

Supplemental Appendix for “Coordination with Cognitive Noise”

by Cary Frydman and Salvatore Nunnari

A Alternative Performance Objective in Model with Endogenous Efficient Coding from Appendix A

Here we revisit the assumption about efficient coding in the model from Appendix A. The specific performance objective that we assume there is only one of several plausible specifications (Ma and Woodford, 2020). In particular, there are other possible objective functions that players may have, besides minimizing the mean squared error of the estimate of θ . For example, a prominent alternative efficient coding objective from the literature on sensory perception is to maximize the mutual information between the state and its noisy internal representation. In the proof of Proposition 3, we confirm that the coding rule we use in our model is robust to this alternative objective.

Yet another alternative objective that has been examined in the economics literature is maximization of expected reward. In this section, we show that the result in Proposition 3 is robust to using this alternative objective function. Specifically, we maintain the constraints in Assumption 4 and we analyze a two-stage game. In the first stage, each player optimally chooses, simultaneously and independently, the parameters of the encoding function. In the second stage, players choose strategies in the simultaneous move game, conditional on their chosen encoding function from the first stage. We show that the optimal encoding function still takes the form characterized in Proposition 3. Thus, our theoretical predictions from Appendix A are robust to three performance objectives: (i) minimizing mean squared error of the estimate of θ , (ii) maximizing mutual information between the noisy internal representation and θ and (iii) maximizing expected reward.

Assumption 6 (Alternative Performance Objective) *Players choose the encoding function which maximizes their expected reward in the simultaneous move game.*

Consider the following two-stage game: in stage 1, each player $i = \{1, 2\}$ chooses simultaneously and independently the parameters of his encoding function, (ξ_i, ψ_i) , to maximize the performance objective in Assumption 6 under the constraints in Assumption 4; in stage 2, players participate to the simultaneous move game endowed with the encoding functions chosen in the previous stage. We solve this game by backward induction.

Stage 2: Simultaneous move Game (with Exogeneous Encoding Functions)

For each player $i = \{1, 2\}$, we have $S_i|\theta \sim N(m_i(\theta), \sigma_S^2)$, where $m_i(\theta) = \xi_i + \psi_i\theta$.

Consider the transformed internal representation $Z_i = (S_i - \xi_i)/\psi_i$. We have:

$$Z_i|\theta \sim N(\theta, \beta_i^2)$$

where $\beta_i = (\sigma_S/\psi_i)$.

Proposition 6 *Suppose Assumptions 3, 4, 6, and $\mu_\theta = (a + b)/2$. Regardless of σ_θ , σ_S , (ξ_1, ψ_1) , and (ξ_2, ψ_2) , there exists an equilibrium of the game where each player invests if and only if $Z_i \leq \mu_\theta$. Moreover, if $\frac{\sigma_\theta^2 \sqrt{\beta_i^2(2\sigma_\theta^2 + \beta_i^2)}}{(b-a)\beta_i^2 \sqrt{\sigma_\theta^2 + \beta_i^2}} > \frac{1}{\sqrt{2\pi}}$ for all $i = \{1, 2\}$, this is the unique monotone equilibrium of the game.*

Proof. Since the likelihood function of Z_i is conjugate to the prior distribution of θ , we have a closed form solution for the distribution of player i 's posterior beliefs over θ . In particular, player 1's posterior distribution of θ given Z_1 is

$$\theta|Z_1 \sim \mathcal{N}\left(\frac{\beta_1^2 \mu_\theta + \sigma_\theta^2 Z_1}{\sigma_\theta^2 + \beta_1^2}, \frac{\sigma_\theta^2 \beta_1^2}{\sigma_\theta^2 + \beta_1^2}\right)$$

Thus, we have:

$$EU[\text{Not Invest}|Z_1] = E[\theta|Z_1] = \frac{\beta_1^2 \mu_\theta + \sigma_\theta^2 Z_1}{\sigma_\theta^2 + \beta_1^2}$$

On the other hand, player 1's expected utility from investing is

$$EU[\text{Invest}|Z_1] = a + (b - a)\Pr[\text{Opponent Invests}|Z_1]$$

Assume player 1 believes his opponent uses a monotone strategy with threshold k_2 . In this case, player 1's expectation that the opponent invests is $\Pr[Z_2 \leq k_2|Z_1]$. Player 1's belief over the distribution of Z_2 conditional on Z_1 is:

$$Z_2|Z_1 \sim \mathcal{N}\left(\frac{\beta_1^2 \mu_\theta + \sigma_\theta^2 Z_1}{\sigma_\theta^2 + \beta_1^2}, \frac{\sigma_\theta^2 (\beta_1^2 + \beta_2^2) + \beta_1^2 \beta_2^2}{\sigma_\theta^2 + \beta_1^2}\right)$$

Thus, we have:

$$\Pr[Z_2 \leq k_2|Z_1] = \Phi\left(\frac{k_2 (\sigma_\theta^2 + \beta_1^2) - \beta_1^2 \mu_\theta - \sigma_\theta^2 Z_1}{\sqrt{\sigma_\theta^2 + \beta_1^2} \sqrt{\sigma_\theta^2 (\beta_1^2 + \beta_2^2) + \beta_1^2 \beta_2^2}}\right)$$

where $\Phi(\cdot)$ is the cumulative distribution of the standard normal.

Player 1's best response is to invest if and only if

$$\frac{\beta_1^2 \mu_\theta + \sigma_\theta^2 Z_1}{\sigma_\theta^2 + \beta_1^2} \leq a + (b - a) \Phi \left(\frac{k_2 (\sigma_\theta^2 + \beta_1^2) - \beta_1^2 \mu_\theta - \sigma_\theta^2 Z_1}{\sqrt{\sigma_\theta^2 + \beta_1^2} \sqrt{\sigma_\theta^2 (\beta_1^2 + \beta_2^2) + \beta_1^2 \beta_2^2}} \right)$$

Assume $k_2 = \mu_\theta$. We want to show that player's best response is to use the same cutoff. In this case, player 1's best response is to invest if and only if

$$E \frac{\beta_1^2 \mu_\theta + \sigma_\theta^2 Z_1}{\sigma_\theta^2 + \beta_1^2} \leq a + (b - a) \Phi \left(\frac{\sigma_\theta^2 (\mu_\theta - Z_1)}{\sqrt{\sigma_\theta^2 + \beta_1^2} \sqrt{\sigma_\theta^2 (\beta_1^2 + \beta_2^2) + \beta_1^2 \beta_2^2}} \right)$$

First, note that the LHS is a convex combination of μ_θ and Z_1 and that, thus, it is a) equal to μ_θ when $Z_1 = \mu_\theta$, b) smaller than μ_θ when $Z_1 < \mu_\theta$, and c) larger than μ_θ when $Z_1 > \mu_\theta$. Second, remember that $\mu_\theta = (a + b)/2$ and note that the RHS is a) equal to μ_θ when the argument of $\Phi(\cdot)$ is 0 (that is, when $Z_1 = \mu_\theta$, since the denominator is strictly positive); b) larger than μ_θ when the argument of $\Phi(\cdot)$ is strictly positive (that is, when $Z_1 < \mu_\theta$), and c) smaller than μ_θ when the argument of $\Phi(\cdot)$ is strictly negative (that is, when $Z_1 > \mu_\theta$). This means that, when player 2 invests if and only if $Z_2 \leq k_2 = \mu_\theta$, then player 1's best response is to invest if and only if $Z_1 \leq \mu_\theta$. This proves that there exists an equilibrium where both players use a monotone strategy with cutoff equal to μ_θ for any value of (ξ_1, ψ_1) , (ξ_2, ψ_2) , σ_S and σ_θ . Finally, to show that, when the condition in the statement of the proposition is satisfied, this is the unique equilibrium of the game, we can use the same steps in the proof of Proposition 4 to show that the best response mapping is a contraction (and that, thus, we can apply the contraction mapping theorem). In particular, it is sufficient to show that the derivative of the best response function of player 1 with respect to k_2 and the derivative of the best response function of player 2 with respect to k_1 have both an absolute value strictly smaller than 1. ■

Stage 1: Encoding Function Choice

When deriving the optimal choice of the encoding function in stage 1, we assume that, in stage 2, players use the cutoff strategy in the (unique) equilibrium from Proposition 6.

Proposition 7 *Suppose Assumptions 3, 4, 6, and $\mu_\theta = (a + b)/2$. The optimal encoding function is the same for both players and is given by $m^*(\theta) = \xi^* + \psi^* \theta = -\frac{\Omega \mu_\theta}{\sigma_\theta} + \frac{\Omega}{\sigma_\theta} \theta$.*

Proof. In stage 2, each player $i = \{1, 2\}$ invests if and only if $Z_i \leq \mu_\theta$. Given the conditional

distribution of Z_i , the probability player i invests for a given θ and encoding function is

$$\mathbb{P}_i(\text{Invest}|\theta, \psi_i) = \Phi\left(\frac{\mu_\theta - \theta}{\sigma_S/\psi_i}\right)$$

Thus, the expected utility player i gets from the game with a given value of θ is

$$\begin{aligned} EU_i(\theta, \psi_i) &= \mathbb{P}_i(\text{Invest}|\theta, \psi_i) (a + \mathbb{P}_{-i}(\text{Invest}|\theta, \psi_{-i})(b - a)) + (1 - \mathbb{P}_i(\text{Invest}|\theta, \psi_i))\theta \\ &= \theta + \Phi\left(\frac{\mu_\theta - \theta}{\sigma_S/\psi_i}\right) \left(a + \Phi\left(\frac{\mu_\theta - \theta}{\sigma_S/\psi_{-i}}\right) (b - a) - \theta\right) \end{aligned}$$

where we use $-i$ to denote i 's opponent. How does this expected utility change with ψ_i (taking ψ_{-i} as given)?

$$\frac{\partial EU_i(\theta, \psi_i)}{\partial \psi_i} = \phi\left(\frac{\mu_\theta - \theta}{\sigma_S/\psi_i}\right) \left(\frac{\mu_\theta - \theta}{\sigma_S}\right) \left(a + \Phi\left(\frac{\mu_\theta - \theta}{\sigma_S/\psi_{-i}}\right) (b - a) - \theta\right) \quad (12)$$

Since $\phi(\cdot)$ is strictly positive for any argument, the sign of equation (12) is determined by the product of its second and third term. First, note that the second term is a) equal to 0 when $\theta = \mu_\theta$, b) strictly positive when $\theta < \mu_\theta$ and c) strictly negative when $\theta > \mu_\theta$. Second, note that — since $\mathbb{P}_{-i}(\text{Invest}|\theta, \psi_{-i})$ is greater than 1/2 if and only if $\theta < \mu_\theta$ and $\mu_\theta = (a + b)/2$ — the third term is a) strictly positive when $\theta < \mu_\theta$ and b) strictly negative when $\theta > \mu_\theta$. This means that the product of the second and third term of equation (12) is always positive, with the exception of the case when $\theta = \mu_\theta$, in which case it is 0.

We have shown that the expected payoff in a game with a given θ is strictly increasing in ψ_i for any value of $\theta \neq \mu_\theta$ and it is constant in ψ_i for $\theta = \mu_\theta$. This means that, from an ex-ante perspective (that is, when a player knows the distribution of θ but does not know its actual realization), each player's expected reward from the simultaneous move game — that is, $EU_i(\psi_i) = \int EU_i(\theta, \psi_i) f(\theta) d\theta$ — is strictly increasing in ψ_i . Therefore, it is desirable to make ψ_i as large as possible consistent with the power constraint. When the distribution of θ is normal, we have

$$E[m^2] = \xi^2 + \psi^2 E[\theta^2] + 2\xi\psi E[\theta] = (\xi + \psi\mu_\theta)^2 + \psi^2 \sigma_\theta^2 \leq \Omega^2$$

The largest value of ψ consistent with this constraint in Assumption 2 is achieved when

$$\xi = -\psi\mu_\theta, \psi = \frac{\Omega}{\sigma_\theta}$$

Thus, $m^*(\theta) = -\frac{\Omega}{\sigma_\theta}\mu_\theta + \frac{\Omega}{\sigma_\theta}\theta$. ■

B Experiment on Awareness of Cognitive Noise

Here we report results from an additional experiment that is designed to investigate whether subjects are aware of their own imprecision and the imprecision of others. If subjects are not aware of the cognitive noise of others, then this would shut down the channel that generates strategic uncertainty in our model, which is key to generating the unique threshold equilibrium.

Experimental Design

Our method for studying awareness of imprecision is to create a simplified version of the coordination game experiment in the main text, but one that retains the core individual decision-making prediction that subjects play a threshold strategy. We employ a task from the numerical cognition literature where subjects are incentivized to quickly and accurately classify whether a two-digit number is larger or smaller than the number 55. Note that this threshold strategy is identical to the equilibrium strategy in Experiment 1; the important difference is that here, we exogenously impose the strategy on subjects without any strategic considerations or equilibrium requirements. We then incentivize subjects to report beliefs about errors in their own classification and in the classification of others. These beliefs are the main object of study in this experiment.

We recruit 300 subjects from Prolific who did not participate in Experiment 1 or Experiment 2. We pay subjects 1 GBP for completing the study, in addition to earnings from three phases of the experiment. In Phase 1, on each of 150 rounds, subjects are incentivized to quickly and accurately classify whether a two-digit Arabic numeral on the experimental display screen is larger or smaller than 55. Subjects earn $(1.5 \times \text{accuracy} - 1 \times \text{speed})$ GBPs, where ‘accuracy’ is the percentage of trials where the subject classifies the number correctly, and ‘speed’ is the average response time in seconds.¹ As in Experiment 1, there are two conditions, and the only difference across conditions is the distribution from which the two-digit Arabic numeral (which we again denote by θ) is drawn. We use the same two distributions as in Experiment 1: in the high volatility condition, $\theta \sim \mathcal{N}(55, 400)$, and in the low volatility condition, $\theta \sim \mathcal{N}(55, 20)$. We then round each value of θ to the nearest integer and re-draw if the rounded integer is less than 11 or greater than 99 (again, to ensure that each number contains exactly two digits).

We note that one difference in incentives compared to those in Experiment 1 involves decision speed. Here, we penalize subjects for the time it takes them to respond. The reason we impose the speed incentive comes from the well known “speed-accuracy tradeoff”

¹The experimental instructions are available in Supplemental Appendix C.

in perceptual decision-making: one can obtain higher accuracy in classification as decision speed slows down. Thus, in order to increase statistical power to detect how accuracy differs for values of θ close and far from the threshold, we jointly reward speed and accuracy.

In Phase 2 of the experiment, we incentivize subjects to report beliefs about others' performance in the task. Furthermore, we collect data on whether subjects believe that others are more imprecise when the number on screen is closer to the reference level of 55, compared to when the number is farther from the reference level. This feature of beliefs is important because the equilibrium predictions from our previous experiment depend on the noise structure in perception. In particular, recent theoretical work has shown that an important property of the noise structure for determining equilibrium is that discriminating between nearby states is harder than discriminating between far away states (Morris and Yang, 2022; Hébert and Woodford, 2021). We ask subjects to consider the 149 other participants in their experimental condition of the study, who also just completed Phase 1. We then ask subjects the following two questions:

1. Consider only trials where the number on screen was equal to 47. In what percentage of these trials do you think the other participants gave a correct answer, that is, they correctly classified whether the number was smaller or larger than 55?
2. Consider only trials where the number on screen was equal to 54. In what percentage of these trials do you think the other participants gave a correct answer, that is, they correctly classified whether the number was smaller or larger than 55?

For each of the two questions, we pay the subject 0.5 GBP if their forecast is within 1% of the true percentage.² Question 1 elicits beliefs about others' imprecision when the distance between the number is far from the threshold (47 vs. 55), whereas Question 2 elicits beliefs about others' imprecision when the distance is close (54 vs. 55). While we could have asked subjects about their beliefs about others' imprecision for a range of numbers — rather than the single numbers 47 and 55 — this would have introduced a confound, since the distribution of numbers is different across conditions.

In Phase 3, we ask subjects about their own performance on the number classification task (that they completed in Phase 1). This question is not trivial because we do not provide subjects with feedback after any round in Phase 1 (nor after the end of Phase 1). Here, we are also interested in subjects' awareness of their own imprecision for numbers that are close and far from the threshold. Specifically, we ask subjects the following two questions:

²Following Hartzmark, Hirshman and Imas (2021), we choose this elicitation procedure as opposed to a more complex mechanism such as the Binarized Scoring Rule (BSR) due to recent evidence showing that the BSR can systematically bias truthful reporting (Danz, Vesterlund and Wilson, 2022).

1. Consider only trials where the number on screen was between 52 and 58. In what percentage of these trials do you think you correctly classified whether the number was smaller or larger than 55?
2. Consider only trials where the number on screen was less than 52 or greater than 58. In what percentage of these trials do you think you correctly classified whether the number was smaller or larger than 55?

For each of these two questions, we again reward subjects with 0.50 GBP if they provide an answer that is within 1% of their true accuracy. All subjects first go through Phase 1, and the order of Phase 2 and Phase 3 is randomized across subjects. We note that one potential concern with our design, is that when asking subjects about their performance in Phase 1, we are testing memory, not *ex-ante* beliefs. This is a reasonable concern, and an alternative is to have subjects forecast their performance before undertaking the classification task. However, under this alternative design, subjects' classification performance would be endogenous to their beliefs, and would invalidate the incentive compatibility of our belief elicitation procedure. For this reason, we opt to implement Phase 1 first for all subjects.

Experimental Results

The upper panel of Figure A1 replicates the classic result from previous experiments on number discrimination, whereby subjects exhibit errors, and these errors increase as the number on screen approaches the threshold (Dehaene, Dupoux and Mehler, 1990). Moreover, we see that, for numbers between 47 and 63, errors are systematically higher in the high volatility condition (Frydman and Jin, 2022). Similar patterns are reflected in the response times shown in the lower panel of Figure A1: response times increase as the number approaches the threshold of 55, and response times are systematically longer in the high volatility condition.

The purpose of Phase 1 is to create a dataset about performance, over which we can ask subjects about their beliefs in Phases 2 and 3. In the left panel of Figure A2, we see that subjects believe their behavior in the classification task exhibits imprecision (that is, beliefs about accuracy are less than 100%). Moreover, we see that subjects are aware that mistakes are more likely for numbers closer to the threshold (greater than 52 and less than 58) than for numbers farther from the threshold (less than 52 or greater than 58; $p < 0.001$).

The results in the middle panel of Figure A2 help validate a crucial assumption in our model. Specifically, we see that subjects are aware of other subjects' imprecision. Moreover, subjects believe that others are less accurate when discriminating 54 vs. 55 compared with discriminating 47 vs. 55 ($p < 0.001$). When embedded in a game, these beliefs are sufficient




Figure A1: **Accuracy and Response Times in the Classification Task.** Note: Upper panel shows the proportion of rounds on which subjects correctly classify θ as greater than or less than the reference level of 55. Lower panel shows the average response time on rounds where subjects correctly classify θ . Inside each data point, the length of the vertical bar denotes two standard errors of the mean. Standard errors are clustered by subject.

to generate strategic uncertainty: if player i believes that player j perceives θ with error, then player i is uncertain about player j 's perception. The data in the middle panel of Figure




Figure A2: **Beliefs about Own and Others' Accuracy in the Classification Task.**

Note: Left panel shows the average belief about own accuracy for values of θ that are far ($\theta < 52$ or $\theta > 58$) and close ($51 < \theta < 59$) to the threshold 55. Middle panel shows the average belief about others' accuracy for values of θ that are far ($\theta = 47$) and close ($\theta = 54$) to the threshold 55. Right panel shows the average belief about others' accuracy when $\theta = 54$, split by experimental condition. In each panel, the vertical bars denote two standard errors of the mean.

A2 therefore provide support for the mechanism that generates strategic uncertainty in our model.

Finally, our data also enable us to test one other feature of beliefs about others' imprecision. As outlined in our pre-registration, we test whether beliefs about others' accuracy on rounds when $\theta = 54$ is higher for those subjects who experience the low volatility distribution in Phase 1.³ Such a test investigates the hypothesis that subjects are aware that others' perception of a given number varies as a function of the experienced distribution. Indeed, the right panel of Figure A2 shows that, for $\theta = 54$, subjects who experience the high volatility distribution in Phase 1 report that others make more errors, compared to those subjects who experience the low volatility distribution in Phase 1 ($p = 0.048$).

³Pre-registration document is available at https://aspredicted.org/OGG_XNK.

C Experimental Instructions

Experiment 1 (Coordination Game)

Welcome!

You will earn £2 for completing this study and will have the opportunity to **earn more money** depending on your decisions during the study.

Specifically, at the end of the study, the computer will randomly select one question. You will receive points from the randomly selected question and the number of points depends on your decision and the decision of another participant. Points will be converted to pounds using the rate 20 points = £1. For example, if you earned 60 points for the selected question, you would then earn $60/20 = £3$ (in addition to the completion fee).

All questions are equally likely to be selected so **make all choices carefully**.

The next pages give detailed instructions. Following the instructions, you will take a quiz on them. You will be allowed to continue and will be entitled to payment **only if you answer all questions on the quiz correctly**.

Instructions (1/2)

The study is separated into 6 parts of 50 rounds each.

In each round, you are randomly matched with another participant, who we call your **opponent**.

In each round, both you and your opponent will be asked to choose between **two options**:

"Option A" or "Option B"

Here is how to earn points:

- If you choose Option A, the number of points you receive does not depend on whether your opponent chooses Option A or B. The amount of points you receive for choosing Option A can be different in different rounds and will be displayed on your screen.
- If you choose Option B, the number of points you receive depends on your opponent's decision: if your opponent chooses Option A, you will receive 47 points; if your opponent also chooses Option B you will receive 63 points.

Importantly, your opponent is reading these same exact instructions. This means that:

- If your opponent chooses Option A, his/her payoff does not depend on your decision and the number of points he/she earns are those given by Option A.
- If your opponent chooses Option B, the number of points he/she receives depends on your decision: if you choose Option A, your opponent will receive 47 points; if you also choose Option B, your opponent will receive 63 points.

Instructions (2/2)

Below is an example screen from the study:

Option A

53

Option B

47 if other participant chooses A

63 if other participant chooses B

In this example, Option A is on the LEFT side of the screen and Option B is on the RIGHT.

In each round, you will choose one of the two options by pressing either the "A" key on your keyboard for the LEFT option or the "L" key on your keyboard for the RIGHT option. On some rounds, Option A will be on the LEFT, and in other rounds it will be on the RIGHT.

In the example above:

- Option A pays you 53 points regardless of your opponent's decision, while Option B pays you 47 points if your opponent chooses Option A and 63 points if your opponent chooses Option B.
- Note also that, if your opponent chooses Option A, he/she earns 53 points regardless of your decision. If your opponent, instead, chooses Option B, he/she earns 47 points if you choose Option A and 63 points if you choose Option B.

Experiment 2 (Human vs. Algorithm)

Welcome!

You will earn £2 for completing this study.

You will also have the opportunity to **earn more money** depending on your decisions.

Specifically, at the end of the study, the computer will randomly select one question. You will receive points from the randomly selected question and the number of points depends on your decision and the decision of an opponent. Points will be converted to pounds using the rate 20 points = £1. For example, if you earned 60 points for the selected question, you would then earn $60/20 = £3$ (in addition to the completion fee).

All questions are equally likely to be selected so **make all choices carefully**.

The next pages give detailed instructions.

Following the instructions, you will take a quiz on them. You will be allowed to continue and will be entitled to payment **only if you answer all questions on the quiz correctly**.

Instructions (1/2)

The study is separated into 6 parts of 50 rounds each.

In each round, you will play a game against the computer, who we call your **opponent**.

In each round, both you and your opponent will choose between two options:

"Option A" or "Option B"

Importantly, **your computerized opponent chooses according to a rule** that you will learn below.

Here is how to earn points:

- If you choose Option A, the number of points you receive does not depend on whether your opponent chooses Option A or B. The amount of points you receive for choosing Option A can be different in different rounds and will be displayed on your screen.
- If you choose Option B, the number of points you receive depends on your opponent's decision: if your opponent chooses Option A, you will receive 47 points; if your opponent also chooses Option B you will receive 63 points.

Your Opponent's Rule

You will not have to guess how your computerized opponent will behave.

This is because **your opponent is programmed to choose according to the following rule**:

- If Option A delivers 55 points or more, your opponent chooses Option A.
- If Option A delivers less than 55 points, your opponent chooses Option B.

To be clear: your computerized opponent follows the rule above without exceptions. So, you can be certain about how your opponent's choice depends on the number of points that Option A delivers.

Instructions (2/2)

Below is an example screen from the study:

Option A

53

Option B

47 if opponent chooses A

63 if opponent chooses B

In this example, Option A is on the LEFT side of the screen and Option B is on the RIGHT.

In each round, you will choose one of the two options by pressing either the "A" key on your keyboard for the LEFT option or the "L" key on your keyboard for the RIGHT option. On some rounds, Option A will be on the LEFT, and in other rounds it will be on the RIGHT.

In the example above:

- Since Option A pays 53 points, the rule says your opponent chooses Option B.
- Therefore, if you choose Option A, you get 53 points.
- If, instead, you choose Option B, you get 63 points.

Experiment 3 (Awareness of Cognitive Noise)

Thank you for participating in this study!

Before we begin, please close all other applications on your computer and put away your cell phone. This study will last approximately **10 minutes**. During this time, we ask your complete and undistracted attention. You will earn £1 for completing the study and you will have the opportunity to **earn more money** depending on your answers during the study.

This study consists of **two phases**. The instructions for Phase 1 are given in the next page. After you go through Phase 1, you will be given a new set of instructions for Phase 2.

When you are ready to continue, press ENTER.

In **Phase 1**, you will see a series of numbers and will be asked to classify whether each number is **larger or smaller than 55**. If the number displayed is smaller than 55, press the "A" key on your keyboard. If the number displayed is larger than 55, press the "L" key.

Your bonus payment will depend on the speed and accuracy of your classification. Specifically:

$$\text{Bonus Payment} = \text{£} (1.5 \times \text{accuracy} - 1 \times \text{speed})$$

where "accuracy" is the percentage of trials where you correctly classified the number as larger or smaller than 55, and "speed" is the average amount of time it takes you to classify the number on all trials throughout the study, in seconds.

Thus, you make the most money by answering as **quickly** and as **accurately** as possible.

For example, if you correctly classified the number on all trials and it took you 0.3 seconds to respond to each question, you would earn $\text{£}(1.5 \times 100\% - 10 \times 0.3) = \text{£}1.20$. If instead you only classified 70% of the numbers correctly and took 0.8 seconds to respond to each question, you would earn $\text{£}(1.5 \times 70\% - 10 \times 0.8) = \text{£}0.25$.

Phase 1 will be separated into 3 parts of 50 trials each. In between, you can take a short break.

Before starting with the classification task, you will be asked a question to check your understanding of the instructions. You will be allowed to continue **only if you answer this question correctly**.

When you are ready to continue with the comprehension question, press ENTER.

This is **Phase 2** of the study.

Phase 2 consists of four questions, two on this page and two on the next one.

There are 99 other participants in this study.

Consider the task completed by **the other participants** in Phase 1.

Question 1

Consider only trials where the number on the screen was **equal to 47**. In what percentage of these trials do you think **the other participants** gave a correct answer, that is, they correctly classified whether the number was smaller or larger than 55? Give us your forecast on a scale between 0% and 100%, where 0% means you believe no answer in these trials was correct and 100% means you believe all answers in these trials were correct. If your forecast is within plus or minus 1% of the true percentage, you will earn £0.5.

Question 2

Consider only trials where the number on the screen was **equal to 54**. In what percentage of these trials do you think **the other participants** gave a correct answer, that is, they correctly classified whether the number was smaller or larger than 55? Give us your forecast on a scale between 0% and 100%, where 0% means you believe no answer in these trials was correct and 100% means you believe all answers in these trials were correct. If your forecast is within plus or minus 1% of the true percentage, you will earn £0.5.

Press ENTER to confirm your answers.

Consider the task **you** completed in Phase 1.

Question 3

Consider only trials where the number on the screen was **between 52 and 58**. In what percentage of these trials do you think **you** gave a correct answer, that is, you correctly classified whether the number was smaller or larger than 55? Give us your forecast on a scale between 0% and 100% where 0% means you believe no answer in these trials was correct and 100% means you believe all answers in these trials were correct. If your forecast is within plus or minus 1% of your true accuracy, you will earn £0.5.

Question 4

Consider only trials where the number on the screen was **smaller than 52 or larger than 58**. In what percentage of these trials do you think **you** gave a correct answer, that is, you correctly classified whether the number was smaller or larger than 55? Give us your forecast on a scale between 0% and 100% where 0% means you believe no answer in these trials was correct and 100% means you believe all answers in these trials were correct. If your forecast is within plus or minus 1% of your true accuracy, you will earn £0.5.

Press ENTER to confirm your answers.

D Response Times from Experiment 1




Figure A3: **Average Response Time as a Function of θ .** Note: Response times are averaged across subjects and across rounds. Inside each data point, the length of the vertical bar denotes two standard errors of the mean. Standard errors are clustered by subject.

Here we analyze the distribution of response times in both the high volatility and low volatility conditions from Experiment 1. The response time variable is defined at the round level, and measures how long it takes a player to execute a decision after the game is presented on the screen. As outlined in our pre-registration, we test two hypotheses regarding the distribution of response times. First, response times should peak at the unique equilibrium cutoff level of 55. Second, conditional on θ , response times should be longer in the high volatility condition. Our hypotheses are motivated by the literature on sequential sampling models (Ratcliff, 1978; Bogacz, Brown, Moehlis, Holmes and Cohen, 2006), which robustly predict that response times become longer as the values of two items under comparison become closer together. Thus, the tests we present in this section are joint tests of cognitive noise, which predicts that subjects use a unique threshold strategy, and sequential sampling models, which predicts how long it takes to implement the threshold strategy on each round.

In many sequential sampling models (see, e.g., Krajbich, Armel and Rangel 2010), the agent will execute a decision as soon as a stream of incoming signals has reached a pre-defined reliability threshold. Because signals are sampled sequentially, response times increase with the number of signals drawn. While the model we present in Section 2 only allows the agent to draw a single noisy signal, S_i , one could generalize the model to allow a sequence of independent noisy signals. For every additional noisy signal that the player collects, her posterior will become narrower, and, thus, the entire stream of signals provides more reliable evidence about whether θ is less than 55. As signals become more informative about whether θ is below the (equilibrium) threshold, the agent will reach the pre-defined reliability threshold with fewer signals, and thus response times will be shorter.

In our setting, there are two particular ways in which a signal can provide more information about whether θ is less than 55. First, recall that in our model, the mean of S_i varies monotonically with θ . Thus, S_i provides cardinal information about θ , and not just ordinal information about whether θ is below 55. It follows that as $|\theta - 55|$ increases, S_i provides a more informative signal about whether $\theta < 55$. Second, as the precision of S_i increases, this naturally provides more information about whether $\theta < 55$. Taken together, sequential sampling models predict that, when a player is tasked with implementing a cutoff strategy (which is derived as the equilibrium strategy under cognitive noise), response times should decrease as (i) $|\theta - 55|$ increases and (ii) the precision of S_i increases. We can test the first prediction by relying on variation in θ within an experimental condition. We can test the second prediction by relying on the variation in signal precision across conditions, which is endogenously generated by efficient coding.

Figure A3 plots the average response time, conditional on θ , for each of the two experimental conditions. We highlight two features of the figure. First, in the high volatility condition, the peak response time is at $\theta = 55$; in the low volatility condition, the peak is not far away, at $\theta = 54$. Moreover, response times fall almost monotonically as θ moves away from the equilibrium threshold of 55 ($p = 0.001$ in a mixed-effects regression of response time on $|\theta - 55|$ for each of the two conditions). Second, there is a clear separation of the curves across conditions: conditional on θ , response times are longer in the high volatility condition compared to the low volatility condition (unconditionally, the average response time is significantly longer in the high volatility condition, $p < 0.001$). These two features of the data are roughly consistent with the predictions outlined above.

One caveat to this analysis is that the player chooses the precision of S_i according to efficient coding, but under the assumption that she can only draw one signal. Predictions may change if we endogenized the signal precision and the number of signals to be drawn (or the reliability threshold). That said, the data from Figure A3 provide suggestive evidence

that subjects are implementing threshold strategies in a manner that is consistent with core predictions of sequential sampling models. Thus, the response time data help validate our assumptions about the cognitive constraints that subjects face when playing the game.

E Heterogeneity in Experiment 2




Figure A4: **Heterogeneity in % Correct Decisions in Experiment 2.** Note: Each point represents a single subject from the Algorithm condition. A “near to threshold round” is one where θ takes on one of the following values: 51, 52, 53, 54, 56, 57, 58, or 59. A “far from threshold round” is one where θ takes on one of the following values: 47, 48, 49, 50, 60, 61, 62, or 63.

Figure 5 in the main text shows that, on average, subjects in the Algorithm condition from Experiment 2 exhibit noise in their behavior. However, Figure 5 only shows the data at the aggregate level. Here, we investigate individual-level data. We demonstrate that a majority of subjects exhibit noise in their behavior, and that the noise structure is consistent with our model assumptions.

For each of the 100 subjects in the Algorithm condition, we compute the proportion of decisions that are “correct” in the sense that subjects best respond to the computer’s exogenous strategy. In particular, we define a decision as “correct” if the subject (chooses invest and $\theta < 55$) or (chooses not invest and $\theta > 55$). Otherwise, we define a decision as “incorrect”. We analyze the sample used to create Figure 5: all rounds for which $46 < \theta < 64$ and subjects exhibit a response time greater than 0.5 seconds.

Before presenting the results, note that our model of cognitive noise induces incorrect decisions because the subject does not observe θ , and instead only observes $S_i = \theta + \sigma \varepsilon_i$.

Thus, even if a subject perfectly follows the optimal threshold strategy of choosing invest if and only if $S_i < 55$, there may still be “incorrect” decisions, and the frequency of these incorrect decisions is increasing in σ .

Figure A4 shows that nearly all subjects (91%) exhibit incorrect decisions, consistent with them using a noisy signal to implement a threshold strategy. We also note that a majority of subjects (78%) are located above the 45-degree line, indicating that the frequency of incorrect decisions is higher when θ is near the threshold of 55, compared to when it is far from the threshold of 55. Taken together, these results are consistent with: 1) subjects exhibiting noise at the individual level and 2) our assumption of the noise specification, whereby incorrect decisions become more likely as θ approaches the threshold of 55.

References

- Bogacz, Rafal, Eric Brown, Jeff Moehlis, Philip Holmes, and Jonathan D Cohen**, “The Physics of Optimal Decision Making: A Formal Analysis of Models of Performance in Two-Alternative Forced-Choice Tasks,” *Psychological Review*, 2006, *113* (4), 700.
- Danz, David, Lise Vesterlund, and Alistair J Wilson**, “Belief Elicitation and Behavioral Incentive Compatibility,” *American Economic Review*, 2022, *112* (9), 2851–83.
- Dehaene, Stanislas, Emmanuel Dupoux, and Jacques Mehler**, “Is Numerical Comparison Digital? Analogical and Symbolic Effects in Two-Digit Number Comparison,” *Journal of Experimental Psychology: Human Perception and Performance*, 1990, *16* (3), 626.
- Frydman, Cary and Lawrence J Jin**, “Efficient Coding and Risky Choice,” *Quarterly Journal of Economics*, 2022, *137* (1), 161–213.
- Hartzmark, Samuel M, Samuel Hirshman, and Alex Imas**, “Ownership, Learning, and Beliefs,” *Quarterly Journal of Economics*, 2021, *136* (3), 1665–1717.
- Hébert, Benjamin M and Michael Woodford**, “Neighborhood-Based Information Costs,” *American Economic Review*, 2021, *111* (10), 3225–55.
- Krajbich, Ian, Carrie Armel, and Antonio Rangel**, “Visual Fixations and the Computation and Comparison of Value in Simple Choice,” *Nature Neuroscience*, 2010, *13* (10), 1292–1298.
- Ma, Wei Ji and Michael Woodford**, “Multiple Conceptions of Resource Rationality,” *Behavioral and Brain Sciences*, 2020, *43*.

Morris, Stephen and Ming Yang, “Coordination and Continuous Stochastic Choice,”
Review of Economic Studies, 2022, 89 (5), 2687–2722.

Ratcliff, Roger, “A Theory of Memory Retrieval.,” *Psychological Review*, 1978, 85 (2), 59.