Contents

1 Pilot experiment **1**
Methods . . . 1
Results and Discussion . . . 2

2 Supplementary Methods **4**
Task instructions: Experiment 1 . . . 4
Task instructions: Experiments 2 - 3 . . . 4
Experiment 1: Model fitting for Both and None conditions . . . 5

3 Experiment 1 **5**
Behavioral results: Main task . . . 5
Post-test . . . 6

4 Experiment 2 **6**
Behavioral results: Main task . . . 6
Post-test . . . 6

5 Experiment 3 **7**
Behavioral results: Main task . . . 7
Post-test . . . 7

6 Experiments 2–3: Likelihood Estimation model **8**

7 Supplementary figures **9**

1 Pilot experiment

In this pilot experiment, we tested whether participants perform better with advice based on incomplete information than they would on their own.

Methods

Participants

50 participants were recruited for an online study on Amazon Mechanical Turk. All participants were paid $2 and a performance bonus of up to $2 proportional to their

final score. All participants had US IP addresses and provided informed consent in accordance with the requirements of the institutional review board.

Procedure

Participants played a card game where they chose whether to stay with a visible card or switch to a hidden card. There were two within-subject conditions. In the Solo condition, participants played alone; in the Advisor condition, participants received advice from an advisor who could only see the hidden card (same as Exp.1 Hidden condition in the main text). As in the main experiment, the advisor's responses were sampled using a logistic function:

$$P(A = \text{Switch}|C_H) = \frac{1}{1 + e^{(-m(C_H - C_{med}))}} \tag{1}$$

Where A is the advisor's advice, C_H is the value of the hidden card, and $m = 1.5$ is the steepness of the logistic curve. The median card value, $C_{med} = 4.5$, was subtracted from C_H to center the logistic function in the middle of the card range.

Participants played 4 blocks of 28 trials (2 Solo blocks and 2 Advisor blocks). Each trial was composed of two parts: participants first saw the visible card and (in the Advisor blocks) the advice and decided whether to switch or stay; second, participants were given feedback on how many points they had earned (Figure S1).

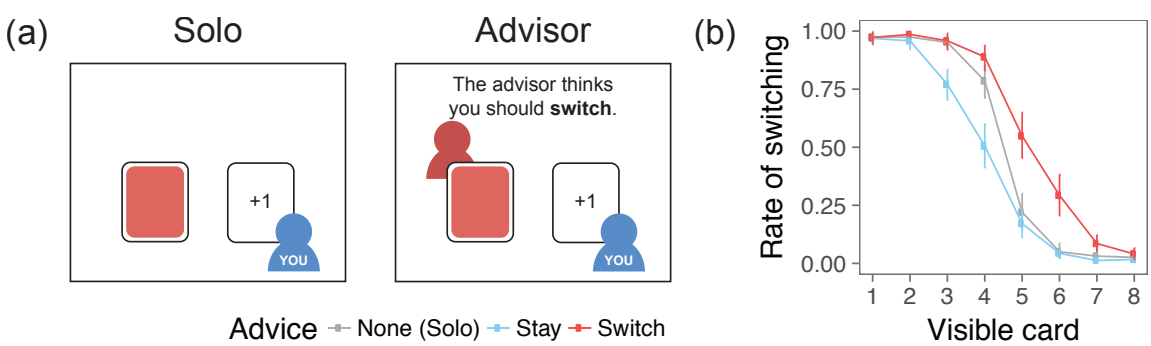

Figure S1: Pilot experiment results. (a) Schematic of choice screen. (b) Proportion of trials where participants chose to switch. Error bars denote bootstrapped 95% CI.

Results and Discussion

First, we asked whether participants followed the advisor's advice at all. We tested the effect of the advisor's advice on participants' responses using a logistic mixed-

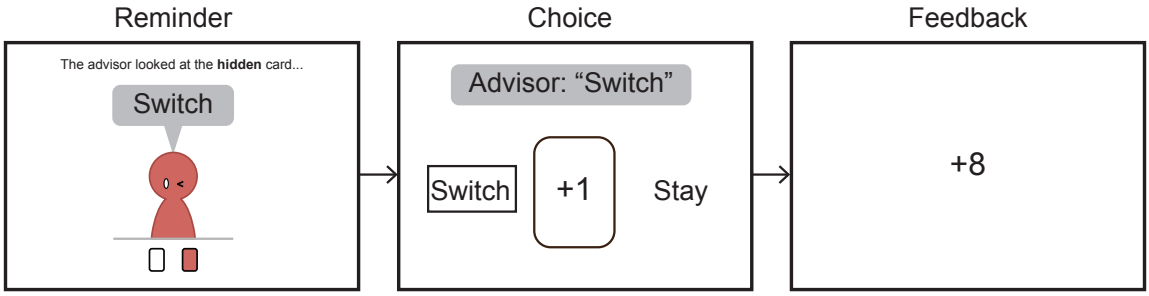

Figure S2: Experiment 1 procedure. Each trial consisted of three parts: first, participants were shown the advisor's advice and a picture to remind them of which cards the advisor could see (Reminder); participants were trained in advance on what each reminder picture meant. Next, as in Exp.2–3, participants saw the visible card and made their Choice, and they received Feedback on how many points they had earned.

effects model (response ~ visible_card + advice + 1|participant). In order to compare advice to the Solo blocks, we coded advice with three levels: "none" (our reference level, for None blocks), "stay" and "switch". Compared to their responses in the Solo blocks, participants switched more often when the advisor suggested to switch (mixed-effects logistic regression: $\beta = 9.66, z = 13.67, p < 0.001$) and switched less when the advisor suggested to stay ($\beta = -7.15, z = -10.13, p < 0.001$; Figure S1b).

Second, we asked whether participants earned more points, overall, in the Advisor condition than in the Solo condition. On average, participants earned 631.68 (SE = 2.65) points in the Advisor blocks, compared to 620.78 (SE = 2.90) in the Solo blocks (paired Wilcoxon test: $V = 1039.5, p < 0.001$).

These results provide preliminary evidence that participants used and benefited even from advice based on incomplete, non-overlapping knowledge. However, in the Advisor condition, participants could have benefited from a mere match between their own and advisor's advice, which was unavailable in the Solo condition because participants played alone without advice. In the main text, we compare conditions where participants did receive advice, and systematically vary the informational value of the advice by manipulating which cards the advisor can see (Exp.1) and how the advisor selects advice (Exp.2-3).

2 Supplementary Methods

Task instructions: Experiment 1

Links to the full task are available at osf.io/w37up/. Before the main task, participants first completed a brief tutorial to familiarize themselves with the task. After introducing the card game and keyboard controls for 'stay' and 'switch' choices (key mappings counterbalanced across participants), participants were introduced to the advisor. In Experiment 1, this portion of the tutorial emphasized that the advisor's access to information would vary from block to block. In three separate screens, participants were introduced to each of the conditions and asked a brief attention check question ("Which cards can the advisor see?"). We explained each condition in a fixed order (None, Hidden, Both). In each condition, they saw a prompt with a schematic of the advisor.

> None: "Sometimes, the advisor will be shown no cards. In these cases, the advisor will have to say something, but they have not seen any of the cards before giving you advice."
>
> Hidden: "Sometimes, the advisor will be shown the hidden card. In these cases, the advisor can't see your card, but they can see the card that is hidden to you. They will give you advice after having seen the card that you can't see."
>
> Both: "Sometimes, the advisor will be shown both cards. In these cases, the advisor will give you advice after having seen both your card and the hidden card. In these cases, the advisor will always know which card is best."

Task instructions: Experiments 2 - 3

In Experiments 2-3, the tutorial emphasized that the advisor could only see the hidden card. In these experiments, the advisor was given a name, Alice, to simplify the wording of the task. Participants were told which cards Alice could see using the following prompt:

> "Alice cannot see the card that is face-up to you. She can only see the hidden card. Based on what's on the hidden card, she will suggest whether you should stay with your own card, or switch and take the hidden card."

As in Exp.1, participants were asked to indicate the card that Alice can (and cannot) see.

In all experiments, after the main task, participants completed a brief survey where they rated their use of the advice ("How often did you follow the advisor's advice?"), trust in the advice ("I trusted the advisor's advice, given the information available to them."), and benefit from the advice ("Thanks to the advisor's advice, I scored higher than I would have if I had been playing on my own."; 1 = Strongly disagree, 5 = Strongly agree). Due to an error in data collection, some participants' responses to this question were lost in Experiment 1 (total n: None = 50, Hidden = 44, Both = 39) and Experiment 2 (total $n = 149$).

Experiment 1: Model fitting for Both and None conditions

In Experiment 1, both models assumed that the advisor was perfectly accurate in the Both condition and completely random in the None condition. Thus, we first inferred the posterior distribution over model parameters (Mental-State Reasoning: m_A, b_A; Accuracy Heuristic: acc; both models: m_L) using participants' responses in only the Hidden condition.

We then sampled from the posterior distribution over model parameters to generate predicted responses across all three conditions. In the Mental-State Reasoning model, we generated predicted responses for None and Both conditions by assuming that the likelihood of the observed advice varied, depending on the advisor's access to information. In the None condition, $P(A|C_H) = .5$ for all values of C_H. In the Both condition, $P(A|C_H) = 1$ if $C_H > C_V$ and is 0 otherwise. To generate predicted responses of the Accuracy Heuristic model, we set fixed values of acc for each condition (None: acc = .5, Both: acc = 1).

3 Experiment 1

Behavioral results: Main task

We tested the effect of advice on participants' responses ("stay" or "switch") in each condition using a mixed-effects logistic regression (response $\sim$ advice $\times$ condition + card + (1|participant)). Conditions were coded using orthogonal polynomial coding, where each condition was ordered according to the amount of information available to the advisor (None < Hidden < Both).

Collapsing across all conditions, participants were more likely to switch when the advisor told them to switch ($\beta_{\text{advice}} = 3.52, z = 30.03, p < 0.001$) and less likely to

switch as the value of the visible card increased ($\beta_{\text{card}} = -1.53, z = -39.00, p < 0.001$). Critically, we observed an interaction between the advisor's access to information and the advice ($\beta_{\text{advice}\times\text{advisor}} = 5.14, z = 24.86, p < 0.001$): participants' choices were more strongly influenced by the advisor's advice as the advisor's access to information increased.

Post-test

Participants' responses to the post-test survey verified that they used advice differently across conditions (Figure S4a). We modeled participants' ratings on each question as a mixed-effects linear regression (rating condition+(1|participant)). Condition was coded using an orthogonal polynomial contrast, where each condition was ordered according to the advisor's access to information (None < Hidden < Both). Participants reported that they followed advice more often (Use of advice: $\beta = 1.93, t(93.74) = 16.84, p < 0.001$), trusted the advice more (Trust in advice: $\beta = 1.89, t(91.50) = 13.13, p < 0.001$), and scored more points (Benefit from advice: $\beta = 19.94, t(89.55) = 12.95, p < 0.001$) when the advisor could see more cards.

4 Experiment 2

Behavioral results: Main task

We tested the effect of advice on participants' responses ("stay" or "switch") in each condition using a mixed-effects logistic regression (response advice + card + (1|participant)). In the Random condition, participants largely ignored the advice ($\beta_{advice} = 0.14, z = 1.64, p > 0.10$), mirroring results from Exp.1 None condition. In the Helpful condition, participants' decisions were influenced by the advice ($\beta_{advice} = 1.99, z = 21.02, p < 0.001$), despite her partial knowledge and imperfect advice, mirroring results from Exp.1 Hidden condition. In the Opposite condition, participants showed a small but significant bias to choose the opposite of the advisor's advice ($\beta_{advice} = -0.37, z = -4.19, p < 0.001$, Figure 3, main text).

Post-test

Participants' responses to the post-test survey suggested that they consistently viewed the Helpful advisor as helpful, and that they had a weak tendency to value the Opposite advisor's advice more than the Random advisor's (Figure S4b). We modeled participants' ratings on each question using linear regression (rating $\sim$ advisor +

(1|participant)). We used dummy coding for different levels of advisor, such that the Helpful and Opposite advisor were both compared to the Random advisor. Participants consistently reported that they followed the Helpful advisor's advice more often ($\beta = 0.84, t = 5.38, p < 0.001$), trusted the advice more ($\beta = 1.38, t = 7.03, p < 0.001$), and reported that they did better with the Helpful advisor's advice than they would have on their own ($\beta = 1.72, t = 7.36, p < 0.001$). By contrast, participants reported that they trusted advice from the Opposite advisor slightly less ($\beta = -0.34, t = -2.16, p = 0.03$) and followed it slightly less often ($\beta = -0.41, t = -2.07, p = 0.04$) than advice from the Random advisor; in addition, participants did not report that they had done better on the task with the Opposite advisor's advice ($\beta = 0.16, t = 0.67, p > 0.40$).

5 Experiment 3

Behavioral results: Main task

In Exp.3, the advisors differed in their tendency to recommend switching (conservative or risky). If participants take the advisor's strategy into account, then they would be more likely follow the Conservative advisor's advice when it suggested to switch, because the advisor only suggested to switch when the hidden card had a very high value (+7 or +8); similarly, participants would be more likely to follow the Risky advisor's advice when it suggested to stay. In addition to the predictions of our cognitive models (see main text), we tested for the presence of an interaction between the advisor and the advice using a mixed-effects logistic regression (follow $\sim$ advisor $\times$ advice + card + (1|participant)). Participants' responses (follow) were coded as a "success" if their action matched the advisor's advice.

As predicted, in the behavioral data, we observed an interaction between the advisor and the advice (advisor $\times$ advice: $\beta = 1.15, z = 14.78.p < 0.001$; Figure S6a). Only the Mental-State Reasoning model predicted this crossover interaction (Figure S6b–c). Conversely, the Accuracy Heuristic model predicts that participants would ignore the Risky advisor's advice altogether and go against the Conservative advisor's advice.

Post-test

Although both advisors were equally accurate, participants' responses suggested that they viewed the Conservative advisor as more reliable than the Risky advisor. Compared to the Risky advisor, participants reported that they followed the

Conservative advisor's advice more often ($\beta = 0.42, t = 3.03, p = 0.002$) and trusted the advice more ($\beta = 0.51, t = 2.68, p = 0.089$). In addition, participants had a weak tendency to report that they did better with the Conservative advisor's advice ($\beta = 0.49, t = 2.18, p = 0.03$).

6 Experiments 2–3: Likelihood Estimation model

To generate predictions of the Mental-State Reasoning model reported in the main text, we estimated the two parameters that describe participants' beliefs about how the advisor selects advice (m_A and b_A) based on participants' choices.

In a separate analysis reported below, we estimated the same parameters based on participants' responses in the Likelihood Estimation task. The goal was to verify that the parameter estimates from the choice data indeed reflect participants' beliefs about the advisor.

We estimated m_A and b_A from participants' responses in each condition (Exp.2: Helpful, Random, Opposite; Exp.3: Conservative, Risky) and each phase (Before, After) separately. As described in the main text, parameter values for each participant were sampled from a group-level distribution that was shared across participants ($m_A \sim \mathcal{N}(\mu_m, \sigma_m)$, $b_A \sim \mathcal{N}(\mu_b, \sigma_b)$). We used the same priors over group-level hyperparameters as in the Mental-state Reasoning model ($\mu_b, \mu_m \sim \text{Uniform}(-20, 20)$; all other parameters $\sim \text{Uniform}(0, 40)$). We estimated the posterior distribution over model parameters by collecting 50,000 samples from an MCMC chain (discarding the first 2,000 samples) using the Metropolis-Hastings algorithm.

The results of this model are shown in Figure S8. In Experiment 2, the advisors varied in the true value of m_A (Helpful: 1.5, Random: 0, Opposite: -1.5) but not b_A (all conditions: 0). Before playing the card game, participants had similar beliefs about the advisor across all condition; after the task, the posterior distributions of μ_m diverged across conditions. In Experiment 3, the advisors varied in the true value of b_A (Risky: 10, Conservative: -10) but not m_A (all conditions: 5). Because we provided a hint to participants about the advisor's tendency to say "switch" or "stay", the posterior distributions μ_b do not overlap even before participants played the card game. By the end of the task, participants' beliefs about the advisor have become even more distinct.

Taken together, we recovered similar parameter values based on participants' choices in the main card game and their explicit report in the Likelihood Estimation task (see also Figure S7). These results provide additional evidence that participants' choices were consistent their beliefs about how the advisor selects advice.

7 Supplementary figures

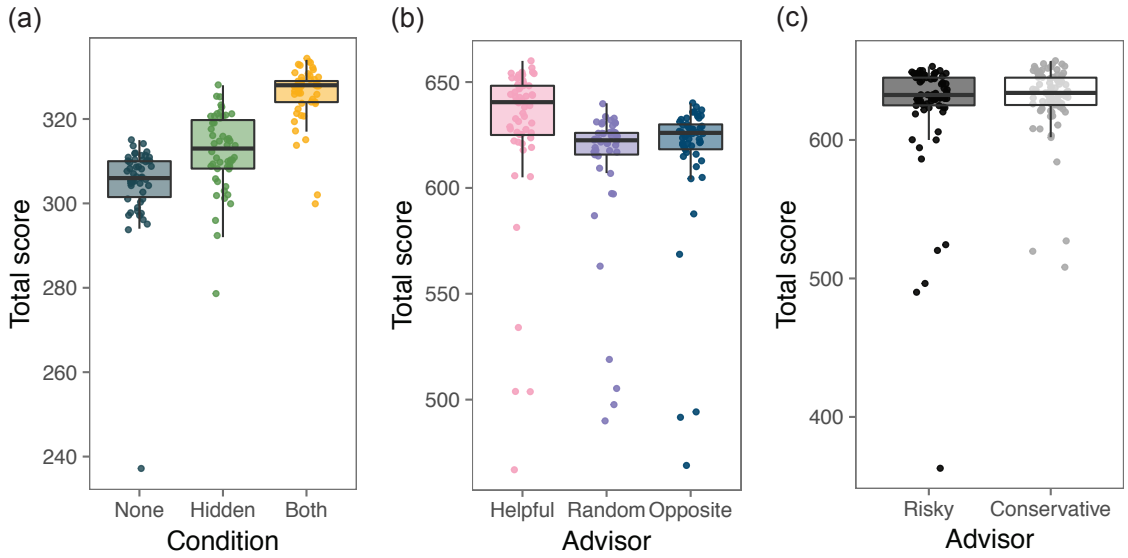

Figure S3: Total points earned in Experiments (a) 1 (b) 2 and (c) 3.

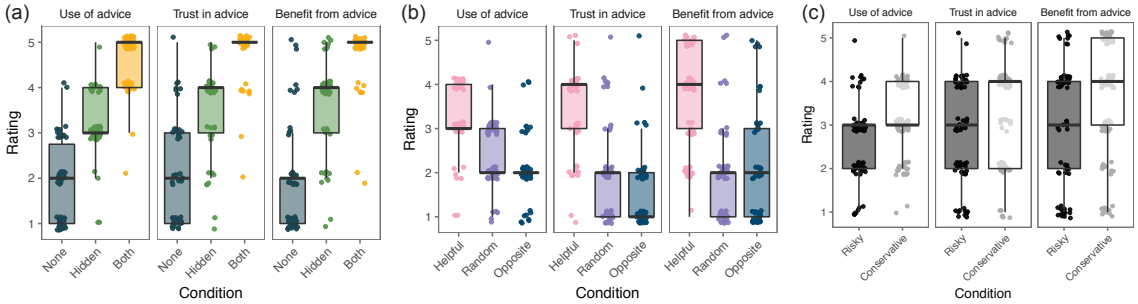

Figure S4: Post-test advisor ratings in (a) Exp.1, (b) Exp.2, and (c) Exp.3.

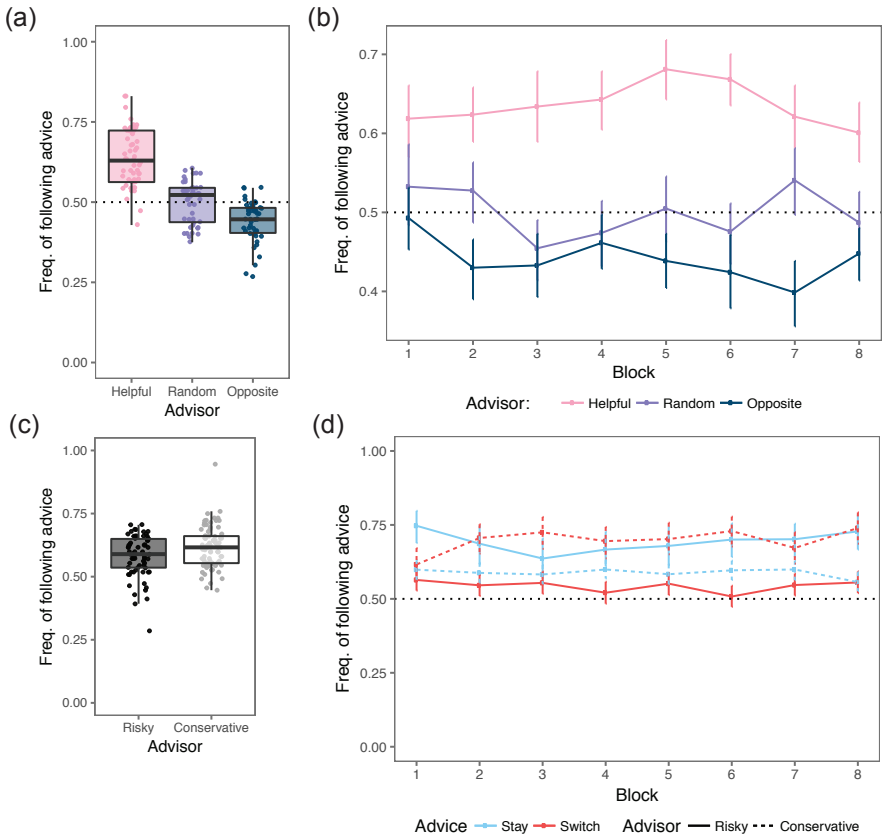

Figure S5: Rates at which participants followed advice in Experiments 2–3, (a,c) collapsed across blocks and (b,d) split by block.

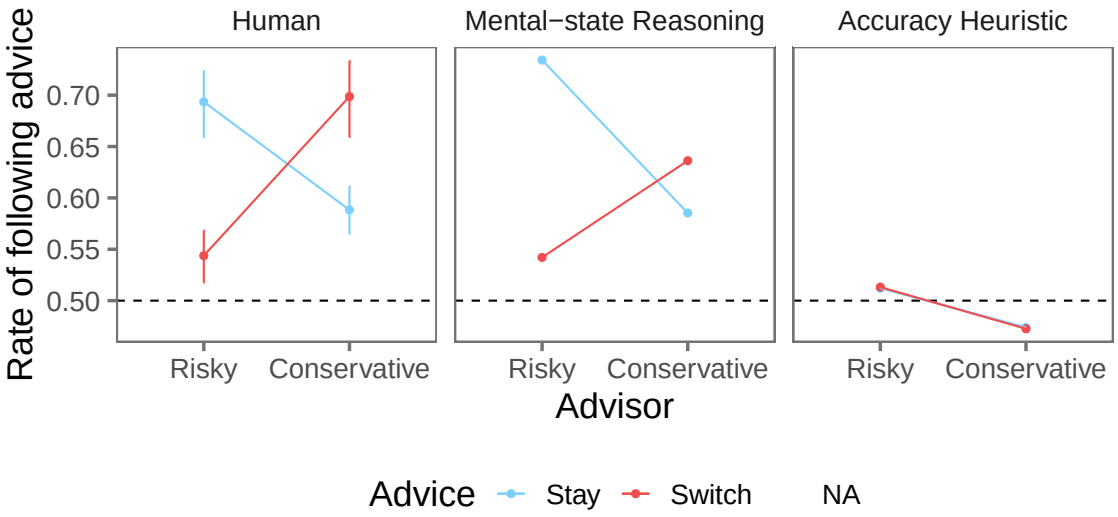

Figure S6: Interaction effects in Experiment 3. Left: Participants were more likely to follow the Risky advisor's advice when she suggested to stay and the Conservative advisor's advice when she suggested to switch. Middle, Right: The Mental-State Reasoning model predicts this crossover interaction, while the Accuracy Heuristic does not.

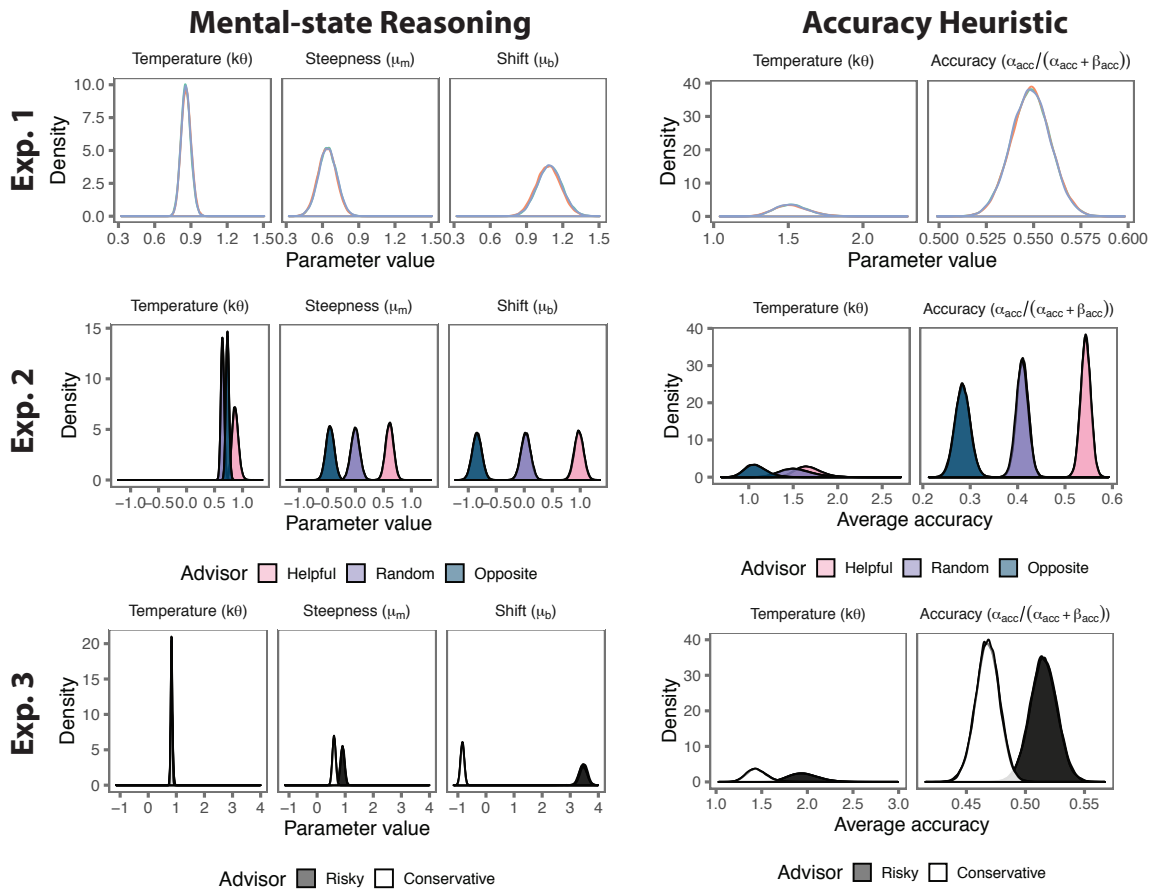

Figure S7: Posterior distribution over group-level parameters. Left: Mental-state Reasoning model. Plots show the posterior distribution over the group-level parameters μ_m, μ_b, and average temperature (k/θ). Right: Accuracy Heuristic model. Plots show the posterior distribution over average temperature (k/θ) and accuracy $(= \alpha_{\text{acc}}/(\alpha_{\text{acc}} + \beta_{\text{acc}}))$

.

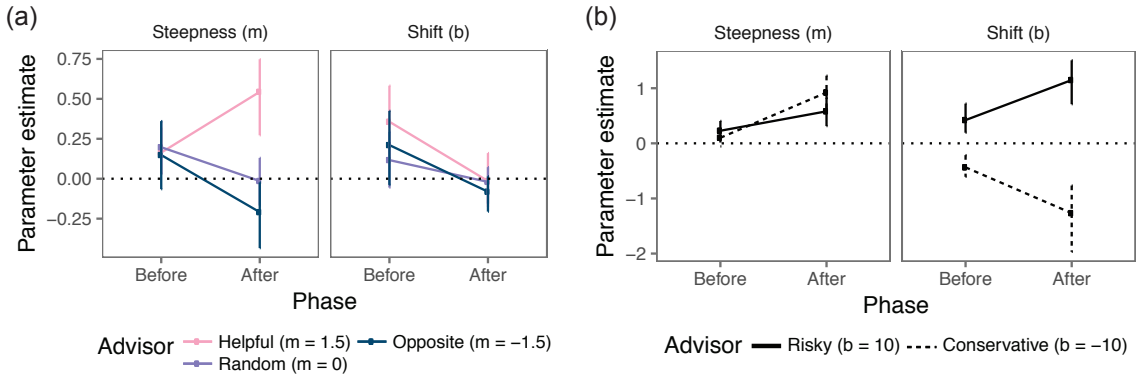

Figure S8: Posterior distribution over the group-level parameters μ_m and μ_b, based on participants' responses on the likelihood estimation task before playing the card game (Before) and after all trials of the card game (After). Error bars indicate 95% credible intervals.