p<1/5000	p<1/5000

Table 2: Effects of path length on internal attributions and attributions to Maze

	Unusual orientations			
	Self		Other	
	Win	Loss	Win	Loss
Internal	m=0.04	m=-0.01	m=0.05	m=-0.02
	p<1/5000	p<1/5000	p<1/5000	p<1/5000
Rotations	m=-0.03	m=0.04	m=-0.02	m=0.04
	p<1/5000	p<1/5000	p<1/5000	p<1/5000

Table 3: Effects of proportion of unusual orientations on internal attributions and attributions to Rotations

	Reported skill			
	Self		Other	
	Win	Loss	Win	Loss
Maze	m=-0.01	m=0.01	m=-0.03	m=0.01
	p=0.0014	p<1/5000	p<1/5000	p=0.04
Rotations	m=-0.02	m=0.01	m=-0.02	m=0.02
	p<1/5000	p<1/5000	p<1/5000	p<1/5000

Table 4: Effects of reported skill on external attributions

	Key press accuracy			
	Self		Other	
	Win	Loss	Win	Loss
Maze	m=0.03	m=0.02	m=0.03	m=0.04
	p<1/5000	p<1/5000	p<1/5000	p<1/5000
Rotations	m=-0.01	m=-0.07	m=-0.01	m=-0.04
	p=0.0088	p<1/5000	p=0.02	p<1/5000

Table 5: Effects of key press accuracy on external attributions