Rational Inattention Meets Cognitive Activation: Experimental Evidence*

Andreas Hefti[†] Tobias Müller[‡] Peiyao Shen[§] Ernst Fehr[¶]

Abstract

This paper examines how external stimuli influence choice performance in a sequential dual-task choice experiment, shedding light on the interplay between attention allocation and cognitive efficiency. We find that sudden, neutral perturbations in decision parameters trigger immediate and persistent improvements in choice quality, driven by enhanced cognitive processing. While these findings cannot be rationalized by standard rational inattention models, which predict stable decision quality across a sequence of similar decision problems, they nevertheless need not contradict the broader principles of optimal information acquisition. Once external activation is accounted for, behavior evolves in a manner consistent with a rational inattention framework in which individuals respond optimally to information costs and choice quality trade-offs. Our results highlight the importance of integrating rational inattention with neurophysiological mechanisms, demonstrating how stimulus-driven cognitive activation influences attention allocation and enhances decision-making efficiency. Our insights suggest that highlighting relevant changes – without altering information – can activate cognitive processing and enhance decision quality by activating dormant cognitive capacity, fostering endogenous deliberation without directing choices.

Conference Version (February 2025)

Keywords: Limited Attention, Rational Inattention, Cognitive Activation, Choice Quality

JEL Classification: D83; D81; C91

*We would like to especially thank Andrew Caplin as well as the numerous participants of the Viznau 2020 meeting on Attention and Choice for valuable comments.

[†] I) School of Management and Law, Zurich University of Applied Sciences and II) Department of Economics, University of Zurich. Email: ahefti@heftynomics.com.

[‡] Bern University of Applied Sciences (BFH). Email: tobias.mueller@bfh.ch.

[§] School of Economics, University of Queensland. Email: peiyao.shen@uq.edu.au.

[¶] Department of Economics, University of Zurich. Email: ernst.fehr@econ.uzh.ch.

1 Introduction

In today's digitized economy, individuals are inundated with information, from conducting financial investments to selecting insurance policies, all while contending with limited cognitive resources. Classical economic theory often assumes that decision-makers are endowed with unlimited cognitive processing abilities, enabling them to make optimal choices. However, a growing body of research demonstrates that real-world decision-making is systematically shaped by cognitive constraints, leading to departures from mere rational behavior. Evidence shows that individuals simplify complex choices, rely on heuristics rather than exhaustive information processing, and often make errors due to limited cognitive capacity (Kahneman and Tversky, 1979; Kahneman, 2003). These constraints influence economic decisions in domains ranging from consumer choice (Hossain and Morgan, 2011) to financial markets (Gabaix, 2019) and policy design (Chetty et al., 2009).

Devoting more attention to a choice task under uncertainty enables individuals to obtain more precise information, improving their ability to make refined choices and ultimately achieve better decision outcomes. However, because attention is scarce, allocating more to one task reduces the attention available for others, with a potentially lower decision quality elsewhere. Consider, for instance, the selection of a healthcare plan, where individuals must evaluate multiple options with varying costs and coverage structures under uncertainty. Greater attention to plan evaluation improves cost estimates, allowing for a more informed selection that minimizes expected expenditures given individual needs and preferences. However, this reallocation of attention comes at the expense of other important decisions, such as financial planning, workplace tasks, or family responsibilities, potentially diminishing decision quality in those areas. The optimal allocation of attention thus resolves a fundamental trade-off between improving accuracy in one domain and preserving decision effectiveness across multiple competing tasks.

The rational inattention framework, first formalized by Sims (2003), acknowledges the cognitive constraints that limit individuals' ability to process information and assumes that individuals allocate attention in a rational, goal-driven manner to maximize expected utility. While rational inattention provides a central normative benchmark for economic analysis, it does not incorporate the possibility that external attention stimuli – beyond an individual's direct control – may not only redirect attention but also actively enhance cognitive processing. Psychological and evolutionary perspectives suggest that certain stimuli, indicating some form of change by their suddenness or salience, can arouse attention and activate a heightened state of information processing. This mechanism may reflect an adaptive evolutionary function: for instance, the sound of a predator in the wild would have historically triggered an immediate, enhanced processing response, enabling individuals to respond effectively to changes in the environment.

If stimulus-driven activation has played a vital role throughout human evolution, it is highly unlikely that it would simply vanish in the modern age. Instead, it would remain embedded in human behavior, manifesting as a natural and relevant part of decision-making processes also in today's complex environments. In modern decision-making contexts, such stimulus-driven activation could similarly "switch on" heightened awareness and information processing capabilities. Importantly, however, such an activation would strongly depend on external, salient changes to initiate it. Under the rational inattention framework, individuals are presumed capable of autonomously triggering enhanced processing whenever it would improve their choices. Yet, if stimulus-driven mechanisms play a significant role, it suggests that these capabilities may be dormant or underutilized unless activated by an external change.

It is precisely this type of external change that may be challenging to reconcile within the existing rational inattention framework. A central benchmark in the rational inattention literature, established by Caplin and Dean (2015), characterizes the rationalizability of stochastic choice data under costly information acquisition. Their framework provides a revealed preference test that identifies all patterns of choice consistent with optimal costly information acquisition, without imposing parametric assumptions on the structure of information costs. According to Caplin and Dean (2015), the NIAC ("No Improving Attention Cycles") and NIAS ("No Improving Action Switches") conditions are necessary and sufficient for a dataset to be consistent with a model of optimal costly information acquisition, given a fixed information cost function. A key implication is that when a decision-maker is exposed to a sequence of decision problems that are consistent in economic incentives and cognitive difficulty, the rational inattention framework predicts a stable behavioral pattern across the sequence. This follows directly from the model's tenet that individuals optimize attention allocation autonomously and consistently under a given information cost function.

In this paper, we follow up on this observation by studying a new experimental design that introduces a sequence of incentivized decision problems, identical in structure and cognitive challenge, that require subjects to allocate attention and make choices. We are particularly interested whether the occurrence of a sudden and salient external change may affect the intensity of cognitive deliberation, and the quality of choice, through an alluded activation effect. Our experiment employs a series of choice problems, where participants engage in two sequentially ordered decision tasks: a contract-choice problem that mirrors real-world insurance selection, and a secondary quiz task designed to simulate competing cognitive demands. Participants must allocate their attention between these tasks, with their performance reflecting how cognitive resources are distributed. A key feature of our design is that the contract-choice task can change abruptly and infrequently over time, where participants are saliently alerted once a change has occurred. These alerts, however, exclusively indicate that a change in the contract task has occurred, without providing any hints or cues about how to best solve the task. Moreover, the contract task was designed such that the changes that occurred did not alter the cognitive challenges associated with identifying the ideal contract. These alternations serve as a test for the activation hypothesis: if decision-makers operate entirely within the rational inattention framework, the presence of such changes should not systematically alter their performance. Our alternative hypothesis is that the sudden changes act as a stimulus that alers subjects' cognition, thereby elevating their *intensive margin of attention* – their capability of efficiently processing perceived information per unit of time. We thus expect the subject to raise their overall decision quality due to this *activation effect*. In the rational inattention model, there is no scope for such external activation effects, because rational inattention exactly means that such activation itself should be part of the optimal information acquisition behavior.

Our approach is significant for several reasons. First, it enables us to test key predictions of the rational inattention model in a controlled decision-making context, contributing to a deeper understanding of its applicability and possible boundaries. Second, by examining how external changes might trigger enhanced cognitive processing, our experiment sheds light on mechanisms that operate alongside the rational allocation of attention, offering a broader perspective on how individuals manage cognitive constraints. Finally, this experiment provides empirical insights into how attention is allocated and how external factors shape decision quality – a topic of central importance in information-rich environments. By addressing these questions, our study seeks to enrich the understanding of attention and decision-making under cognitive constraints. A more comprehensive understanding of how human attention is allocated and how it influences choices has broad implications. It can inform the design of decision environments, policies, and tools that better accommodate the ways individuals process information in complex settings. By examining the factors that shape attention and cognitive engagement, our findings contribute to a deeper understanding of the cognitive processes underlying decision-making, with potential relevance for fields ranging from behavioral economics to public policy and technology design.

2 Related Literature

Our study contributes to the rich literature on attention allocation and decision-making under cognitive constraints. Attention is widely recognized as a gating mechanism that filters information, enabling individuals to focus on relevant aspects of their environment while disregarding less critical details. This process becomes particularly important in contexts of information overload, where individuals must allocate scarce cognitive resources effectively. Theoretical models have explored attention allocation through frameworks such as sequential search, information acquisition with varying signal precision, and entropy-based costs for state-dependent stochastic choices (e.g., McCall (1970); Sims (2003); Caplin and Dean

(2015); Matějka and McKay (2015)). These models provide foundational insights into the trade-offs between information costs and decision quality.

The rational inattention framework has been central to this line of inquiry, offering a formal approach to understanding how individuals allocate cognitive resources optimally in the face of uncertainty and informational constraints (Sims, 2003; Caplin and Martin, 2015). While these models emphasize rational, goal-directed attention allocation, they typically abstract from the influence of external stimuli that can activate attention in ways not fully captured by endogenous optimization. Our work extends this literature by explicitly testing the impact of non-informative external stimuli on attention allocation and decision quality, addressing a gap in the understanding of how exogenous changes in the informational environment affect cognitive engagement.

Empirical studies on attention often employ experimental methods, such as eye-tracking or Mouselab, to measure information acquisition processes (Reutskaja et al., 2011a; Gabaix and Laibson, 2003). However, while these techniques provide valuable insights into how attention is deployed, they do not directly address whether the observed patterns of attention allocation lead to optimal decision-making. Our experiment bridges this gap by simultaneously measuring attention allocation and evaluating the resulting decision quality in a controlled dual-task environment.

Finally, our study complements experimental efforts to test models of inattention. Prior work has examined sequential search (Caplin and Dean, 2011), status quo bias (Geng, 2016), and sharp predictions of rational inattention in state-dependent stochastic choices (Dean and Neligh, 2017). We build on this foundation by introducing external, non-informative stimuli as potential triggers for cognitive activation, allowing us to investigate how such stimuli interact with endogenous attention mechanisms. By focusing on the interplay between stimulus-driven and rational attention, our study provides new insights into the cognitive processes that shape decision-making in uncertain and complex environments.

2.1 Activation Effects in the Literature

Empirical and neuroscientific evidence suggests that human attention is governed by two distinct but interacting systems: a top-down, goal-directed mechanism and a bottom-up, stimulus-driven mechanism Corbetta and Shulman (2002). The top-down system, associated with the dorsal attention network, allows individuals to allocate cognitive resources voluntarily, based on internal goals and anticipated benefits. In contrast, the bottom-up system, mediated by the ventral attention network, is designed to detect and respond to unexpected, salient stimuli, ensuring that important environmental changes do not go unnoticed. While these two systems are anatomically distinct, they are functionally intertwined: top-down attention can suppress distractions, but bottom-up signals can override this suppression when sufficiently salient. This suggests that attention should not be modeled as a purely endogenous process but rather as an integrated system where external stimuli can modulate attention allocation, even when the individual has no prior intention of reorienting focus. Such interactions highlight the need for a more comprehensive model of rational inattention—one that incorporates stimulus-driven activation effects into endogenous attention allocation. A key implication of this two-system model is that bottom-up attention does not merely redirect focus but may also enhance cognitive processing. By interrupting a lower-arousal attentional state, external stimuli may increase cognitive efficiency, leading to higher-quality information processing in complex decision tasks. This idea—that bottom-up attention may serve as a cognitive activator—is central to our theoretical framework, where we formalize how exogenous stimuli can enhance processing efficiency beyond what standard rational inattention models predict.

Empirical evidence suggests that externally driven attention shifts can significantly influence economic choices. For instance, Reutskaja et al. (2011b) use eye-tracking to demonstrate that objects placed in salient positions receive disproportionate visual attention and are more likely to be chosen, independent of their intrinsic value. This finding highlights the role of stimulus-driven attention in economic decision-making, suggesting that exogenous factors can shape choices beyond purely endogenous optimization mechanisms.

A well-documented mechanism demonstrating the power of stimulus-driven attention is the orienting response Sokolov (1963), an involuntary shift in cognitive focus triggered by unexpected external stimuli. This response, which evolved as a survival mechanism, redirects attention toward salient environmental changes, enhancing sensory processing and cognitive efficiency.

Petersen and Posner (2012) provide neuroscientific evidence that external stimuli enhance cognitive processing. Specifically, they show that phasic activation of the brain's alerting network, triggered by transient external cues such as warning signals or unexpected sensory changes, leads to a temporary increase in alertness and readiness to process information. This activation is mediated by the locus coeruleus-norepinephrine (LC-NE) system, which facilitates faster reaction times and improved sensory sensitivity. While these mechanisms are well-documented in attentional control, they suggest a broader implication: external stimuli can serve as cognitive activators, increasing processing efficiency without increasing deliberation time.

The influence of external stimuli on attention may have a strong evolutionary basis. Carretié (2014) argues that exogenous (stimulus-driven) attention evolved as an adaptive mechanism to ensure rapid detection and response to critical environmental changes, particularly those related to survival. This perspective suggests that, despite the increasing complexity of modern decision-making environments, human cognition may still retain an inherent sensitivity to external stimuli, shaping attention allocation in ways that extend beyond purely endogenous, utility-maximizing considerations.

Psychological research has demonstrated that external arousal can increase cognitive engagement captured by physiological measures. For example, Kahneman (1973) tracked pupil dilation as an indicator of mental effort and attentional load in a series of experiments. His findings revealed that externally induced arousal – such as environmental stressors like noise – can enhance or impair cognitive engagement in tasks like tone discrimination and arithmetic problem-solving. By contrast, a key implication of rational inattention is that external stimuli can alter decision-making only if they change the costs or benefits of acquiring information.

If a stimulus interferes with information processing – such as distracting noise impairing cognitive performance in complex tasks as in Kahneman (1973) – it can be incorporated into existing RI frameworks as an increase in the cost of attention. This reasoning explains why Kahneman found evidence that such distractors can decrease performance. However, this perspective struggles to reconcile another key finding: Kahneman also observed cases where environmental stressors, such as noise, elevated arousal levels and enhanced performance.[1]

A particularly relevant class of external stimuli in this context are informationally neutral stimuli – stimuli that induce arousal and alertness while not altering the cognitive aspects of the decision task. Unlike distracting stimuli that increase the cost of attention, informationally neutral stimuli do not alter the structure of the decision problem or the complexity of information acquisition. RI models predict that such stimuli should have no effect on decision-making, since they do not change the trade-off between attention costs and benefits. However, if external stimuli can activate cognitive engagement independently of cost-based attention allocation, then informationally neutral stimuli provide a crucial empirical test of whether stimulus-driven activation effects exist. If such stimuli can improve decision quality, this would suggest that attention is not solely governed by endogenous optimization, but can, and in some cases needs to be, exogenously activated.

3 Behavioral Predictions

We consider an agent facing multiple decision tasks, where the objective is to decide how much attention to invest into each task and determine a task-wise optimal choice based on available information. In each task j, the most desirable response is represented by an unknown quantity X_j, and before making a choice, the agent receives a signal S_j about X_j. Consistent with models of rational inattention, the informativeness of S_j depends on the attention devoted to task j. The agent seeks to optimally allocate attention between the tasks, balancing the cost of acquiring better information against its expected benefits, where greater attention leads to a

[1]Kahneman's evidence overall points towards an inverted-U-shaped relation between arousal and effort, consistent with the Yerkes-Dodson Law.

better understanding of the task. We extend this framework by introducing an external activation mechanism: the agent's ability to extract and process information from attention depends on their arousal state, which can be modulated by external stimuli.

3.1 A Reduced-Form Model of Attention and Choice

To exemplify the above notions with two tasks $j = 1, 2$, suppose that the agent seeks to choose task-specific actions κ_1, κ_2 to minimize the expected quadratic loss $U = c_1 E[(\kappa_1 - X_1)^2] + c_2 E[(\kappa_2 - X_2)^2]$, with decision weights c_1, c_2, representing the relative importance of the two tasks to the agent. A rational agent updates via Bayes' Rule, where the task-wise optimal response is given by the mean posterior $\kappa_j(S_j) = E[X_j|S_j]$, and expected utility by the mean conditional variance. Let $X_j \sim N(0, \sigma_j^2)$, $S_j = X_j + \varepsilon_j$, $\varepsilon_j \sim N(0, \sigma_{\varepsilon_j}^2)$, where σ_j^2 represents the agent's ex-ante information about task j, and $\sigma_{\varepsilon_j}^2$ captures how revealing the signal is about the optimal action. Task-specific states X_1, X_2 and errors $\varepsilon_1, \varepsilon_2$ are independent.

Consider the important benchmark case of an infrequent task, like choosing a healthcare plan or buying a new car, where familiarity and routine are absent; the agent holds an uninformative prior. In such a case, the agent's posterior belief is determined solely by the signal: $X_j|S_j \sim S_j$, or equivalently $\sigma_j^2 \to \infty$. Thus, by Bayesian updating, ex-post information relies entirely on the precision of the signal, $q_j \equiv 1/\sigma_{\varepsilon_j^2}$, which we will interpret as the *quality of choice* in task j.

This quality depends on the amount of attention $\alpha_j \geq 0$ directed to the task, following $f_j = \alpha_j \tau_j$, where $\tau_j > 0$ captures the marginal productivity of attention in generating knowledge. Consistent with our experimental setup, we suppose an exogenously fixed amount of attention available for both tasks, normalized to unity, such that $\alpha_1 + \alpha_2 \leq 1$.[2] A rational agent chooses qualities q_1, q_2 to minimize $U = c1/q_1 + c_2/q_2$, anticipating how attention α_1, α_2 translates into quality. The fact that attention is scarce ($\alpha_1 + \alpha_2 \leq 1$) implies the existence of an attention

[2] In the experiment, this amounts to the fixed time budget of 120 seconds per period. Although not relevant to our analysis, one could endogenize this total amount by an Envelope-Theorem argument, where the agent also trades offs the benefits of an additional unit of additional total attention available for both tasks against its current costs, e.g., in terms of metabolic energy (Lennie, 2003).

frontier, similar to a budget constraint, reflecting the trade-offs in acquiring a more precise internal representation of the signal in each task. The shape and location of this frontier depend on the marginal productivities of attention, capturing how paying more attention to a task transforms in a better understanding of that task.

Figure 1 depicts the optimal choice of q_1, q_2 subject to the attention frontier. Our

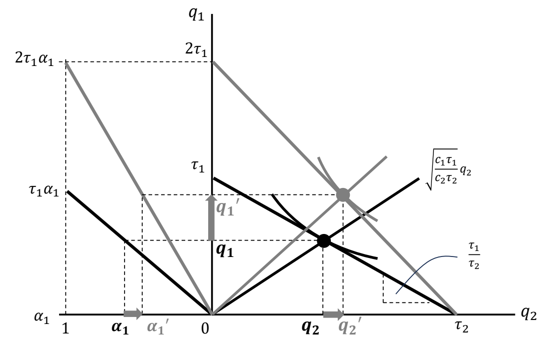

Figure 1: Optimal Choice and Activation Effect

simple model is explicitly solvable and has a unique optimal attention allocation and quality levels, indicated in black in Figure 1, and formally given by

$$q_1 = \frac{\tau_1}{1+\gamma}, \quad q_2 = \frac{\gamma\tau_2}{1+\gamma}, \quad \alpha_1 = \frac{1}{1+\gamma}, \quad \alpha_2 = \frac{\gamma}{1+\gamma}, \quad \gamma \equiv \sqrt{\frac{\tau_1 c_2}{\tau_2 c_1}}. \tag{1}$$

The optimal solution fully exhausts attention – reflecting monotonicity of the preferences for quality – and has further notable properties. E.g., situations with $q_1 > q_2$ but $\alpha_1 < \alpha_2$ can emerge as a consequence of rational attention, reflecting an inherent propensity to use attention to smooth qualities across different tasks.[3] This smoothing property also is relevant for the activation effect and intertemporal considerations pertaining to the use of attention.

3.1.1 The Activation Effect and its Implications

The *activation effect* represents an exogenous increase in task-specific cognitive processing efficiency triggered by an external stimulus. Let $A_j \in \{0, 1\}$ denote the state of arousal in task j. If arousal occurs, this manifests in a heightened *intensive*

[3]Such a situation arises, e.g., if $\tau_1 > \tau_2$ and $c_1 = c_2$. If one is indifferent between both tasks, it is optimal to exploit some of the higher cognitive productivity in task $j = 1$ to free additional attention for task $j = 2$ and re-balance decision quality accordingly.

margin of attention and a correspondingly improved ability to convert attention into decision quality, formalized as $\tau_j' = \tau_j + \Delta_j A_j$, with $\Delta_j > 0$. The main effects of arousal in task $j = 1$ are as follows:

Proposition 1 *Activation in task $j = 1$ implies higher optimal qualities in both tasks, with $\frac{dq_1}{q_1} > \frac{dq_2}{q_2}$, and an increase in the attention α_2 directed to task $j = 2$.*

Proof: Solution (1) is continuously differentiable in τ_1, and all claims follow from calculating the respective derivatives. $\square$

An upward jump in the processing rate τ_1 leads to a rotation of the attention frontier, as depicted in Figure 1 (gray lines). The key implication is that a given level of q_1 now requires less attention, freeing up attention for task 2. As a result, the agent can achieve higher quality in task 2 without sacrificing quality in task 1, via a reallocation of attention along the *extensive margin of attention.* Proposition 1 then shows that it is indeed optimal for the agent to use some of the freed attention capacity to improve choice quality in task 2. However, the increase in quality in task 2 occurs at a smaller rate, reflecting the localized nature of the activation effect with its primary impact on the activated task.

3.2 Information Processing: Neuroscientific Foundations

We now relate our reduced-form model to neurophysiological evidence of cognitive information processing. Neural decision-making is based on the accumulation of information over time, where sensory inputs are processed through coordinated neural activity. A well-established principle in neuroscience is that task-relevant information is encoded in the firing activity of populations (ensembles) of neurons (Shadlen and Newsome, 1998).[4] The following principles are fundamental to neurophysiological information processing (Gold and Shadlen, 2007). Neurons process sensory inputs through *spike trains* – stochastic sequences of spikes (electrical pulses) fired by neurons. The informativeness of these trains depends primarily on two key dimensions: *rate coding* (average firing rate of neurons) and *spike vari-*

[4]Neural information processing is typically studied at the level of neural ensembles, where functionally related neurons work together to encode and integrate sensory signals.

ability (fluctuations in the spike trains due to synaptic and cellular noise).[5] When accumulating evidence for a decision, the brain integrates information from neuron ensembles over time through *population averaging*, a process in which multiple individual spike trains are combined to reduce noise and improve the reliability of the encoded information. Besides firing rate and variability, the accuracy of these representations therefore depends on the duration of evidence accumulation. Longer integration periods enhance signal stability by averaging out stochastic fluctuations in neural activity (Shadlen and Newsome, 2001; Ratcliff and McKoon, 2008).

The above notions suggest a neurophysiological foundation for our reduced-form model from Section 3.1. The spike trains generated by the neuron ensembles relevant for processing task j contribute evidence, which we model as a sequence of information signals ε_{ji}. Each neuron in the ensemble fires at an average rate of λ_j, and over a duration (attention span) of α_j, resulting in the accumulation of an expected $n_j = \lambda_j \alpha_j$ noisy signals per neuron. Population averaging combines responses from multiple neurons, leading to a reduction in the variability of the encoded information as the relative impact of noise diminishes. To capture this key effect, it is common to approximate neural variability as independent across the ensemble.[6] Accordingly, we assume that each spike train constitutes an iid signal $\varepsilon_{ji} \sim N(0, m_j\sigma^2)$. The variance $m_j\sigma^2$ captures neural variability arising from synaptic noise, cellular fluctuations, and experience-dependent refinements in neural circuits. This formulation accounts for the fact that different tasks may engage distinct neuronal populations, which exhibit varying levels of variability in their firing activity. Population averaging then suggests an aggregate signal $\varepsilon_j = 1/n_j \sum_i^{n_j} \varepsilon_{ji}$, which is zero-normal with precision $\frac{\lambda_j}{m_j\sigma^2}\alpha_j$.

From this perspective, the intensive margin of attention in our rational decision model, $\tau_j \equiv \frac{\lambda_j}{m_j\sigma^2}$, should depend on the firing rate λ_j relative to neural noise

[5] While the timing of the spikes in a train may also carry information (so-called temporal coding), evidence indicates that rate coding is the primary mechanism for encoding information in perceptual decision tasks (Gold and Shadlen, 2007).

[6] For example, independence is a common approximation in decision accumulation models, where moment-to-moment evidence is treated as independent noisy increments for analytical tractability (Ratcliff and McKoon, 2008). While neuronal signals can exhibit correlations (**?**), these tend to be weak in higher-level decision areas. Moreover, in many cases, ignoring them leads to little or no loss of information, supporting the use of independence as a reasonable simplification (Eyherabide and Samengo, 2013).

$m_j\sigma^2$ in a task.[7] In fact, this expression clarifies that there are two key ways how the intensive margin of attention may improve: either via a higher arrival rate of neural signals (λ_j) or by reducing stochastic variability due to gaining incremental experience. Neurophysiology has provided evidence for both, showing important differences, which we incorporate in the following.

3.2.1 Neural Firing rates and the Activation Effect

The activation effect from Section 3.1.1, has a direct neurophysiological basis in the modulation of neural firing rates. Empirical evidence surveyed by Gold and Shadlen (2007) shows that firing rates are dynamically modulated by external factors such as stimulus salience, urgency, or sudden environmental changes, facilitating faster evidence accumulation. One key aspect is that external stimuli can induce abrupt and substantial increases in neuronal firing rates Rotstein (2013); Vardi et al. (2013). Another is that neuronal activation task-specific external stimuli selectively enhance the firing rates of the specific neurons engaged in the relevant task, without broadly affecting other neural populations (Rischka et al., 2021). This aligns with our task-specific activation mechanism, which can be represented as $\lambda_j' = \lambda_j + \Delta_j A_j$ in the current refinement.

3.2.2 Learning Dynamics

While external activation leads to a sudden increase in the efficiency of attention, learning operates through a gradual and endogenous process of refinement. Learning is best understood as an experience-driven reduction in trial-to-trial variability, stabilizing neural responses over time. In their review, Gold and Shadlen (2007) show that learning and experience primarily reduce trial-to-trial variability in spike trains rather than increasing average firing rates. As tasks are repeated, neurons become more consistent in their responses, leading to lower baseline noise in spike trains but not necessarily higher overall activity. Moreover, task-evoked activity has been shown to reduce neural variability across multiple cortical areas through

[7] In reality, neurons fire in stochastic rather than uniform patterns, making a Poisson-distributed spike count, $n_j \sim Poisson(\lambda_j\alpha_j)$, a more accurate representation. The above solution remains approximately correct in this more realistic case, as $Var(\varepsilon_j) = E[Var(\varepsilon_j|n_j)] = m_j\sigma_2 E[1/n_j]$, where $E[1/n_j] \approx 1/E[n_j]$ holds for sufficiently large $\lambda_j\alpha_j$.

a mechanism that stabilizes neural responses during task engagement. Ito et al. (2020) demonstrate that task engagement globally quenches neural variability and correlations across different cortical areas, leading to more stable and consistent neural responses. This stabilization process reduces variability and shared noise across cortical areas, indicating a general mechanism by which repeated task engagement refines neural coding over time.

In our model, we capture this process through a gradual reduction in baseline neural variability $\sigma^2(t)$ over time, formalized by the condition $\frac{d}{dt}\sigma^2(t) < 0$, as a consequence of repeated task exposure. This reflects the idea that neural representations become more precise with experience, reducing noise in evidence accumulation. We refer to this time-dependent refinement of neural variability as *learning dynamics.*

While both learning and activation enhance the intensive margin of attention τ_j, they differ in fundamental ways (Gold and Shadlen, 2007). Activation is an externally induced and discontinuous surge in the firing rate, whereas learning is endogenous and gradual. Activation is a discontinuous surge in the firing rate for the exposed task in response to sudden, task-specific external stimuli. Learning, in contrast, results from repeated exposure to tasks, leading to a steady refinement of neural responses through a reduction in spike train variability. This distinction is crucial for understanding how experience-driven improvements in attention allocation differ from externally driven activation effects. It also has a metabolic rationale: since maintaining high firing rates is energy-intensive, the brain optimizes decision-making by improving signal reliability through reduced variability rather than increased spike output (Moujahid et al., 2014).

We now derive the key predictions of our optimal attention model under the assumption that learning follows the above gradual process over time, holding the activation level λ_j constant.

Proposition 2 *As learning progresses, the optimal quality levels q_1 and q_2 both increase over time, while the allocation of attention α_1 and α_2 remains constant.*

<u>Proof</u>: Follows from (1) and $\tau_j(t) = \frac{\lambda_j}{m_j \sigma^2(t)}$. $\square$

The intuition is that the main effect of learning – the reduction in baseline neural variability – is an outward parallel shift of the attention frontier in Figure 1. By homogeneity of preferences over quality, this implies a proportional increase in both quality levels while maintaining a constant allocation of attention across tasks.

4 The Attention-Choice-Experiment

4.1 Experimental Design

This study employs a laboratory experiment designed to empirically test predictions of the rational inattention model. Unlike previous studies that consider fairly artificial decision-making tasks, such as choosing the most frequent color or number, our experiment introduces a complex, state-dependent environment meant to mirrors real-world consumer decision-making. By observing how participants allocate limited cognitive resources between competing tasks, we provide a rigorous test of the hypothesis that individuals optimally allocate their attention based on expected utility.

4.1.1 Sequential Dual-Taks Setting

The core of our design is a dual-task setting that simulates real-world choices faced by consumers, particularly in insurance markets. Participants engage in two incentivized tasks: (a) a contract-choice problem and (b) a general-knowledge quiz task.

In the **contract-choice problem**, participants are presented with an exogenous risk profile indicating the probabilities of three potential adverse events (A, B, and C) (Panel (a) in Figure 2). Each event results in a different monetary loss, which participants can mitigate by selecting one of four insurance contracts. The contracts vary in terms of premium costs and coverage levels, requiring participants to evaluate trade-offs between upfront costs and potential losses. For example, Contract 1 features a premium of 2,000 points, with losses of 50 points for event A, 150 points for event B, and 690 points for event C. Determining the optimal contract necessitates mental effort, as participants must integrate information about premiums, probabilities, and potential losses to rank the contracts based on expected net benefits.

Simultaneously, participants can allocate time to a **quiz task** consisting of general-knowledge questions (Panel (b) in Figure 2). Each correctly answered question contributes a fixed piece-rate payment to the participant's overall payoff. The quiz questions are designed to be cognitively demanding, requiring similar levels of mental effort as the contract-choice problem. This ensures that both tasks are comparable in terms of cognitive load, allowing us to observe genuine trade-offs in attention allocation.

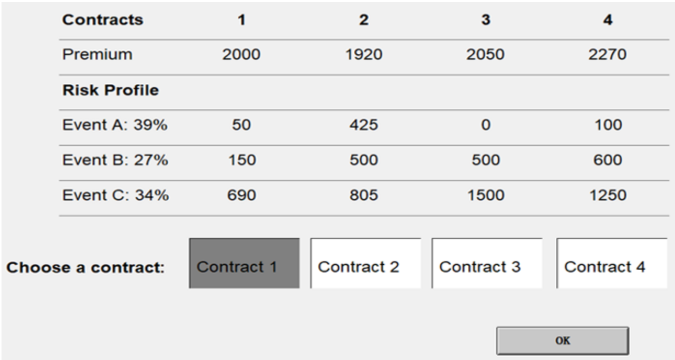

(a) Contract Choice

(b) Quiz Task

Figure 2: Example of Contract Task and Quiz Task

4.1.2 Measuring Attention and Decision Quality

Each period of the experiment is divided into two stages (Figure 3). In the first stage, participants are given a risk profile with probabilities for the three adverse events and have 20 seconds to decide whether to allocate any time to the contract-choice problem. In the second stage, participants receive a fixed time budget of 120 seconds, which they can freely allocate between the contract-choice task (X seconds) and the quiz task (120 - X seconds). If participants choose to allocate zero seconds to the contract-choice problem, they spend the full 120 seconds on the quiz task. The primary measure of attention is the time spent on the contract-choice problem, which reflects the participant's allocation of cognitive resources. The opportunity cost of focusing on the contract problem is the potential earnings from the quiz task, creating a clear trade-off in attention.

At the end of each period, a random draw based on the given risk profile determines the occurrence of one of the three adverse events, establishing the monetary consequences of the selected contract. This randomization ensures that the risks are exogenous, allowing for a clean evaluation of decision quality based on the chosen contract.

4.1.3 Treatment Design

Our experiment employs a **2x2 treatment design** to examine how changes in cognitive demand and monetary incentives affect attention allocation. The first dimension of the design is a within-subject treatment that varies the risk profiles across periods. While the cognitive difficulty of the contract-choice problem remains constant, the optimal contract may change, allowing us to observe shifts in decision quality that result from variations in attention rather than changes in task complexity. Additionally, the expected net benefit of the optimal contract is designed to first-order stochastically dominate the second-best contract, ensuring a consistent monetary incentive across periods.

The second dimension of the design is a between-subject treatment that manipulates the piece-rate reward for correctly answered quiz questions. In one condition, the reward is set at a high level, increasing the relative incentive to allocate time to the quiz task. In the other condition, the reward is lower, making the contract-choice task relatively more attractive. This variation allows us to test whether participants adjust their attention allocation in response to changes in relative payoffs, as predicted by the rational inattention model.

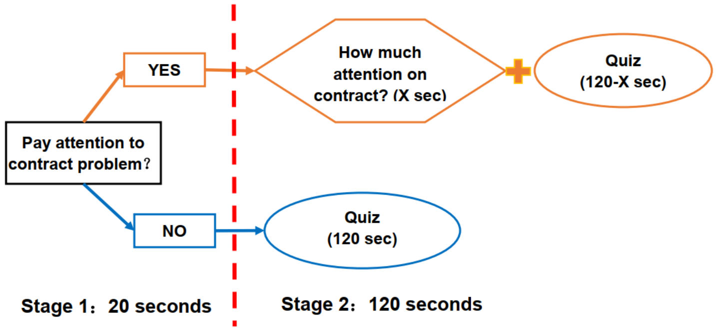

Figure 3: Decision Problem in One Period

4.1.4 Implementation and Procedure

The experiment was conducted over six sessions at the University of Zurich using the z-Tree software (Fischbacher, 2007). A total of $n = 137$ subjects participated in the experiment, with each subject only participating in one session. Before a session started, participants had to answer a series of control questions correctly to ensure their understanding of the instructions. A session lasted approximately 1.5 hours.

5 Experimental Results

This section examines key aspects of participants' behavior, focusing on attention allocation, decision quality, and quiz performance. We begin by analyzing attention allocation, assessing how participants distribute their cognitive resources between stimulus and non-stimulus periods. Next, we evaluate choice quality, examining whether exposure to stimuli leads to persistent improvements in decision outcomes, consistent with our mental activation hypothesis. We then shift our focus to quiz performance, exploring how participants' cognitive effort allocation affects their ability to answer quiz questions accurately and efficiently. Finally, we categorize participants into distinct cognitive types, allowing us to assess heterogeneity in decision-making responses and its relationship with individual cognitive traits.

5.1 Attention Allocation

In the first step of our analysis, we examine how participants allocate their attention to the contract-choice task in periods when the risk profile changes (hereafter referred to as "stimulus periods") compared to periods without any change in the risk profile ("non-stimulus periods"). Figure 4 presents two key visualizations of attention allocation across the 16 experimental rounds. The left panel shows the proportion of participants who accessed the contract-choice screen in each round, while the right panel displays the average time (in seconds) spent on the contract-choice task.

Overall, the results indicate a clear pattern: participants allocate significantly more attention to the contract-choice problem during stimulus periods. For example, in rounds 4 and 8, which correspond to stimulus periods, 96% of participants chose to enter the contract-choice screen. This behavior aligns with the predictions of rational inattention theory, suggesting that participants increase their cognitive effort when the task environment becomes more uncertain. However, a substantial share of participants also chose to engage with the contract-choice task during non-stimulus periods. Both the likelihood of accessing the contract screen and the time spent on the task in non-stimulus periods are significantly greater than zero, posing a challenge to the strict predictions of rational inattention. One possible explanation for this behavior is cognitive uncertainty, as suggested by recent studies (Enke and Graeber, 2023). Cognitive uncertainty may lead participants to overestimate the potential importance of the contract-choice task, even in periods when no change in the risk profile occurs, resulting in a more consistent engagement with the task across all periods.

Result 1 *Participants allocate significantly more attention to the contract-choice task during stimulus periods, consistent with a baseline prediction of rational inattention. Nonetheless, a notable share of participants also engage with the contract task during non-stimulus periods. The presence of cognitive uncertainty may explain this pattern, as participants potentially overestimate the task's importance even in stable periods, leading to sustained attention allocation across all rounds.*

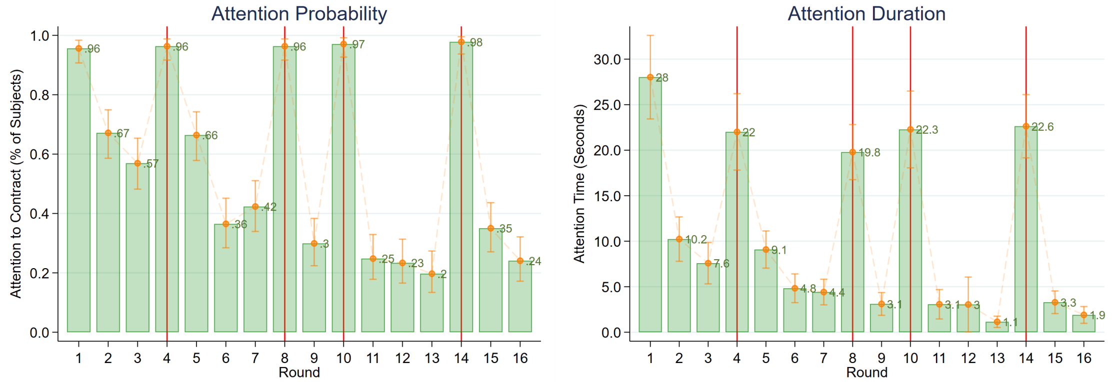

(a) Share of Participants Entering Contract Screen

(b) Average Duration Spent on Contract Screen

Figure 4: Attention on Contract Task

Note: Panel (a) shows the share of individuals paying attention to the insurance choice problem. Panel (b) shows the amount of time (in seconds) spent on the insurance choice screen. Error bars represent 95% confidence intervals. Vertical red lines indicate the stimulus rounds, i.e., periods when the risk profile has changed.

5.2 Choice Quality

In this section, we evaluate whether the quality of choices in the contract task improves following the stimulus appearance. The primary question is whether the stimulus induces a persistent enhancement in decision-making performance in this task. To assess decision quality, we mainly focus on *Expected Excess Cost (EEC)*, which quantifies the expected loss incurred from selecting a suboptimal contract. This measure provides a direct and interpretable metric of economic decision-making performance. Additionally, we consider two alternative measures: *Rank*, an ordinal variable ranging from 1 (best) to 4 (worst), and *Best*, a binary indicator equal to 1 if the optimal contract is chosen.

Figure 5(a) shows the average EEC over periods, highlighting a sharp reduction in excess costs during the first stimulus period ($t = 4$).[8] Similarly, Figure 5(b) reveals that the probability of choosing the optimal contract doubles during the first stimulus period. Importantly, decision quality remains persistently high in subsequent periods without further improvements from additional stimuli. Likewise,

[8]Excess costs are computed as the expected cost of the selected contract divided by the expected costs of the best contract and then multiplied by 100.

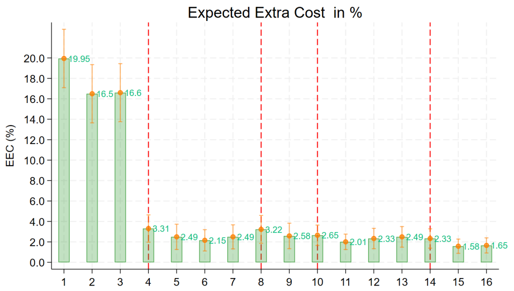

(a) Expected Excess Costs

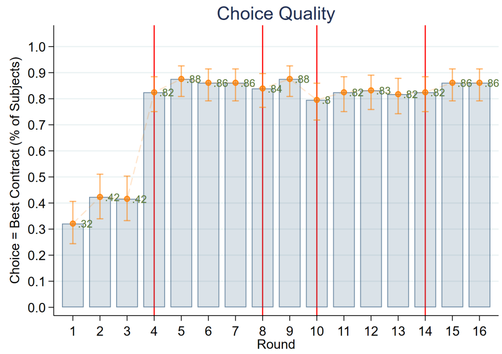

(b) Prob to chose Best Contract

Figure 5: Contract Task: Choice Quality

Note: Figure (a): EEC, over periods. Figure (b): Probability to chose the optimal contract, over periods. Error bars represent 95% CI's. Vertical red lines indicate stimulus periods.

Figure 6(a) reveals a prominent difference in the distribution of *Rank* between the first three periods (prior to any stimulus) and the stimulus periods, also indicating an evident improvement in choice quality. Figure 6(b) confirms that this finding is closely related to attention, as mainly participants who respond immediately to the stimulus achieve better outcomes.

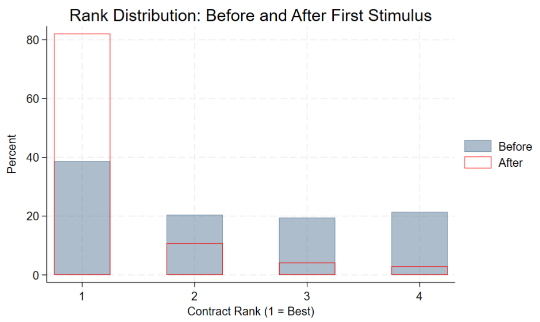

(a) Before vs. After First Stimulus

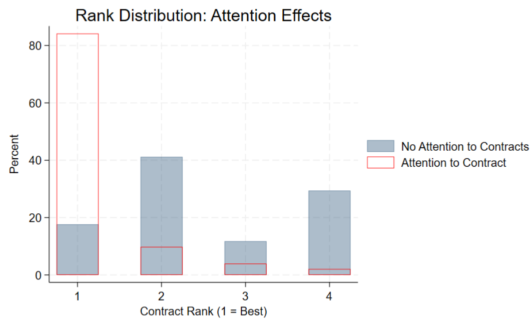

(b) Attention Effects

Figure 6: Quality of Contract Choices

Note: Share of participants by contract rank, whereas rank 1 corresponds to the best contract. Figure (a): Compares Rank in the first three periods to Rank in all subsequent stimulus periods. Figure (b): Compares Rank for subjects who entered the contract task in stimulus periods to those who have not.

5.2.1 Statistical Analysis and Results

To substantiate these visual findings, we conduct formal statistical analyses. First, we perform a formal structural break test for panel data based on Bai and Perron (2003). The test identifies a significant break at $t = 4$ ($p < 0.01$), coinciding with the first stimulus period. By contrast, when applied to periods $t > 4$, the test fails to reject the null hypothesis of a structural break.

Second, we estimate panel regressions using a random-effects specification.[9] Our model regresses EEC on a dummy variable, *Treat4*, which captures the effect of the first stimulus period, and control variables *Period* and $Period^2$ to account for potential learning effects. The effect of *Treat4* is large and significant ($\beta = -6.83$, $p < 0.01$), indicating a discontinuous improvement in decision quality in the first stimulus period. The time trend reveals a steady decrease in excess costs ($\beta_t = -3.9$, $p < 0.01$), with diminishing returns ($\beta_{t^2} = 0.17$, $p < 0.01$) as participants gain more experience. Restricting the sample to post-stimulus periods ($t > 4$) results in no significant time trends or additional stimulus effects.

Taken together, the break analysis and panel regressions show that decision quality remains persistently high after the first stimulus period, with no further substantial gains induced by subsequent stimuli or continued learning. These patterns persist when we include subject-specific linear time trends to account for individual learning effects and a lagged measure of *EEC* to control for potential serial correlation in decision quality across periods.[10]

To further address subject heterogeneity, we further include two subject-specific variables: *Math_Abil* (math ability score) and *Initial_EEC* (EEC in the first period) in our panel regressions. Both measures predict lower EEC over time, indicating better decisions by more capable individuals. However, only *Initial_EEC* interacts significantly and negatively with *Treat4*, suggesting that participants with poorer initial decision-making benefit more from the stimulus. In other words, individuals

[9] We use random effects because all regressors represent common time trends or vary identically across subjects, leaving no meaningful between-subject variation. In fact, both the fixed and random effects estimations yield identical coefficients and standard errors in this case.

[10] Such serial correlation may stem from individual ability, behavioral inertia, or learning dynamics. When incorporating lagged EEC, we employ the GMM estimator developed by Arellano and Bond (1991).

with poorer initial decision-making quality benefit more from the initial stimulus.[11]

5.2.2 Mere Learning Effects?

The fact that decision quality in repeated decision tasks improves over time is a well-documented regularity in the literature on cumulative and reinforcement learning. These theories have long been central to understanding how individuals adapt their behavior over time in response to experience and feedback. Reinforcement learning, formalized by Roth and Erev (1998), describe and experimentally validate a trial-and-error process wherein agents incrementally adjust their strategies based on the observed outcomes of past actions. Similarly, cumulative learning models emphasize the refinement of behavior as individuals gain familiarity with a task. Empirical evidence indicates that decision quality improves gradually over repeated exposures, with competence building steadily over time (Heath and Tversky, 1991). Experimental studies support this notion, showing that participants refine their strategies progressively through incremental updates (Nagel, 1995).

The common thread of both frameworks is that learning evolves incrementally. Performance improvements are continuous and predictable, driven by accumulated experience or the reinforcement of previously successful actions. However, the sudden and dramatic improvement observed at the first stimulus period is inconsistent with these gradualist perspectives. Such abrupt shifts suggest a mechanism beyond conventional learning, potentially arising from contextual triggers or external shocks that instantaneously enhance cognitive engagement. This insight, together with the above findings, is summarized as follows:

Result 2 *Decision quality exhibits a sharp and significant improvement during the first stimulus period, inconsistent with the gradual progression predicted by conventional learning theories. This abrupt shift suggests an external trigger that instantaneously enhances cognitive engagement rather than a slow adaptation process. Furthermore, subjects with poorer initial decision quality benefit the most, highlighting the potential of external triggers to unlock latent cognitive potential.*

[11] An examination of the marginal effects indicates that this heterogeneous effect is not just reflecting mere ceiling effect, where able choosers simply cannot improve their choices: we find that the marginal effects of all possible values of *Initial_EEC* are significantly different from zero.

5.2.3 Mental Activation

To further investigate the nature of the improvement in decision quality, we distinguish between two potential drivers: an increase in time allocated to decision-making (extensive margin) and improved cognitive efficiency (intensive margin). Specifically, we test the mental activation hypothesis by which exposure to the stimulus enhances mental processing.

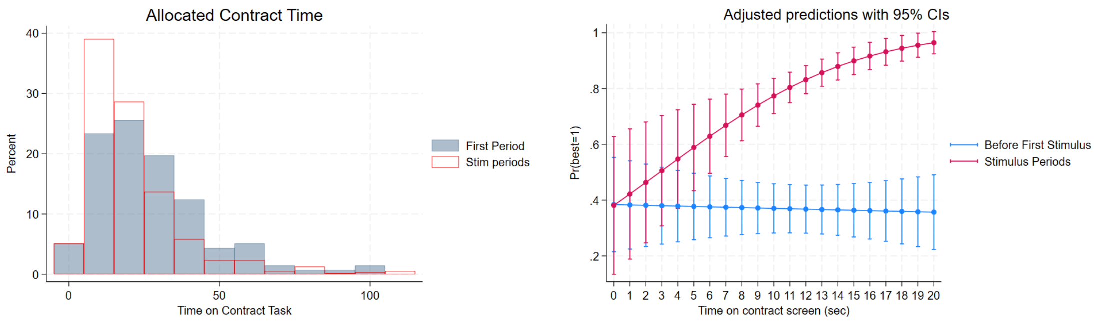

(a) Extensive Margin of Attention (b) Intensive Margin of Attention

Figure 7: Quality of Contract Choices

Note: Share of participants by contract rank, whereas rank 1 corresponds to the best contract. Figure (a): Compares Rank in the first three periods to Rank in all subsequent stimulus periods. Figure (b): Compares Rank for subjects who entered the contract task in stimulus periods to those who have not.

As Figure 5(b) showed, the likelihood to choose the best contract increases substantially after the first stimulus period. However, Figure 7(a) presents the distribution of time spent on the contract task, showing that subjects allocate less time to the contract in stimulus periods compared to the initial period. In other words, subjects reduce their extensive margin of attention to the contract but still manage to choose better contracts after the first stimulus periods. These facts indicate that their intensive margin of attention has likely improved upon the first stimulus period. To examine intensive margin improvements, we estimate a random-effects probit model for the probability of selecting the best contract, conditional on time spent.[12] Figure 7(b) shows that additional seconds increase the likelihood of choosing the best contract, but only in stimulus periods. Moreover, there is no noticeable

[12]We constrain this regression to subjects allocating $\tau \in \{1, 2, ..., 20\}$ seconds to the contract task. The upper bound $\tau = 20$ includes most subjects, and excludes subjects with very long durations. Adding these subjects results in noisier estimates.

difference in this probability for people who barely look at the contracts ($< 6s$). In sum, this picture is strongly consistent with the hypothesis that the stimulus activates an enhanced information processing, increasing the effectiveness of analyzing the contract information.

Result 3 *Following the initial improvement, decision quality remains persistently high across subsequent periods without additional gains from further stimuli. Our analysis indicates that this improvement is driven by an enhanced intensive margin of attention – participants process information more effectively – and not by allocating more time to the task.*

5.3 Quiz Performance

In the next step of our analysis, we shift our focus from the contract-choice task to the performance in the quiz task. We assess quiz performance using two key metrics: (i) quiz speed, defined as the number of seconds spent per quiz attempt (regardless of whether the answer is correct), and (ii) quiz efficiency, defined as the number of seconds spent per correct quiz answer.

5.3.1 Quiz Speed & Efficiency

To characterize quiz performance, Figure 8 presents two panels. Panel (a) displays the average quiz speed, measured by the number of seconds participants spent per quiz question attempt, while panel (b) shows the average time spent per correct answer, which we refer to as "quiz efficiency". The results indicate a smoothly declining pattern in "quiz speed" over time, suggesting a learning effect. Participants become faster at answering quiz questions as they progress through the experiment, reducing the time per attempt. In contrast, participants significantly increase quiz efficiency as indicated by the reduction in time spent per correct answer between periods 1-3 and period 4 onwards.

5.3.2 Hit Rate and Opportunity Costs

In this part of the analysis, we present two additional metrics to further evaluate participants' performance in the quiz task: (i) the hit rate of correct answers and

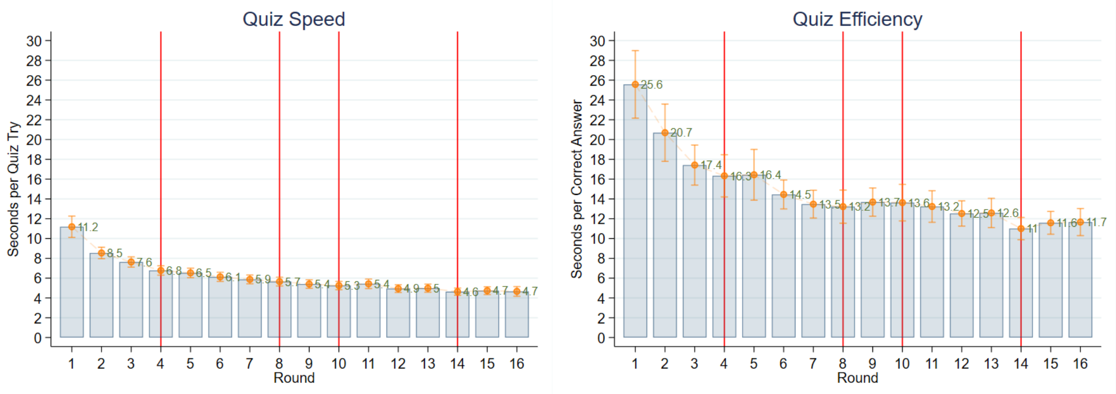

(a) Seconds per Quiz Try (b) Seconds per Correct Quiz Answer

Figure 8: Mental Productivity of Quiz questions

Note: Figure (a) plots the average number of seconds spent per quiz question (tries). Figure (b) plots the average number of seconds per correct quiz answer.

(ii) the opportunity cost of time spent on the contract task, measured by the slope of correct answers over time spent on the contract task.

The first measurement of quiz performance is the hit rate, calculated as the average proportion of correctly answered quiz questions relative to the total number of questions attempted. Figure 9, panel (a), shows the hit rate for both stimulus and non-stimulus periods. The results indicate no significant difference in the hit rate between stimulus and non-stimulus periods. On average, participants maintain a consistent level of accuracy across all periods, suggesting that their ability to correctly answer quiz questions does not vary based on the presence of a stimulus.

The second metric examines the opportunity cost of attention allocated to the contract task. We measure this as the slope of the number of correctly answered quiz questions plotted against the time spent on the contract task. Figure 9, panel (b), displays the relationship between the time allocated to the contract task and the corresponding reduction in correctly answered quiz questions. The results indicate a higher opportunity cost during stimulus periods compared to non-stimulus periods, meaning that each additional second spent on the contract task is associated with a greater reduction in the number of correct quiz answers in stimulus periods.

Result 4 *Participants exhibit a gradual improvement in quiz speed, reducing the time spent per attempt, while significantly enhancing quiz efficiency over time.*

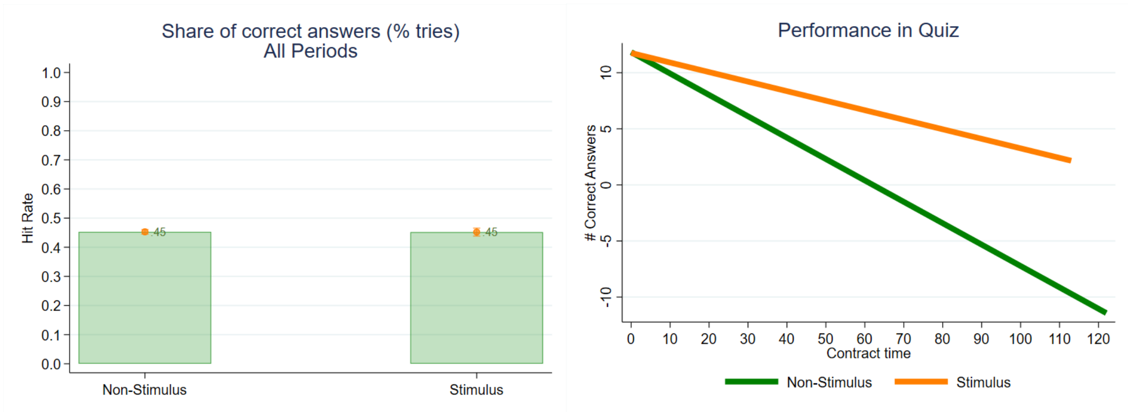

(a) Hit Rate (full sample)

(b) Number of Correct Answers and Contract Time

Figure 9: Performance in Quiz

Note: The hit rate is computed as the number of correct answers in comparison to the number of quiz tries (i.e., share of correct answers). Figure (a) plots the average hit rate between non-stimulus and stimulus periods. Figure (b) plots the predicted nubmer of correct answers (fitted values) by contract time in the non-stimulus and stimulus periods.

However, the accuracy of quiz responses remains stable across both stimulus and non-stimulus periods, indicating that cognitive performance in the quiz task is not influenced by external stimuli.

5.3.3 Do Higher Stakes Change Attention and Choice?

A key theoretical prediction is that, under a higher-stakes incentive structure of the quiz, we should observe a shift in decision quality along the attention frontier – with subjects reallocating cognitive effort towards higher quiz performance at the expense of contract decision quality. Table 1 reports the estimated effects of the high-stakes treatment on key outcome variables. While the high-stakes incentive marginally reduces contract time by 1.8 seconds, it has no statistically significant effect on expected excess costs (EEC), quiz efficiency, or quiz speed. This suggests that subjects are not substantially reallocating attention in response to the higher quiz incentives.

This presents a potential puzzle. If subjects respond optimally to incentives, why do we not observe the expected shift in behavior? Does this contradict rationality in their decision at a fundamental level? To address this central question, we

Table 1: The Effects of Stakes

	Contract Time (1)	EEC (2)	Quiz Efficiency (3)	Quiz Speed (4)
High Stake	-1.839**	-0.093	0.382	0.275
	(0.932)	(0.772)	(1.309)	(0.411)
Period	-1.847***	-3.834***	-2.059***	-0.943***
	(0.217)	(0.341)	(0.245)	(0.072)
Period^2	0.065***	0.171***	0.080***	0.037***
	(0.011)	(0.016)	(0.012)	(0.003)
Constant	19.998***	21.945***	24.697***	10.590***
	(1.212)	(1.650)	(1.490)	(0.438)
Observations	2192	2192	2179	2188

Notes: Random-effect regressions. High Stake is a dummy variable equal to one if in the high stake treatment. Contract Time is the time spent on the contract task; EEC = Expected Excess Cost; Quiz Efficiency = time spent per correct quiz answer; Quiz Speed = time spent per quiz attempt. Robust standard errors in parenthesis are clustered at the individual level. Significance levels are indicated as * 0.10, ** 0.05, *** 0.01.

assess the effective opportunity cost of improving quiz performance in the low-stakes treatment and compare it to the potential gains from shifting effort toward the quiz task in the high-stakes treatment. Specifically, we estimate the expected opportunity cost of answering one additional quiz question. Using a random-effects model, we then predict the expected loss in contract quality (measured in excess costs) when reducing contract time to accommodate additional quiz attempts.

We first estimate the time productivity in the contract task under the low-stakes treatment. This allows us to predict the expected earnings from the contract if the time typically required to answer correctly a quiz question – 15 seconds on average – was instead spent on the contract task. We use a random-effect approach to estimate the following regression model

$$y_{it} = \beta_0 + \alpha t_{it} + \gamma T + \beta t_{it} \times T + \delta_t + \delta_t^2 + \epsilon_{it}, \tag{2}$$

where the dependent variable y_{it} is the excess cost of subject i's selected contract compared to the cost of the best contract in Period t. This is calculated as the difference between the expected costs of the selected contract and the best contract in the period. t_{it} denotes the contract time duration of subject i in Period t. T is the indicator variable, equal to 1 in the stimulus periods and 0 in the non-stimulus periods. δ_t is the period fixed effect, which is included to control for the learning effect. β captures the average effect of a stimulus on contract costs as the contract duration increases by 1 second. Using the above estimation, we then predict the average excess cost of selected contracts as contract time increases from 1 second to 20 seconds[13].

Figure 10 plots the estimated relationship between contract time and excess costs in both stimulus and non-stimulus periods. Looking at the stimulus periods (red line), we observe a steep increase in excess costs when contract time is shortened. Specifically, reducing contract time by 15 seconds (the average time required to answer a quiz question) increases expected excess costs from 32 to 132 points. This means that the true average opportunity cost of answering an additional quiz question is *twice* the value of the high-stakes incentive.

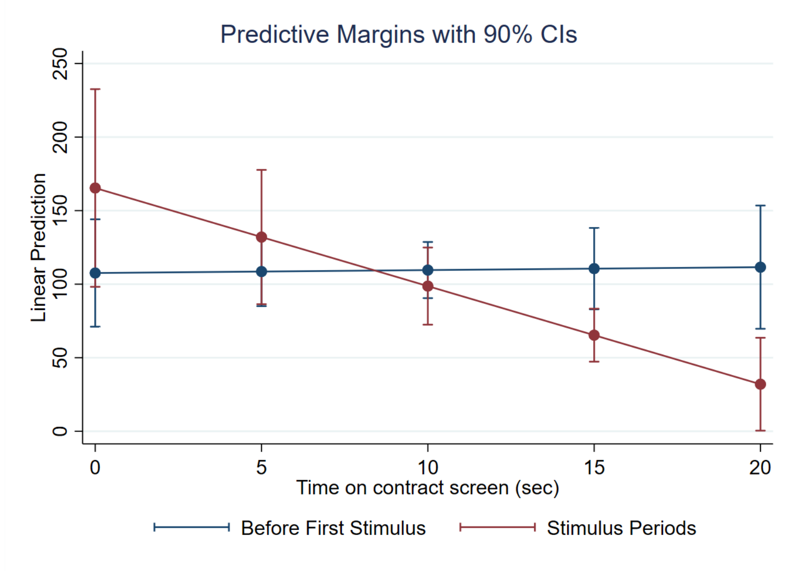

Figure 10: Contract Productivity

Thus, rather than contradicting rational decision-making, the absence of a shift in behavior is actually consistent with optimal behavior. The theoretical prediction of a behavioral shift depends on the assumption that the change in incentives is

[13]We analyze the observations where contract time duration ranges between 0 and 20 seconds to minimize estimation bias caused by outliers. In addition, we use only observations from the low-stakes treatment to establish the productivity baseline for the high-stakes treatment.

sufficiently strong to justify reallocating attention. However, whether the incentive change is “strong” in this sense is ultimately an empirical question – one that could not have been known in advance. Our findings reveal that, in this experimental setting, the incentive change was not large enough to make shifting attention worthwhile.

Result 5 *Despite higher quiz incentives, we find no significant reallocation of attention or improvement in quiz performance. Our analysis reveals that the opportunity cost of shifting effort is too high, making the absence of behavioral change nevertheless consistent with optimal choice.*

5.4 Cognitive Types

Going forward, we classify participants into ”cognitive types” to further characterize the interplay between choice behavior and the stimulus. We apply the following rough classification:

1. Hyper-rational: Subjects who have chosen the best contract at least twice in Periods 1-3 and have chosen the best contract at least 12 times in Periods 4-16. At the same time, they do not respond to the first stimulus period significantly.

2. Rational: Subjects who respond to the first stimulus period significantly.

3. Non-stable: Subjects who do not respond to the first stimulus period significantly and only sometimes select the best choice either before or after stimulus periods.

The time path of their likelihood of choosing the best contract is displayed in Figure 11.

What are the differences in choosing the best contract between different types? Math ability and the quiz strategy clicker are significantly correlated with the hyper-rational type. In total, there are 35 hyper-rational type subjects, and 102 are either rational or non-stable type. Among the 35 hyper-rational subjects, 24 are the math high type (68.6%) which is higher than the percentage of high math type in the

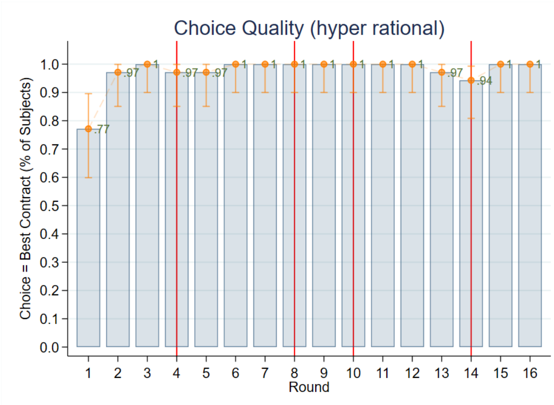

(a) Hyper-rational

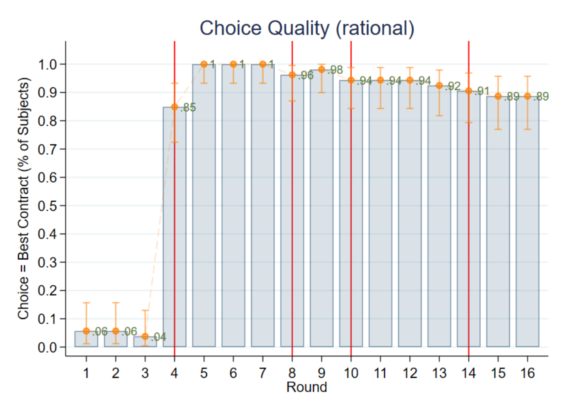

(b) Rational

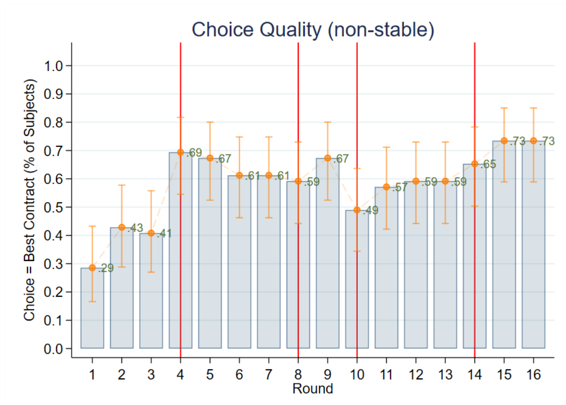

(c) Non-stable

Figure 11: Probability of Best contract choice by rational types

Note: Figure (a) plots the probability of hyper-rational subjects choosing the best contract over periods. Figure (b) plots the probability of rational subjects choosing the best contract over periods. Figure (c) plots the probability of non-stable subjects choosing the best contract over periods.

other types (i.e., 43%). Similarly, 11 out of the 35 (i.e., 31.4%) hyper-rational type are clickers, while only 16 out of the 102 (15.7%) other types are clickers. Although hyper-rational are more likely to be a high math type or a clicker, the coefficient of "Clicker × Post Period4" and "Above-median Math × Post Period4" indicate no significant difference between the high and low math type, and between the clicker and thinker in terms of the quality of their contract choices. We summarize our empirical results as follows:

Result 6 *Hyper-rational individuals, a minority of the sample, exhibit negligible response to the first stimulus, aligning with rational inattention theory. In contrast,*

Table 2: Cognitive Types

	Contract Hard (1)	Quiz Hard (2)	Gender (3)	**Clicker** (4)	Above-median Age (5)	Above-median Education (6)	**Above-median Math Score** (7)	Above-median Math Time (8)
Hyper-rational	-0.132*	0.017	-0.212**	0.157*	-0.097	0.036	0.254***	-0.014
	(0.079)	(0.087)	(0.100)	(0.087)	(0.093)	(0.084)	(0.093)	(0.099)
Rational	-0.040	-0.206**	-0.073	-0.075	-0.139	-0.199**	-0.102	0.052
	(0.078)	(0.080)	(0.088)	(0.068)	(0.084)	(0.079)	(0.088)	(0.088)
Non-stable	0.151*	0.198***	0.247***	-0.053	0.224**	0.175**	-0.106	-0.042
	(0.082)	(0.072)	(0.082)	(0.069)	(0.087)	(0.071)	(0.089)	(0.090)

Notes: Each cell reports the coefficient from a separate OLS regression. The dependent variables are: a dummy variable equal to one if the subject's perception of the contract difficulty is above the sample's median (column 1, Contract Hard), a dummy variable equal to one if the subject's perception of the quiz difficulty is above the sample's median (column 2, Quiz Hard), gender equal to one if female (column 3), clicker equal to one if the subject has more than three rounds where his hit rate is below 0.3 and his number quiz tries is above median (column 4), a dummy variable equal to one if the subject's age is above the sample's median (column 5), a dummy variable equal to one if the subject's education level is above the sample's median (column 6), a dummy variable equal to one if the subject's math score is above the sample's median (column 7), a dummy variable equal to one if the subject's math time is above the sample's median (column 8). Robust standard errors in parenthesis are clustered at the individual level. Significance levels are indicated as * 0.10, ** 0.05, *** 0.01.

Table 3: Contract quality and Subject Types I

	Best Contract		Extra Cost		Best Contract		Extra Cost	
	Full (1)	Period $\geq$ 4 (2)	Full (3)	Period $\geq$ 4 (4)	Full (5)	Period $\geq$ 4 (6)	Full (7)	Period $\geq$ 4 (8)
Clicker	0.416**	0.413**	-2.005**	-1.448***				
	(0.173)	(0.204)	(0.844)	(0.551)				
Math Score					0.533***	0.537***	-2.608***	-1.889***
Above-median					(0.130)	(0.159)	(0.718)	(0.642)
Constant	-0.551***	0.867***	20.342***	3.592***	-0.747***	0.702***	21.242***	4.244***
	(0.115)	(0.132)	(1.445)	(0.715)	(0.126)	(0.145)	(1.441)	(0.801)
Observations	2192	1781	2192	1781	2192	1781	2192	1781
(Pseudo) R^2	0.1513	0.0163	0.304	0.014	0.1699	0.0392	0.313	0.029
cmd	probit	probit	OLS	OLS	probit	probit	OLS	OLS

Notes: All regressions control for period. Robust standard errors in parentheses are clustered at the individual level. Significance levels: * 0.10, ** 0.05, *** 0.01.

rational participants respond significantly to the stimulus, while non-stable individuals display inconsistent decision patterns. Furthermore, we detect no significant variation in contract choice quality between high and low math ability participants or between strategic and non-strategic quiz takers.

6 Conclusion

This study examines the interaction between attention allocation and decision-making under cognitive constraints. Our findings provide evidence that sudden, non-informative perturbations in the decision environment trigger heightened cognitive engagement, leading to persistent improvements in decision performance. These results highlight the role of externally aroused cognitive processing shifts as

Table 4: Contract quality and Subject Types II

	Best Contract (1)	Extra Cost (2)	Best Contract (3)	Extra Cost (4)
Post Period 4	1.576***	-18.879***	1.607***	-20.199***
	(0.184)	(1.667)	(0.219)	(1.957)
Clicker	0.423*	-4.420		
	(0.238)	(3.153)		
Clicker × Post Period4	-0.009	2.972		
	(0.279)	(3.030)		
Above-median Math			0.522***	-5.727**
			(0.200)	(2.534)
Above-median Math× Post Period4			0.015	3.838
			(0.252)	(2.592)
Constant	-0.552***	20.818***	-0.741***	22.790***
	(0.125)	(1.594)	(0.153)	(1.836)
Observations	2192	2192	2192	2192
(Pseudo) R^2	0.151	0.306	0.170	0.318

Notes: Robust standard errors in parenthesis are clustered at the individual level. Significance levels are indicated as * 0.10, ** 0.05, *** 0.01.

a complement to endogenous attention allocation, challenging the assumption that individuals always regulate cognitive effort optimally.

Our experiment isolates this activation effect by introducing salient external changes in a contract-choice task. Participants exhibited a marked improvement in decision quality following the first such stimulus period, reflected in lower excess costs and more frequent selection of the optimal contract. Crucially, these effects arose without additional informational content or changes in cognitive difficulty, underscoring that the mechanism at play is an externally induced enhancement of cognitive processing, rather than an adjustment in information acquisition incentives or costs. These results challenge the tenet that individuals are consistently capable of self-regulating their attention and decision effort as propagated by the rational inattention framework. Instead, they point to the critical interplay between internal cognitive strategies and external activating factors as a fundamental component of decision-making.

The implications extend beyond the laboratory. In real-world decision contexts, like insurance marketplaces or investment platforms, designing interventions that incorporate attention-activating stimuli could enhance decision quality without overwhelming individuals with excessive information nor coercing them into non-autonomous ways of thinking and choosing.

By providing empirical evidence on the interplay between endogenous and exogenous attention mechanisms, this study contributes to a more comprehensive understanding of attention and choice. It underscores the need for economic models to account for externally induced attention shifts, which have a neurophysiological foundation and likely play a crucial role in shaping decision efficiency in information-rich environments. As modern economies continue to grapple with information overload, our findings offer a new directions for designing decision-support systems and policy interventions that enhance cognitive engagement without coercion, paternalism or information manipulation.

References

Arellano, M. and Bond, S. (1991). Some tests of specification for panel data: Monte carlo evidence and an application to employment equations. *Review of Economic Studies*, 58(2):277–297.

Bai, J. and Perron, P. (2003). Computation and analysis of multiple structural change models. *Journal of Applied Econometrics*, 18(1):1–22.

Caplin, A. and Dean, M. (2011). Search and satisficing. *American Economic Review*, 101(7):2899–2922.

Caplin, A. and Dean, M. (2015). Revealed preference, rational inattention, and costly information acquisition. *American Economic Review*, 105(7):2183–2203.

Caplin, A. and Martin, D. (2015). Testable implications of models of intertemporal choice: Exponential discounting and its generalizations. *Econometrica*, 83(5):1849–1890.

Carretié, L. (2014). Exogenous (automatic) attention to emotional stimuli: a review. *Cognitive, Affective, and Behavioral Neuroscience*, 14(4):1228–1258.

Chetty, R., Looney, A., and Kroft, K. (2009). Salience and taxation: Theory and evidence. *American Economic Review*, 99(4):1145–1177.

Corbetta, M. and Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, 3(3):201–215.

Dean, M. and Neligh, N. (2017). Experimental tests of rational inattention models. *Journal of Economic Theory*, 169:67–100.

Enke, B. and Graeber, T. (2023). Cognitive uncertainty. *The Quarterly Journal of Economics*, 138(4):2021–2067.

Eyherabide, F. and Samengo, I. (2013). When and why noise correlations are important in neural decoding. *The Journal of Neuroscience*, 33(45):17921–17936.

Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental economics*, 10(2):171–178.

Gabaix, X. (2019). Behavioral inattention. *Handbook of Behavioral Economics: Applications and Foundations 1*, 2:261–343.

Gabaix, X. and Laibson, D. (2003). Allocation of attention: Theory and evidence. *Journal of Economic Perspectives*, 17(3):1–23.

Geng, S. (2016). Decision time, consideration time, and status quo bias. *Economic Inquiry*, 54(1):433–449.

Gold, J. I. and Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, 30:535–574.

Heath, C. and Tversky, A. (1991). Preference and belief: Ambiguity and competence in choice under uncertainty. *Journal of Risk and Uncertainty*, 4(1):5–28.

Hossain, T. and Morgan, J. (2011). Shrouded attributes and information suppression: Evidence from field experiments. *Quarterly Journal of Economics*, 126(3):1403–1440.

Ito, T., Brincat, S. L., Siegel, M., Mill, R. D., He, B. J., Miller, E. K., Rotstein, H. G., and Cole, M. W. (2020). Task-evoked activity quenches neural correlations and variability across cortical areas. *PLOS Computational Biology*, 16(8):e1007983.

Kahneman, D. (1973). *Attention and Effort.* Prentice-Hall, Englewood Cliffs, NJ.

Kahneman, D. (2003). Maps of bounded rationality: Psychology for behavioral economics. *American Economic Review*, 93(5):1449–1475.

Kahneman, D. and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–292.

Lennie, P. (2003). The cost of cortical computation. *Current biology*, 13(6):493–497.

Matějka, F. and McKay, A. (2015). Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review*, 105(1):272–98.

McCall, J. J. (1970). Economics of information and job search. *The Quarterly Journal of Economics*, 84(1):113–126.

Moujahid, A., D'Anjou, A., and Graña, M. (2014). Energy demands of diverse spiking cells from the neocortex, hippocampus, and thalamus. *Frontiers in Computational Neuroscience*, 8:41.

Nagel, R. (1995). Unraveling in guessing games: An experimental study. *American Economic Review*, 85(5):1313–1326.

Petersen, S. E. and Posner, M. I. (2012). The attention system of the human brain: 20 years after. *Annual Review of Neuroscience*, 35(1):73–89.

Ratcliff, R. and McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural Computation*, 20(4):873–922.

Reutskaja, E., Nagel, R., Camerer, C. F., and Rangel, A. (2011a). Search dynamics in consumer choice under time pressure: An eye-tracking study. *American Economic Review*, 101(2):900–926.

Reutskaja, E., Nagel, R., Camerer, C. F., and Rangel, A. (2011b). Search dynamics in consumer choice under time pressure: An eye-tracking study. *American Economic Review*, 101(2):900–926.

Rischka, L., Furtner, M., Härtl, W., Hahn, M., Müller, F. M., Windischberger, C., Schmidt, M. H., Lanzenberger, R., Stocker, G., Ernst, T., Rupprecht, R., Wadsak, W., Hacker, M., and Rami-Mark, C. (2021). Reliability of task-specific neuronal activation assessed with functional pet, asl, and bold imaging. *Journal of Cerebral Blood Flow & Metabolism*, 41(10):2734–2749.

Roth, A. E. and Erev, I. (1998). Predicting how people play games: Reinforcement learning in experimental games. *American Economic Review*, 88(4):848–881.

Rotstein, H. G. (2013). Abrupt and gradual transitions between low and hyper-excited firing frequencies in neuronal models with fast synaptic excitation: A comparative study. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 23(4):046104.

Shadlen, M. N. and Newsome, W. T. (1998). The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *Journal of Neuroscience*, 18:3870–3896.

Shadlen, M. N. and Newsome, W. T. (2001). Neural basis of a perceptual decision in the parietal cortex (area lip) of the rhesus monkey. *Journal of Neurophysiology*, 86(4):1916–1936.

Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665–690.

Sokolov, E. N. (1963). *Perception and the Conditioned Reflex*. Pergamon Press, Oxford, UK. Translated by Stefan W. Waydenfeld.

Vardi, R., Goldental, A., Guberman, S., Kalmanovich, A., Marmari, H., and Kanter, I. (2013). Sudden synchrony leaps accompanied by frequency multiplications in neuronal activity. *Frontiers in Neural Circuits*, 7:176.