

Efficient coding and incentive effects: cognitive precision and effort in risky choice

Abdelaziz Alsharawy* Sheryl Ball† Alec Smith‡

October 21, 2026

Abstract

The brain has limited computational resources with which to make decisions, necessitating efficient representations of information. The principle of efficient coding posits that value representations are context dependent, such that decision-makers perceive values more or less accurately according to the distribution of values in the decision environment. However, it remains unclear whether the level of precision is a fixed biological constraint or is dynamically modulated by motivational factors such as the level of incentives. To address this, we test whether incentives modulate the fidelity of value representations in risky choice across a narrow range payoff environment (low volatility) and a wide range payoff environment (high volatility), with each participant facing real payoffs in one environment and hypothetical payoffs in the other. We find that pairing hypothetical payoffs with the wide range environment exacerbates the loss of payoff sensitivity, amplifying the behavioral signature of efficient coding. The reverse assignment, with real payoffs in the wide range environment, does not significantly attenuate it. The pairing of hypothetical payoffs with the wide range environment is characterized by lower self-reported arousal and faster response times in that environment, consistent with the withdrawal of cognitive effort. Our results suggest that the precision of value representations responds to experienced changes in incentives.

*The Department of Management, Policy, and Community Health, School of Public Health, The University of Texas Health Science Center at Houston.

†The Department of Economics, Virginia Polytechnic Institute and State University.

‡The Department of Economics, Virginia Polytechnic Institute and State University. Corresponding author: alecsmith@vt.edu.

1 Introduction

Understanding how people evaluate risky outcomes is central to behavioral economics and decision science. A growing body of work shows that choices are context dependent and are shaped by the statistical structure of the environment (Bhui and Gershman, 2018; Frydman and Jin, 2022; Gershman et al., 2015; Louie and Glimcher, 2012; Woodford, 2020). This work emphasizes that because neuronal firing is stochastic and the brain’s computational resources are limited, values are represented imprecisely, and choices computed from those representations are noisy (Marois and Ivanoff, 2005; Woodford, 2020). The efficient coding hypothesis addresses this limitation, positing that the brain optimizes resource use by minimizing redundancy in sensory perception (Barlow, 1961).

In the context of risky choice, an implication of efficient coding is that the representation of values is influenced by the distribution of stimuli in the environment. The decision-maker’s sensitivity to marginal changes in objective values depends on how frequently these stimuli occur. While efficient coding offers a compelling account of imprecise valuation in risky choice (Frydman and Jin, 2022; Heng et al., 2020; Khaw et al., 2021), these models assume a fixed capacity constraint. An open question is whether the precision of value representations is determined solely by such biological constraints or whether it is also influenced by motivation and the decision to exert cognitive effort.

In this paper, we study how motivation affects cognition in risky choice. Our question is whether the precision of value representations responds to changes in incentives. We show that the efficient coding models of Heng et al. (2020, henceforth, HWP) and Frydman and Jin (2022, henceforth, FJ) possess a *scale invariance* property: choice probabilities are unchanged when payoffs and the prior to which the decision-maker has adapted are rescaled together (Proposition 1). This property holds for any prior and under any of the three performance objectives HWP and FJ consider (maximizing mutual information, expected financial gain, or accuracy). Under the uniform priors of FJ’s experimental design, the three objectives coincide in a single coding rule.¹ We test this property by extending FJ’s design to vary payoff consequentiality (real versus hypothetical). This manipulation lowers the effective stakes of a decision without altering the displayed payoffs. Because the coding rule depends on neither the scale of payoffs nor the objective being maximized, the model predicts no change in behavior.

Existing evidence on the effect of incentives on lottery choice is nuanced. Higher stakes and real incentives were shown to generally increase risk aversion (Holt and Laury, 2002;

¹The scale-invariance property here is analogous to that of Khaw et al. (2021, henceforth, KLW), though in their model logarithmic encoding renders choice probabilities invariant to a proportional rescaling of payoffs even at a fixed prior.

Bouchouicha and Vieider, 2017; Alsharawy et al., 2021), whereas some other work reported weak or null effects of incentives on risk preferences (Wiseman and Levin, 1996; Taylor, 2013; Hackethal et al., 2023). One interpretation of these findings is that incentives need not uniformly alter preferences, but may instead affect the cognitive processes underlying choice. Camerer and Hogarth (1999) argue that incentives are most likely to matter when additional effort can improve performance, such as in settings involving prediction, judgment, and memory recall. Consistent with this view, Horn and Freund (2022) find that real rewards improve decision quality and reduce decision noise without affecting risk preferences. Moreover, real and hypothetical payoffs can produce different psychological and neural responses (Camerer and Mobbs, 2017), and real incentives elicit stronger activity in reward-encoding regions (Kang et al., 2011). More recently, Enke et al. (2023) show that very high incentives substantially increase response times while only weakly reducing cognitive biases. Similarly, Gneezy et al. (2024) find that increasing stakes can amplify, rather than attenuate, deviations from expected utility, and suggest that differences between hypothetical and real choices may partly reflect differences in cognitive effort and attention. While this earlier work establishes that incentives operate partly through cognitive processes rather than preferences alone, we ask how the precision of risky choice responds to incentives.

In motor and cognitive control tasks, reward recruits effort that suppresses intrinsic neural noise (Manohar et al., 2015). If incentives modulate the precision of value computations in the same way, choice sensitivity should depend on consequentiality, contrary to the scale-invariance prediction. Our design also elicits self-reports of arousal and valence at intervals during the task. We use arousal and response times as markers of cognitive engagement. Arousal reflects the mobilization of cognitive resources (Howells et al., 2010). Response times index the depth of deliberation; faster responses indicate less processing effort (Bhui and Jiao, 2023; Caplin et al., 2020; Konovalov and Krajbich, 2019; Krajbich and Rangel, 2011). These measures provide process-level evidence on effort allocation in decision-making.

We find that cognitive precision is determined by both contextual and motivational factors: low-motivation contexts with hypothetical payoffs amplify the loss of choice sensitivity associated with wider outcome distributions, accompanied by reduced arousal and faster response times, consistent with the withdrawal of cognitive effort. The remainder of the paper proceeds as follows. Section 2 presents the efficient-coding framework and the scale-invariance proposition. Section 3 describes the experimental design and hypotheses. Section 4 reports the behavioral and process-level results. Section 5 discusses implications.

2 Theory

2.1 Efficient coding and value representation in risky choice

Efficient coding originates in sensory neuroscience: perceptual systems allocate limited representational capacity to maximize information about frequently encountered stimuli (Barlow, 1961). A growing literature extends this principle to the representation of economic value (Heng et al., 2020; Frydman and Jin, 2022; Khaw et al., 2021; Woodford, 2020; Louie and Glimcher, 2012; Polania et al., 2019; Webb et al., 2021). To make these concepts more concrete, we review the efficient coding theory of HWP as adapted to risky choice by FJ.

A decision maker (DM) must choose between a risky option, which pays $\$X$ with probability p and $\$0$ with probability $1 - p$, and a certain amount of money $\$C$. In expected utility theory, the DM deterministically selects the risky option if $p \cdot U(X) + (1 - p)U(0) > U(C)$, is indifferent at equality, and otherwise chooses the certain payoff.

Under the efficient coding theory of HWP/FJ, the DM’s decision is determined by the probabilistic output of a population of neurons that imprecisely encode numerical values. Assume the values X and C are drawn from independent probability distributions f_X and f_C , with cumulative distribution functions (CDFs) F_X and F_C , respectively. The nervous system encodes X imprecisely via an internal representation R_X , generated by n binary processing units (“neurons” or “perceptrons”) that each output 1 with probability $\theta_X(X)$ and 0 otherwise. The function $\theta_X(\cdot)$ is the encoding rule for X . The certain payoff C is encoded by a similar process with representation R_C and encoding rule $\theta_C(\cdot)$. We focus below on the derivation here for the risky payoff; the representation of the certain payoff is analogous (see FJ).

Conditional on X , the total number of high-state units R_X is distributed $\text{Binomial}(n, \theta_X(X))$, with likelihood

$$p(R_X|X) = \binom{n}{R_X} \theta_X(X)^{R_X} (1 - \theta_X(X))^{n-R_X}. \quad (1)$$

In general, the optimal encoding rule depends on the objective being maximized. HWP and FJ consider three different objectives: maximizing mutual information between X and its representation R_X , maximizing the probability of an accurate choice in a two-alternative forced choice, and maximizing the expected value of the chosen option. Maximizing mutual information (equivalently, maximizing accuracy) yields

$$\theta_X(X) = \sin^2\left(\frac{\pi}{2} F_X(X)\right). \quad (2)$$

Maximizing expected financial gain yields instead

$$\theta_X(X) = \left[\sin \left(\frac{\pi \int_{-\infty}^X f_X(x)^{\frac{2}{3}} dx}{2 \int_{-\infty}^{\infty} f_X(x)^{\frac{2}{3}} dx} \right) \right]^2. \quad (3)$$

Both rules depend on the prior so each adapts to changes in the statistical structure of the environment.

Given the noisy representations R_X and R_C , the DM applies Bayes' rule to form posterior distributions for X and C . The posterior mean for X takes the form

$$\mathbb{E}[\tilde{X} \mid R_X] = \frac{\int p(R_X \mid X) f_X(X) X dX}{\int p(R_X \mid X) f_X(X) dX}, \quad (4)$$

where the integrals are over the support of f_X (FJ, eq. 13); the posterior mean for C , $\mathbb{E}[\tilde{C} \mid R_C]$ takes the analogous form. Equation (4) expresses Bayes' rule applied to the binomial likelihood in equation (1) and an arbitrary prior; it does not require any restriction on the shape of f_X or the encoding rule θ_X .

Under linear utility, the DM chooses the risky option whenever $p \cdot \mathbb{E}[\tilde{X} \mid R_X] > \mathbb{E}[\tilde{C} \mid R_C]$. Because the representations R_X and R_C are random, the predicted probability of risky choice, conditional on the realized payoffs (X, C) , is obtained by summing over (R_X, R_C) weighted by their binomial likelihoods:

$$\begin{aligned} \Pr(\text{risky} \mid X, C) &= \sum_{R_X=0}^n \sum_{R_C=0}^n \mathbf{1}\left\{p \cdot \mathbb{E}[\tilde{X} \mid R_X] > \mathbb{E}[\tilde{C} \mid R_C]\right\} p(R_X \mid X) p(R_C \mid C) \\ &\quad + \frac{1}{2} \sum_{R_X=0}^n \sum_{R_C=0}^n \mathbf{1}\left\{p \cdot \mathbb{E}[\tilde{X} \mid R_X] = \mathbb{E}[\tilde{C} \mid R_C]\right\} p(R_X \mid X) p(R_C \mid C), \end{aligned} \quad (5)$$

where the second term implements a uniform tie-breaking rule (FJ, eq. 17). Like equation (4), equation (5) is general: it does not depend on the shape of the prior or the specific form of the encoding rule.

Frydman and Jin (2022) show that when pX and C are i.i.d. uniform (specifically, f_X is uniform on $[X_l, X_u]$, f_C is uniform on $[C_l, C_u]$, with $C_l = pX_l$ and $C_u = pX_u$), the optimal coding rule under each of the three objectives (mutual information, accuracy, and expected financial gain) reduces to the same closed form:

$$\theta_X(X) = \sin^2 \left(\frac{\pi}{2} \cdot \frac{X - X_l}{X_u - X_l} \right), \quad \theta_C(C) = \sin^2 \left(\frac{\pi}{2} \cdot \frac{C - C_l}{C_u - C_l} \right). \quad (6)$$

This is FJ's equation (10); they prove this *objective invariance* property in their Online

Appendix A. Under more general priors, the three objectives generally imply different rules.

2.2 Scale invariance

In psychophysics, the discriminability of two magnitudes is thought to depend on their ratio rather than their absolute difference, a regularity known as Weber’s law. Applied to models of risky choice, this implies scale invariance: multiplying every payoff by a common factor $t > 0$ leaves the probability of each choice unchanged, so that only the relative magnitudes of payoffs matter. Khaw et al. (2021, henceforth, K LW) develop a model with this property by design. For a binary lottery that pays X with probability p and 0 otherwise, with C again a certain amount of money, K LW derive internal representations via continuous noisy signals $R_X \sim \mathcal{N}(\log X, \nu^2)$ and $R_C \sim \mathcal{N}(\log C, \nu^2)$, involving log encoding with constant-variance Gaussian noise, combined with log-normal priors over the payoffs. Under this specification, the probability of choosing the risky lottery depends only on the ratio X/C . K LW test the theory via an experiment, finding support and rejecting additive random utility. We next show that the HWP/FJ framework possesses an analogous scale invariance property, for any prior without point masses.

Proposition 1 (Scale invariance). *Fix any one of the three performance objectives of Section 2.1, with optimal encoding rules θ_X, θ_C given by equation (2) (mutual information or accuracy) or equation (3) (expected financial gain). Consider the HWP/FJ model of Section 2.1 with that objective, arbitrary priors f_X, f_C without point masses, capacity $n \in \mathbb{N}$, and linear utility. For any scalar $t > 0$, let the rescaled payoffs be $Y = tX$ and $D = tC$, with rescaled priors $f_Y(y) = (1/t) f_X(y/t)$ and $f_D(d) = (1/t) f_C(d/t)$. Then for every (X, C) ,*

$$\Pr(\text{risky} \mid tX, tC) = \Pr(\text{risky} \mid X, C), \quad (7)$$

where $\Pr(\text{risky} \mid \cdot, \cdot)$ is the choice probability in equation (5).

Proof sketch. The DM’s priors scale accordingly, and the optimal coding rules for the rescaled priors satisfy $\theta_Y(tX) = \theta_X(X)$ and $\theta_D(tC) = \theta_C(C)$, so the implied signal distributions are unchanged. The conditional expectations scale linearly, and the condition $p \cdot \mathbb{E}[\tilde{Y} \mid R_X] > \mathbb{E}[\tilde{D} \mid R_C]$ is therefore equivalent to $p \cdot \mathbb{E}[\tilde{X} \mid R_X] > \mathbb{E}[\tilde{C} \mid R_C]$. See Appendix B for the full proof. $\square$

Remark. Proposition 1 holds objective by objective: for each of the three performance objectives, taken separately, choice probabilities are invariant to a common rescaling of payoffs and priors, for arbitrary priors. Whether choice probabilities also coincide across objectives

is a separate question, settled by FJ’s objective invariance: under the i.i.d. uniform priors of FJ’s design, which we adopt, the three objectives yield the same coding rule, equation (6) (FJ, Online Appendix A), though not under general priors.

In KIW, scale invariance is built into the encoding. A common rescaling of payoffs shifts both log-scale signals by the same constant, which cancels from the choice rule even at a fixed prior. In the HWP/FJ framework the code is not invariant at a fixed prior, since a rescaled payoff occupies a different quantile of the unadapted prior. Proposition 1 establishes invariance when payoffs and the adapted prior are rescaled together. For our design the distinction does not matter. Participants see the same payoff distributions in the real and hypothetical conditions and are told in advance which condition they are in. A participant who treated a hypothetical payoff as worth a fixed fraction t of its face value would therefore scale payoffs and the prior together.

Importantly, Proposition 1 assumes that the encoding capacity n is fixed and exogenous, and therefore the same in the rescaled environment. If n instead responds to the effective stakes of the choice problem, scale invariance fails: lower stakes induce a flatter encoding rule and reduced sensitivity to payoff differences.

3 Experimental design and hypotheses

Section 2 establishes two invariance properties of the efficient coding model with fixed capacity n . Under uniform priors, FJ show that choice probabilities are robust to changing the objective that the coding rule serves, whether mutual information, accuracy, or expected financial gain (objective invariance, Section 2.1). We establish that choice probabilities are also unchanged by a proportional rescaling of payoffs and the adapted prior (scale invariance, Proposition 1).

The model’s primitives are the displayed payoff distributions, the capacity n , and the performance objective. Payment is not among them, so applying the theory to our manipulation requires a position on what making payoffs hypothetical does. It might lower the stakes of a decision, change the objective the decision-maker pursues, or both. If hypothetical payoffs lower the stakes, a displayed payoff of X is worth tX to the participant for some $t \in (0, 1)$: incentives are dampened but not eliminated. Since participants know the condition in advance, their prior over what the payoffs are worth is the displayed distribution scaled by the same t , which is the rescaled prior of Proposition 1. Choice probabilities over displayed pairs are then the same as under real payment, for any $t > 0$. If instead hypothetical payoffs change the objective, a participant with no money at stake aims to choose accurately. By objective invariance the coding rule, and hence behavior, is again unchanged. If they do

both, the two results apply in sequence, since a rescaled uniform prior is again uniform. If the displayed payoffs cease to matter to the participant at all ($t = 0$), the model is degenerate. The quantity encoded is then a constant, so every coding rule is optimal under each of the three objectives. Since both options are worth zero, the decision rule in equation (5) is a tie on every trial, and the model predicts that the risky option is chosen with probability one half regardless of pX , C , or the volatility condition. That prediction is inconsistent with long-standing evidence that hypothetical choices are orderly and payoff-sensitive, and that behavior under hypothetical payoffs resembles behavior under low real stakes (Holt and Laury, 2002; Camerer and Hogarth, 1999). We rule out $t = 0$ by assumption, though the prediction is easily tested, and we return to it in Section 4.

With the displayed payoff distributions held fixed across conditions, the fixed-capacity model predicts the same behavior in real and hypothetical blocks, in both volatility conditions. In particular, the loss of payoff sensitivity from the low- to the high-volatility condition does not depend on whether payoffs are real or hypothetical. This is the null hypothesis our design tests. If instead the capacity n responded positively to incentives, lower effective stakes would induce a flatter encoding rule and reduced sensitivity to payoff differences (Section 2.2), yielding the directional hypotheses of Section 3.3.

3.1 Sample

We recruited a total of 70 participants to complete our risky choice task, which is based on FJ’s task (2022) but is preceded by a brief training phase for participants to report their emotional experiences (as discussed further below). Participants were recruited across 17 experimental sessions, with 3 to 5 participants per session. After completing the choice task, participants also responded to a survey that captures individual characteristics. The experiment was conducted at the Virginia Tech Economics Laboratory in spring 2021. All procedures were approved by Virginia Tech Institutional Review Board and participants provided informed consent prior to their participation. The study was not preregistered. The hypotheses and treatment design were fixed before data collection. The average payment per participant was \$20.8, which includes a \$10 show-up payment. All parts of the experiment were conducted on a computer terminal with the task and survey being programmed via oTree (Chen et al., 2016).

Participants were recruited from the laboratory’s participant pool and were all affiliated with Virginia Tech as undergraduate students, graduate students, or university staff: 31 women/39 men, with a median age of 21.5 years (mean = 24.0 years), 79% reported being at the university for between one and four years, and only 14% reported working full-time. The

FJ (2022) benchmark sample, which we also incorporate in the pooled analyses below, was recruited from University of Southern California students; participant-level demographic data from that sample are not available to us. Accordingly, pooled analyses should be interpreted with this difference in subject pools in mind. In addition, our sample size was chosen to approximate the sample size of the original working-paper version of Frydman and Jin (2022), which included 34 participants and already documented evidence for efficient coding in risky choice; we recruited 35 participants per incentive treatment, for a total of 70 participants. As shown in the Results section, our data exhibit the key pattern predicted by efficient coding: payoff sensitivity is lower in the high volatility condition relative to the low volatility condition. However, our experiment is less powered to detect smaller effects, so statistically insignificant findings should be interpreted with caution.

3.2 FJ Task

We administered a version of the risky choice task used to test the efficient coding hypothesis (Frydman and Jin, 2022). This task involves a series of 600 decisions to choose between a sure payoff and a lottery with 50% chance of a positive payoff and a 50% of a zero payoff. The 600 decisions are divided into two conditions that differ in terms of the underlying distribution from which the (positive) payoffs are sampled (see Figure 1A–C). In the low volatility (LV) condition, the payoffs are sampled independently for each participant from uniform narrow distributions (sure payoff range: 8 to 12; risky option’s positive payoff: 16 to 24). On the other hand, the high volatility (HV) condition payoffs are sampled independently from uniform broad distributions (sure payoff range: 4 to 16; risky option’s positive payoff: 8 to 32). Note that the mean values across the volatility conditions’ payoffs are constant (10 for the sure payoff and 20 for the risky positive payoff). The order of the conditions was randomized across participants.

An important feature of this task is the common trials, in which 30 trials were selected a priori (matching those in FJ) to fall in the support of both the low and high volatility conditions and to be presented across both conditions. These identical trials were sampled to span the low volatility payoff distributions. All our analyses are based upon these common trials, which enable us to exploit within-subject variation in behavior across volatility conditions. To maintain the statistical properties of the underlying payoff distributions, the sampled payoffs in the HV condition were re-drawn if they fell in the support of LV condition. After a short adaptation phase of 30 trials with the payoffs sampled independently from the underlying distributions, the common trials were presented in a random order in the remaining 270 trials of each volatility condition across participants.




Figure 1: Risky Choice and Emotion Classification Tasks. The distributions of the (A) risky payoffs and (B) sure payoffs across low and high volatility conditions. Note that the spikes in the narrow (low volatility) range are due to the common trials presented to all participants and in both volatility conditions (Frydman and Jin, 2022). (C) An example of a trial from the risky choice task. Note that the location of the risky option varied across trials and participants. (D) Procedure for elicitation of emotional experiences. (E) The list of emotions presented in the Emotion Classification Task (Heffner et al., 2021). (F) The average coordinate of each of the 20 emotions across participants.

3.3 Hypothesis and modifications to the FJ task

Our hypotheses are motivated by two threads of literature. First, existing models from neuroscience and psychology positing that the brain’s representation of value and sensory information is relative to the available context and is governed by a cognitive constraint. The divisive normalization framework assumes a fixed range of neural activity (Louie et al., 2015; Rustichini et al., 2017), the efficient coding and noisy coding models assume a fixed capacity constraint (Frydman and Jin, 2022; Khaw et al., 2021), and the decision by sampling model assume a finite number of samples (Bhui and Gershman, 2018; Stewart et al., 2006). Second, rational inattention and resource-rational theories are instead centered around agent’s flexibility in deploying cognitive resources and in-response to task demands (Sims, 2003; Van den Berg and Ma, 2018). As Van den Berg and Ma (2018, p. 2) put it, “... *why the brain would not simply recruit more resources whenever it has more items to remember? After all, your heart does something similar by beating faster whenever you increase your physical activity.*” Such flexibility would imply a top-down mechanism that can upregulate or downregulate value perception and sensitivity to stimuli in risky choice; we conjecture that incentives can play that role.

To test this hypothesis, we adapted FJ’s risky choice paradigm in two key ways. First, participants were randomly assigned to one of two incentive conditions (see Table 1). One group received real payment based on a randomly selected trial from the LV condition, while payoffs from the HV condition were hypothetical. This incentive structure, labelled Enhance, is hypothesized to increase the difference in payoff sensitivity between the low- and high-volatility conditions, because fewer encoding resources are devoted to the hypothetical HV payoffs. On the other hand, the real payment for the second group of participants is tied to the HV condition, with hypothetical payoffs in the LV condition. This condition, labelled Suppress, is hypothesized to decrease the difference in payoff sensitivity between the low- and high-volatility conditions, because fewer encoding resources are devoted to the hypothetical LV payoffs. These hypotheses are consistent with decision rules in resource-rational theories of decision-making (Sims, 2003; Van den Berg and Ma, 2018). Our two hypotheses are:

Hypothesis 1 (Enhance condition). *Weaker incentives in high volatility (vs. low volatility) will reduce payoff sensitivity.*

Hypothesis 2 (Suppress condition). *Stronger incentives in high volatility (vs. low volatility) will increase payoff sensitivity.*

Note that our incentivized condition involved paying only one randomly determined trial that was selected at the end of the session as the paying trial for each volatility condition.

Thus, participants only realized real payment from either the HV or LV conditions depending on treatment assignment. This can be contrasted with FJ’s task that paid for one randomly selected trial out of both conditions, which neither enhances nor suppresses precision across the HV and LV conditions (Benchmark).

Second, we measured participants’ emotional experiences at regular intervals using the dynamic Affective Representation Mapping (dARM) (Heffner et al., 2021). Prior to the risky choice task, participants were trained to classify emotional experiences on an unpleasant/pleasant and low/high arousal dimensions (Russell et al., 1989) simultaneously on a 200x200-pixel grid. Then, participants were asked to report their feelings (valence and arousal) regularly during the risky choice task (Figure 1D–F).

3.4 Procedures

The experimental session proceeded as follows. After providing informed consent and completing emotion classification training, participants completed the risky choice task starting with either the low volatility condition or the high volatility one (order counter-balanced across participants; see Table 1). Participants were not informed of the underlying distribution of payoffs prior to making their decisions, but were told whether the upcoming decisions would result in real or hypothetical payment. In the hypothetical-payoff blocks, participants were explicitly informed that their choices would not affect their earnings; however, they were asked to think carefully about each decision and indicate what they would choose if faced with the decision. A complete copy of the experimental instructions is provided in Appendix C. Each condition consisted of 300 trials, divided into blocks of 50 decisions. The online appendix identifies the 30 common trials that were presented in both volatility conditions (see section 3.2). Participants were informed that they could take short breaks in between blocks and were reminded at the start of each block of the underlying incentives that pertained to one’s assigned treatment group (real or hypothetical). Participants also reported their feelings in the beginning of each block and every 10 trials thereafter. During each trial, two keys on the keyboard were used for making decisions between the left and right options that corresponded to either the sure or risky options (the presentation format (left vs. right) of the options was randomized across trials). A 2-second fixation preceded the presentation of the options in each trial, and response time was recorded in addition to which option was selected.

Table 1: Analysis data based on FJ’s risky choice task and order of experimental treatments.

		Second 300 trials			
		HV Hypothetical	HV Real	LV Hypothetical	LV Real
First 300 trials	High Volatility (HV) Hypothetical	—	—	—	Enhance $N = 18$ (ABS)
	High Volatility (HV) Real	—	—	Suppress $N = 17$ (ABS)	Benchmark $N = 75$ (FJ)
	Low Volatility (LV) Hypothetical	—	Suppress $N = 18$ (ABS)	—	—
	Low Volatility (LV) Real	Enhance $N = 17$ (ABS)	Benchmark $N = 74$ (FJ)	—	—

ABS = data collected by authors of this paper. FJ = data collected and shared by Frydman and Jin (2022).

4 Results

We examine whether our data reproduce the core efficient coding pattern documented by Frydman and Jin (2022). FJ test for efficient coding using separate interaction terms on X and C (their Table I); the single-slope version of this result, in which risk-taking responds more steeply to $pX - C$ under low volatility, appears in their Figure V. Because our primary objective is to test whether incentive manipulations amplify or dampen the efficient coding effect, we characterize payoff sensitivity using $pX - C$. The coefficient on $pX - C$ measures payoff sensitivity; its interaction with HV measures how volatility shifts that sensitivity; and the triple interaction with the incentive condition measures how our manipulation modifies that effect. The coefficient on $HV \times (pX - C)$ provides a parsimonious measure of the efficient coding effect, and the higher-order interactions allow us to test whether incentives modify its magnitude. This approach is also used in related work examining whether an experimental manipulation alters the efficient coding effect (Frydman and Jin, 2025). Tables 2 and 3 report results from this specification estimated separately on our data and on the Frydman and Jin (2022) dataset.

First, we replicate FJ’s finding supporting the efficient coding hypothesis in risky choice (Frydman and Jin, 2022). This replication uses our pooled sample that includes participant

choices from the common trials in our two experimental treatments (Enhance and Suppress). As shown in Table 2, sensitivity to the difference between the expected value of the gamble and the sure payoff ($pX - C$) is lower in the HV condition relative to LV; a finding that was robustly replicated across between- and within-subject tests. Following FJ, column (1) includes data from the first 300 trials only and provides a between-subjects test while column (2) focuses on a subset of trials that exclude excessively fast trials with responses times of less than 0.5 seconds (less than 1.5% of the data in ABS dataset). Column (3) includes data from all 600 trials and instead provides a within-subject test of efficient coding, with column (4) excluding the first half of the second condition to allow for longer adaptation and columns (5) and (6) disaggregating further by order of HV and LV conditions.

Across columns 1-5, risk taking was more responsive to changes in payoff differences when the underlying sampling distribution was narrow, where the negative and statistically significant coefficient of $HV \times (pX - C)$ indicating steeper slope under LV ($p < 0.05$). In column (6), however, the coefficient of $HV \times (pX - C)$ is not statistically significant and includes data from participants completing the HV condition first only. This again replicates FJ (2022) results, who conjecture that “it is easier for subjects to detect a change in the environment when moving from a low-volatility to high-volatility condition, because “outlier payoffs” that are never experienced in the first condition begin to appear in the second condition” (p. 186). Our incentive manipulations, presented next, will offer novel insights on this adaptation process and how transitioning from real to hypothetical payoffs can yield comparable declines in sensitivity to $pX - C$ to changes in the underlying volatility in the environment. In summary, perception is more optimized for payoffs that were more commonly presented, providing confirmatory evidence that payoff sensitivity depends on the distribution of payoffs, as posited by efficient coding. Similar to findings based on FJ’s dataset, which we present in Table 3, such evidence is stronger in magnitude and statistical significance in (a) between-subjects tests and in within-subject tests that (b) enable longer adaptation or (c) among participants assigned to complete the LV condition first.

Table 2: Frequency of choosing the risky option (ABS dataset)

	Between-subjects tests		Within-subject tests			
	(1) Unrestricted (first 300)	(2) Restricted (first 300)	(3) Restricted	(4) Restricted (w/o 301–450)	(5) LV first (w/o 301–450)	(6) HV first (w/o 301–450)
Dependent variable: chose risky						
High Volatility (HV)	0.141*	0.142*	0.013	0.024	0.044	-0.007
	(0.059)	(0.059)	(0.034)	(0.034)	(0.043)	(0.052)
$pX - C$	0.166***	0.166***	0.131***	0.138***	0.166***	0.098***
	(0.009)	(0.009)	(0.010)	(0.010)	(0.009)	(0.017)
$HV \times (pX - C)$	-0.076***	-0.076***	-0.025*	-0.042***	-0.068***	-0.009
	(0.016)	(0.016)	(0.010)	(0.011)	(0.016)	(0.014)
Constant	0.417***	0.417***	0.488***	0.485***	0.417***	0.567***
	(0.039)	(0.039)	(0.031)	(0.029)	(0.039)	(0.044)
Observations	2100	2089	4155	3244	1634	1610
Participants	70	70	70	70	35	35

Notes: Mixed-effects linear regressions. Standard errors of the fixed effects are in parentheses, clustered at the participant level. High volatility takes a value of 1 for the condition with distributions sampled from the wider distributions. The models include a random intercept and a random effect for $pX - C$. Restricted sample excludes fast response trials (<0.5 seconds). * $p < .05$, ** $p < .01$, *** $p < .001$.

Table 3: Frequency of choosing the risky option (FJ dataset)

	Between-subjects tests		Within-subject tests			
	(1) Unrestricted (first 300)	(2) Restricted (first 300)	(3) Restricted	(4) Restricted (w/o 301–450)	(5) LV first (w/o 301–450)	(6) HV first (w/o 301–450)
Dependent variable: chose risky						
High Volatility (HV)	0.020 (0.046)	0.032 (0.046)	-0.001 (0.013)	0.006 (0.015)	0.017 (0.021)	-0.008 (0.020)
$pX - C$	0.149*** (0.010)	0.161*** (0.009)	0.142*** (0.007)	0.145*** (0.007)	0.160*** (0.009)	0.127*** (0.011)
$HV \times (pX - C)$	-0.033* (0.014)	-0.043** (0.013)	-0.011* (0.005)	-0.018** (0.006)	-0.024** (0.008)	-0.008 (0.010)
Constant	0.417*** (0.028)	0.406*** (0.027)	0.429*** (0.023)	0.426*** (0.022)	0.409*** (0.027)	0.447*** (0.036)
Observations	4470	4170	8257	6411	3125	3286
Participants	149	148	149	149	74	75

Notes: Mixed-effects linear regressions. Standard errors of the fixed effects are in parentheses, clustered at the participant level. High volatility takes a value of 1 for the condition with distributions sampled from the wider distributions. The models include a random intercept and a random effect for $pX - C$. Restricted sample excludes fast response trials (<0.5 seconds). * $p < .05$, ** $p < .01$, *** $p < .001$.

We next investigate the impact of incentives on payoff sensitivity. Participants were informed before each 300-trial block whether the upcoming decisions would be real or hypothetical, but they were not told at the outset that the experiment would contain both real and hypothetical blocks, nor that incentives would switch at trial 301. Thus, participants in the first block knew only the incentives of the current block. A direct between-subjects test of whether these stakes affect payoff sensitivity classifies each participant by the incentive status of the current block. Table 4 reports mixed-effects regressions of risky choice on the first 300 trials under this classification, with the indicator Hypothetical taking a value of 1 when payoffs in the participant’s current block were hypothetical; all such observations come from the ABS dataset. Columns 1 and 2 pool the ABS and FJ datasets, and columns 3 and 4 restrict the sample to ABS participants, for whom recruitment and procedures are held fixed. The coefficient on $\text{Hypothetical} \times \text{HV} \times (\text{pX} - \text{C})$ is negative but not statistically significant in the pooled estimates and virtually zero within the ABS sample. Thus, we find no evidence that the level of stakes alone modulates payoff sensitivity in the first block. Nor is payoff sensitivity absent under hypothetical payoffs: the $\text{Hypothetical} \times (\text{pX} - \text{C})$ interaction in Table 4 is 0.009 (SE 0.019) in columns 3 and 4, so payoff sensitivity under hypothetical payoffs is indistinguishable from that under real payoffs, consistent with the assumption $t > 0$ in Section 3. The within-subject tests, reported next, examine whether payoff sensitivity instead responds to the change in incentives that ABS participants experience between blocks.

Within-subject tests examine whether the change in incentives between blocks modulates payoff sensitivity. Since participants faced hypothetical payoffs in exactly one volatility condition that we designed to either enhance or suppress the efficient coding effect (see Table 1), we examine within-subject variation through these treatment assignments directly. Our treatment conditions are thus defined by how incentives are paired with volatility across both blocks of the experiment. Table 5 reports mixed-effects linear regressions that pool our dataset with FJ’s dataset, enabling us to pinpoint differences in sensitivity to $\text{pX} - \text{C}$ across the Enhance, Suppress, and Benchmark treatments; Figure 2 visualizes the estimated coefficients from the models in Table 5 that moderate the effect of $\text{pX} - \text{C}$ on the frequency of choosing the risky option. In these within-subject tests, which account for participants being exposed to the change in incentive structure across the two blocks, responsiveness to changes in $\text{pX} - \text{C}$ differed across treatment groups.

The omitted category in Table 5 is Suppress, so the coefficient $\text{HV} \times (\text{pX} - \text{C})$ reports the efficient coding effect for that group, and the Enhance and Benchmark interactions report differences from it. Columns 1–3 (Figure 2 panels 1–3) demonstrate that participants assigned to complete the HV condition with hypothetical payoffs (Enhance) were less respon-

Table 4: Frequency of choosing the risky option in the first block (between-subjects tests), by payoff type (real versus hypothetical)

	Pooled ABS and FJ		ABS only	
	(1) Unrestricted	(2) Restricted	(3) Unrestricted	(4) Restricted
Dependent variable: chose risky				
High Volatility (HV)	0.045 (0.041)	0.055 (0.041)	0.157 (0.089)	0.157 (0.089)
$pX - C$	0.151*** (0.008)	0.161*** (0.007)	0.161*** (0.014)	0.161*** (0.014)
$HV \times (pX - C)$	-0.041*** (0.012)	-0.049*** (0.012)	-0.075** (0.024)	-0.075** (0.024)
Hypothetical	-0.038 (0.055)	-0.029 (0.055)	-0.066 (0.077)	-0.065 (0.077)
Hypothetical $\times$ HV	0.081 (0.085)	0.073 (0.085)	-0.031 (0.116)	-0.029 (0.116)
Hypothetical $\times$ ($pX - C$)	0.019 (0.015)	0.009 (0.015)	0.009 (0.019)	0.009 (0.019)
Hypothetical $\times$ HV $\times$ ($pX - C$)	-0.036 (0.025)	-0.028 (0.025)	-0.001 (0.032)	-0.002 (0.033)
Constant	0.423*** (0.025)	0.415*** (0.025)	0.451*** (0.059)	0.451*** (0.059)
Observations	6570	6259	2100	2089
Participants	219	218	70	70

Notes: Mixed-effects linear regressions. Standard errors of the fixed effects are in parentheses, clustered at the participant level. High volatility takes a value of 1 for the condition with distributions sampled from the wider distributions. Hypothetical takes a value of 1 for participants whose first-block payoffs were hypothetical. All such observations are from the ABS dataset; the omitted category in columns 1–2 therefore pools ABS participants with real payoffs and the FJ sample, and columns 3–4 restrict the sample to ABS participants. The models include a random intercept and a random effect for $pX - C$. Restricted sample excludes fast response trials (<0.5 seconds). * $p < .05$, ** $p < .01$, *** $p < .001$.

sive to changes in $pX - C$ than participants assigned to complete the HV condition with real payoffs (Suppress). This difference is captured by the coefficient $HV \times (pX - C) \times Enhance$, which ranges between -0.044 and -0.063 across columns 1–3 ($p < 0.05$). This negative effect is also large in magnitude, about 2.4 to 4.1 times the efficient coding effect observed under Benchmark (given by the sum of $HV \times (pX - C)$ and its Benchmark interaction). We also report in Table 5 comparisons between payoff sensitivity across the Enhance and Benchmark conditions, with statistically significant differences (χ^2 tests, $p = 0.014$, 0.009 , and 0.002 in columns 1–3). In contrast, we find no evidence that responsivity to changes in $pX - C$ among participants assigned to Suppress differed from Benchmark ($p > 0.10$ in all columns). Thus, we find support for Hypothesis 1: hypothetical payoffs in the HV condition (the Enhance treatment) amplify the loss of payoff sensitivity relative to both the Suppress treatment and the Benchmark. The evidence for Hypothesis 2 is more limited. Even though the difference between the Suppress and Enhance conditions is statistically significant in columns 1–3, we do not find a significant difference in payoff sensitivity between the Suppress and Benchmark conditions. Taken together, these results suggest that encoding precision can respond to incentives rather than being determined by a fixed biological capacity.

The findings from the full sample do not replicate in a subsample analysis that focuses only on participants who completed the HV condition first (Table 5 column 4 and Figure 2 panel 4). Nonetheless, the switch in incentives across our treatment groups reveals more nuanced responses to presented payoffs that go beyond distribution-dependent payoff sensitivity. Participants assigned to Suppress who completed the high volatility condition first show significantly lower responsivity to changes in $pX - C$ throughout both volatility conditions than Benchmark ($Benchmark \times (pX - C) = 0.056$, $p < .05$); the corresponding comparison with Enhance is similar in magnitude but not statistically significant. Even though participants still seem unable to detect the change in volatility when HV is completed first (FJ’s conjecture outlined earlier), this finding highlights that overall responsivity to $pX - C$ can diminish with specific pairings of incentive and direction of change in volatility: HV first with real payoffs followed by LV with hypothetical payoffs. We conjecture that this combination is not only more difficult for detecting the change in environment from HV to LV, but it also dampens participants’ motivation to be responsive to changes in payoffs in the LV condition.

Several features of the data weigh against a fatigue-based interpretation of the incentive-modulated payoff-sensitivity effects. The treatment effects are identified from counterbalanced changes in volatility and incentives, not from monotonic trial progression. Controlling directly for trial number and trial number squared leaves the main estimates unchanged. Moreover, extremely fast responses remain rare throughout the task, increasing only mod-

estly from Block 1 to Block 2 (4.7% to 6.4%, pooling ABS and FJ data). These checks suggest that the treatment effects reflect condition-specific changes in payoff sensitivity, rather than general deterioration in engagement over the session.

Table 5: Frequency of choosing the risky option across incentive treatments, within-subject tests (pooled ABS and FJ datasets)

	(1) Restricted	(2) Restricted (w/o 301–450)	(3) LV first (w/o 301–450)	(4) HV first (w/o 301–450)
Dependent variable: chose risky				
High Volatility (HV)	-0.109** (0.036)	-0.094** (0.035)	-0.089* (0.041)	-0.120* (0.059)
$pX - C$	0.121*** (0.015)	0.130*** (0.015)	0.170*** (0.013)	0.071** (0.026)
$HV \times (pX - C)$	-0.002 (0.013)	-0.017 (0.013)	-0.034 (0.019)	0.015 (0.019)
Enhance	-0.129* (0.061)	-0.113 (0.059)	0.065 (0.076)	-0.304*** (0.072)
Benchmark	-0.124* (0.051)	-0.118* (0.049)	0.024 (0.056)	-0.281*** (0.058)
Enhance $\times$ High Volatility	0.244*** (0.060)	0.227*** (0.061)	0.266*** (0.073)	0.212* (0.097)
Benchmark $\times$ High Volatility	0.108** (0.038)	0.100** (0.038)	0.107* (0.046)	0.112 (0.062)
Enhance $\times pX - C$	0.020 (0.019)	0.015 (0.019)	-0.009 (0.019)	0.048 (0.033)
Benchmark $\times pX - C$	0.022 (0.017)	0.016 (0.017)	-0.010 (0.015)	0.056* (0.028)
Enhance $\times HV \times (pX - C)$	-0.047* (0.020)	-0.044* (0.020)	-0.063* (0.030)	-0.042 (0.027)
Benchmark $\times HV \times (pX - C)$	-0.010 (0.014)	-0.001 (0.015)	0.010 (0.021)	-0.024 (0.022)
Constant	0.553*** (0.046)	0.544*** (0.044)	0.385*** (0.049)	0.728*** (0.045)
$\beta_E = \beta_B$ (χ^2 test), $p =$	0.014*	0.009**	0.002**	0.390
Observations	12412	9655	4759	4896
Participants	219	219	109	110

Notes: Mixed-effects linear regressions. Standard errors of the fixed effects are in parentheses, clustered at the participant level. Omitted category is Suppress. β_E and β_B are the Enhance $\times$ HV $\times$ (pX - C) $\times$ and Benchmark $\times$ HV $\times$ (pX - C) coefficients, respectively. High volatility takes a value of 1 for the condition with distributions sampled from the wider distributions. The models include a random intercept and a random effect for pX - C. Restricted sample excludes fast response trials (<0.5 seconds). * $p < .05$, ** $p < .01$, *** $p < .001$.




Figure 2: Estimated coefficients for the effect of interaction with pX - C on frequency of choosing the risky option (full results for models (1)–(4), with Suppress as the omitted category and including statistical significance, are presented in Table 5). Error bars show 95% confidence intervals; comparisons across coefficients are based on the tests reported in Table 5.

We turn next to examining differences in process-level measures that offer additional insights on potential mechanisms underlying differences in decision-making across volatility conditions and incentive treatments. First, we focus on self-reported emotional experiences, which were periodically assessed throughout the experimental session (60 times total split across volatility conditions) among ABS participants who were assigned to either Suppress or Enhance. As shown in Figure 3A, participants in the Enhance treatment, who presented the strongest evidence for efficient coding (Table 5 and Figure 2), reported significantly lower levels of arousal in HV relative to LV, compared to Suppress participants (Wilcoxon rank-

sum test; $N=70$, $p<0.001$). This finding was specific to the arousal dimension and did not extend to the valence dimension (Figure 3B; $P>0.10$). This suggests that dampened arousal levels under high volatility environments may serve as a mechanistic pathway facilitating the loss of payoff sensitivity during risky decision-making. In Figure 4, we visualize the evolution of self-reported arousal over the 600 trials in the experimental session (60 measurements in total), which demonstrates a general decreasing trend throughout the session with a notable spike in arousal after the change in volatility environment (starting at trial 301) coinciding with the transition from hypothetical to real payoffs (panels A and D). This pattern did not extend to self-reports of valence (see Figure 7 in the supplement).




Figure 3: Difference in self-reported (A) arousal and (B) valence between HV and LV across treatment assignments (Suppress and Enhance). Arousal and valence measures were z-scored at the participant level using repeated measurements throughout the HV and LV trials. Data includes arousal/valence elicitation before either common or noncommon trials. ABS dataset only as elicitation of arousal and valence was not part of FJ 2022 experimental design. Error bars represent 95% confidence intervals for each group mean and are not intended to be used for visual inference about differences across treatment conditions. Statistical significance, based on Wilcoxon rank-sum tests ($N=70$), are indicated by asterisks: *** p -value < 0.001 ; N.S. denotes difference not statistically significant.

Second, we examine response times to gauge changes in processing speed across volatility conditions and incentive treatments. As shown in Figure 5, we find that average response time was significantly lower during HV relative to LV for the Enhance treatment relative to the other treatment groups. When we trace down the evolution of response times across the 600 trials in the experimental session (Figure 6), we find that the spike at trial 301 that FJ identifies going from LV to HV is masked by changes in underlying incentives: it is virtually absent when the HV condition involved hypothetical payoffs. Instead, we again find that such spike is linked to changes in incentives more broadly and extends to participants assigned to




Figure 4: Self-reported arousal across trials by treatment assignments and order of volatility conditions. Arousal measures were z-scored at the participant level using repeated measurements throughout the HV and LV trials. ABS dataset only as elicitation of arousal and valence was not part of FJ 2022 experimental design. The red vertical line marks the change in environment from HV to LV (or vice versa), and data includes common and noncommon trials. Grey intervals represent 95% confidence bounds around locally smoothed estimates (shown by the black line, with bandwidth chosen to minimize mean integrated squared error).

complete the HV condition first. Thus, we find converging evidence that Enhance induces differences in process-level measures between the HV and LV conditions that are consistent with changes in risk-taking behavior. In other words, the condition showing the largest decrease in payoff sensitivity was characterized by reduced self-reported arousal and faster decision times in the high volatility environment. Although arousal and response times decline over the session, the patterns shown in Figures 4 and 6 vary with the pairing of incentives and volatility, suggesting that these process-level results are not simply a fatigue effect.

To further examine the relationship between arousal and payoff sensitivity, we conducted an individual-level analysis. For each participant, we estimated a regression of risky choice on HV, $(pX - C)$, and their interaction, and used the coefficient on $HV \times (pX - C)$ as an individual-level measure of volatility-dependent payoff sensitivity. We then related this measure to the participant-specific difference in z-scored arousal between the HV and LV conditions. We find a positive and statistically significant association ($\rho=0.27$, $p=0.023$). In other words, participants who exhibit larger declines in arousal in HV relative to LV (which we interpret as being consistent with the withdrawal of cognitive effort) also tend to show stronger reductions in sensitivity to payoff variation. Confirming our aggregate analysis results (Table 5), the distribution of this individual-level volatility-dependent payoff-sensitivity measure also differed significantly across the Enhance and Suppress treatments, with participants in the former condition exhibiting lower sensitivity than the latter (mean difference = 0.050; rank-sum test $p=0.017$). These findings provide additional evidence linking variation in arousal to heterogeneity in payoff sensitivity; nonetheless, this analysis is correlational and establishing a causal channel will require further work.

5 Discussion

We find evidence consistent with the efficient coding hypothesis in risky choice and show that incentives can modulate the distribution-dependence of payoff sensitivity across volatility environments. Participants in the Enhance treatment, who received real payment in the low-volatility condition and hypothetical payoffs in the high-volatility condition, showed a wider gap in payoff sensitivity between the low- and high-volatility conditions than participants in either the Suppress treatment, who received the reverse assignment, or the benchmark condition, who received real payment in both. In our data, this modulation emerges only after participants experience the change in incentives between blocks. The Suppress treatment itself did not significantly attenuate the gap relative to the benchmark. Consistent with previous findings on the relationship between incentives and arousal (Anderson and




Figure 5: Difference in response time between HV and LV across treatment assignments (Suppress, Enhance, and Benchmark). Data includes common trials only. Error bars represent 95% confidence intervals for each group mean and are not intended to be used for visual inference about differences across treatment conditions. Statistical significance, based on Wilcoxon rank-sum tests ($N = 35, 35$, and 149 participants in the Suppress, Enhance, and Benchmark conditions, respectively), are indicated by asterisks: $*p < 0.05$, $**p < 0.01$; N.S. denotes difference is not statistically significant.

Brown, 1984; Richter and Gendolla, 2009), we observed increased arousal in the real stakes condition; but this pattern did not extend to the valence dimension. Moreover, response times, often used as indicators of variation in the decision-making process (Cooper et al., 2019; Konovalov and Krajbich, 2019), were faster under HV relative to LV condition for the Enhance treatment, which showed the largest loss of payoff sensitivity.

Taken together, our results indicate that shifting between hypothetical and real incentives can alter the precision of value representation underlying risk-taking behavior and induce corresponding changes in the decision-making process (arousal and response time). The reduced sensitivity to changes in payoff values in the high volatility condition under hypothetical stakes is consistent with prior work demonstrating greater brain activation of value encoding regions in real rather than hypothetical consumer good choices (Kang et al., 2011). Moreover, our findings are consistent with prior research linking arousal to mental effort (Howells et al., 2010). We demonstrate that choice precision is not only sensitive to the statistical properties of the environment, as shown by KLW and FJ, but also to the motivational context, where incentives produce multi-faceted shifts that extend to response time and arousal.

Broadly, our results suggest that the fidelity of the value representation is not a fixed




Figure 6: Average response time across trials by treatment assignments and order of volatility conditions. Similar to FJ (2022), we display datapoints with average response time between 1-4 seconds. The red vertical line marks the change in environment from HV to LV (or vice versa), and data includes common and noncommon trials. Grey intervals represent 95% confidence bounds around locally smoothed estimates (shown by the black line, with bandwidth chosen to minimize mean integrated squared error).

biological constraint but a dynamic decision variable responsive to experienced changes in incentives. While prior work demonstrates that choice sensitivity naturally adapts to the statistical range of stimuli (Frydman and Jin, 2022; Khaw et al., 2021), our incentive manipulation suggests that this phenomenon also reflects the decision-maker’s allocation of cognitive effort. Because the optimal code depends on neither the objective being maximized (objective invariance, under the uniform priors of our design) nor the scale of payoffs (scale invariance), a model with fixed capacity predicts no effect of consequentiality on behavior (for any scaling factor $t > 0$). The natural mechanism by which incentives might affect behavior is through capacity, which in the efficient coding model is given by n , the number of encoding units. We interpret the incentive effect as a change in these resources.

What mattered was the pairing of incentives across blocks, which every participant in our experiment experienced as a change. Pairing hypothetical payment with the high-volatility condition and real payment with the low-volatility condition (Enhance) amplified the loss of sensitivity between the two conditions, whereas the reverse pairing (Suppress) left that sensitivity loss no smaller than that of FJ’s participants, who were paid throughout.

A plausible explanation is that effort responds to a change in incentives relative to recent experience rather than to their level. Removing real payment leads to a withdrawal of effort (resulting in less precise value representations and hence lower payoff sensitivity), while adding payment raises effort only to the level of participants who were paid throughout. The lower self-reported arousal and faster response times we observe are consistent with this kind of effort withdrawal. Whether incentives expand capacity globally or reallocate a fixed capacity toward stimulus representation from other task demands is not something our design distinguishes; Frydman and Jin (2022, Section V.B) suggest the latter interpretation, citing the models of Woodford (2012a,b) and Dewan (2020).

Our results are consistent in spirit with theories of rational inattention (Matějka and McKay, 2015), in which the precision of processing responds to the stakes of the decision, though in our data only to changes in those stakes. Such models typically frame this as the costly acquisition of information about an uncertain state; in our setting, the payoffs are observed directly, so the relevant imprecision lies in how accurately they are internally represented rather than in what is learned about them. Our conjecture is that when incentives are negligible, reduced motivation leads to effort withdrawal (Camerer and Hogarth, 1999; Dickhaut et al., 2009) and lower precision of the representation. Supporting this interpretation, participants in unincentivized high volatility contexts exhibited faster response times and lower arousal, consistent with a rational withdrawal of effort.

A Supplementary Material

A.1 Additional figures




Figure 7: Self-reported valence across trials by treatment assignments and order of volatility conditions. Valence measures were z-scored at the participant level using repeated measurements throughout the HV and LV trials. ABS dataset only as elicitation of arousal and valence was not part of FJ 2022 experimental design. The red vertical line marks the change in environment from HV to LV (or vice versa), and data includes common and noncommon trials. Grey intervals represent 95% confidence bounds around locally smoothed estimates (shown by the black line, with bandwidth chosen to minimize mean integrated squared error).

B Proof of Proposition 1

Proof. Let θ_Y and θ_D denote the optimal coding rules for the rescaled priors f_Y and f_D under the same objective as θ_X and θ_C . We proceed in four steps.

Step 1: Encoding invariance. We show that $\theta_Y(tX) = \theta_X(X)$ for every X . The argument splits depending on which CDF-based encoding rule is in force.

Step 1a (mutual information / accuracy rule). Under equation (2), $\theta_X(X) = \sin^2(\frac{\pi}{2} F_X(X))$. The change of variable $v = u/t$ in the CDF of f_Y gives

$$F_Y(tX) = \int_{-\infty}^{tX} \frac{1}{t} f_X(u/t) du = \int_{-\infty}^X f_X(v) dv = F_X(X).$$

Therefore $\theta_Y(tX) = \sin^2(\frac{\pi}{2} F_Y(tX)) = \sin^2(\frac{\pi}{2} F_X(X)) = \theta_X(X)$.

Step 1b (expected financial gain rule). Under equation (3), $\theta_X(X) = \sin^2(\frac{\pi}{2} G_X(X))$, where

$$G_X(X) \equiv \frac{\int_{-\infty}^X f_X(x)^{2/3} dx}{\int_{-\infty}^{\infty} f_X(x)^{2/3} dx}.$$

The same change of variable $v = u/t$ gives

$$\int_{-\infty}^{tX} f_Y(u)^{2/3} du = \int_{-\infty}^{tX} \frac{1}{t^{2/3}} f_X(u/t)^{2/3} du = t \cdot \frac{1}{t^{2/3}} \int_{-\infty}^X f_X(v)^{2/3} dv = t^{1/3} \int_{-\infty}^X f_X(v)^{2/3} dv,$$

and the identical computation produces the same $t^{1/3}$ factor in the denominator $\int_{-\infty}^{\infty} f_Y(u)^{2/3} du$. The two $t^{1/3}$ factors cancel, so $G_Y(tX) = G_X(X)$ and hence $\theta_Y(tX) = \theta_X(X)$.

The same arguments applied to f_C and f_D yield $\theta_D(tC) = \theta_C(C)$.

Step 2: Likelihood invariance. The binomial likelihood in equation (1) depends on X only through $\theta_X(X)$. By Step 1, $\theta_Y(tX) = \theta_X(X)$, so

$$p(R | tX) = \binom{n}{R} \theta_Y(tX)^R (1 - \theta_Y(tX))^{n-R} = p(R | X)$$

for every $R \in \{0, 1, \dots, n\}$. The analogous identity $p(R | tC) = p(R | C)$ holds for the certain payoff.

Step 3: Posterior scaling. Substituting $y = tx$ (so $dy = t dx$) in the posterior-mean formula (4) applied to the rescaled environment, and using Step 2 to replace $p(R | tx)$ with $p(R | x)$:

$$\mathbb{E}[\tilde{Y} | R] = \frac{\int p(R | y) f_Y(y) y dy}{\int p(R | y) f_Y(y) dy} = \frac{\int p(R | x) \cdot \frac{1}{t} f_X(x) \cdot (tx) \cdot t dx}{\int p(R | x) \cdot \frac{1}{t} f_X(x) \cdot t dx} = t \cdot \mathbb{E}[\tilde{X} | R].$$

The $1/t$ from f_Y and the t from dy cancel, and the factor tx in the numerator integrand supplies the surviving t . The same calculation gives $\mathbb{E}[\tilde{D} | R] = t \cdot \mathbb{E}[\tilde{C} | R]$.

Step 4: Choice probability invariance. Evaluate the choice probability (5) at (tX, tC) .

For any (R_X, R_C) , Step 3 implies

$$p \cdot \mathbb{E}[\tilde{Y} \mid R_X] > \mathbb{E}[\tilde{D} \mid R_C] \iff p \cdot t \mathbb{E}[\tilde{X} \mid R_X] > t \mathbb{E}[\tilde{C} \mid R_C] \iff p \cdot \mathbb{E}[\tilde{X} \mid R_X] > \mathbb{E}[\tilde{C} \mid R_C],$$


since $t > 0$. The indicator in (5) is therefore unchanged for every (R_X, R_C) , and by Step 2 the weights $p(R_X \mid tX)p(R_C \mid tC) = p(R_X \mid X)p(R_C \mid C)$ also coincide. The same two equalities hold for the tie-breaking term. Hence the double sum in (5) takes the same value at (tX, tC) as at (X, C) , establishing (7). $\square$

C Experimental instructions

[Screen 1]

Instructions

The purpose of this task is to study how people classify feelings. In this task, you will use a diagram like the one below to describe how you are feeling. It is in the form of a square - a kind of map for feelings. The center of the square represents a neutral, average, everyday feeling. It is neither positive nor negative.



The right half of the grid represents pleasant feelings. The farther to the right, the more pleasant. The left half represents unpleasant feelings. The farther to the left, the more unpleasant

The vertical dimension of the map represents arousal. Arousal has to do with how wide awake, alert, or activated a person feels - independent of whether the feeling is positive or negative. The top half is for feelings that are higher in arousal, and the lower half for feelings that are lower in arousal. The very bottom represents sleep, and the higher you go, the more awake a person feels. At the top is maximum arousal. If you imagine a state we might call frantic excitement (remembering that it could be either positive or negative), then this feeling would define the top of the grid.


Click **Next** when ready to begin the task.

Next

[Screen 2]

In this part of the study, you will be shown words that describe certain feelings. Think about what the feeling is like in terms of the diagram. Some people think about situations, or draw on memories of situations, that have made them feel that way in the past. Please **drag and drop** each **word** on the diagram to rate that particular feeling. You can still shuffle the words around on the diagram after dropping them if you like. The last position for each word will be recorded as your final rating of that feeling.

neutral	still
surprised	quiet
aroused	sluggish
peppy	sad
enthusiastic	disappointed
happy	disgusted
satisfied	annoyed
relaxed	angry
calm	afraid
sleepy	nervous




Next

[Screen 3]

Instructions

After getting familiar with how feelings can be classified on the low arousal/high arousal and unpleasant/pleasant grid, we are interested in knowing how you currently feel at different stages of the session. On the next page, you will see a diagram similar to the one below. You will be asked to indicate how you are feeling using the diagram. Note that you can click on the diagram any number of times but be aware that the last click will be recorded as your final rating of your own feeling.




Click **Next** when ready to begin the task.

Next

[Screen 4]

How do you feel?

Click on the graph.




[Screen 5]

Instructions

For today's session, you will receive \$10 for your participation, and you will have the opportunity to earn more money depending on your answers during the experiment.

In this experiment, you will be asked to make a series of decisions about choosing a "risky gamble" or a "sure thing". The risky gamble will pay a positive amount with 50% chance, and \$0 with 50% chance. The amount shown for the sure thing will be paid with 100% chance, if chosen. Below is an example screen from the experiment.



In the example, the risky gamble is on the LEFT side of the screen, and the sure thing is on the RIGHT side of the screen. The risky gamble pays \$22.85 with 50% chance, and \$0 with 50% chance. The sure thing pays \$10.78 with 100% chance.

For each question, you will be asked to select one of the two options for each question by pressing either the 'a' key for the LEFT option or the 'k' key for the RIGHT option on the keyboard.

On some questions, the risky gamble will be on the LEFT, and in other questions it will be on the RIGHT.

Today's session is divided into two parts of 6 blocks each. We will next present you the instructions for first part.

Click **Next** to continue instructions.

Next


[Screen 6: randomized to view either Hypothetical Payment or the Real Payment]

[If randomized to complete Hypothetical Condition first; otherwise, the participant will view this screen after they complete the first 300 questions]

Hypothetical Payment

This part of the experiment is broken down into 6 blocks, and each block contains 50 questions. Thus, in this part of the experiment, you will be presented with 300 questions.

All payoffs from this part of the experiment are hypothetical. **You will not receive a payment from this part.**



At the end of the experiment, one trial from this part of the experiment will be randomly selected for **hypothetical** payment. For example, if the trial on the left was chosen and you selected the sure thing, your **hypothetical** payoff would be \$10.78.

If instead you chose the risky gamble, your **hypothetical** payoff would be either \$0 or \$22.85, depending on which outcome the computer randomly selects.

To reiterate, your choices in this part of the experiment **have no effect on your earnings** and have no effect on what options will be shown to you subsequently. We are interested in what you would do if you actually faced these choices, so please think about them carefully.

Next, we will ask you a comprehension question on this instruction.

Click **Next** to start practice.


Next

[If randomized to complete Real Condition first; otherwise, the participant will view this screen after they complete the first 300 questions]

Real Payment

This part of the experiment is broken down into 6 blocks, and each block contains 50 questions. Thus, in this part of the experiment, you will be presented with 300 questions.

At the end of the experiment, **one trial from this part of the experiment** will be randomly selected, and you **will be paid** according to your decision on that trial.



For example, if the trial on the left was chosen and you selected the sure thing, you would be paid \$10.78.

If instead you chose the risky gamble, you would be paid either \$0 or \$22.85, depending on which outcome the computer randomly selects.

Therefore, you should choose the option on each question that you prefer, since it may end up being the question that you are actually paid for. Remember that your earnings will be added to your \$10 participation fee.

Next, we will ask you a comprehension question on this instruction.

Click **Next** to start practice.

Next

[Screen 7]

Comprehension Check

Please answer the following question to the best of your ability.

True or False: One of the 300 upcoming questions will be randomly selected by the computer, and you will be paid a bonus depending on your answer to the randomly selected question.

- ☐ True
☐ False

When you have answered the above question, press **Next** to proceed.

Next


[Screen 8]

Before starting the task, let's get you familiar with the how you make decisions. First practice round: **Press "a" key on the keyboard to choose the LEFT option.**



[Screen 9]


Second practice round: Press 'k' key on the keyboard to choose the RIGHT option.



[Screen 10]

Good job! You are now ready to begin the experiment. Click **Next** to begin the first part of the experiment.

Next



[Screens 11+]

[Example of the separator screen between the blocks of 50 questions]

Block 1

This block will contain **50** questions.

You can now take a short break before beginning the task.

Click **Next** when ready to begin the task.

Next

[At the beginning of each block of 50 questions in the Hypothetical Condition]

Reminder: Hypothetical Payment

All payoffs from this block are hypothetical. **You will not receive a payment from any decisions in this block.**



Click **Next** to proceed to the block's decisions.

Next

[At the beginning of each block of 50 questions in the Real Condition]

Reminder: Real Payment

Your decision from one trial from this block **may be randomly selected to determine your payment today.**




Click **Next** to proceed to the block's decisions.

Next


[At the beginning of each block, and once every 10 questions thereafter]

How do you feel?

Click on the graph.



[Example of a risky choice question; the location of the risky option was randomized across screens; between screens, a brief fixation cross was presented]



[At the end of experimental session; after completing all 600 decisions; the screenshots below are for a participant randomized to complete the hypothetical condition first; the participant chose the gamble but ended up with the zero outcome]

Your hypothetical payoff:

Hypothetical bonus round (selected by computer)	278 / 600
Hypothetical bonus payment (you chose the certain amount)	\$11.48

This is the hypothetical payoff from the unpaid part of the session (Part 1).

Next

Your payoff:

Completion	\$10.00
Real bonus round (selected by computer)	305 / 600
Real bonus payment (you chose the gamble)	\$0.00
Total	\$10.00

You will receive the bonus payment from part 2 in addition to the fixed participation fee.

Next

D Common trial payoffs

Common Trial Payoffs

Common trials across the low and high volatility conditions and in both ABS and FJ tasks. These payoff values were presented once in each volatility condition, and the order of the common trials within a volatility condition were randomized across participants.

X (risky option's payoff)	C (Sure option's payoff)
17.30	11.20
17.31	9.62
17.32	8.04
17.34	8.75
17.38	10.38
18.63	8.79
18.64	11.16
18.68	8.03
18.68	9.61
18.72	10.44
19.96	9.62
19.98	8.01
19.99	8.81
20.03	10.4
20.04	11.22
21.29	8.82
21.30	10.39
21.34	9.58
21.34	11.23
21.38	8.00
22.62	9.62
22.66	8.79
22.67	8.04
22.69	11.21
22.71	10.42
23.96	9.55
23.97	11.80
23.98	8.03
23.98	10.37
23.99	8.84

Note that the bolded payoff is for a trial that differed between ABS and FJ datasets. ABS used the C payoff of **11.80** paired with the X of 23.97, which was presented in FJ's (2022) supplementary material document. However, in FJ's dataset, the associated C payoff was **11.18**.

E Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this work, the authors used generative AI tools (Claude, Anthropic; Gemini, Google) to assist with editing and formatting the manuscript. All content and any errors are solely the responsibility of the authors.

References

- Alsharawy, A., Zhang, X., Ball, S. B., and Smith, A. (2021). Incentives affect the process of risky choice. *Available at SSRN 3943681*.
- Anderson, G. and Brown, R. I. F. (1984). Real and laboratory gambling, sensation-seeking and arousal. *British Journal of Psychology*, 75(3):401–410.
- Barlow, H. B. (1961). Possible principles underlying the transformation of sensory messages. *Sensory communication*, 1:217–234.
- Bhui, R. and Gershman, S. J. (2018). Decision by sampling implements efficient coding of psychoeconomic functions. *Psychological review*, 125(6):985.
- Bhui, R. and Jiao, P. (2023). Attention constraints and learning in categories. *Management Science*, 69(9):5394–5404.
- Bouchouicha, R. and Vieider, F. M. (2017). Accommodating stake effects under prospect theory. *Journal of Risk and Uncertainty*, 55(1):1–28.
- Camerer, C. and Mobbs, D. (2017). Differences in behavior and brain activity during hypothetical and real choices. *Trends in Cognitive Sciences*, 21(1):46–56.
- Camerer, C. F. and Hogarth, R. M. (1999). The effects of financial incentives in experiments: A review and capital-labor-production framework. *Journal of risk and uncertainty*, 19(1-3):7–42.
- Caplin, A., Csaba, D., Leahy, J., and Nov, O. (2020). Rational inattention, competitive supply, and psychometrics. *The Quarterly Journal of Economics*, 135(3):1681–1724.
- Chen, D. L., Schonger, M., and Wickens, C. (2016). oTree—An open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9:88–97.

- Cooper, D. J., Krajbich, I., and Noussair, C. N. (2019). Choice-process data in experimental economics. *Journal of the Economic Science Association*, 5(1):1–13.
- Dewan, A. (2020). Costly Multidimensional Information. Working paper, New York University.
- Dickhaut, J., Rustichini, A., and Smith, V. (2009). A neuroeconomic theory of the decision process. *Proceedings of the National Academy of Sciences*, 106(52):22145–22150.
- Enke, B., Gneezy, U., Hall, B., Martin, D., Nelidov, V., Offerman, T., and Van De Ven, J. (2023). Cognitive biases: Mistakes or missing stakes? *Review of Economics and Statistics*, 105(4):818–832.
- Frydman, C. and Jin, L. J. (2022). Efficient coding and risky choice. *The Quarterly Journal of Economics*, 137(1):161–213.
- Frydman, C. and Jin, L. J. (2025). Adaptation through experience or description? evidence from an experiment on efficient coding. Working Paper 34214, National Bureau of Economic Research.
- Gershman, S. J., Horvitz, E. J., and Tenenbaum, J. B. (2015). Computational rationality: A converging paradigm for intelligence in brains, minds, and machines. *Science*, 349(6245):273–278.
- Gneezy, U., Halevy, Y., Hall, B., Offerman, T., and van de Ven, J. (2024). How real is hypothetical? a high-stakes test of the allais paradox. *Harvard Business School Working Paper 25-005*.
- Hackethal, A., Kirchler, M., Laudenbach, C., Razen, M., and Weber, A. (2023). On the role of monetary incentives in risk preference elicitation experiments. *Journal of Risk and Uncertainty*, 66(2):189–213.
- Heffner, J., Vives, M.-L., and FeldmanHall, O. (2021). Emotional responses to prosocial messages increase willingness to self-isolate during the COVID-19 pandemic. *Personality and Individual Differences*, 170:110420.
- Heng, J. A., Woodford, M., and Polania, R. (2020). Efficient sampling and noisy decisions. *Elife*, 9:e54962.
- Holt, C. A. and Laury, S. K. (2002). Risk aversion and incentive effects. *American economic review*, 92(5):1644–1655.

- Horn, S. and Freund, A. M. (2022). Adult age differences in monetary decisions with real and hypothetical reward. *Journal of Behavioral Decision Making*, 35(2):e2253.
- Howells, F. M., Stein, D. J., and Russell, V. A. (2010). Perceived mental effort correlates with changes in tonic arousal during attentional tasks. *Behavioral and Brain Functions*, 6(1):1–15.
- Kang, M. J., Rangel, A., Camus, M., and Camerer, C. F. (2011). Hypothetical and real choice differentially activate common valuation areas. *Journal of neuroscience*, 31(2):461–468.
- Khaw, M. W., Li, Z., and Woodford, M. (2021). Cognitive imprecision and small-stakes risk aversion. *The Review of Economic Studies*, 88(4):1979–2013.
- Konovalov, A. and Krajbich, I. (2019). Revealed strength of preference: Inference from response times. *Judgment & Decision Making*, 14(4).
- Krajbich, I. and Rangel, A. (2011). Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. *Proceedings of the National Academy of Sciences*, 108(33):13852–13857.
- Louie, K. and Glimcher, P. W. (2012). Efficient coding and the neural representation of value. *Annals of the New York Academy of Sciences*, 1251(1):13–32.
- Louie, K., Glimcher, P. W., and Webb, R. (2015). Adaptive neural coding: from biological to behavioral decision-making. *Current opinion in behavioral sciences*, 5:91–99.
- Manohar, S. G., Chong, T. T.-J., Apps, M. A., Batla, A., Stamelou, M., Jarman, P. R., Bhatia, K. P., and Husain, M. (2015). Reward pays the cost of noise reduction in motor and cognitive control. *Current Biology*, 25(13):1707–1716.
- Marois, R. and Ivanoff, J. (2005). Capacity limits of information processing in the brain. *Trends in cognitive sciences*, 9(6):296–305.
- Matějka, F. and McKay, A. (2015). Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review*, 105(1):272–298.
- Polania, R., Woodford, M., and Ruff, C. C. (2019). Efficient coding of subjective value. *Nature neuroscience*, 22(1):134–142.
- Richter, M. and Gendolla, G. H. (2009). The heart contracts to reward: Monetary incentives and pre-ejection period. *Psychophysiology*, 46(3):451–457.

- Russell, J. A., Weiss, A., and Mendelsohn, G. A. (1989). Affect grid: a single-item scale of pleasure and arousal. *Journal of personality and social psychology*, 57(3):493.
- Rustichini, A., Conen, K. E., Cai, X., and Padoa-Schioppa, C. (2017). Optimal coding and neuronal adaptation in economic decisions. *Nature communications*, 8(1):1–14.
- Sims, C. A. (2003). Implications of rational inattention. *Journal of monetary Economics*, 50(3):665–690.
- Stewart, N., Chater, N., and Brown, G. D. (2006). Decision by sampling. *Cognitive psychology*, 53(1):1–26.
- Taylor, M. P. (2013). Bias and brains: Risk aversion and cognitive ability across real and hypothetical settings. *Journal of Risk and Uncertainty*, 46(3):299–320.
- Van den Berg, R. and Ma, W. J. (2018). A resource-rational theory of set size effects in human visual working memory. *ELife*, 7:e34963.
- Webb, R., Glimcher, P. W., and Louie, K. (2021). The normalization of consumer valuations: Context-dependent preferences from neurobiological constraints. *Management Science*, 67(1):93–125.
- Wiseman, D. B. and Levin, I. P. (1996). Comparing risky decision making under conditions of real and hypothetical consequences. *Organizational behavior and human decision processes*, 66(3):241–250.
- Woodford, M. (2012a). Inattentive Valuation and Reference-Dependent Choice. Working paper, Columbia University.
- Woodford, M. (2012b). Prospect Theory as Efficient Perceptual Distortion. *American Economic Review (Papers and Proceedings)*, 102(3):41–46.
- Woodford, M. (2020). Modeling imprecision in perception, valuation, and choice. *Annual Review of Economics*, 12:579–601.