

Frydman, Cary D.; Nunnari, Salvatore

Working Paper

Cognitive Imprecision and Strategic Behavior

CESifo Working Paper, No. 9483

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Frydman, Cary D.; Nunnari, Salvatore (2021) : Cognitive Imprecision and Strategic Behavior, CESifo Working Paper, No. 9483, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/249028>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Cognitive Imprecision and Strategic Behavior

Cary Frydman, Salvatore Nunnari

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

<https://www.cesifo.org/en/wp>

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: <https://www.cesifo.org/en/wp>

Cognitive Imprecision and Strategic Behavior

Abstract

We propose and experimentally test a theory of strategic behavior in which players are cognitively imprecise and perceive a fundamental parameter with noise. We focus on 2×2 coordination games, which generate multiple equilibria when perception is precise. When adding a small amount of cognitive imprecision to the model, we obtain a unique equilibrium where players use a simple cutoff strategy. The model further predicts that behavior is context-dependent: players implement the unique equilibrium strategy with noise, and the noise decreases in fundamental volatility. Our experimental data strongly support this novel prediction and reject several alternative game-theoretic models that do not predict context-dependence. We also find that subjects are aware of other players' imprecision, which is key to generating strategic uncertainty. Our framework has important implications for the literature on global games and, more broadly, illuminates the role of perception in generating both random and context-dependent behavior in games.

JEL-Codes: C720, C920, D910, E710.

Keywords: perception, efficient coding, coordination, global games.

Cary Frydman
University of Southern California
Marshall School of Business
Los Angeles / CA / USA
cfrydman@marshall.usc.edu

Salvatore Nunnari
Bocconi University
Department of Economics
Milan / Italy
salvatore.nunnari@unibocconi.it

October 9, 2021

We are grateful to Benjamin Enke, Evan Friedman, Nicola Gennaioli, Duarte Goncalves, Maxim Goryunov, Lawrence Jin, Chad Kendall, Stephen Morris, Rosemarie Nagel, Nicola Pavoni, Alexandros Rigos, Hyun Shin, Ajalavat Viriyavipart, Roberto Weber, Michael Woodford, Ming Yang and audiences at Princeton University, University of Glasgow, Indiana University, University of Zurich, University of Alicante, University of Insubria, Florida State University, MiddExLab Virtual Seminar, the HBS Workshop on Cognitive Noise and the NBER Behavioral Macro SI for helpful comments. Luca Congiu and Massimiliano Pozzi provided excellent research assistantship. We acknowledge financial support from the NSF (Frydman, Grant No. 1749824) and the ERC (Nunnari, Grant No. 852526).

1 Introduction

Over the past decade, economists have begun investigating the effects of *cognitive imprecision* on individual decision-making (see Woodford 2020 for a review). This agenda proposes that the decision-maker’s perception of the economic environment is noisy, and does not coincide with the objective environment. To date, the emphasis of cognitive imprecision in economics has largely been in the domains of choice under risk and intertemporal choice. For example, theory has shown that even when an agent has linear utility and perfectly patient time preferences, noisy perception of payoffs can lead agents to behave as if they are risk averse (Khaw, Li and Woodford, 2021) and as if they discount the future (Gabaix and Laibson, 2017). Initial experimental tests of these theories have produced encouraging results (Gershman and Bhui, 2020; Khaw, Li and Woodford, 2021; Enke and Graeber, 2021; Frydman and Jin, Forthcoming).

Motivated by the evidence from individual decision-making studies, it is natural to ask whether cognitive imprecision also affects strategic behavior. This question is important not only to test whether cognitive imprecision extends into other environments, but also because noisy perception can fundamentally affect equilibrium predictions. Indeed, it is well known that, in games often used to model bank runs, currency attacks, and revolutions, there can be multiple equilibria.¹ This multiplicity largely comes from the assumption that players can precisely perceive a fundamental parameter, which then serves as a coordination device. However, as pointed out by Woodford (2020), if one adds a small amount of cognitive imprecision to the game so that each player perceives a slightly different fundamental, then coordination becomes more difficult and multiplicity breaks down. This line of reasoning follows directly from the vast literature on global games (Carlsson and Van Damme, 1993; Morris and Shin, 2003; Angeletos and Lian, 2016) with the important distinction that, here, we interpret the noise as arising from perceptual errors rather than from traditional sources of asymmetric information (e.g., different opportunities that agents have to acquire information about an uncertain state of the world).

In this paper, we theoretically develop and experimentally test the hypothesis that perceptual noise systematically affects strategic behavior. The hypothesis enables us to apply standard results from the global games literature without explicitly introducing private information. Specifically, we analyze a 2×2 simultaneous-move game where players can choose to invest or not invest in an asset. Each player’s payoff depends on the value of a fundamental and on the action of the other player. While theory predicts multiple equilibria for a range

¹For examples of game theoretic models in these domains, see Diamond and Dybvig (1983), Obstfeld (1996), Morris and Shin (1998), Atkeson (2000), Goldstein and Pauzner (2005), and Edmond (2013).

of fundamental values in the complete information version of the game, a small amount of perceptual noise generates a unique equilibrium: each player invests if and only if their noisy perception of the fundamental crosses a threshold. This is the classic conclusion from the global games literature, which typically interprets the noisy observation of fundamentals as capturing asymmetric information.²

Because we adopt the view that noise arises from perceptual error, we can leverage principles from psychology to generate even sharper equilibrium predictions. Specifically, we draw on the principle of *efficient coding*, which states that the decision-maker’s perceptual system optimally reallocates resources as the statistics of the environment change. In our setting, efficient coding implies that perception of a given fundamental value will be more precise when fundamental volatility is lower. The intuition is that each player has a limited set of cognitive resources, and she optimally allocates these resources towards perceiving those fundamental values that she expects to occur more frequently.³ Thus, when the distribution of fundamentals becomes more volatile, cognitive resources are dispersed more broadly, which leads to larger perceptual errors (for values near the center of the distribution). Efficient coding therefore predicts *context-dependent* behavior: each player implements the equilibrium threshold strategy with more precision as fundamental volatility decreases.

To test this prediction, we conduct a pre-registered experiment in which subjects play a simultaneous-move game in each of three hundred rounds. The game is characterized by the value of a fundamental parameter, which is clearly displayed to both subjects on each round as a two-digit Arabic numeral, such as “45”. We rely on subjects’ inherent cognitive imprecision to transform this “public” signal into a private signal, owing to idiosyncratic perceptual errors. The perceptual error induces uncertainty over the fundamental. More importantly, perceptual error also induces *strategic uncertainty* over the other player’s action, which is key to breaking the multiplicity of equilibria. Our key experimental treatment is to manipulate the volatility of the fundamental across a high volatility and a low volatility condition, in order to test for context-dependent behavior.

Our data strongly support the novel prediction that equilibrium outcomes are context-dependent. Specifically, we observe that the probability of investing is monotone in the

²Several previous experimental studies have found empirical support for the global games prediction that the probability of investing is monotonic in the fundamental (Heinemann, Nagel and Ockenfels, 2004, 2009; Cabrales, Nagel and Armenter, 2007; Avoyan, 2019; Szkup and Trevino, 2020; Goryunov and Rigos, 2020). As we describe further below, one feature that differentiates our study from the existing experimental literature on global games is that we invoke the global games arguments without explicitly endowing subjects with private information

³Such an assumption has been validated in many papers on sensory perception (Girshick, Landy and Simoncelli, 2011; Wei and Stocker, 2015; Payzan-LeNestour and Woodford, Forthcoming) and in economic decision making (Polania, Woodford and Ruff, 2019; Frydman and Jin, Forthcoming).

fundamental, and we find that this monotonic relationship is significantly stronger in the low volatility condition than in the high volatility condition. In light of our model, we interpret the observed treatment effect as a consequence of more accurate perception of fundamentals in the low volatility condition. Further evidence comes from the distribution of response times, which indicates that subjects choose their action significantly faster when they are adapted to the low volatility distribution of fundamentals.

We emphasize that, in both experimental conditions, the strong monotonic relationship that we observe between fundamentals and investing is not predicted under the complete information version of the game. As such, our data suggest that even when subjects receive no explicit private signals from the experimenter, private information is inherent in the game because subjects encode the fundamental with idiosyncratic perceptual noise. Our framework therefore provides a new explanation for earlier experimental papers that find a high correlation between behavior and fundamentals, regardless of whether subjects receive explicit private signals (Heinemann, Nagel and Ockenfels, 2004; Van Huyck, Viriyavipart and Brown, 2018).⁴ Such a result may initially appear puzzling because, in the absence of private information, there are multiple equilibria and behavior should not vary smoothly with fundamentals. However, one can explain the correlation between fundamentals and behavior by taking a broader view of the potential sources of private information to also include perceptual errors.⁵

The data from our experiment can also separate between cognitive imprecision and several alternative models from game theory. Perhaps the closest model to cognitive imprecision is Quantal Response Equilibrium (QRE; McKelvey and Palfrey 1995, 1998). In QRE, each player stochastically best responds to their opponent. Cognitive imprecision and QRE share the prediction that, in equilibrium, behavior is random (even when players are not indifferent between actions). However, a unique prediction generated by cognitive imprecision, that is not shared by QRE, is that behavior depends on the distribution from which the fundamental is drawn. In fact, this context dependence further separates cognitive imprecision from a broader class of theories including level-k thinking (Stahl and Wilson, 1994, 1995; Nagel, 1995). We discuss these differences across theories in more detail in Section 5.

⁴Moreover, when discussing an experiment where there is no explicit private information about payoffs, Heinemann, Nagel and Ockenfels (2009) argue that “Of course, players know the true payoff. Their uncertainty about others’ behavior makes them behave *as if* they are uncertain about payoffs” (p. 203). Our results indicate that it may well be the case that subjects do not know the true payoff, because of perceptual error.

⁵Van Huyck and Stahl (2018) conduct an experiment by simultaneously varying both the range and the mean of payoffs in a stag hunt game. However, because the fundamental in their experiment never takes on values in the “dominance regions”, one cannot interpret the experimental results through the theory of global games.

In our model, perceptual error generates a unique equilibrium because each player is uncertain about their opponent’s perception of the fundamental. The model therefore relies on the assumption that each player is aware of their opponent’s imprecision. To investigate the validity of this assumption, we conduct a second experiment, where subjects are asked to classify whether a two-digit number is greater than a reference level of 55 (which is chosen to be the same as the threshold in the unique equilibrium of the game in our first experiment). We then incentivize subjects to report their beliefs about (i) the average accuracy of all other subjects in the experiment and (ii) their own accuracy.

We find that subjects are aware of their own errors and, more importantly, they are aware of others’ errors in the classification task. Subjects also report beliefs that discriminating between a number close to the threshold, say “54”, is harder than discriminating between a number far from the threshold, say “47.” This property has been shown theoretically to have important implications for equilibrium selection (Morris and Yang, 2021), and has recently been formalized in the rational inattention literature using a so-called “neighborhood cost function” (Hébert and Woodford, Forthcoming). The data from our second experiment therefore provide novel evidence supporting the assumption that subjects are aware that others’ ability to discriminate between two states depends on the “distance” between states.

Our results build directly on a set of papers that has begun testing whether principles of cognitive imprecision are active in individual economic decision-making (Enke and Graeber, 2021; Gershman and Bhui, 2020; Khaw, Li and Woodford, 2021; Frydman and Jin, Forthcoming; Polania, Woodford and Ruff, 2019). There has, however, not yet been a similar set of tests in strategic settings. Our results provide clear evidence that cognitive imprecision does indeed extend into strategic settings, and that subjects are aware of this feature of their own and others’ perception. By testing whether similar cognitive mechanisms apply in individual and strategic decision-making, our work is related to recent brain imaging evidence from Nagel, Brovelli, Heinemann and Coricelli (2018) who show that common neural circuits are activated during lottery choice and games with strategic uncertainty.

The remainder of the paper proceeds as follows: Section 2 presents the model and derives the theoretical predictions for our experimental manipulation. Sections 3 and 4 describe the experimental design and report the experimental results for Experiment 1 (the simultaneous-move game) and Experiment 2 (the number classification task), respectively. Section 5 discusses assumptions of our theoretical framework and draws connections with the global games literature and alternative game theory models. Section 6 concludes.

	Not Invest	Invest
Not Invest	θ, θ	θ, a
Invest	a, θ	b, b

Figure 1: **The Game**

2 Model

In this section, we present a model in which players imprecisely perceive their strategic environment. We assume that each player forms a noisy perception of the “fundamental” payoff in the game. As a consequence, each player’s perception of the fundamental payoff will, in general, differ from the true fundamental and from their opponent’s perception of the fundamental. We illustrate the strategic implications of noisy perception in the setting of a 2×2 simultaneous-move game. We focus our analysis on those parameter values that generate the essential features of a coordination game.

Consider the game in Figure 1, where $b > a$. In what follows, we always assume that a and b are perceived precisely (i.e., without any noise) by both players, and we are interested in the effect of imprecise perception of θ .⁶ As a benchmark, we first consider the predictions of a model in which θ is perceived precisely, and then we relax this assumption to investigate the implications of cognitive imprecision.

2.1 Benchmark: Precise Perception

When both players perceive θ precisely, the game is one of complete information and its Nash equilibria depend on the true value of θ , as outlined below:

- If $\theta > b$, then Invest is a strictly dominated action for each player, and (Not Invest, Not Invest) is the unique Nash (and dominant strategy) equilibrium.
- If $\theta < a$, then Not Invest is a strictly dominated action for each player, and (Invest, Invest) is the unique Nash (and dominant strategy) equilibrium.
- If $a \leq \theta \leq b$, then there are two Nash equilibria in pure strategies: (Not Invest, Not Invest) and (Invest, Invest). There also exists one Nash equilibrium in mixed strategies.

Thus, when θ takes on values in the intermediate range $[a, b]$, there are multiple pure strategy Nash equilibria. This prediction relies on the assumption that each player precisely perceives θ . Precise perception generates common knowledge about θ , which enables

⁶Our assumption that a and b are perceived without noise can be justified, for example, through a learning mechanism. In our experiment, we keep a and b constant across all rounds, so the amount of noise in perceiving a and b is arguably minimal.

coordination and gives rise to multiple self-fulfilling equilibria. The predictions change dramatically, however, when we relax the assumption that players can perceive θ precisely.

2.2 Cognitive Imprecision

Suppose now that players perceive θ with noise. This assumption is backed up by a large literature in numerical cognition, which finds that people encode numerical quantities with noise, even when the quantities are presented symbolically (see Dehaene 2011 for a review). To model the imprecision, we assume that all players have a common prior that θ is distributed normally: $\theta \sim \mathcal{N}(\mu_\theta, \sigma_\theta^2)$. Our key assumption is that, instead of precisely observing the realized value of θ , each player only has access to a noisy internal representation of θ .

Assumption 1 (Cognitive Imprecision) *Each player i , $i = \{1, 2\}$, observes a noisy internal representation of θ , $S_i = m(\theta) + \epsilon_i$, where each ϵ_i is independently and normally distributed: $\epsilon_i \sim \mathcal{N}(0, \sigma_S^2)$, with $\sigma_S^2 > 0$.*

Assumption 1 says that each player’s internal representation is conditionally independent and depends on θ through an encoding function, $m(\theta)$. We follow Khaw, Li and Woodford (2021) and assume a linear encoding function with a “power constraint”:⁷

Assumption 2 (Encoding Function) *The encoding function is linear: $m(\theta) = \xi + \psi\theta$. In addition, there is a power constraint, $E[m^2] \leq \Omega^2 < \infty$.*

The power constraint ensures that the encoded value, $m(\theta)$, does not vary too much, which captures the idea that the brain cannot encode an arbitrarily large set of values. The parameters of the encoding function, (ξ, ψ) , are allowed to vary with the player’s environment – which we characterize by the prior distribution of θ . Thus, the conditional distribution of noisy signals can vary across environments. This assumption is built on substantial empirical evidence, mainly from the literature on sensory perception, which demonstrates that the distribution of noisy internal representations is *optimally* adapted to the statistical regularities of the environment. This principle is called *efficient coding*, and recent work has empirically documented effects of efficient coding in economic choices (Polania, Woodford and Ruff, 2019; Frydman and Jin, Forthcoming). To close the efficient coding model, we need to specify the performance objective which drives the players’ optimal choice of the encoding function parameters.

⁷Khaw, Li and Woodford (2021) assume a slightly different specification of the encoding function, which is linear in the logarithm of a payoff value. See their Appendix C for details.

Assumption 3 (Performance Objective) *Payers choose the encoding function which minimizes the mean squared error between θ and its conditional mean, $E[\theta|s_i]$.*

The previous three assumptions describe the perceptual constraints and the objective function of each player. Given these constraints and objectives, we can derive the efficient encoding function that each player optimally chooses.⁸

Proposition 1 (Efficient Coding) *Given Assumptions 1-3, the optimal encoding function features $\xi^* = -\frac{\Omega}{\sigma_\theta}\mu_\theta$ and $\psi^* = \frac{\Omega}{\sigma_\theta}$. Consider the transformed internal representation, $Z_i \equiv (S_i - \xi^*)/\psi^*$. The conditional distribution of this transformed internal representation is $N(\theta, \omega\sigma_\theta^2)$, where $\omega = \sigma_S^2/\Omega^2$. The variance of the transformed internal representation is proportional to the variance of θ .*

At an intuitive level, efficient coding implies that perceptual resources are allocated so as to better discriminate between different values of θ that are expected to occur more frequently under the players' prior beliefs. Specifically, as the volatility of the prior decreases, perceptual resources are reallocated towards the center of the distribution.

Given the optimal encoding function in Proposition 1, we can now solve for the equilibria of the game. We restrict our analyses to monotone equilibria of the incomplete information game, that is, equilibria in which actions are monotonic in the transformed internal representation, Z_i . In a monotone equilibrium, players' mutual best response is to choose Invest if and only if their transformed internal representation is below a threshold k^* . Adapting results from the global games literature (Carlsson and Van Damme, 1993; Morris and Shin, 2003; Morris, 2010) to the game in Figure 1, with the further assumption that $\mu_\theta = (a + b)/2$ (as in our experiment in the next section), we can establish there exists a monotone equilibrium such that player i invests if and only if $Z_i \leq \mu_\theta$, for any value of σ_θ , σ_S and Ω . Furthermore, if the noise in the transformed internal representation is sufficiently small, then this is the unique monotone equilibrium.

Proposition 2 (Equilibrium Existence and Uniqueness) *Suppose Assumptions 1-3 and $\mu_\theta = (a + b)/2$. There exists an equilibrium of the game where each player invests if and only if $Z_i \leq \mu_\theta$ (or, equivalently, $E[\theta|Z_i] \leq \mu_\theta$). Moreover, if $\sqrt{\omega(1 + \omega)} < \frac{\sqrt{6\pi}}{(b-a)}\sigma_\theta^2$, this is the unique monotone equilibrium of the game.*

In deriving Proposition 2, we assume common knowledge of the distribution of internal representations. However, precise knowledge of the underlying information structure is not

⁸In Section 5, we discuss the robustness of our main theoretical results to different assumptions about the players' performance objective.




Figure 2: **Predicted Probability of Investing as a Function of θ .** Note: The solid line denotes the prediction for a low volatility distribution with $\theta \sim N(55, 20)$; the dashed line denotes the prediction for a high volatility distribution with $\theta \sim N(55, 400)$; we set the following parameter values: $a = 47$, $b = 63$, and $\omega = 1.5$.

necessary for this equilibrium to arise. As we show in Appendix C, in a simpler model where $m(\theta) = \theta$, it is enough to assume that (i) $\mu_\theta = (a + b)/2$, (ii) $E[\epsilon_i] = 0$, (iii) the distribution of ϵ_i is symmetric, quasiconcave and independent of the realized value of θ , and (iv) the distribution of θ is symmetric and continuous on \mathbb{R} .

Proposition 2 implies a particular set of comparative statics with respect to θ . The probability of investing is pinned down by the distribution of the transformed internal representation: $Pr[\text{Invest}|\theta] = Pr[Z_i \leq \mu_\theta|\theta] = \Phi\left(\frac{\mu_\theta - \theta}{\omega\sigma_\theta}\right)$, where $\Phi(\cdot)$ is the cumulative density function of the standard normal. This result indicates that, in the unique monotone equilibrium, the probability of investing is monotonically decreasing in θ . Moreover, the probability of investing depends not only on θ but also on the prior distribution from which θ is drawn: σ_θ modulates the optimal encoding rule and, thus, the noise in implementing the unique monotone equilibrium strategy of the game. Thus, if we experimentally manipulate the volatility of the prior, we should see that, in the unique equilibrium of the game obtained when ω is sufficiently small, the probability of investing is more sensitive to θ when the prior volatility is smaller. This prediction is summarized in the following proposition.

Proposition 3 (Comparative Statics) *Suppose Assumptions 1-3, $\mu_\theta = (a + b)/2$, and $\sqrt{\omega(1 + \omega)} < \frac{\sqrt{6\pi}}{(b-a)}\sigma_\theta^2$. In the unique monotone equilibrium of the game, the probability that each player invests for a given value of θ is $Pr[Invest|\theta] = \Phi\left(\frac{\mu_\theta - \theta}{\omega\sigma_\theta}\right)$. Decreasing the variance of θ will increase the sensitivity of choices to θ , (that is, the rate at which $Pr[Invest|\theta]$ decreases with θ) for values of θ close to μ_θ .*

We illustrate the implications of Proposition 3 in Figure 2. The figure shows that the prior variance of θ strongly affects choice sensitivity to θ . This dependence of equilibrium behavior on the distribution of θ motivates our experimental design.

3 Experiment 1: Simultaneous-Move Game

3.1 Experimental Design

In this experiment, we test the model by incentivizing subjects to play a simultaneous-move game, and we manipulate the distribution that generates the fundamental payoff, θ . We pre-register the experiment and recruit 300 subjects from the online data collection platform, Prolific.⁹ We restrict our sample to subjects who, at the time of data collection, (i) were UK nationals and residents, (ii) did not have any previous “rejected” submissions on Prolific, and (iii) answered all comprehension quiz question correctly. Subjects are paid 2 GBP (~ 2.8 USD) for completing the experiment, and they have the opportunity to receive additional earnings based on their choices and the choices of other participants.

The experiment consists of 300 rounds, and each subject participates in all rounds. In each round, a subject is randomly matched with another subject and, together, they play the simultaneous-move game in Figure 1. In all rounds, we set the payoff parameters $a = 47$ and $b = 63$. The only feature of the game that varies across rounds is the value θ , which is drawn from the condition-specific distribution $f(\theta)$. In each round, both subjects observe the same realization of θ . In order to shut down learning about other participants’ behavior, we choose not to provide subjects with feedback about their payoff or their opponent’s choice in a given round. At the end of the experiment, one round is selected at random, and subjects are paid according to the number of points they earned in that round, which in turn, depends on their action, their opponent’s action, and the (round-specific) value of θ . Points are converted to GBPs using the rate 20:1. The average duration of the experiment was ~ 25 minutes and average earnings, including the participation fee, were ~ 5.5 GBP (~ 7.7 USD).

Subjects are randomized into one of two experimental conditions: a *high volatility* condition or a *low volatility* condition, which differ only based on the distribution of θ . In the

⁹The pre-registration document is available at https://aspredicted.org/IHU_KCE.

high volatility condition, $f(\theta)$ is normally distributed with mean 55 and variance 400. In the low volatility condition, $f(\theta)$ is normally distributed with mean 55 and variance 20. In both conditions, after drawing θ from its respective distribution, we round θ to the nearest integer, and we re-draw θ if the rounded value is less than 11 or greater than 99. We implement these modifications to the normal distribution to control complexity and ensure that θ is a two-digit number on each round. We do not give subjects any explicit information about $f(\theta)$ in the instructions, as our intention is to test whether a subject’s perceptual system can adapt to the statistical properties of the environment without explicit top-down information. Moreover, we believe that such a design is more natural than explicitly telling subjects the distribution of parameters they will experience, as this could artificially direct their attention to the distribution (a similar design feature is used in Frydman and Jin (Forthcoming)). Each condition contains an identical set of instructions and comprehension quiz.¹⁰

Recall that, in the complete information version of the game, there are multiple equilibria when θ is in the range $[47, 63]$. We therefore focus our analyses on games for which θ lies in this range, which occurs on 93% of rounds in the low volatility condition and on 31% of rounds in the high volatility condition. We pre-register that our main analyses are restricted to those rounds for which $\theta \in [47, 63]$, which we call “common rounds.” This is a crucial feature of our design, because it allows us to compare behavior across conditions using the exact same set of games and varying only the context — which is characterized by the distribution of games. In choosing the variance of $f(\theta)$ for each condition, we thus strike a balance between generating a substantial number of common rounds to analyze and creating a large difference in prior variance across conditions. As outlined in our pre-registration, we also exclude the first 30 rounds from our analyses, in order to allow subjects time to adapt to the distribution of θ .

Figure 3 provides a screenshot of a single round shown to subjects. In order to avoid framing effects, we label the two options “Option A” and “Option B”, and the left-right location of each option is randomized across rounds. The number “45” is the realized value of θ on the specific round shown in Figure 3. We emphasize that — while the number is clearly displayed to all subjects and, thus, would traditionally be interpreted as public information — here we rely on cognitive imprecision to transform the fundamental value into private information. In other words, we assume that the constraints on a subject’s perceptual system make it impossible to perfectly perceive the fundamental value (Assumption 1). Furthermore, efficient coding implies that the amount of imprecision varies endogenously with the distribution of θ across conditions.

¹⁰The experimental instructions are available in Appendix D.

Option A

45

Option B

47 if other participant chooses A
63 if other participant chooses B

Figure 3: **Sample Screenshot Shown to Participants in Experiment 1.** Note: In this round, the realized value of θ is 45, which is clearly and explicitly displayed to both subjects. Subjects choose either “Option A” or “Option B” by pressing one of two different keys on the keyboard.

3.2 Experimental Results

Choice Behavior

Following our pre-registration, we restrict our analysis to common rounds in which subjects execute a decision with a response time greater than 0.5 seconds, which generates a sample of 50,129 decisions. Across both conditions, subjects choose to invest on 58.9% of rounds and exhibit an average response time of 1.64 seconds.

In Figure 4, we plot the probability of investing as a function of the fundamental, separately for the two experimental conditions. One can see that, in both conditions, the aggregate data are consistent with the hypothesis that subjects implement strategies that are monotone in θ . This finding is in line with previous experimental results on coordination games (Heinemann, Nagel and Ockenfels, 2004, 2009; Szkup and Trevino, 2020). Importantly, the smooth decreasing relationship between θ and the probability of investing obtains even without explicitly endowing subjects with any private signals about θ . We interpret this result as subjects implementing a noisy threshold strategy in both experimental conditions, which is consistent with the comparative static between θ and the probability of investing in Proposition 3.

In order to provide a more targeted test of cognitive imprecision, we exploit the variation in the distribution of θ across our two experimental conditions. Efficient coding predicts context-dependent behavior, where the threshold strategy from the unique monotone equilibrium is implemented with more precision in the low volatility condition. Figure 4 provides




Figure 4: **Empirical Frequency of Investing as a Function θ** . Note: For each value of θ between 47 and 63, we plot the proportion of rounds on which a subject chooses to invest, separately for each of the two experimental conditions. Data are pooled across subjects and are shown for rounds 31-300, after an initial 30-round adaptation period. Vertical bars inside each data point denote two standard errors of the mean. Standard errors are clustered by subject.

evidence consistent with this prediction: the frequency of investing is more sensitive to the fundamental in the low volatility condition, compared to the high volatility condition.

To formally test the difference in slope, we estimate a series of mixed effects logistic regressions. Column (1) of Table 1 confirms our main result: the coefficient on the interaction term $(\theta - 55) \times Low$ is significantly negative, indicating that the probability of investing decreases in the fundamental more rapidly when a subject is adapted to the low volatility condition. Columns (2) and (3) show that this result holds in both early (first 70 trials after adaptation) and late (last 70 rounds of the session) subsamples. Column (4) indicates that the treatment effect becomes moderately stronger over the course of the experiment. The strengthening of the treatment effect over the course of the experiment suggests that subjects have not fully adapted to the distribution by round 100, and additional rounds of play provide the opportunity for further adaptation.

While subjects do not receive feedback after each round, it is still possible that they learn

Dependent Variable: Pr(Invest)	(1)	(2)	(3)	(4)
$(\theta - 55)$	-0.458*** (0.033)	-0.467*** (0.039)	-0.577*** (0.051)	-0.481*** (0.037)
$(\theta - 55) \times \text{Low}$	-0.499*** (0.063)	-0.351*** (0.059)	-0.487*** (0.076)	-0.374*** (0.061)
Low	-0.182 (0.386)	-0.335 (0.343)	-0.170 (0.423)	-0.275 (0.360)
Late				-0.022 (0.121)
$(\theta - 55) \times \text{Late}$				0.013 (0.021)
Low \times Late				0.092 (0.154)
Low $\times (\theta - 55) \times \text{Late}$				-0.065* (0.034)
Constant	1.351*** (0.221)	1.316*** (0.224)	1.465*** (0.222)	1.292*** (0.229)
Observations	50,129	13,196	12,861	25,864
Rounds	31-300	31-100	231-300	(31-100) & (231-300)

Table 1: **Treatment Effect Estimates.** Note: Table displays results from mixed effects logistic regressions. Observations are at the subject-round level. The dependent variable takes the value 1 if the subject chooses to Invest and 0 otherwise. The variable *Low* takes the value 1 if the round belongs to the low volatility condition and 0 otherwise. The variable *Late* takes the value 1 if the round number is 231 or greater, and 0 otherwise. Only data from rounds where $46 < \theta < 64$ are included in the regressions. There are random effects on $(\theta - 55)$ and the intercept. Standard errors of the fixed effect estimates are clustered at the subject level and shown in parentheses. ***, **, * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

about the strategic environment through repeated exposure to the game, as in Weber (2003) and Rick and Weber (2010). Moreover, our experimental design implies that subjects in different conditions will experience the same game, characterized by θ , a different number of times (e.g., games characterized by a value of θ close to 55 will occur more frequently in the low volatility condition). This raises the potential concern that our observed treatment effect is due to the differential ability to learn, rather than to cognitive imprecision. To investigate this alternative explanation, Table 2 presents results when we restrict to subsamples where subjects have identical experience with a given game in both conditions. In particular, the first column restricts to those rounds on which subjects in both conditions have previously observed 3 games with the same value of θ as in the current round. Columns (2) – (4) further restrict the data based on more and more experience with a given game. The regression results indicate that our treatment effect obtains among each of the different subsamples. Thus, learning cannot explain the entire treatment effect we observe.

Dependent Variable: Pr(Invest)	(1)	(2)	(3)	(4)
$(\theta - 55)$	-0.384*** (0.039)	-0.389*** (0.041)	-0.364*** (0.036)	-0.447*** (0.047)
$(\theta - 55) \times \text{Low}$	-0.266*** (0.051)	-0.275*** (0.054)	-0.285*** (0.052)	-0.299*** (0.063)
Low	-0.317 (0.276)	-0.356 (0.285)	-0.129 (0.283)	-0.207 (0.317)
Constant	1.067*** (0.181)	1.202*** (0.199)	0.993*** (0.192)	1.174*** (0.217)
Observations	4,263	4,053	3,677	3,255
Rounds of Experience with Game (θ)	3	4	5	6

Table 2: **Controlling for Experience with θ .** Note: Table displays results from mixed effects logistic regressions. Observations are at the subject-round level. The dependent variable takes value 1 if the subject chooses to Invest and 0 otherwise. The variable *Low* takes value 1 if the round belongs to the low volatility condition and 0 otherwise. Only data from rounds where $46 < \theta < 64$ are included in the regressions. There are random effects on $(\theta - 55)$ and the intercept. Standard errors of the fixed effect estimates are clustered at the subject level and shown in parentheses. ***, **, * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

It is important to point out that the results in Figure 4 are aggregated across subjects. Therefore, while the data are consistent with the prediction that, at the individual subject level, signals are drawn from a noisier distribution in the high volatility condition, there is another potential explanation based on aggregation. Specifically, suppose that subjects

perceive θ perfectly and that they use a potentially non-equilibrium cutoff strategy. Further suppose that there is heterogeneity with respect to the cutoff that each subject adopts. If some subjects use low cutoffs, while others use high cutoffs, then this heterogeneity would give rise to the decreasing relationship observed in both aggregate curves in Figure 4. In addition, if the variance in cutoff strategies across subjects is larger in the high volatility condition, then this alternative hypothesis could explain the weaker relationship between θ and the probability of investing in the high volatility condition. To investigate this alternative hypothesis, based on heterogeneity of cutoff strategies, we structurally estimate the model to obtain subject-specific cutoffs and measures of perceptual noise.

Structural Estimation

According to the model described in Section 2, subject i chooses the parameters of the encoding rule, $m_i(\theta) = \xi_i + \psi_i\theta$. She then observes a noisy internal representation, $S_i = m_i(\theta) + \epsilon_i$. If we define a transformed version of the noisy internal representation as $Z_i = (S_i - \xi_i)/\psi_i$, then, for a cutoff Z_i^* , our model predicts that she invests if and only if $Z_i \leq Z_i^*$. In the unique monotone equilibrium of the game with cognitive imprecision, all subjects in the same treatment choose the same (ξ_i, ψ_i, Z_i^*) . Here, we allow subjects to make heterogeneous (non-equilibrium) choices and we structurally estimate these parameters using behavior observed in the experiment.

Consider subject i who adopts a cutoff value of Z_i^* and, in round t , receives a noisy internal representation $S_{it} = \xi_i + \psi_i\theta_t + \epsilon_{it}$. The probability that subject i invests in round t is the probability that her transformed noisy internal representation is below her cutoff:

$$\text{IP}(\text{Invest}|\theta_t, \sigma_S, \psi_i, Z_i^*) = \Phi\left(\frac{Z_i^* - \theta_t}{\sigma_S/\psi_i}\right) \quad (1)$$

We structurally estimate the model using maximum likelihood estimation. In particular, for each subject, we estimate the standard deviation of the transformed noisy internal representation, $\sigma_i = \sigma_S/\psi_i$, and the cutoff Z_i^* .¹¹ We maximize the following log-likelihood

¹¹We cannot separately identify σ_S and ψ_i since these two parameters are perfect substitutes in the conditional density of Z_i . At the same time, while ψ_i is an endogenous variable, we interpret σ_S as an exogenous parameter, capturing the degree of a subject's perceptual acuity. In Section 2, we assumed σ_S is homogeneous. Even if we allowed for heterogeneity across subjects, the randomization into experimental conditions would guarantee a similar distribution of σ_S in the two sub-populations. For this reason, we attribute any difference in the distribution of the estimated σ_i 's across conditions to differences in ψ_i .

function over (σ_i, Z_i^*) , using behavior in rounds 31 – 300:

$$LL(\sigma_i, Z_i^*, \mathbf{y}_i) = \sum_{t=31}^{300} y_{it} \cdot \log(\mathbb{P}(\text{Invest}|\theta_t, \sigma_i, Z_i^*)) + (1 - y_{it}) \cdot \log(1 - \mathbb{P}(\text{Invest}|\theta_t, \sigma_i, Z_i^*)), \quad (2)$$

where $\mathbf{y}_i \equiv \{y_{it}\}_{t=31}^{300}$ and y_{it} denotes the subject’s choice in round t , with $y_{it} = 1$ if the subject chooses to invest and $y_{it} = 0$ if the subject chooses not to invest. We maximize the log-likelihood function in (2) by searching over grid values of $[\sigma_i, Z_i^*] \in [0.1, 50.1] \times [11, 99]$, in increments of 0.5 along each dimension.

Figure 5 plots the distribution of estimated parameters for the 300 subjects (150 in each condition). Beginning with the upper panel, we see that, for most subjects, the estimated cutoff lies between 50 and 60. The mean cutoff in the high volatility condition is 58.5 and the mean cutoff in the low volatility condition is 57.2. These means are not significantly different from one another ($p = 0.15$). The average cutoff in each condition is, however, significantly greater than 55. As can be seen from the figure, this difference relative to 55 is driven mainly by the right tail of the distribution, which captures a small fraction of subjects who almost always choose to invest.

More importantly, we find that the standard deviation of estimated cutoffs is not significantly different across conditions (8.4 in high volatility vs. 7.5 in low volatility, $p = 0.43$ Levene’s test). This suggests that heterogeneity in cutoffs is *not* driving our main result. If it were, we would have observed a more concentrated distribution of cutoffs in the low volatility condition and, thus, a significantly lower standard deviation of estimated cutoffs in the low volatility condition.

Instead, the lower panel of Figure 5 reveals that the difference in behavior across conditions stems from the standard deviation of the noisy internal representations. The mean estimated value of σ_i is significantly higher in the high volatility condition (14.4 vs. 5.9, $p < 0.001$). One can easily see from the figure that this effect holds not only on average, but across the whole distribution. In summary, while the aggregate data in Figure 4 are consistent with subjects in the high volatility condition exhibiting (i) a wider dispersion of cutoffs or (ii) a higher amount of noise in the internal representation of the fundamental, our structural estimation indicates that the effect is coming only through the second channel, as predicted by the theory developed in Section 2.

Response Times

Here we analyze the distribution of response times in both conditions. As outlined in our pre-registration, we hypothesize that if subjects are implementing cutoff strategies, then

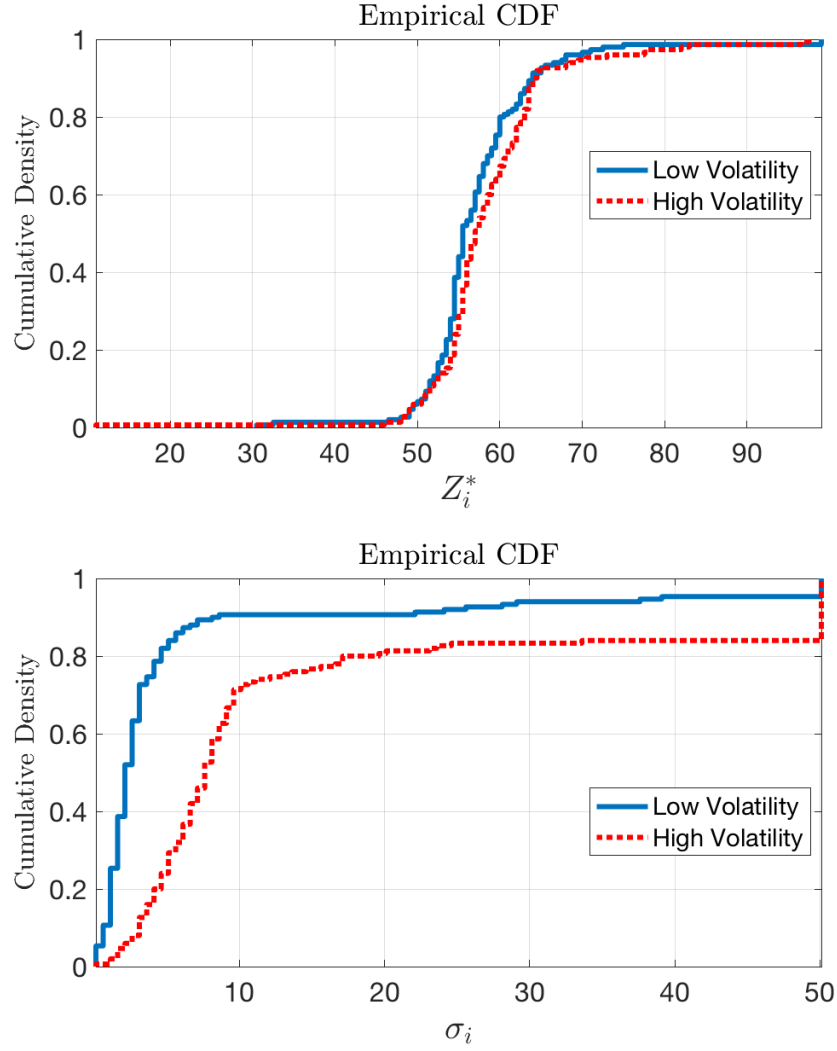


Figure 5: **Empirical CDFs of Subject-Level Structural Estimates.** Note: Upper panel is empirical CDF of estimated cutoffs. Lower panel is empirical CDF of estimated standard deviations of noisy internal representations.

response times should peak at the cutoff value. Figure 6 shows that response times are longer in the high volatility condition. Moreover, in the high volatility condition, the peak response time is at 55, whereas in the low volatility condition, the peak is at 54. If subjects are implementing the unique equilibrium threshold strategy, which involves discriminating whether the fundamental is above or below 55, then models of sequential sampling from the mathematical psychology literature (Ratcliff, 1978; Krajbich, Armel and Rangel, 2010) would predict that response times should peak at the predicted threshold of 55, since these are the most “difficult” discrimination problems. The response time data provide some support for




Figure 6: **Average Response Time as a Function of θ in Experiment 1.** Note: Response times are averaged across subjects and across rounds. Vertical bars denote two standard errors of the mean. Standard errors are clustered by subject.

this prediction.

4 Experiment 2: Number Classification Task

Here we report results from a second 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 imprecision 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.

4.1 Experimental Design

Our method for studying awareness of imprecision is to create a simplified version of the previous experiment, 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 the previous experiment; the main 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. 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.¹² 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 between Experiment 1 and Experiment 2 involves decision speed. Here, in Experiment 2, we penalize subjects for the time it takes them to respond. The reason we impose the speed incentive in Experiment 2 comes from the well known “speed-accuracy tradeoff” 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, 2021; Hébert and Woodford, Forthcoming)¹³. We ask subjects to consider the 149

¹²The experimental instructions are available in Appendix D.

¹³For example, an alternative model of imperfect perception that does not feature the property that

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.¹⁴ 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:

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

nearby states are harder to distinguish than far away states is proposed in Gul, Pesendorfer and Strzalecki (2017).

¹⁴Following 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, 2020).

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.

4.2 Experimental Results

The upper panel of Figure 7 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, Forthcoming). Similar patterns are reflected in the response times shown in the lower panel of Figure 7: 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 8, 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 8 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 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 8 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 1.¹⁵ Such a test investigates the hypothesis that subjects are aware that

¹⁵Pre-registration document is available at https://aspredicted.org/OGG_XNK.

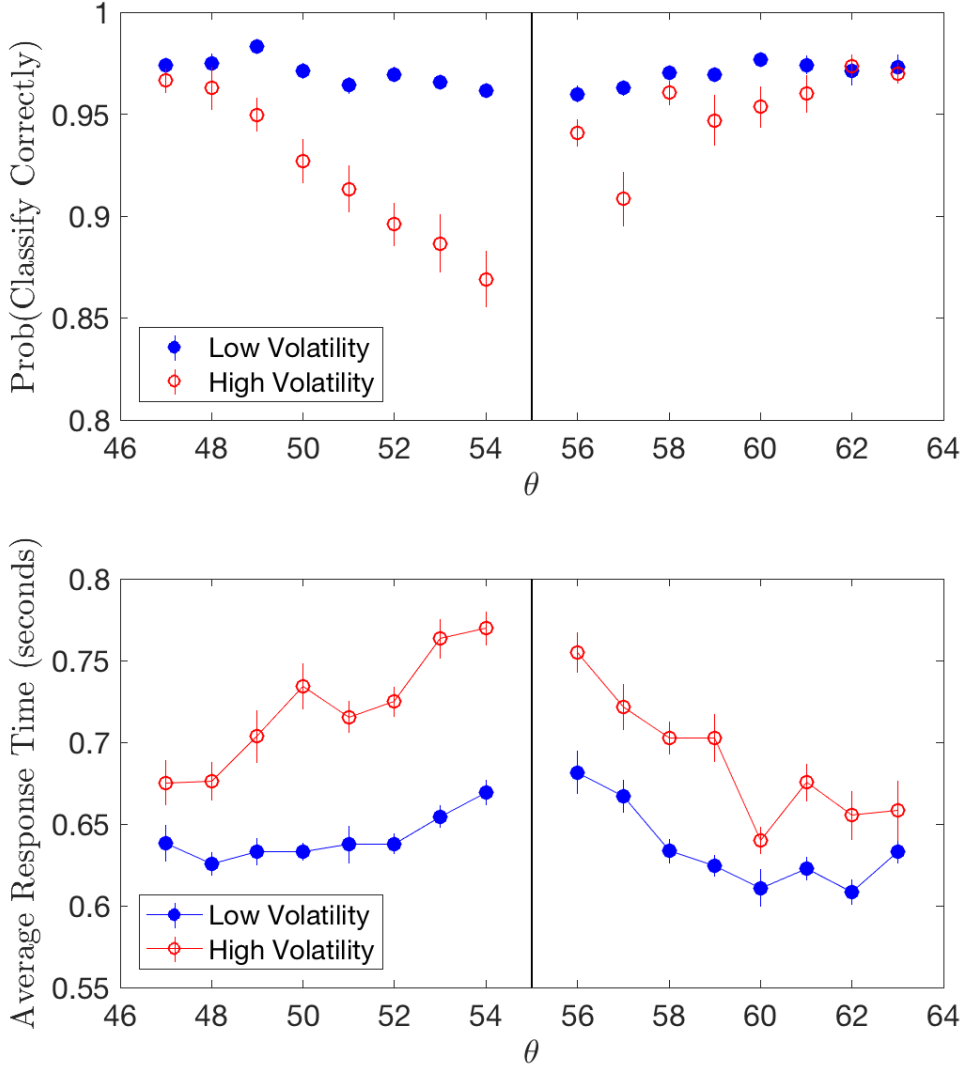


Figure 7: **Accuracy and Response Times in the Classification Task in Experiment 2.** Note: Upper panel shows the proportion of rounds on which subjects correctly classify θ as greater than or less than the reference level of 55. Lower panel shows the average response time on rounds where subjects correctly classify θ . In both panels, the vertical bars denote two standard errors of the mean. Standard errors are clustered by subject.

others' perception of a given number varies as a function of the experienced distribution. Indeed, the right panel of Figure 8 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 ($p = 0.048$).



Figure 8: **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 all panels, vertical bars denote two standard errors of the mean.

5 Discussion

5.1 Interpretation of Signals in Global Games

One theme that emerges from both our theoretical and experimental analyses is that the noise that is assumed in models of global games can be interpreted literally as irreducible error stemming from perceptual constraints. This theme is related to the idea from Heinemann, Nagel and Ockenfels (2009), that behavior in a complete information coordination game can be interpreted *as if* players are observing a fundamental parameter with noise. Like us, Heinemann, Nagel and Ockenfels (2009) structurally estimate a global games model, and find a sizable standard deviation of private signals. However, Heinemann, Nagel and Ockenfels (2009) argue that the (only) source of the estimated standard deviation of private signals is strategic uncertainty. In contrast, we argue that the standard deviation of private signals is driven by cognitive imprecision. By adopting an “as is” interpretation of noise in private signals, we are able to generate and test novel hypotheses about how the standard deviation of private signals varies across environments.

Another important implication for the literature on global games has to do with the role of public vs. private signals. A series of papers has argued that when an institution like the government or a financial market can generate public signals, then a unique equilibrium may no longer obtain in a global games model (Angeletos and Werning 2006, Atkeson 2000,

Hellwig, Mukherji and Tsyvinski 2006). The argument is that a sufficiently precise public signal can act as a coordination device, and thus restore multiple equilibria. However, our theory and experimental results suggest that there is an important difference between *access* to a public signal and precise perception of public signal. Specifically, even if all players have access to the public signal, each player’s perceptual system will process the same public signal slightly differently. This perceptual friction therefore transforms the public signal into private information, thus making it difficult to use the public signal as a coordination device. Our results therefore imply that the provision of a public signal is not enough to overturn the classic global games result. The ability to precisely perceive public information is also necessary and, as we have shown, this cannot be taken for granted.

5.2 Efficient Coding Assumption

Here we revisit the assumption about efficient coding in our model. The specific performance objective that we assume in Section 2 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 1, 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 Appendix B, we show that the result in Proposition 1 is robust to using this alternative objective function. Specifically, we maintain the constraints in Assumption 2 and we analyze a two-stage game. In the first stage, each player optimally chooses 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 1. Thus, our theoretical predictions 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.

5.3 Comparison with Quantal Response Equilibrium

In our model, stochastic perception of θ immediately generates stochastic strategic behavior. As such, our model is related to Quantal Response Equilibrium (QRE; McKelvey and Palfrey 1995, 1998, Goeree, Holt and Palfrey 2016), which is a workhorse model of stochastic

behavior in experimental game theory.¹⁶ In QRE, the utility associated with a given action is subject to a random shock, and each player best responds given the distribution of these shocks. For some parameter values, the models of QRE and cognitive imprecision deliver similar predictions, in that both theories predict that the probability of investing decreases monotonically in θ .

However, there is one fundamental difference in the assumptions of the two theories, which generates distinguishing predictions. When applying QRE to our experimental games, player i 's utility from not investing is $\theta + \epsilon_i$, where ϵ_i is mean zero noise. Crucially, QRE assumes that player i knows both θ and his private shock ϵ_i . This implies that player i believes that his opponent's utility from not investing is centered at θ (even if player i does not know player j 's shock, ϵ_j). In this sense, θ is common knowledge. In contrast, in our model of cognitive imprecision, the stochastic element of the theory is embedded in the *perception* of θ , rather than in the *utility* of not investing. Therefore, in cognitive imprecision, player i does not observe θ , but a noisy signal $Z_i = \theta + \epsilon_i$. This implies that player i believes that his opponent's perceived utility of not investing is centered at $\theta + \epsilon_i$. Thus, in contrast to QRE, there is no common knowledge of θ in cognitive imprecision.

This difference in assumptions leads to two important distinguishing predictions. The first difference is that, in QRE, each player has precise perception of θ , and thus there is no role for a prior distribution over θ . The prior does, however, play a key role in cognitive imprecision, since it impacts the perception of θ and, more importantly, player i 's beliefs about player j 's perception of θ . Our main result, displayed in Figure 4, clearly shows that the prior distribution has a systematic effect on behavior, which supports the prediction of cognitive imprecision.

While the noise structure in QRE is almost always taken to be exogenous, efficient coding may offer one such source for noisy responses in QRE. Interestingly, in the original QRE paper, McKelvey and Palfrey (1995) propose that “...to the extent that we can find observable independent variables that a priori one would expect to be correlated with the precision of these [expected payoff] estimates, one can make predictions about the effects of different experimental treatments that systematically vary these independent variables” (pg. 7-8). Efficient coding provides one such independent variable, which is the volatility of the payoff distribution.¹⁷

The second difference between QRE and cognitive imprecision involves the theoretical

¹⁶For other models of strategic interaction with stochastic choice, see Goeree and Holt (2004), Friedman and Mezzetti (2005), Gonçalves (2020), and Goeree and Louis (Forthcoming).

¹⁷In related work, Friedman (2020) proposes a model that endogenizes the precision parameter in QRE, though it is the set of payoffs in the *current game* that determine the precision parameter — rather than the distribution of payoffs within a class of games, as in our model.

conditions that are sufficient to generate a unique equilibrium. As shown in Proposition 3, cognitive imprecision generates a unique equilibrium when the noise in perception is sufficiently small. One interpretation of this condition, is that when players pay sufficient attention to the coordination game, so that the variance of the internal representation Z_i is below a threshold, then uniqueness obtains under our theory of cognitive imprecision. In contrast, QRE delivers a unique equilibrium when the variance of the shock to utility is sufficiently *large* (Ui, 2006). While our data do not enable us to test between this difference in conditions for uniqueness, one implication is that when players devote a substantial amount of attention to the coordination game, the multiplicity of equilibria is more likely to be eliminated under cognitive imprecision, compared with QRE.

5.4 Comparison with Level-k Thinking

Our results also relate to another behavioral theory of games called Level-k Thinking (Stahl and Wilson, 1994, 1995; Nagel, 1995). In one prominent version of this theory, there are different types of players, and each type best responds to another type who exhibits one less degree of strategic sophistication. For example, a Level-0 type would be characterized by no strategic sophistication and, thus, would exhibit purely random behavior. A Level-1 type would then best respond to a Level-0 player, and a Level-2 player would best respond to a Level-1 player, and so on. What are the predictions of Level-k Thinking for the game in our first experiment? Given that Level-0 players randomize, the expected utility of a Level-1 player from Invest is

$$EU_{L1}(\text{Invest}) = \frac{1}{2}a + \frac{1}{2}b$$

Thus, $EU_{L1}(\text{Invest}) > EU(\text{Not Invest})$ if and only if $\theta < (a + b)/2$. Next, under the assumption that Level-2 players believe they are facing a Level-1 opponent, the expected utility from Invest for a Level-2 player is

$$EU_{L2}(\text{Invest}) = \begin{cases} b & \text{if } \theta < (a + b)/2 \\ a & \text{if } \theta > (a + b)/2 \end{cases}$$

When $\theta < (a + b)/2$, then $EU_{L2}(\text{Invest}) = b > \theta$. Conversely, when $\theta > (a + b)/2$, then $EU_{L2}(\text{Invest}) = a < \theta$. Thus, Level-2 players choose Invest if and only if $\theta < (a + b)/2$. Using the same logic, we obtain the same prediction for all upper levels.

In sum, the fraction of subjects who choose Invest is:

$$\Pr[\text{Invest}] = \begin{cases} \Pr[L_0]\frac{1}{2} + (1 - \Pr[L_0]) & \text{if } \theta < (a + b)/2 \\ \Pr[L_0]\frac{1}{2} & \text{if } \theta > (a + b)/2 \end{cases}$$

where $\Pr[L_0]$ is the fraction of Level-0 players in the population. The theory therefore predicts that, in the aggregate, the probability of investing is monotone in θ and exhibits a sharp decrease at $\theta = (a + b)/2$. However, Level-k Thinking does not predict any difference across our experimental treatments; thus the theory would need to be augmented with some extra feature in order to explain the clear context-dependence we observe in our data.

6 Conclusion

We have provided and experimentally validated a framework for analyzing strategic behavior when players have imprecise perception of a fundamental parameter. Our results are in line with previous experiments on global games, which find evidence consistent with an equilibrium where all players invest once the fundamental crosses a threshold (Heinemann, Nagel and Ockenfels, 2004, 2009; Cabrales, Nagel and Armenter, 2007; Szkup and Trevino, 2020; Goryunov and Rigos, 2020). At the same time, our experimental data suggest that the predictions from the global games literature may be more applicable than previously thought: even when there is no explicit private information given to players, imprecise perception can serve as a source of private information. Interestingly, the particular manner in which we model imprecise perception is closely connected to the noise structure used in the global games literature to select a unique equilibrium.

We also find empirical evidence of *context-dependent* strategic behavior, which is consistent with efficient coding. This context-dependence rules out several alternative game-theoretic models. We argue that the unstable strategic behavior that we observe across experimental conditions is a consequence of the efficient use of cognitive resources. In our setting of a 2×2 coordination game, efficient coding provides a mechanism that modulates the probability that two players coordinate and play the same action. One important direction for future research is to understand the implications of cognitive imprecision in more general games. The idea that “public” payoffs may be perceived as private information, is likely to have important implications for other strategic behavior outside the coordination games we study here.

References

- Angeletos, George-Marios and Chen Lian**, “Incomplete Information in Macroeconomics: Accommodating Frictions in Coordination,” in John B. Taylor and Harald Uhlig, eds., *Handbook of Macroeconomics Vol. 2*, Elsevier, 2016.
- **and Iván Werning**, “Crises and Prices: Information Aggregation, Multiplicity, and Volatility,” *American Economic Review*, 2006, *96* (5), 1720–1736.
- Atkeson, Andrew**, “Rethinking Multiple Equilibria in Macroeconomic Modeling: Comment,” *NBER Macroeconomics Annual*, 2000, *15*, 162–171.
- Avoyan, Ala**, “Communication in Global Games: Theory and Experiment,” 2019.
- Cabrales, Antonio, Rosemarie Nagel, and Roc Armenter**, “Equilibrium Selection through Incomplete Information in Coordination Games: An Experimental Study,” *Experimental Economics*, 2007, *10* (3), 221–234.
- Carlsson, Hans and Eric Van Damme**, “Global Games and Equilibrium Selection,” *Econometrica*, 1993, *61* (5), 989–1018.
- Chambers, Christopher P and Paul J Healy**, “Updating Toward the Signal,” *Economic Theory*, 2012, *50* (3), 765–786.
- Danz, David, Lise Vesterlund, and Alistair J Wilson**, “Belief Elicitation: Limiting Truth Telling with Information on Incentives,” 2020. NBER Working Paper No. 27327.
- Dehaene, Stanislas**, *The Number Sense: How the Mind Creates Mathematics*, Oxford University Press, 2011.
- **, 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.
- Diamond, Douglas W and Philip H Dybvig**, “Bank Runs, Deposit Insurance, and Liquidity,” *Journal of Political Economy*, 1983, *91* (3), 401–419.
- Edmond, Chris**, “Information Manipulation, Coordination, and Regime Change,” *Review of Economic Studies*, 2013, *80* (4), 1422–1458.
- Enke, Benjamin and Thomas Graeber**, “Cognitive Uncertainty,” 2021.

- Friedman, Evan**, “Endogenous Quantal Response Equilibrium,” *Games and Economic Behavior*, 2020, *124*, 620–643.
- Friedman, James W and Claudio Mezzetti**, “Random Belief Equilibrium in Normal Form Games,” *Games and Economic Behavior*, 2005, *51* (2), 296–323.
- Frydman, Cary and Lawrence J Jin**, “Efficient Coding and Risky Choice,” *Quarterly Journal of Economics*, Forthcoming, *Forthcoming*.
- Gabaix, Xavier and David Laibson**, “Myopia and Discounting,” 2017. NBER Working Paper No. 23254.
- Gershman, Samuel J and Rahul Bhui**, “Rationally Inattentive Intertemporal Choice,” *Nature Communications*, 2020, *11* (1), 1–8.
- Girshick, Ahna R, Michael S Landy, and Eero P Simoncelli**, “Cardinal Rules: Visual Orientation Perception Reflects Knowledge of Environmental Statistics,” *Nature Neuroscience*, 2011, *14* (7), 926–932.
- Goeree, Jacob K and Charles A Holt**, “A Model of Noisy Introspection,” *Games and Economic Behavior*, 2004, *46* (2), 365–382.
- **and Philippos Louis**, “M Equilibrium: A Theory of Beliefs and Choices in Games,” *American Economic Review*, Forthcoming.
- **, Charles A Holt, and Thomas R Palfrey**, *Quantal Response Equilibrium: A Stochastic Theory of Games*, Princeton University Press, 2016.
- Goldstein, Itay and Ady Pauzner**, “Demand–Deposit Contracts and the Probability of Bank Runs,” *The Journal of Finance*, 2005, *60* (3), 1293–1327.
- Gonçalves, Duarte**, “Sequential Sampling and Equilibrium,” 2020.
- Goryunov, Maxim and Alexandros Rigos**, “Discontinuous and Continuous Stochastic Choice and Coordination in the Lab,” 2020.
- Gul, Faruk, Wolfgang Pesendorfer, and Tomasz Strzalecki**, “Coarse competitive equilibrium and extreme prices,” *American Economic Review*, 2017, *107* (1), 109–37.
- 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*, Forthcoming.
- Heinemann, Frank, Rosemarie Nagel, and Peter Ockenfels**, “The Theory of Global Games on Test: Experimental Analysis of Coordination Games with Public and Private Information,” *Econometrica*, 2004, 72 (5), 1583–1599.
- , —, and —, “Measuring Strategic Uncertainty in Coordination Games,” *Review of Economic Studies*, 2009, 76 (1), 181–221.
- Hellwig, Christian, Arijit Mukherji, and Aleh Tsyvinski**, “Self-Fulfilling Currency Crises: The Role of Interest Rates,” *American Economic Review*, 2006, 96 (5), 1769–1787.
- Huyck, John Van, Ajalavat Viriyavipart, and Alexander L Brown**, “When Less Information is Good Enough: Experiments with Global Stag Hunt Games,” *Experimental Economics*, 2018, 21 (3), 527–548.
- and **Dale O Stahl**, “Conditional Behavior and Learning in Similar Stag Hunt Games,” *Experimental Economics*, 2018, 21 (3), 513–526.
- Khaw, Mel Win, Ziang Li, and Michael Woodford**, “Cognitive Imprecision and Small-Stakes Risk Aversion,” *Review of Economic Studies*, 2021, 88 (4), 1979–2013.
- 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.
- McKelvey, Richard D and Thomas R Palfrey**, “Quantal Response Equilibria for Normal Form Games,” *Games and Economic Behavior*, 1995, 10 (1), 6–38.
- and —, “Quantal Response Equilibria for Extensive Form Games,” *Experimental Economics*, 1998, 1 (1), 9–41.
- Morris, Stephen**, “Global Games,” in Steven Durlauf and Lawrence E. Blume, eds., *Game Theory*, Plagrave Macmillan, 2010.
- and **Hyun Song Shin**, “Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks,” *American Economic Review*, 1998, pp. 587–597.

- **and** –, “Global Games: Theory and Applications,” in Mathias Dewatripont, Lars Peter Hansen, and Stephen J. Turnovsky, eds., *Advances in Economics and Econometrics (Proceedings of the Eighth World Congress of the Econometric Society)*, Cambridge University Press, 2003.
- **and Ming Yang**, “Coordination and Continuous Stochastic Choice,” 2021.
- Nagel, Rosemarie**, “Unraveling in Guessing Games: An Experimental Study,” *The American Economic Review*, 1995, 85 (5), 1313–1326.
- , **Andrea Brovelli, Frank Heinemann, and Giorgio Coricelli**, “Neural Mechanisms Mediating Degrees of Strategic Uncertainty,” *Social Cognitive and Affective Neuroscience*, 2018, 13 (1), 52–62.
- Obstfeld, Maurice**, “Models of Currency Crises with Self-Fulfilling Features,” *European Economic Review*, 1996, 40 (3-5), 1037–1047.
- Payzan-LeNestour, Elise and Michael Woodford**, “Outlier Blindness: A Neurobiological Foundation for Neglect of Financial Risk,” *Journal of Financial Economics*, Forthcoming.
- Polania, Rafael, Michael Woodford, and Christian C Ruff**, “Efficient Coding of Subjective Value,” *Nature Neuroscience*, 2019, 22 (1), 134–142.
- Ratcliff, Roger**, “A Theory of Memory Retrieval,” *Psychological Review*, 1978, 85 (2), 59.
- Rick, Scott and Roberto A Weber**, “Meaningful Learning and Transfer of Learning in Games Played Repeatedly without Feedback,” *Games and Economic Behavior*, 2010, 68 (2), 716–730.
- Stahl, Dale O and Paul W Wilson**, “Experimental Evidence on Players’ Models of Other Players,” *Journal of Economic Behavior and Organization*, 1994, 25 (3), 309–327.
- **and** –, “On Players’ Models of Other Players: Theory and Experimental Evidence,” *Games and Economic Behavior*, 1995, 10 (1), 218–254.
- Szkup, Michal and Isabel Trevino**, “Sentiments, Strategic Uncertainty, and Information Structures in Coordination Games,” *Games and Economic Behavior*, 2020, 124, 534–553.
- Ui, Takashi**, “Correlated Quantal Responses and Equilibrium Selection,” *Games and Economic Behavior*, 2006, 57 (2), 361–369.

- Weber, Roberto A**, “Learning with No Feedback in a Competitive Guessing Game,” *Games and Economic Behavior*, 2003, 44 (1), 134–144.
- Wei, Xue-Xin and Alan A Stocker**, “A Bayesian Observer Model Constrained by Efficient Coding Can Explain ‘Anti-Bayesian’ Percepts,” *Nature Neuroscience*, 2015, 18 (10), 1509.
- Woodford, Michael**, “Modeling Imprecision in Perception, Valuation, and Choice,” *Annual Review of Economics*, 2020, 12, 579–601.

A Proofs

Proof of Proposition 1

Here we adapt the theoretical derivation of efficient coding from Khaw, Li and Woodford (2021) to our framework where the distribution of θ is normal rather than lognormal. According to Assumption 1, the internal representation S of θ is drawn from

$$S|\theta \sim N(m(\theta), \sigma_S^2) \quad (3)$$

where the *encoding rule*, $m(\theta)$, is a linear transformation of θ , $m(\theta) = \xi + \psi\theta$, which satisfies the power constraint in Assumption 2. Parameters ξ and ψ are endogenous while the precision parameter σ_S is exogenous. The efficient coding hypothesis requires that the encoding rule $m(\theta)$ is chosen (among all linear functions satisfying the constraint) so as to maximize the system’s objective function, for a given prior distribution of θ . As in Khaw, Li and Woodford (2021), we assume that the system produces an estimate of θ on the basis of S , $\tilde{\theta}(S)$, and that the goal of the design problem is to have a system that achieves as low as possible a mean squared error of this estimate. Given a noisy internal representation, the estimate which minimizes the mean squared error is $E[\theta|S]$ for all S . The goal of the design problem is, thus, to minimize the variance of the posterior distribution of θ .

Consider the transformed internal representation, $Z \equiv (S - \xi)/\psi$. The distribution of the transformed internal representation conditional on θ is $Z|\theta \sim N(\theta, \sigma_S^2/\psi^2)$. Thus, the distribution of θ given the transformed internal representation is

$$\theta|Z \sim N\left(\mu_\theta + \frac{\sigma_\theta^2}{\sigma_\theta^2 + (\sigma_S^2/\psi^2)}(Z - \mu_\theta), \frac{\sigma_\theta^2(\sigma_S^2/\psi^2)}{\sigma_\theta^2 + (\sigma_S^2/\psi^2)}\right) \quad (4)$$

The variance of the posterior distribution of θ is strictly increasing in the variance of

Z , σ_S^2/ψ^2 . Thus, it is desirable to make ψ as large as possible (in order to make the mean squared error of the estimate as small as possible) consistent with the power constraint. When the distribution of θ is normal, we have

$$E[m^2] = \xi^2 + \psi^2 E[\theta^2] + 2\xi\psi E[\theta] = (\xi + \psi\mu_\theta)^2 + \psi^2\sigma_\theta^2 \leq \Omega \quad (5)$$

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

$$\xi = -\psi\mu_\theta, \quad \psi = \frac{\Omega}{\sigma_\theta} \quad (6)$$

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

$$Z|\theta \sim N\left(\theta, \frac{\sigma_S^2}{\Omega^2}\sigma_\theta^2\right) \quad (7)$$

The same optimal coding rule obtains under an alternative goal of the system. Consider the more conventional hypothesis from sensory perception literature, whereby the encoding rule is assumed to maximize the Shannon mutual information between the objective state θ and its subjective representation S . Denote with ρ_θ the precision of θ and with ρ_S the precision of S . We have $\theta \sim N\left(\mu_\theta, \frac{1}{\rho_\theta}\right)$, $S|\theta \sim N\left(\xi + \psi\theta, \frac{1}{\rho_S}\right)$, $Z|\theta \sim N\left(\theta, \frac{1}{\rho_Z}\right)$, and $\theta|Z \sim N\left(\frac{\rho_\theta\mu_\theta + \rho_Z Z}{\rho_\theta + \rho_Z}, \frac{1}{\rho_\theta + \rho_Z}\right)$, where $Z = \frac{S-\xi}{\psi}$ and $\rho_Z = \psi^2/\sigma_S^2$. The Shannon mutual information between θ and Z is

$$I(\theta, Z) = \frac{1}{2}\log_2\left(\frac{\sigma_\theta^2}{\sigma_{\theta|Z}^2}\right) = \frac{1}{2}\log_2\left(1 + \frac{\rho_Z}{\rho_\theta}\right) \quad (8)$$

which is strictly increasing in ρ_Z and, thus, strictly decreasing in σ_Z^2 . This means that, as for the previous goal, it is desirable to make ψ as large as possible (consistent with the power constraint).

Proof of Proposition 2

First, we show that, when the conditions in the statement of the Proposition are satisfied, there exists a unique monotone equilibrium of the game. Remember that $Z_i \sim N(\theta, \sigma_Z^2)$, where $\sigma_Z^2 = \omega\sigma_\theta^2 = (\sigma_S^2/\Omega^2)\sigma_\theta^2$. Thus, player 1's posterior distribution of θ given Z_1 is

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

Therefore, we have:

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

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 . In this case, player 1's expectation that the opponent invests is $\Pr[Z_2 \leq k|Z_1]$. Player 1's belief about the distribution of Z_2 given Z_1 is:

$$Z_2|Z_1 \sim \mathcal{N}\left(E[\theta|Z_1] = \frac{\sigma_Z^2}{\sigma_\theta^2 + \sigma_Z^2}\mu_\theta + \frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_Z^2}Z_1, \frac{2\sigma_\theta^2\sigma_Z^2 + \sigma_Z^4}{\sigma_\theta^2 + \sigma_Z^2}\right)$$

Thus, we have:

$$\Pr[Z_2 \leq k|Z_1] = \Phi\left(\frac{(\sigma_\theta^2 + \sigma_Z^2)k - \sigma_Z^2\mu_\theta - \sigma_\theta^2 Z_1}{\sqrt{2\sigma_\theta^2\sigma_Z^2 + \sigma_Z^4}\sqrt{\sigma_\theta^2 + \sigma_Z^2}}\right)$$

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

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

$$\frac{\sigma_Z^2 \mu_\theta + \sigma_\theta^2 Z_1}{\sigma_\theta^2 + \sigma_Z^2} \leq a + (b - a)\Phi\left(\frac{(\sigma_\theta^2 + \sigma_Z^2)k - \sigma_Z^2\mu_\theta - \sigma_\theta^2 Z_1}{\sqrt{2\sigma_\theta^2\sigma_Z^2 + \sigma_Z^4}\sqrt{\sigma_\theta^2 + \sigma_Z^2}}\right) \quad (9)$$

If we write $\bar{Z}(k)$ for the unique value of Z_1 such that player 1 is indifferent between investing and not investing (this is well defined since player 1's expected payoff from not investing is strictly increasing in Z_1 and player 1's expected payoff from investing is strictly decreasing in Z_1), the best response of player 1 is to follow a monotone strategy with threshold equal to $\bar{Z}(k)$, that is, to invest if and only if $Z_1 \leq \bar{Z}(k)$.

Observe that as $k \rightarrow -\infty$ (that is, player 2 never invests), $EU[\text{Invest}|Z_1, k]$ tends to a , so $\bar{Z}(k)$ tends to $\frac{(\sigma_\theta^2 + \sigma_Z^2)a - \sigma_Z^2\mu_\theta}{\sigma_\theta^2}$. As $k \rightarrow \infty$ (that is, player 2 always invests), $EU[\text{Invest}|Z_1]$ tends to b , so $\bar{Z}(k)$ tends to $\frac{(\sigma_\theta^2 + \sigma_Z^2)b - \sigma_Z^2\mu_\theta}{\sigma_\theta^2}$. A fixed point of $\bar{Z}(k)$ — that is a value k^* such that $\bar{Z}(k^*) = k^*$ — is a monotone equilibrium of the game where each player invests if and only if his signal is below k^* . Since $\bar{Z}(k)$ is a mapping from \mathbb{R} to itself and is continuous in k , there exists $k \in \left[\frac{(\sigma_\theta^2 + \sigma_Z^2)a - \sigma_Z^2\mu_\theta}{\sigma_\theta^2}, \frac{(\sigma_\theta^2 + \sigma_Z^2)b - \sigma_Z^2\mu_\theta}{\sigma_\theta^2}\right]$, such that $\bar{Z}(k) = k$ and a threshold equilibrium of this game exists.

When is there a unique equilibrium? Define $W(\bar{Z}(k), k)$ as

$$W(\bar{Z}(k), k) = \frac{\sigma_Z^2 \mu_\theta + \sigma_\theta^2 \bar{Z}(k)}{\sigma_\theta^2 + \sigma_Z^2} - a - (b - a) \Phi \left(\frac{(\sigma_\theta^2 + \sigma_Z^2) k - \sigma_Z^2 \mu_\theta - \sigma_\theta^2 \bar{Z}(k)}{\sqrt{2\sigma_\theta^2 \sigma_Z^2 + \sigma_Z^4} \sqrt{\sigma_\theta^2 + \sigma_Z^2}} \right)$$

At a fixed point, $\bar{Z}(k^*) = k^*$. Thus, we have:

$$W(k^*) = \frac{\sigma_Z^2 \mu_\theta + \sigma_\theta^2 k^*}{\sigma_\theta^2 + \sigma_Z^2} - a - (b - a) \Phi \left(\frac{\sigma_Z^2}{\sqrt{2\sigma_\theta^2 \sigma_Z^2 + \sigma_Z^4} \sqrt{\sigma_\theta^2 + \sigma_Z^2}} (k^* - \mu_\theta) \right)$$

Then,

$$\frac{\partial W(k^*)}{\partial k^*} = \frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_Z^2} - \phi \left(\frac{\sigma_Z^2}{\sqrt{2\sigma_\theta^2 \sigma_Z^2 + \sigma_Z^4} \sqrt{\sigma_\theta^2 + \sigma_Z^2}} (k^* - \mu_\theta) \right) \frac{\sigma_Z^2 (b - a)}{\sqrt{2\sigma_\theta^2 \sigma_Z^2 + \sigma_Z^4} \sqrt{\sigma_\theta^2 + \sigma_Z^2}}$$

And there is a unique fixed point if and only if $\frac{\partial W(k^*)}{\partial k^*} > 0$ at the fixed point. When $\frac{\partial W(k^*)}{\partial k^*} < 0$, there are at least three fixed points. Since $\phi(y) \leq \frac{1}{\sqrt{2\pi}}$, this is a sufficient condition for $\frac{\partial W(k^*)}{\partial k^*} > 0$:

$$\begin{aligned} \frac{\sigma_\theta^2}{\sigma_\theta^2 + \sigma_Z^2} &> \frac{1}{\sqrt{2\pi}} \frac{\sigma_Z^2 (b - a)}{\sqrt{2\sigma_\theta^2 \sigma_Z^2 + \sigma_Z^4} \sqrt{\sigma_\theta^2 + \sigma_Z^2}} \\ \frac{\sigma_\theta^2 \sqrt{2\sigma_\theta^2 \sigma_Z^2 + \sigma_Z^4}}{(b - a) \sigma_Z^2 \sqrt{\sigma_\theta^2 + \sigma_Z^2}} &> \frac{1}{\sqrt{2\pi}} \\ \sqrt{2\pi} &> \frac{(b - a) \sigma_Z^2 \sqrt{\sigma_\theta^2 + \sigma_Z^2}}{\sigma_\theta^2 \sqrt{2\sigma_\theta^2 \sigma_Z^2 + \sigma_Z^4}} \end{aligned}$$

The condition $\sqrt{\omega(1 + \omega)} < \frac{\sqrt{6\pi}}{(b-a)} \sigma_\theta^2$ is obtained by replacing $\sigma_Z = \omega \sigma_\theta$ in the condition above and re-arranging terms. Thus, this shows that, when the conditions in the statement of the Proposition are satisfied, there exists a unique monotone equilibrium of the game.

Second, we show that, when $\mu_\theta = \frac{(a+b)}{2}$, there exists a monotone equilibrium of the game where $k^* = \mu_\theta$ for any value of σ_θ , σ_S and ω (or, equivalently, for any value of σ_θ and σ_Z). Assume player 2 uses a threshold strategy where he invests if and only if $Z_2 \leq k = \mu_\theta$. Is this an equilibrium, that is, is $\bar{Z}(\mu_\theta) = \mu_\theta$? $\bar{Z}(\mu_\theta)$ is the value of Z_1 such that the following equation is satisfied with equality:

$$\frac{\sigma_Z^2 \mu_\theta + \sigma_\theta^2 Z_1}{\sigma_\theta^2 + \sigma_Z^2} = a + (b - a) \Phi \left(\frac{(\sigma_\theta^2 + \sigma_Z^2) k - \sigma_Z^2 \mu_\theta - \sigma_\theta^2 Z_1}{\sqrt{2\sigma_\theta^2 \sigma_Z^2 + \sigma_Z^4} \sqrt{\sigma_\theta^2 + \sigma_Z^2}} \right)$$

$$\frac{\sigma_Z^2 \mu_\theta + \sigma_\theta^2 Z_1}{\sigma_\theta^2 + \sigma_Z^2} = a + (b - a) \Phi \left(\frac{\sigma_\theta^2 \mu_\theta - \sigma_\theta^2 Z_1}{\sqrt{2\sigma_\theta^2 \sigma_Z^2 + \sigma_Z^4} \sqrt{\sigma_\theta^2 + \sigma_Z^2}} \right)$$

If we set $Z_1 = \mu_\theta$, we get:

$$\mu_\theta = a + (b - a) \Phi(0)$$

$$\mu_\theta = \frac{(a + b)}{2}$$

which is true by one of the assumptions in the statement of the Proposition.

Proof of Proposition 3

From Proposition 2 and the condition in the statement of Proposition 3, we know that there exists a unique monotone equilibrium of the game where each player invests if and only if his transformed internal representation is smaller than μ_θ . In this equilibrium, $Pr[Invest|\theta] = Pr[Z_i \leq \mu_\theta|\theta] = \Phi\left(\frac{\mu_\theta - \theta}{\omega\sigma_\theta}\right)$ and $\frac{\partial Pr[Invest|\theta]}{\partial \theta} = -\phi\left(\frac{\mu_\theta - \theta}{\omega\sigma_\theta}\right)\left(\frac{1}{\omega\sigma_\theta}\right)$. Thus, $Pr[Invest|\theta]$ grows with σ_θ if $\theta < \mu_\theta$ and it decreases with σ_θ if $\theta > \mu_\theta$. Moreover, the sensitivity of choices to θ decreases with σ_θ for values of θ around the cutoff.

Indeed, we have

$$\begin{aligned} \frac{\partial Pr[Invest|\theta]}{\partial \theta \partial \sigma_\theta} &= \phi\left(\frac{\mu_\theta - \theta}{\omega\sigma_\theta}\right)\left(\frac{1}{\omega\sigma_\theta^2}\right) + \phi'\left(\frac{\mu_\theta - \theta}{\omega\sigma_\theta}\right)\left(\frac{\mu_\theta - \theta}{\omega\sigma_\theta^2}\right)\left(\frac{1}{\omega\sigma_\theta}\right) \\ &= \phi\left(\frac{\mu_\theta - \theta}{\omega\sigma_\theta}\right)\left(\frac{1}{\omega\sigma_\theta^2}\right) - \left(\frac{\mu_\theta - \theta}{\omega\sigma_\theta}\right)\phi\left(\frac{\mu_\theta - \theta}{\omega\sigma_\theta}\right)\left(\frac{\mu_\theta - \theta}{\omega\sigma_\theta^2}\right)\left(\frac{1}{\omega\sigma_\theta}\right) \\ &= \phi\left(\frac{\mu_\theta - \theta}{\omega\sigma_\theta}\right)\left(\frac{1}{\omega\sigma_\theta^2}\right) - \phi\left(\frac{\mu_\theta - \theta}{\omega\sigma_\theta}\right)\left(\frac{(\mu_\theta - \theta)^2}{\omega^3\sigma_\theta^4}\right) \\ &= \phi\left(\frac{\mu_\theta - \theta}{\omega\sigma_\theta}\right)\left(\frac{\omega^2\sigma_\theta^2 - (\mu_\theta - \theta)^2}{\omega^3\sigma_\theta^4}\right) \end{aligned}$$

which is positive if and only if $(\mu_\theta - \theta)^2 < \omega^2\sigma_\theta^2$.

(In the second line, we used the fact that $\phi'(x) = -x\phi(x)$.)

B Alternative Model of Efficient Coding

Assumption 4 (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, (x_i, ψ_i) , to maximize the performance objective in Assumption 4 under the constraints in Assumption 2; 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 4 *Suppose Assumptions 1, 2, 4 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 2 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 4.

Proposition 5 *Suppose Assumptions 1, 2, 4, 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)?

$$\frac{\partial EU_i(\theta, \psi_i)}{\partial \psi_i} = \phi\left(\frac{\mu_\theta - \theta}{\sigma_S/\psi_i}\right) \left(\frac{\mu_\theta - \theta}{\sigma_S}\right) \left(a + \Phi\left(\frac{\mu_\theta - \theta}{\sigma_S/\psi_{-i}}\right) (b - a) - \theta\right) \quad (10)$$

Since $\phi(\cdot)$ is strictly positive for any argument, the sign of equation (10) 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 (10) 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

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

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

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

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

C Robustness of Monotone Equilibrium with $k^* = \mu_\theta$

Let us introduce the following definitions from Chambers and Healy (2012):

Definition 1 A random variable with cumulative density function F and mean μ is **symmetric** if, for every $a \geq 0$, $F(\mu + a) = 1 - \lim_{x \rightarrow a^-} F(\mu - a)$.

Definition 2 A random variable is **quasiconcave** (or unimodal) if it has a density function f such that for all $x, x' \in \mathbb{R}$ and $\lambda \in (0, 1)$, $f(\lambda x + (1 - \lambda)x') \geq \min\{f(x), f(x')\}$.

Definition 3 The error term ϵ_i satisfies **symmetric dependence** with respect to the random variable θ if, for each realization of θ , $\epsilon_i|\theta$ has a continuous density function $f_{\epsilon_i|\theta}$ satisfying $f_{\epsilon_i|\theta}(\epsilon_i|\mu_\theta + a) = f_{\epsilon_i|\theta}(\epsilon_i|\mu_\theta - a)$ for almost every ϵ_i and a in \mathbb{R} . (Note that error terms that are independent of θ satisfy this definition).

Consider the following assumptions:

$$(A1) \quad S_i = \theta + \epsilon_i$$

$$(A2) \quad E[\theta] = \mu_\theta < \infty$$

$$(A3) \quad \theta \text{ is a symmetric random variable and its density is continuous on } \mathbb{R}$$

$$(A4) \quad E[\epsilon_i|\theta] = 0 \text{ for each } \theta$$

$$(A5) \quad \epsilon_i \text{ is a symmetric and quasiconcave random variable}$$

$$(A6) \quad \epsilon_i \text{ satisfies symmetric dependence with respect to } \theta$$

Lemma 1 (Chambers and Healy 2012, Proposition 2) *Assume A1-A6. A Bayesian agent updates his beliefs over θ in the direction of the signal, that is, for almost every $S_i \in \mathbb{R}$, there exists some $\alpha \geq 0$ such that $E[\theta|S_i] = \alpha S_i + (1 - \alpha)\mu_\theta$.*

Proposition 6 *Assume A1-A6 and $\mu_\theta = (a + b)/2$. There exists a monotone equilibrium of the game where $k^* = \mu_\theta$.*

Proof of Proposition 6 The proof can be carried out with general values for a and b (such that $b > a$). For ease of exposition, we focus on the experimental parameters: $a = 47$, $b = 63$, $\mu_\theta = 55$. Assume that player j uses threshold $k_j = 55$, that is, he invests if and only if $S_j < 55$. We want to show that player i 's best response is to use the same threshold, $k_i = 55$. Player i prefers to invest if and only if $EU[\text{Not Invest}|S_i] < EU[\text{Invest}|S_i, k_j]$. Thus, we want to show that (1) when $S_i = 55$, $EU[\text{Not Invest}|S_i] = EU[\text{Invest}|S_i, k_j = 55]$; (2) when $S_i < 55$, $EU[\text{Not Invest}|S_i] < EU[\text{Invest}|S_i, k_j = 55]$; and (3) when $S_i > 55$, $EU[\text{Not Invest}|S_i] > EU[\text{Invest}|S_i, k_j = 55]$.

By Lemma 1, $EU[\text{Not Invest}|S_i] = E[\theta|S_i] = \alpha S_i + (1 - \alpha)\mu_\theta$ where $\alpha \geq 0$. Note also that $EU[\text{Invest}|S_i, k_j = 55] = 47 + (63 - 47)Pr[S_j < k_j = 55|S_i]$. First, we prove (1). Assume $S_i = 55$. We want to show that $EU[\text{Not Invest}|S_i] = EU[\text{Invest}|S_i, k_j = 55]$. By Lemma 1, $EU[\text{Not Invest}|S_i = 55] = \alpha S_i + (1 - \alpha)\mu_\theta = \alpha(55) + (1 - \alpha)(55) = 55$. Thus, the equality we want to show becomes $55 = 47 + (63 - 47)Pr[S_j < k_j = 55|S_i = 55]$. This equality is satisfied if and only if $Pr[S_j < k_j = 55|S_i = 55] = 1/2$. By A1 and A4 (and linearity of expectation), $E[S_j|S_i] = E[\theta|S_i] = 55$. By A5, the density of $S_j|S_i$ is symmetric. Thus, the probability S_j takes a value below its posterior mean (55) is $1/2$. This proves (1).

Second, we prove (2). Assume $S_i < 55$. By Lemma 1, $EU[\text{Not Invest}|S_i] = \alpha S_i + (1 - \alpha)55$. This is smaller than 55 for any positive α . This also means that, by A1 and A4, $E[S_j|S_i] = E[\theta|S_i] < 55$. The probability that the opponent invests is the posterior probability that his signal is below 55 (given S_i). Since the conditional distribution of the opponent's signal is symmetric around its mean (by A5), the median is equal to the mean. This means that the conditional CDF of the opponent signal equals $1/2$ at the posterior mean, is greater than $1/2$ for values of S_j above the mean and is lower than $1/2$ for values of S_j below the mean. Since the posterior mean of the opponent's signal is lower than 55, the probability that player j 's signal is lower than 55 (conditional on $S_i < 55$) is greater than $1/2$. Thus, $EU[\text{Invest}|S_i, k_j = 55] = 47 + (63 - 47)Pr[S_j < k_j = 55|S_i] > 55$. This proves that $EU[\text{Invest}|S_i, k_j = 55] > 55 > EU[\text{Not Invest}|S_i]$. (3) can be proven analogously.

D Experimental Instructions

Experiment 1

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

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.