Attention and Choices with Multiple States and Actions: A Laboratory Experiment

Andrea Civelli[*] Cary Deck[†] Antonella Tutino[‡]

Abstract

We study how a rationally inattentive decision maker chooses state-contingent actions under uncertainty in complex environments. We explore a series of decision problems by varying the number of states as well as incentive structures. We fully characterize the theoretical solutions and compare them to choices made by subjects facing those problems in a controlled laboratory experiment. Observed behavior is broadly consistent with the theoretical model, with subjects responding to changes in complexity and incentives by varying their level of attention. Nevertheless, some interesting differences emerge from the experimental data. In particular, we find only moderate support for the invariance under compression property, that some aspects of subject behavior may require explanations such as perceptual distance, and that complexity can affect the ratio of expected utility to information gains.

Keywords: Rational Inattention, Laboratory Experiment, Information Processing Capacity

JEL Classification: C91, D11, D8, E20.

[*]University of Arkansas, Walton College of Business, Department of Economics, Business Building 402, Fayetteville, AR 72701. Email: andrea.civelli@gmail.com

[†]University of Alabama and Chapman University. Email: cdeck@cba.ua.edu

[‡]Federal Reserve Bank of Dallas. Email: tutino.antonella@gmail.com

1 Introduction

Understanding the mapping between information acquisition and decisions in rich economic environments, where multiple states expand the scope of choices faced by the decision maker (DM henceforth), is pivotal in economics. Of particular importance is the role that incentives play in the trade off between the attention required to make informed decisions and the benefits of making those choices. For instance, in a monetary policy setting a person may opt to acquire information about inflation at a different level of precision if that inflation is persistent rather than transitory and as a result may make different investment choices. However, there has been little direct analysis of how people actually choose state-contingent actions under uncertainty in complex environments

Starting from Sims (2003, 2006), a tenet of rational inattention (RI henceforth) theory is that observed behavioral choices are rationalized through the joint consideration of utility and costs of information. RI measures the uncertainty in a system through an information cost based on Shannon's entropy, while allowing subjective preferences and perceived incentives to shape a DM's decisions on the optimal information structure. Although RI theory lays out the foundations for investigating how incentives affect information structure and, ultimately, behavioral choices, the unobservability of preferences and subjective costs of information structures make it difficult to test the theoretical predictions with naturally occurring data. As a result, tests of RI have relied heavily on laboratory experiments (Cheremukhin *et al.*, 2015; Caplin and Dean, 2015; Caplin and Martin, 2014; Dean and Neligh, 2019; Dewan and Neligh, 2020; Caplin *et al.*, 2021). The vast majority of this literature, however, has focused on estimating the cost of information and its consistency with Shannon's structure. By contrast, very few papers consider experiments in which a subject faces more than a binary state-action problem and even fewer explicitly study problems with multiple, heterogenous states and actions.[1]

The goal of this paper is to study attention and choices in richer economic environments

[1] For instance, Dean and Neligh (2019) consider multiple states that can be reduced to a binary state by the Invariance Under Compression (IUC) property. The approaches based on relative thinking of Bushong *et al.* (2020) and on focusing of Kőszegi and Szeidl (2013) derive the effects of incentives with varying ranges of available options, which are forms of the extensive margin of complexity. Different from these models, we study the intensive margin of complexity where the number of options are fixed and the states vary within that range.

with more than two states and two actions. Specifically, we build on the approach of Matějka and McKay (2015), Caplin and Dean (2015) and Caplin and Martin (2014) to model discrete choices under RI, but we extend the state and action spaces to accommodate more complex decision environments. The higher complexity of the decision environment is defined in our model by the enlargement of the state space and an increase in the number of available choices for the DM. We then test the theoretical predictions of the model with a laboratory experiment which closely matches this theoretical framework.

In our setting, before taking an action, the DM can choose signals to reduce uncertainty about the realized state of the world. However, the information gain provided by the signal is costly: the signal requires cognitive effort to be processed, and that effort is commensurate with the complexity of the system. The cognitive cost is measured by Shannon's mutual information. The DM chooses the information structure, a mapping from signal to state and from state to actions, that maximizes her expected utility net of the cost associated with the informativeness of the signal about the states. The solution to the RI model is fully characterized by the stochastic choice functions which jointly identify the probability of choosing each action in each state and represent the two mappings corresponding to the optimal information structure.

Starting from this general theoretical setup we develop six decision problems by varying the incentive structure and the complexity of the underlying economic environment. We devise the following approach to model higher complexity. There are N states of the world which are exogenously ranked by nature. The number of possible rankings is $N!$. The DM must indicate which state of the world she believes is first in the ranking and receives a payoff based on the actual ranking of the selected state. Our framework allows for a rich interaction between payoffs and attention choices. We assume the payoff structure is monotonically decreasing in the rank position of the chosen state (i.e., the payoff from selecting the second ranked state is no more than the payoff from selecting the top ranked state). Depending on the payoffs, the DM has different incentives to explore more or less precisely the rankings of the states.

This approach directly accommodates a 2-state and 2-action environment similar to Caplin and Dean (2015), which we label 2s2a for simplicity. In this world only two rankings are possible, and knowledge of the first ranked state trivially defines the entire ranking. We study two decision problems in this basic setup by varying the slope of the payoff gains associated with correctly finding the first ranked state from a case which strongly rewards an accurate differentiation of states (the baseline case) to a problem with very close payoffs.

We then expand the model to a 4-state and 4-action world, labeled as 4s4a, where decisions are based on the 24 possible rankings of the four states. We study four decision problems in this second environment. Two of these problems correspond to higher-dimension versions of the 2s2a problems in which the payoffs of the first and second ranked states are equal to that of a first ranked state in the 2s2a environment, while the payoffs of the third and fourth ranked states are equal to the second ranked state in the 2s2a environment. These problems allow us to directly assess the effects of complexity on attention. The other two problems in the 4s4a environment change the reward structure of the baseline case by either increasing the payoff of the first ranked state or decreasing the payoff of the second ranked state. These problems are used to further examine how attention responds to incentives in a relatively more complex setting.

Comparing the experimental data to the RI theoretical results, we find that our behavioral evidence is generally consistent with the RI model; however, interesting differences between the RI DM and the experimental subjects also emerge. In particular, we can draw four main conclusions.

First, the empirical results moderately support the Invariance Under Compression property of the RI model (Caplin *et al.*, 2021). The theoretical results show that adding states with similar payoffs to the baseline 2s2a problem does not modify the optimal behavior of the RI DM, who continues partitioning the space into paying and non-paying states. However, experimental subjects in the 4s4a environment also separate paying from non-paying actions, as predicted by the theory, but they over-select the first ranked action.

Second, and related to the first conclusion, we find that subjects' decisions respond to changes in incentives in a way consistent with the theory. Subjects gather instrumental information in the 4s4a environment as well, while remaining relatively uninformed about outcomes with smaller payoffs. However, the subjects in the experiment pay excessive attention to the first two ranked actions when payoffs across states are close.

The third result we discuss is related to the perceptual distance theory (Hébert and Woodford, 2020). Two of our decision problems are designed to produce steeper posterior distributions of actions with different implications for the probability of choosing the first action relative to the second one. The subjects in the experiment, however, basically adopt the same flatter than predicted distribution in these two problems, contradicting the RI model. Perceptual distance provides a possible explanation of this result. The similarity between the first two states in the experiment reduces the impact of payoff differentials on the subjects' decisions, as in a 2s2a world, but it does not make the role of the second ranked action completely irrelevant in the more complex environment.

The last result is provided by the comparison of expected utility and information gains across decision problems. With linear utility and mutual information-based costs, the RI model implies a ratio of expected utility to information gains that is approximately constant. The experimental subjects follow a similar rule, with exception of the treatment with the steepest payoff structure, in which they gather less information and obtain a smaller utility gain than predicted by the RI theory. We link this behavior to a large perceived cost of information subjects face to process a more precise signal about the second ranked action.

Our paper contributes to the RI literature in several ways.[2] First, our discrete choice model draws from those of Caplin and Dean (2015) and Caplin and Martin (2014), which resemble Caplin and Dean (2013), Matějka and McKay (2015) and Dean and Neligh (2019). We complement these papers by fully characterizing the solution for each of the decision problems considered. These solutions, expressed as conditional and unconditional optimal posteriors, give sharp theoretical predictions that can be directly compared with the observed empirical frequencies found in our experimental data.

Second, our model also builds upon Caplin and Dean (2015)'s posterior-based approach and the stochastic choice functions for posterior-separable costs of Caplin and Martin (2014) and Caplin *et al.* (2019). These papers show that separable cost functions can rationalize some aspects of state complexity via consideration sets as in Caplin *et al.* (2019), the neighborhood approach as in Hébert and Woodford (2020), and latent variables as in Csaba (2018). These approaches rationalize the concept of state complexity by identifying sets in

[2]A detailed review of the rational inattention literature can be found in Maćkowiak *et al.* (2022).

which differentiating among alternatives is harder the closer the alternatives are. Consistent with previous result, we show that the closeness of the states in terms of payoff differential is key in eliciting changes in attention and, in turn, behavioral reactions finely attuned to the states.

Third, our work also relates to the experimental literature explicitly testing RI to connect how endogenous information acquisition and processing affect expectations about payoffs in complex systems. Examples of these literatures include Cheremukhin *et al.* (2015), Caplin and Dean (2015), Caplin and Martin (2014), Dean and Neligh (2019), Dewan and Neligh (2020), Caplin *et al.* (2021), Duffy and Puzzello (2021), Kryvtsov and Petersen (2021). We complement their contribution by building on Caplin and Dean (2015) and expand their framework to incorporate richer environments. The main way our work differs from this literature is our investigation of the effects of incentives in a setting where multiple states and actions are present.

The remainder of the paper is organized as follows. Section 2 formalizes the theoretical framework, describes the probabilistic environment, and provides numerical solutions for specific decision problems. Section 3 details the experimental design, while the experimental results are presented in Section 4. Our setting applies to a variety of economic applications in which a DM needs to distinguish and rank different states to maximize her utility. As such, Section 5 offers examples of the applicability of our framework to diverse economic environments ranging from the interaction of fiscal and monetary policy to voting behaviors. A final section concludes.

2 Theoretical Framework

To provide a theoretical foundation to evaluate RI, we adapt the model to the tasks we employ in our experiment. The framework builds on those of Caplin and Dean (2015) and Caplin and Martin (2014) for discrete choices. We postulate that the information processing-cost takes on Shannon's functional form as in Matějka and McKay (2015). We extend the state and task spaces to accommodate richer economic environments expanding the scope of the DM's choices. In this section, we present both the model and the theoretical predictions on an abstract level while we discuss the applicability of the framework to different economic environments in Section 5.

2.1 The Probabilistic Environment

We consider a DM who chooses actions where the outcomes depend on the realization of the state. We postulate a state-action matching utility in which the DM's reward is linked to her ability to correctly identify and rank the states from the most to the least likely. Higher ranked states are associated with higher payoffs.

Specifically, we consider a finite set of states $\omega \in \Omega$. The set of outcomes is defined as Z and actions-induced space is represented by Z^{Ω}. For each decision problem, the set of actions is defined as $a \in \mathcal{A}$ taking values in the finite space $\mathcal{A} \subset Z^{\Omega}$. For concreteness, we assume that $\Omega = \{\omega_1, \omega_2, \ldots, \omega_N\}$ describes the N possible realizations of the state. The set Z is the ranking of the states ω's, taking on $N!$ possible values. The set, Z, contains all the rankings that the DM judges worth considering. For a given decision problem, the DM's action a boils down to a particular ranking of the states among the alternatives in $\mathcal{A}$.

The DM's preferences are described by a utility function over outcomes, $u : \mathcal{A} \times \Omega \to \mathbb{R}$. That is, the utility function matches the action corresponding to a particular choice of the ranking of the states to a payoff that reflects the true ranking of the states.

The DM has beliefs over the rankings of the states. The set of all possible beliefs over rankings is defined as $\Gamma := \Delta(\Omega)$. Before processing any information, the DM has a uniform prior belief over rankings defined as $\mu \in \Delta(\Omega)$.

The DM is assumed to be a Bayesian expected utility maximizer. Let u denote the utility associated with each of the ranked states. For a given belief $\gamma \in \Gamma$ and set of actions $\mathcal{A}$, let the best response function be:

$$\phi(\gamma, \mathcal{A}) := \arg\max_{a \in \mathcal{A}} < u_a, \gamma >, \tag{1}$$

with $u_a = u \circ a$. The function in (1) encodes the match between state-actions and payoffs by rewarding the actions according to how they relate to the true ranking of the states.

The associated value function is given by:

$$V(\gamma, \mathcal{A}) := \max_{a \in \mathcal{A}} < u_a, \gamma >, \tag{2}$$

which selects the payoff-maximizing action among alternatives.

Before taking an action, the DM can acquire information about the state through signals $S \subset \mathcal{S}$, where $\mathcal{S}$ is a generic signal space. Signal choices correspond to an information structure that is a Markov kernel, $P : \Omega \to \Delta(S)$ defining a collection of probability distributions over signals $s \in S$ for each realization of the state $\omega \in \Omega$. We denote these conditional distributions as $P(.|\omega) = P_\omega$. Signals can be interpreted as information that the DM decides to acquire and process to better assess the true ranking of the states before selecting a particular ranking. The link between information about rankings (the signals) and true ranking of the states is encoded in the conditional probabilities: signals more attuned to the true ranking have more concentrated conditional distributions.

Let $\mathcal{P}(\Omega)$ be the space of all information structures defined over the state space Ω with finite signal supports. Given a prior distribution over states, the unconditional distribution of signal s under prior μ is:

$$P_\mu(s) = \sum_{\omega} \mu(\omega) P_\omega(s). \tag{3}$$

The corresponding posterior distribution given signal $s \in S$ is:

$$\gamma^s(\omega) := \frac{\mu(\omega) P_\omega(s)}{P_\mu(s)} \tag{4}$$

Together with the prior, the information structure defines a distribution over posteriors. The support of the information structure is defined as $\mathcal{S}(P) := supp(P_\mu)$. Once the signal has been chosen, the DM conditions her actions only on it as a proxy for the true state.[3]

The restriction on signals the DM can potentially acquire is defined through a cost function over information structures. This is the core of RI theory: the cognitive cost of

[3]Following Caplin and Dean (2015) and Caplin and Martin (2014), this implies that redundant signals are disregarded.

processing information prevents the DM from acquiring precise signals about the state. In turn, this assumption implies that the DM's information structure must take into account the trade off between the informativeness of signals and the cognitive effort necessary to process them. As in Matějka and McKay (2015), this cost is based on Shannon's mutual information, an information-theoretic concept that relates informativeness to the joint distribution of states and signals.

Formally, the information cost function is a mapping from information structures and priors, $G : \mathcal{P}(\Omega) \times \Delta(\Omega) \to \mathbb{R}$. We follow the approach of Matějka and McKay (2015) and Caplin and Dean (2015) and assume a posterior separable cost function based on Shannon's mutual information of the form:

$$G(P, \mu) = \sum_{s \in \mathcal{S}(P)} P_\mu(s) \left(\sum_\omega \gamma^s(\omega) \log \frac{\gamma^s(\omega)}{\mu(\omega)} \right) \tag{5}$$

The problem of the rationally inattentive DM is to choose an information structure that maximizes the ex-ante expected utility net of the information cost:

$$\max_{P \in \mathcal{P}(\omega)} \left\{ \sum_{\mathcal{S}(P)} P_\mu(s) V(\gamma^s, \mathcal{A}) - \kappa G(P, \mu) \right\}. \tag{6}$$

where κ is the marginal cost of processing information. Since the problem (6) is strictly concave with strictly convex cost (5), it has a unique solution, P^*, corresponding to the optimal attention strategy. Given the convexity of (5), the support of P^* is constrained by the cardinality of the action set, $|\mathcal{A}|$. This implies that we can take the signal space to be of the same cardinality as the action space and identify each signal as suggestive of a specific action.[4]

In this theoretical framework, we can fully characterize the behavior of the DM in all potential choice sets with a stochastic choice function,i.e. a mapping $\rho : M^\Omega \times Z \to [0, 1]$ such that for $\mathcal{A} \in Z$, $\rho(a, (A)) > 0$ implies $a \in \mathcal{A}$ and $\sum_{a \in \mathcal{A}} \rho(a, \mathcal{A}) = 1$.[5]

We can characterize the stochastic choice functions of the RI model through the

[4]See footnote (3).

[5]Caplin and Martin (2014) and Caplin and Dean (2015) introduced the notion of state-dependent stochastic choice (SDSC) functions and provide detailed characterizations of these notions.

induced distribution over posteriors. Thus, the conditional stochastic choices function for all $\omega \in \Omega$ are given by:

$$\rho_\omega(a, \mathcal{A}) = P^*_\omega\{s \in \mathcal{S}(P^*) : a = \phi(\gamma^s, \mathcal{A})\} \tag{7}$$

and the unconditional stochastic choice functions is the prior weighted average of the unconditional choice probabilities:

$$\rho(a, \mathcal{A}) = P^*_\mu\{s \in \mathcal{S}(P^*) : a = \phi(\gamma^s, \mathcal{A})\} \tag{8}$$

To map the theoretical framework into a setting in which actions, but not signals, are directly observable Caplin and Dean (2015) and Caplin *et al.* (2021) recast the solution P^* in terms of the implied posterior beliefs $\gamma^a(\omega) := \frac{P^*(a|\omega)\mu(\omega)}{P^*(a)}$. This quantity is called the "revealed posterior" in the literature. Since the support of P^* has the same cardinality as $\mathcal{A}$, with each signal conducive to a distinct action, the difference between $\gamma^s(\omega)$ and $\gamma^a(\omega)$ is irrelevant for the theoretical framework. However, in such setting a researcher can only observe $\gamma^a(\omega)$ from the empirical distribution of actions and states.

2.2 Numerical Solution and Theoretical Predictions: Ideal Choice Environment

The optimal attention strategy P^* specifies a collection of stochastic choice functions, one for each possible realization of the state. Under the RI model, the mappings among signals, states and actions for the conditional posterior distribution in (7) and for the unconditional posterior distribution in (8) are fully traceable. The setting in which a researcher can directly map information choices via the observed actions, dubbed the *ideal choice environment* by Caplin and Martin (2014), allows us to derive explicit predictions that can be tested with experimental data.

We begin by replicating the results in 2s2a environment of Caplin and Dean (2015). Note that within this framework there is no difference between states and ranking of the states. Then, we augment the Caplin and Dean (2015) environment by considering decision problems in a 4s4a setting. Note that this framework is comprised of 24 possible rankings of the states, representing the information about the states acquired and processed by the DM. This extension allows us to study the impact of complexity on the DM's behavior in a context in which her payoff depends on how precisely she can identify the state(s) associated with the highest reward.

Overall, we consider six decision problems: two in the 2s2a environment ($\mathcal{DP}'_1$ and $\mathcal{DP}'_2$) and four in the 4s4a environment ($\mathcal{DP}_1$, $\mathcal{DP}_2$, $\mathcal{DP}_3$ and $\mathcal{DP}_4$).

The quantitative solution and theoretical predictions of the model depend on the specific value of the marginal cost parameter κ in (6).[6] This parameter represents the DM's intrinsic cognitive ability of processing information. Thus, it should remain invariant across decision problems. To provide a sharper comparison with the experimental data, we calibrate κ so that the optimal state-dependent stochastic choice functions in (7) and (8) are consistent with the empirical counterpart based on the representative experimental subject in our baseline treatment, corresponding to $\mathcal{DP}_1$. We then derive precise quantitative predictions under this calibrated value.[7] As shown in Figure 1, this exercise suggests that a value of $\kappa = 15$ is a reasonable approximation.

2.2.1 The 2s2a environment

There are 100 balls of two possible colors representing the states $\Omega = \{\omega_1, \omega_2\} = \{blue, red\}$. The possible rankings correspond to the set $Z = \{[blue,\ red], [red,\ blue]\}$. The DM chooses one of two actions $\mathcal{A} = \{a, b\}$, where a is the action corresponding to the belief that the state mostly likely ranked first is blue, and b has the opposite ranking. We assume that the DM has a uniform prior over the ranking of the colored balls. Let $\mathbf{p} = [p_1,\ p_2]$ be the payoff vector associated with the position of a color in the ranking. Note that this payoff vector induces a payoff matrix, u, over which the utility for each ranked state is defined, with dimensionality $|\Omega| \times |\mathcal{A}| = 2 \times 2$.

Given the primitives defining the environment, the conditional and unconditional

[6]The qualitative implications of the theoretical model we discuss would remain valid for most plausible value of κ.

[7]The calibration is obtained by maximum likelihood estimation given the observed data on the revealed posterior distribution of state-contingent actions in the baseline. Numerical predictions are generally robust for values of κ in the range [5, 25].

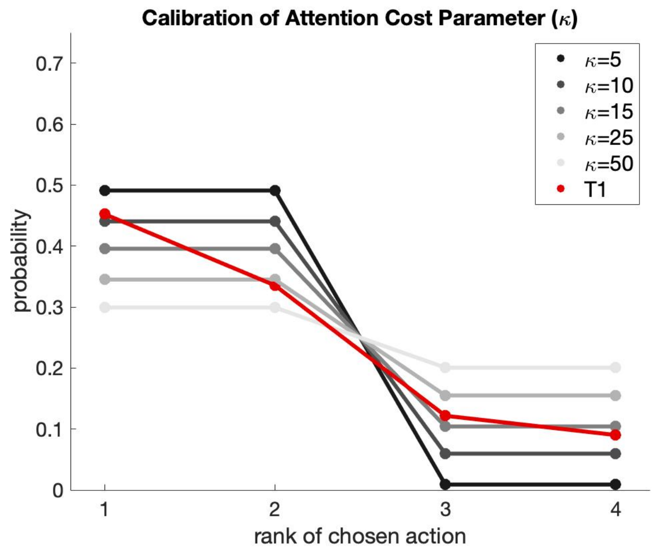

Figure 1: The "revealed posterior" from the experimental data in the baseline treatment (T1 in the notation introduced in Section 3) is plotted against the theoretical conditional distribution of the state contingent actions in $\mathcal{DP}_1$ for different calibrations of the marginal cost parameter κ. Although no perfect match of the experimental data can be obtained with any specific κ, $\kappa = 15$ provides a fair calibration of the model.

stochastic choice functions for the RI model are fully specified. In particular, for a value of the marginal cost of attention $\kappa = 0$ in equation (6), both conditional and unconditional choice probabilities are degenerate, selecting for each state the signals corresponding to actions with the highest payoff. This is the case of full information. As κ approaches infinity, conditional choice probabilities are degenerate and the unconditional choice probability is set as the prior. In this case the DM processes no information. The attention structure varies in a more meaningful way for intermediate values of κ, as the DM considers signals directly conducive to distinguishing states for which the correct actions provide the highest utility.

To make precise quantitative predictions, we characterize $\mathcal{DP}_1^{'}$ and $\mathcal{DP}_2^{'}$ by specifying the payoff vectors $\mathbf{p} = [20,\ 0]$ and $\mathbf{p} = [20,\ 19]$, respectively. $\mathcal{DP}_1^{'}$ is a symmetric problem with ex-ante utility $\mathbb{E}_\mu[u] = 10$ under a uniform prior over the states. The optimal distribution of actions given states derived by the solutions (7)-(8) conditional on *blue* being first ranked is P^*_{blue}= [.8, .2], and it is calculated symmetrically for *red*. The optimal expected value and cost function at the chosen posteriors are $V = 15.83$ and $G = 0.21$ bits. The top panel of Table 1 recapitulates the solution.

The problem $\mathcal{DP}_2^{'}$ allows us to study the effect of a change in the payoff differential on the DM's optimal choice of information structure in this environment. The solution to $\mathcal{DP}_2^{'}$ shows that the DM optimally decides to acquire essentially no information and randomizes across ranked choices, as reported in Table 1.

2.2.2 The 4s4a environment

We now allow for a richer economic environment. In particular, while keeping the number of balls at 100, we assume that the balls can take on four colors, corresponding to the four states $\Omega = \{blue, red, gray, green\}$. Having four states implies that there are $4! = 24$ possible rankings in the set Z. Before acquiring any information on the ranking, the DM has a uniform prior over which color is more common. After acquiring information on the ranking, the DM can choose one action in the expanded space $\mathcal{A} = \{a, b, c, d\}$, where a is the action corresponding to the belief that the state most likely ranked first is blue, b for red, and so on.

The payoff vector now comprises of four payments, $\mathbf{p} = [p_1,\ p_2,\ p_3,\ p_4]$, associated with the position of a color in the actual ranking. As before, different payoff vectors define different decision problems, and we consider four of them, $\mathcal{DP}_1$-$\mathcal{DP}_4$. The bottom panel of Table 1 summarizes all the elements of the solutions of these problems, while we discuss here some of the key aspects of the solutions.

The payoff vector for $\mathcal{DP}_1$ is $\mathbf{p} = [20,\ 20,\ 0,\ 0]$. For illustrative purpose, we fix the ranking at $[blue, red, gray, green]$ and discuss the DM's optimal choices under this ranking, but the same reasoning would apply to the other orderings without loss of generality. The DM optimally places the same probabilities on two pairs of signals by increasing the probabilities of identifying the first two states and reducing the probabilities of identifying the last two states. The probability of the signals suggesting each state-contingent action given *blue* is first in the ranking is given by $P^*_{blue} = [.4, .4, .1, .1]$, and it is defined symmetrically for the other colors and ranking permutations. In turn, this choice makes actions a and b equally profitable given the ranking and makes actions c and d equally profitable as well. The information acquired is used to trade off probabilities of selecting signals conducive to c or d

Environment: 2s2a

	$\mathcal{DP}_1'$	$\mathcal{DP}_2'$
$\mathbf{p}$	[20, 0]	[20, 19]
$\mathbb{E}_\mu$	10	19.5
P^*_ω	[.8, .2]	[.52, .48]
V	15.83	19.5
G	.21	5e-4

Environment: 4s4a

	$\mathcal{DP}_1$	$\mathcal{DP}_2$	$\mathcal{DP}_3$	$\mathcal{DP}_4$
$\mathbf{p}$	[20, 20, 0, 0]	[20, 20,19, 19]	[20, 0,0, 0]	[40, 20,0, 0]
$\mathbb{E}_\mu$	10	19.5	5	15
P^*_ω	[.40, .40, .10, .10]	[.26, .26, .24, .24]	[.55, .15, .15, .15]	[.71, .19, .05, .05]
V	15.83	19.5	11.17	32.27
G	.21	5e-4	.20	.62

Table 1: Theoretical quantitative solutions for the six decisions problems considered in the analysis. The solution is derived under the calibration of $\kappa = 15$. The payoff vectors indicated by $\mathbf{p}$, $\mathbb{E}_\mu$ is the ex-ante utility under uniform prior, P^*_ω is optimal conditional distribution of signal given state, $V =$ is the expected utility at the optimum and G is the cost at the optimum.

for signals suggesting a or b. This information choice allows the DM to increase her expected payoff from an ex-ante value of $\mathbb{E}_\mu[u] = 10$ to an ex-post expected value of $V = 15.83$.

The payoff vector of $\mathcal{DP}_2$ is $\mathbf{p} = [20,\ 20,\ 19,\ 19]$. The expected utility under the flat prior is $\mathbb{E}_\mu[u] = 19.5$, almost double that of $\mathcal{DP}_1$. The optimal information strategy is to remain essentially uninformed. The incentives for gathering information in this environment are low since the payoffs are very similar across alternative actions.

Next, we consider the effects on attention of raising the incentives to distinguishing among states. We implement this change in $\mathcal{DP}_3$, characterized by the payoff vector $p = [20,\ 0,\ 0,\ 0]$. Note that the payoff vector implies that the DM receives payment only when she can correctly identify the first ranked state. Under the example true ranking given above, the optimal solution distinguishes the paying option *blue* from the three other alternatives, assigning a probability of 0.55 to blue and a probability of 0.15 for each of the non-paying state-contingent actions. This information strategy provides a significant lift to expected utility from the prior to the posterior distribution, while keeping the cost close to the same level of $\mathcal{DP}_1$, as shown in Table 1.

Finally, we consider whether changes in attention are triggered by decoy effects in an environment in which there are two states that pay positive but markedly different rewards. We call this decision problem $\mathcal{DP}_4$, with associated payoff vector $p = [40,\ 20,\ 0,\ 0]$. The optimal information strategy is similar to the one for $\mathcal{DP}_1$ in that the optimal choice involves a partition into two sub-spaces. However, in $\mathcal{DP}_4$ the DM now has an incentive to pay additional attention to further differentiate among the first and second ranked states so that she can select the actions corresponding to the highest expected utility. Moreover, the DM does not distinguish between the two non-paying options. This strategy is the most attention demanding across all decision problems analyzed, since it involves an increase in both overall attention (the total cost G_4) and precision in attention's allocation (as shown by the optimal P^*_ω in Table 1). The reason is that the DM's optimal information structure encompasses both vetting several alternatives and differentiating between them in a meaningful way.

3 Experimental Design

3.1 General Description of the Experimental Task

In our experiment, participants face a series of decision tasks. In each decision task subjects observe 1) 100 colored balls comprised on N different colors and 2) an information table. The information table indicates the number of balls of the i^{th} most common color for $i \in [1, \ldots, N]$. The information table also provides the payoff vector $\mathbf{p}$ indicating the payoff the subject will receive if she selects the i^{th} most common color. For each task, the subject has an unlimited amount of time to select a color. Figure 2 provides an example task with $N = 4$.

3.2 Treatments

We consider six treatments as described in Table 2. In treatment T1, there are four colors of balls: red, blue, green, and gray. There are 27 balls of the most common color, 26 balls of the second most common color, 24 balls of the third most common color, and 23 balls of the least common color. For T1 if the color selected by the participant was the most common or the second most common the payoff was \$20 and otherwise it was \$0. That is,

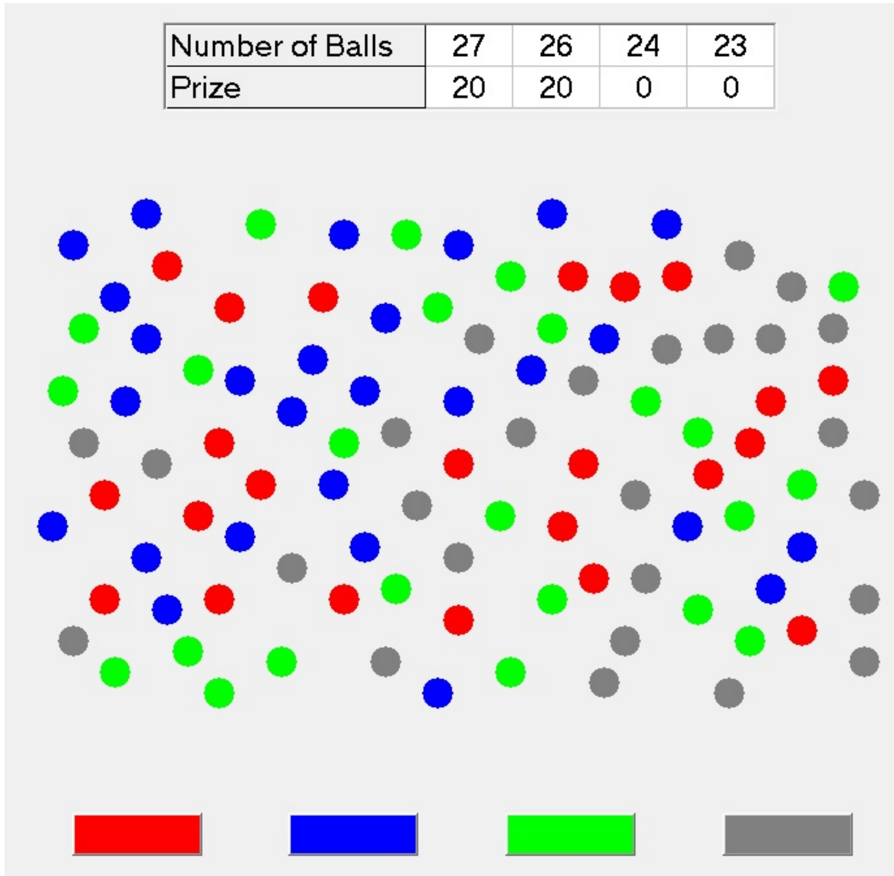

Figure 2: An example of the experimental task in the four-colors world. There are 100 balls: 23 *red* balls, 24 *green* balls, 26 *gray* balls and 27 *blue* balls. For this case, recognizing either of the two most common colors (*blue* and *gray*) yields a payoff of \$20 each. The other two colors yields \$0.

$\mathbf{p} = [20, 20, 0, 0]$ and thus T1 corresponds to $\mathcal{DP}_1$. The example shown in Figure 1 reflects T1. Treatments T2, T3, and T4 use the same colors and the same numbers of balls, but differ by $\mathbf{p}$. In T2 the payoff for the third and fourth most common colors is increased to \$19 so that $\mathbf{p} = [20, 20, 19, 19]$ as in $\mathcal{DP}_2$. T3 is similar to T1 except that the payoff of the second most common color is decreased to \$0 so that $\mathbf{p} = [20, 0, 0, 0]$ consistent with $\mathcal{DP}_3$. T4 is similar to T1 except that the payoff for the most common color is increased to \$40 so that $\mathbf{p} = [40, 20, 0, 0]$ consistent with $\mathcal{DP}_4$. Because T2, T3, and T4 are constructed by modifying T1, for convenience we refer to T1 as the baseline.

The other two treatments, T1$'$ and T2$'$, only involve 2 colors (red and blue) and there are 52 balls of the most common color and 48 balls of the least common color. In T1$'$ the payoff for selecting the most common color is \$20 and the payoff for selecting the least

2 colors

Decision Problems	Experimental Treatments	Color Frequency	Payoff Vector
$\mathcal{DP}_1^{'}$	T1$^{'}$	(52, 48)	[20, 0]
$\mathcal{DP}_2^{'}$	T2$^{'}$	(52, 48)	[20, 19]

4 colors

Decision Problems	Experimental Treatments	Color Frequency	Payoff Vector
$\mathcal{DP}_1$	T1	(27, 26, 24, 23)	[20, 20, 0, 0]
$\mathcal{DP}_2$	T2	(27, 26, 24, 23)	[20, 20, 19, 19]
$\mathcal{DP}_3$	T3	(27, 26, 24, 23)	[20, 0, 0, 0]
$\mathcal{DP}_4$	T4	(27, 26, 24, 23)	[40, 20, 0, 0]

Table 2: Mapping theoretical decision problems, denoted by $\mathcal{DP}$, with experimental treatments, denoted by T. Relative frequencies of balls per color and payoff vectors are reported in the last two columns.

common color is \$0 so that $\mathbf{p} = [20, 0]$ consistent with $\mathcal{DP}_1^{'}$. Thus, T1 and T1$^{'}$ are similar despite the difference in the number of colors. T2$^{'}$ is similar to T1$^{'}$ except that the payoff for the least common color is increased to \$19 so that $\mathbf{p} = [20, 19]$ consistent with $\mathcal{DP}_2^{'}$. Thus T2$^{'}$ is similar to T2 despite the difference in the number of colors.

3.3 Connection Between Tasks and Theoretical Framework

The experimental setting mirrors our theoretical framework. The colors correspond to different possible states. Considering the relative frequency of each color is akin to understanding which state is likely. The relevant space for the participants is given by all possible color rankings and participants are informed in advance that all possible rankings are equiprobable. Further, the cognitive cost of processing information to acquire a signal about the state mimic Shannon's cost in the model. With action and state spaces, prior, cost and payoffs fully specified, the experimental tasks directly align with the decision problems described in Subsection 2.2. That is, T1 aligns with decision problem $\mathcal{DP}_1$; T2, T3 and T4 align with decision problems $\mathcal{DP}_2$, $\mathcal{DP}_3$ and $\mathcal{DP}_4$, respectively, and T1$^{'}$ and T2$^{'}$ align with decision problems $\mathcal{DP}_1^{'}$ and $\mathcal{DP}_2^{'}$, respectively.

Because we do not impose a time limit on the tasks, the color ranking is perfectly observable if a participant exerts sufficient effort. Further, there is no cost associated with a particular strategy of gathering information. As a result, a participant who is successful may have expended sufficient cognitive effort or may have not exerted sufficient cognitive effort and made a lucky guess. What the theoretical framework provides us with is an objective measure of the complexity of the system of colored balls given their relative frequencies based on Shannon's entropy.

A second important thing to note about our design is that we collected choice data in a setting in which participants select their own attention strategies without eliciting beliefs. The choice of not imposing an information structure is deliberate: in the rational inattention framework, the DM faces an intrinsic cognitive cost of acquiring and processing information as opposed to an exogenously imposed information structure. The model posits that the actions-specific posteriors are revealing of the optimally chosen information strategies and can be measured without stated beliefs.

We also note that to have a more precise mapping of each available action to each state of the world, we treat repetitions of the same tasks as multiple, independent realizations of the same decision. Finally, because the focus of the paper is on incentives, we vary the payoff structure while keeping the cost of information fixed by keeping the number of balls of the most common, second most common, and so on fixed given the number of colors used in the task.

3.4 Procedures

The study was conducted at The University of Alabama's TIDE Lab and all of the participants were drawn from the lab's standing pool of research study volunteers. None of the participants had been in any related previous study. Upon arrival at the lab, subjects read and signed a consent form and were then seated at a private workstation where they read computerized instructions. A copy of the instructions is available in Appendix A.

As explained in the instructions, participants completed 160 tasks. The first 40 tasks consisted of 20 T1$'$ tasks and 20 T2$'$ tasks in a completely randomized order. The last 120 tasks consisted of 30 T1 tasks, 30 T2 tasks, 30 T3 tasks, and 30 T4 tasks in a completely randomized order. Blocking the tasks according to the number of colors in the task is meant to avoid confusion and starting with the $N = 2$ color treatments is intended to help the participants better understand the interface when facing the $N = 4$ color treatments, which are the primary focus of the experiment. For each task, the ranking of each color and the order of the colored response buttons were randomized.

Between each task the screen went blank for 1 second. Then the information table appeared for 1 second before the next image of colored balls appeared. After the decision tasks, participants completed a survey that included a question about gender and a question designed to measure their risk intensity.[8]

Once the participant completed the study, one task was chosen at random and the participant was paid based upon their decision in that period. All payoff amounts were in $US. Participants were recruited for one hour study although some finished in as few as 20 minutes and some took up to 75 minutes. The average salient payoff was $18.34 and participants received an additional $5 for completing the study.

4 Experimental Results

We divide the experimental results into two subsections. In the first subsection, we report the results for our 2s2a environment and compare them to those of Caplin and Dean (2015). In the second subsection, we focus on our core 4s4a environment and test the theoretical predictions of the RI model against the laboratory data.

4.1 The 2s2a Benchmark Environment

Figure 3 compares the conditional distributions of choices for problems T1$'$ and T2$'$ in response to a change in incentives from the baseline payoff vector [20 0] to the vector [20 19]. The theory suggests that subjects should expend less attention in T2$'$ as compared to T1$'$ and that the distribution in T2$'$ should become flatter since the incentives to distinguish between the two actions decreases. Caplin and Dean (2015) found that subjects responded to incentives in the predicted way, but that very large changes in payoffs were necessary to

[8]The risk question asks participants what fraction of a larger lottery prize they would be willing to invest in a risk asset that was equally likely to double or halve in value (see Dohmen *et al.* (2011).)

obtain significant effects.[9] We confirm the same type of result in Figure 3 for the aggregate observations across all subjects. In our data, the average probability of choosing the first action goes down from 70.3% to 66.7%, a small negative drop. We test whether this difference is equal to zero against the alternative hypothesis of being negative, and we find the test rejects the null only at 5.9% of significance level.

Evidence consistent with a somewhat larger change of attention allocation between the two treatments is illustrated by Figure 4, where decision times are used as a proxy for attention effort. Subjects spend twice as much time to discern the correct state when they take the first action in T1$'$ than in T2$'$ - 10.4 seconds as compared to 5.2 seconds. The difference between treatments when selecting the second action is only 2 seconds. Further, time spent is virtually identical regardless of the action taken in T2$'$.

4.2 The 4s4a Environment

Having established that factors like the computer interface and subject pool are not impacting behavior given the similarity our results and those previously reported in the literature, we now turn to our main 4s4a environment. The results are summarized in Figure 5, which compares the theoretical optimal conditional distribution for each decision problem with the distributions observed in the experiment for T1, T2, T3, and T4. The theoretical solutions are derived under the calibration of the marginal attention cost parameter $\kappa = 15$ as discussed in Section 2.2. For easy completeness, the bottom panels of Figure 5 report the analogous distributions for T1$'$ and T2$'$ in the 2s2a experiment.[10]

[9]See also Caplin and Martin (2014), Dean and Neligh (2019), Dewan and Neligh (2020) and Caplin *et al.* (2021).

[10]We test for the NIAS (No Improving Action Switches) and NIAC (No Improving Attention Cycles) conditions introduced by Caplin and Martin (2014) and Caplin and Dean (2015) to assess whether the overall data structure generated by our experiment satisfies the RI framework.

The NIAS condition verifies that any chosen action is optimal given the belief implied by the revealed posterior observed in the data. Intuitively, it ensures that actions are optimal given the information acquired by the subjects, ruling out systematic misuse of information. The condition entails a series of tests for each treatment in the experiment and each observed action, which is compared to the others in the action set. The NIAC condition ensures that gross utility cannot be increased by reassigning information structures along any cycle of decision problems. Intuitively, it ensures that attention is allocated efficiently throughout the experiment, ruling out negative response to incentives.

We find that the NIAS and NIAC tests at the aggregate level for the whole sample are significantly non negative. Our data set overall satisfies the conditions. At individual level, one third of the subjects (28 out of 75) fails one or more of the conditions. The results discussed in the section, however, remain robust to

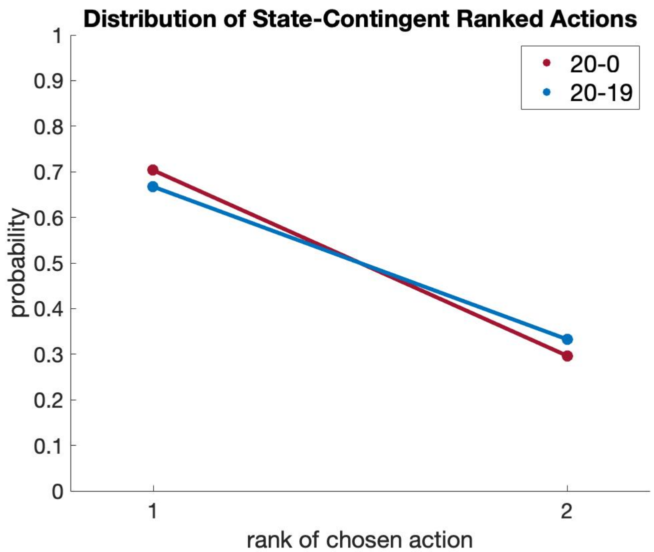

Figure 3: Comparing the experimental distributions of the rank of chosen actions in the two-action problem for T1$'$ and T2$'$. Average probability aggregated across all choices and all subjects.

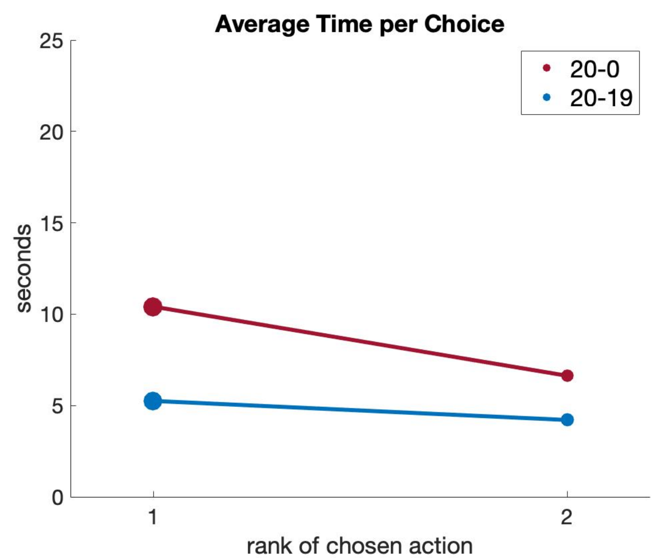

Figure 4: Comparing the decision time of the experimental subjects in the two-action problem for T1$'$ and T2$'$.The time is reported by action taken. The size of the dot is proportional to the number of times the action is observed. Average time aggregated across all action and all subjects.

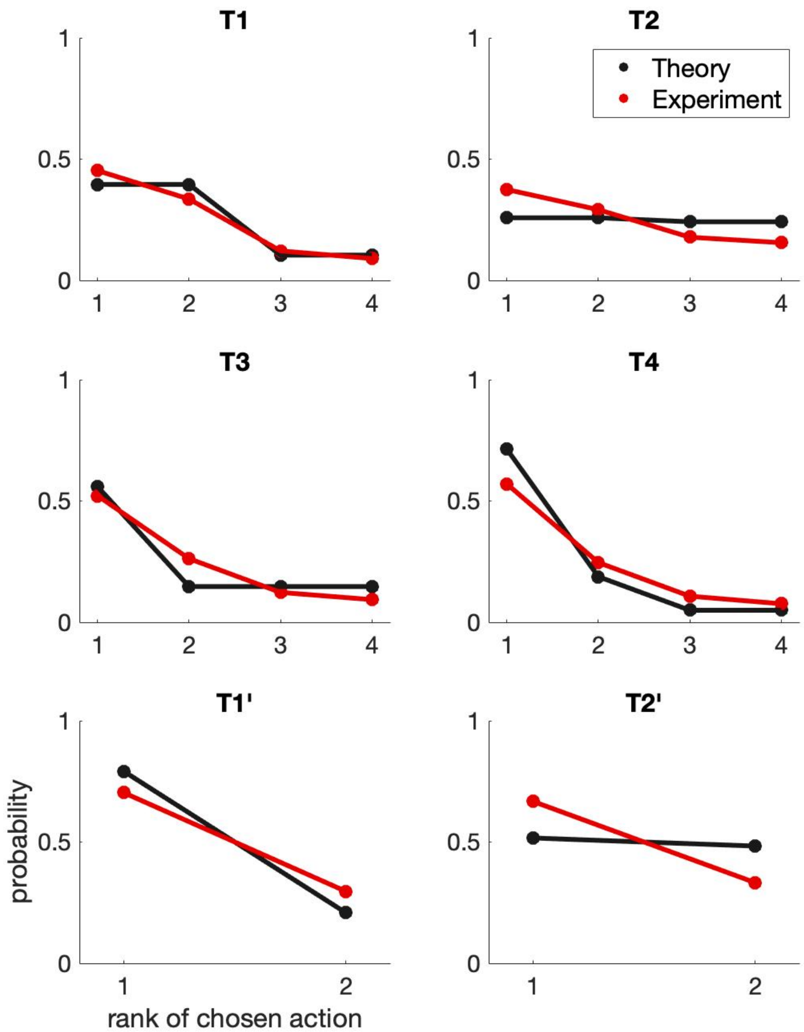

Figure 5: Comparing the theoretical optimal conditional distributions of all the decision problems with their experimental counterparts. Top four panels are for the 4s4a treatments T1, T2, T3, and T4. The bottom two panels are for the 2s2a treatments T1$'$ and T2$'$. The marginal attention cost is set to $\kappa = 15$ in all the treatments.

The experimental choice distributions (the red lines in Figure 5) are broadly consistent with the theoretical distributions (the black lines in Figure 5), at least from a qualitative standpoint. To corroborate the graphic intuition, we calculate the Kullback–Leibler (KL) divergence between observable choice distributions and the theoretical posteriors and report those in Table 3 for each treatment. This quantity, expressed in bits, can be interpreted as the information lost by approximating the experimental data by the RI model. To give a sense of the magnitude of these divergences, we compare them with the corresponding KL measures obtained under the assumption that no information is obtained (second column of Table 3) and the assumption that full information is obtained (third column of Table 3). The last column of the table, then, puts the KL measures into relative terms by dividing them by the entropy levels of the theoretical distributions.

The RI model's KL divergence measures from the observed data are smaller than the no-information and full-information in all six decision tasks. The relatively better fit of the RI model in T1 is due to the fact that this treatment is used to calibrate κ. However, all the other treatments also report small divergences from the RI model. In general the divergence is four to five times smaller than that associated with the no-information strategy, with the exception of T2 in which the two models unsurprisingly have similar values given the payoff vector in that task. In every task the divergence associated with the RI model is far smaller than that associated with the full information model. The ratios reported in the last column confirm that the deviations from the RI theoretical model are relatively small as the ratio is less than 3% for each treatment.

The actual treatments we employ were designed to help identify certain effects. Specifically, we focus on four main comparisons as summarized in Table 4.

Point 1: The theoretical results show that the optimal strategies in $\mathcal{DP}_1$ and $\mathcal{DP}_1'$ should be behaviorally equivalent. This observation reflects the Invariance Under Compression (IUC) property of the model for which making a choice environment more complex by adding states with similar payoffs does not modify the DM's optimal choices (see Caplin *et al.*, 2021, for details of the IUC property). As the first panel of Figure 5 illustrates, the DM's optimal strategy in $\mathcal{DP}_1$ is to acquire signals that partition the space

the exclusion of those subjects. These results are not included here, but are available upon demand.

Treatment	KL Divergences			Ratio to RI Entropy
	RI Model	No Information	Full Information	
T1	.005	.081	.581	.40%
T2	.022	.027	.681	1.62%
T3	.022	.094	.456	1.89%
T4	.024	.123	.389	2.82%
$\mathrm{T1}'$	.009	.037	.197	1.82%
$\mathrm{T2}'$	.020	.025	.239	2.92%

Table 3: Kullback–Leibler divergence of the experimental state-contingent action distributions from the theoretical RI posterior distributions (second column). Kullback–Leibler divergence from a fully-uninformed strategy (third column) and from the full-information model theoretical distribution (fourth column). KL measures are expressed in bits. Since the full-information distributions are degenerate, we approximate the full-information distribution with the function (.97 .01 .01 .01) which provides conservative estimates of the KL divergence in this case.The KL measures for the RI model are normalized by the entropy of the theoretical posterior distributions in the last column of the table.

into paying states (the first and second ranked states) and non-paying states (the third and fourth ranked states). Within each partition, the signals do not differentiate between alternatives. Moreover, the sum of the probabilities of the first and second action is equal to the posterior probability of the first action in $\mathcal{DP}_1'$.

Our data demonstrates a clear separation between the probability placed in the paying actions and non-paying actions in T1. In fact, the sum of the probabilities of the first two ranked actions is very close to the theoretical level. A formal within subject t-test of this hypothesis does not reject the null ($p - value = .80$). But we still observe a greater probability than predicted of the first ranked action relative to the second ranked action. We test whether these two probabilities are statistically different and we find the data do reject the null hypothesis of no difference. Nevertheless, decision times for actions in T1 reported in Figure 6 confirm that subjects exercise very similar efforts to take decisions in favor of the first two actions, which is higher than for the last two actions. Taken together these results moderately support the conclusion that the IUC property is satisfied by the data.

Point 2: Lower incentives mean that less information is processed going from $\mathcal{DP}_1$ to $\mathcal{DP}_2$, as was the case in the 2s2a environment when comparing $\mathcal{DP}_1'$ to $\mathcal{DP}_2'$. In both environments the DM optimally decides to stay uninformed and the optimal conditional

Point	$\mathcal{DP}$ Comparison	Description
1	$\mathcal{DP}_1$ Vs $\mathcal{DP}_1^{'}$	Impact of Complexity
2	$\mathcal{DP}_1$ Vs $\mathcal{DP}_2$	Role of Incentives
3	$\mathcal{DP}_3$ Vs $\mathcal{DP}_4$	Steeper Incentives & Ranking
4	Across all $\mathcal{DP}$'s	Expected Utility & Information Gains

Table 4: Main assessment points for the comparison between experimental data and theoretical implications of the RI model. The second column of the table indicates the decision problems on which they are based.

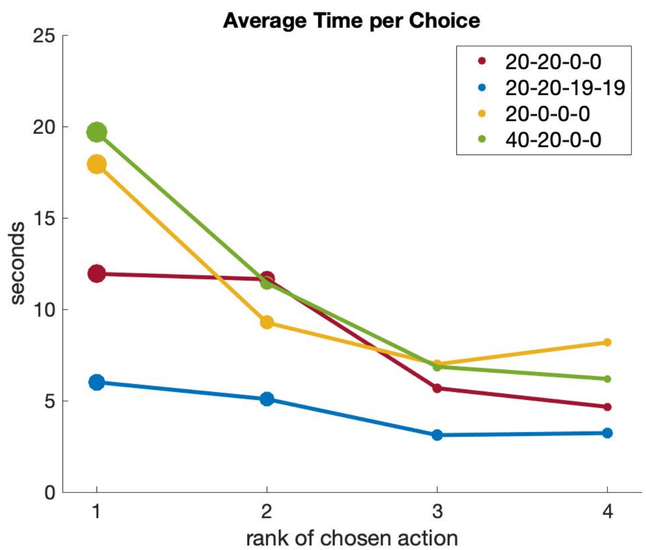

Figure 6: Comparing the decision time of the experimental subjects in the four-action problem for treatments T1, T2, T3, and T4. The time is reported by action taken. The size of the dot is proportional to the number of times the action is observed. Average time aggregated across all action and all subjects.

distribution of actions are almost perfectly flat.

We observe a flattening of the empirical posterior distributions in T2 relative to T1 consistent with the theory. However, the first two ranked choices and the last two are respectively chosen with higher and lower probabilities than predicted. We use a set of within subject t-tests to check whether each point in the distribution is statistically significantly different from the value predicted by the RI model. All tests reject the null at high confidence levels (with $p - values$ in the order of 10^{-4} or smaller), indicating a failure of the model in statistical terms despite the qualitative goodness. The decision times for T2 in Figure 6 depict a similar story. Subjects exercise similar efforts in taking decisions across all actions, but they still spend relatively more time for the two top ranked actions.

Point 3: In presence of a single high-paying state and three states with 0 payoff as in $\mathcal{DP}_3$, under RI the DM abandons the strategy of partitioning the top two ranked actions and the bottom two ranked actions. Instead, the DM optimally focuses on the signal suggesting the first ranked action and remains relatively uninformed about the remaining options. Given the calibration, the posterior probability of the first action is .55, while the three non-paying alternatives take on .15 probability each. By contrast, in T4 the RI DM optimally chooses to differentiate among two top ranked and two bottom ranked actions and then seeks a sharper differentiation between first and second ranked state-contingent actions. The benefit of selecting the top ranked action over the second ranked action, however, is the same in both of these treatments (i.e., the payoff difference between the two choices is 20 both in T3 and T4).

As with T1 and T2, the revealed distributions for T3 and T4 are in line with the theoretical predictions as both distributions get steeper relative to that observed for T1. But the strategies adopted by the subjects in T3 and T4 are remarkably similar, suggesting a new role for the second ranked action in the more complex environment. This is an interesting dimension of analysis that is missing in the less complex 2s2a setting. Three observations are in order.

First, a t-test does not reject the null that the probabilities of the first and second ranked actions selected in T3 and T4 are the same against the alternate hypothesis that the T4 probability is higher, contrary to what implied by the theory ($p-value = .99$). Second, while the probability of the second ranked action being selected in T4 is sizable as predicted, it is chosen too often by the subjects. The observed frequency exceeds the predictive frequency by six percentage points and a one-sided t-test confirms the difference is statistically significant. Finally, subjects in T3 choose the first ranked action with a probability close to the theoretically predicted level. In fact, a t-test fails to reject equality at standard significance levels with $p-value = .09$. However, the gap between the observed and predicted frequency of the second ranked choice being selected is 12 percentage points, which is a significant difference.

Generally speaking, we find that payoffs are effective in modifying attention strategies in the 4s4a environment in line with the theoretical predictions. Figure 7 illustrates this effect by reporting the confidence intervals associated to the t-tests of the within subject differences in the first ranked choice probabilities of each treatment with respect to the baseline T1. All differences are in the direction predicted and statistically significant at the 95% confidence level.[11] Similarly, decision times in Figure 6 also show attention progressively shifts upwards as the incentives to correctly distinguish between actions increase in T3 and T4.

In spite of these effects, however, the greater closeness of first and second ranked action probabilities in T3 and T4 indicates a potential shortcoming of the RI theory. One potential explanation for the divergence between data and predicted probabilities is perceptual distance (see Hébert and Woodford, 2020). The Shannon's function postulates that the cost of differentiating across alternatives depends solely on the number of states in the system, without explicitly taking into account the difficulty in counting the different number of balls representing the states. However, in the experiment some states might be more difficult to tell apart than others. In our results, the flattening of the observed distributions could be due to the similarity between the first two states with 27 and 26 balls.

The effects of perceptual distance may confound, but do not completely offset, the incentives to differentiate between first and second ranked actions in the 4s4a environment though. The second ranked action retains a useful function in a multi-state environment since it still helps subjects separate the first ranked action from the lowest ranked ones, even when it provides a limited payoff. This is an interesting aspect of analysis in this type of environment that deserves further attention in future research.

Point 4: Our last point compares expected utility and information gains thus providing further characterization of the results discussed so far, especially with respect to perceptual distance. We first calculate theoretical expected utility gains as the difference in ex-ante utility under the uniform prior ($\mathbb{E}_\mu$ in Table 1) and the expected utility evaluated at the optimal posterior (V in Table 1). The information gain from an uninformative prior to the optimal distribution implied by the chosen information strategy is given by the difference

[11] t-tests for the null hypothesis that these differences are zero against the alternative that they are positive in T3 and T4 and negative in T1 strongly reject the null (with $p-values$ less than 10^{-4}). The same conclusions are also confirmed by a set of one sample proportion Wald tests for the share of subjects with the predicted directional change in the probability of the first ranked action being selected. The tests for the four-action experiment strongly reject the one-sided $H_0 : p = .5$.

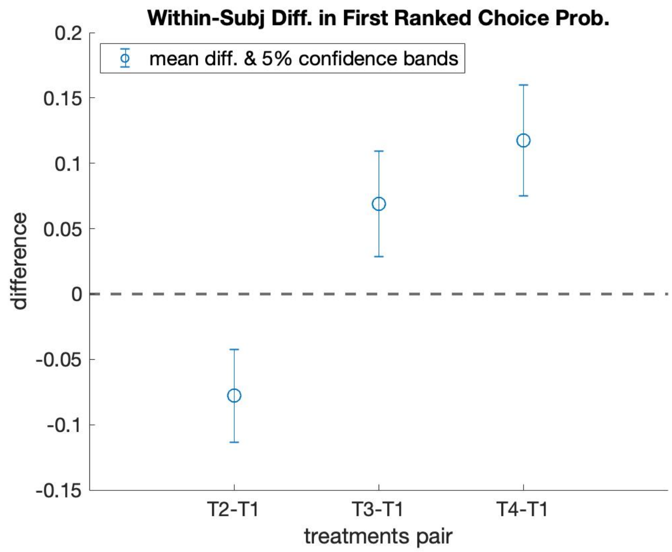

Figure 7: Comparing the within subject difference of the first action probability with respect to the baseline T1. Each candle bar shows the mean difference across all subjects and its confidence interval at 5% significance level for a zero-mean t-test.

between the entropy of the uniform prior over actions and the entropy of the posterior P^*_ω in Table 1, multiplied by $\kappa = 15$.[12] The empirical counterparts of these measures are calculated by replacing the theoretical posteriors with the observable distributions.

Figure 8 displays the expected utilities and information costs for both the theoretical and experimental posterior distributions. On the left panel of the figure, one can see that the ratio between model-implied changes in expected utility and changes in information cost is basically constant and equal to 2 across decision problems, with the exception of T2 where there is little incentive to acquire and process information. For the subjects in the experiment on the right panel, behavior in T1, T2, and T3 is very close to the theoretical level. By contrast, in T4 subjects gather less information than the optimal RI DM and obtain a smaller utility gain as well. The T4 allocation of attention delivers comparatively higher expected utility at a slightly higher cost then in T3. Subjects face a steeper information cost to parcel out the signal for the second ranked action in T4. Because the perceived cost of information is higher than the perceived gain, subjects settled for a sub-optimal payoff gain with a ratio to information cost higher than 2.

[12] More precisely, the information gain corresponds to the cost G in Table 1, i.e., the mutual information of the optimal strategy. The change in entropy is used as a proxy for G to facilitate a direct comparison with the experimental data.

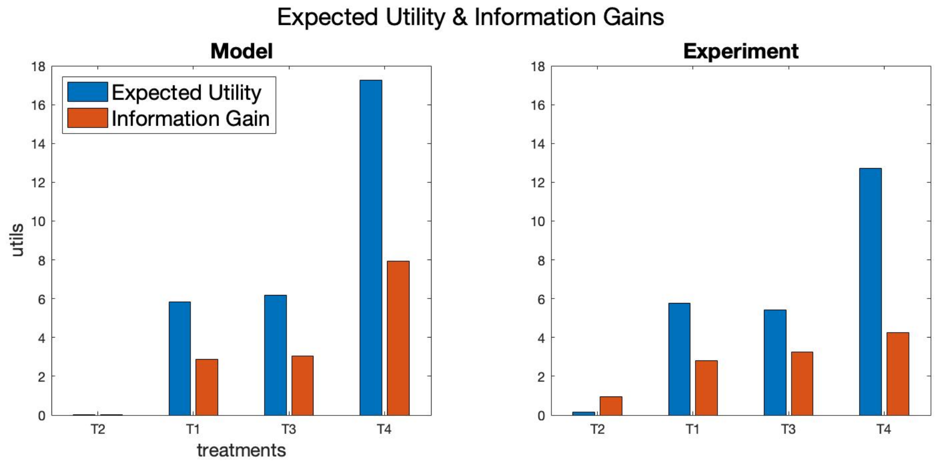

Figure 8: Comparing expected utilities and information gains for model and empirical data. Expected utility gains are obtained as $\mathbb{E}_\mu - V$ from Table 1). Information gains are given by the difference of the entropy of the uniform prior and the posterior P^*_ω in Table 1, multiplied by the calibrated $\kappa = 15$. Their empirical counterparts are calculated by replacing the theoretical posteriors with the observable distributions. The treatments are re-ordered on the x-axis for visual convenience.

5 Discussion

The framework we propose is general enough to allow for a variety of applications in environments where agents need to reduce uncertainty in a complex system to identify which behavior is the most appropriate among multiple alternatives. In this section we discuss three stylized examples of potential applications of our framework.

The first application is in monetary policy. Let us consider an economy in which the states of the world are determined by two features of the inflation rate: the inflation level (high or low) and its duration (temporary or permanent). The Central Bank sets the interest rate following a Taylor-type rule that directly responds to the inflation rate and its duration. The DM in this environment makes decisions about a portfolio of stocks and bonds, the performance of which depends on the policy rate and its time dynamics. The DM, hence, needs to process information about the different inflation outcomes and the associated interest policy to optimize her portfolio management decisions. The optimal attention strategy balances the effort required to reduce uncertainty about the various dimension of inflation and the benefits of selecting a portfolio attuned to the actual policy stance. The implications for monetary policy is that the Central Bank may need to keep into account the scarce attention resources of the private sector when assessing the transmission of policy decisions to the economy. While responding to changes in duration and level of inflation may be desirable, it also puts an additional cognitive burden on private agents to accurately identify their optimal portfolio response which could impair the economic effects of a policy.

The second application is the study of the interactions between monetary and fiscal policies. Consider an economy in which both fiscal and monetary policies can be either active or passive at any given time, controlling tax rates and interest rates respectively. These policy combinations jointly determine a complex economic environment in which a DM needs to process information about multiple aspects of the national economic policy to optimally choose consumption and savings in the short and long term as a function of the prevailing tax and interest rates.

A third example would be the analysis of labor markets with realistically complex compensation schemes. Let us suppose that employers offer a mix of compensation and benefits to attract candidates. The DM decides whether to accept a job offer by analyzing the proposed contracts and ranking them according to their preferences about different dimensions of the compensation scheme. Such an environment may have different predictions on the matching rate in the economy than those from a model that only considers wages.

6 Conclusions

While the rational expectations model has taken on a predominant role in economics, scholars have begun to recognize the limitations of such an approach in modeling behavior. In the rational inattention framework, the DM faces a trade-off between the cognitive effort of processing information relevant for her actions and the benefit arising from making precise state-contingent choices. The core of the theory rests on the interplay between incentives, as described by the DM's utility, and costs based on Shannon's mutual information. This interplay results in the choice of an optimal information structure that relates information about states and the DM's behavior.

While the theory postulates a precise mapping between information and resulting behavior, it is difficult to directly infer the information strategy behind choices with naturally occurring economic data. Thus, there is a growing reliance on data from laboratory experiments to directly test the predictions of the rational inattention model. While the preliminary results are encouraging, the settings behind those results have been relatively simple. Building on the existing literature, we extend the state and action spaces to accommodate richer economic environments expanding the scope of the DM choices in a meaningful way. Specifically, we provide a framework to study discrete choice problems under rational inattention and we characterize the solution of the optimal choices of information structure and behavior in response to complexity driven by changes in incentives and in the number of states and actions available to the decision maker.

Using controlled laboratory experiments, we directly test for the implications of the RI model. The various experimental tasks directly align with the theoretical decision problems we model, allowing us to analyze the empirical frequencies of the subjects' information choices and their behavior in response to changes in the underlying environment and to compare them with their theoretical counterparts.

We find that the experimental subjects generally behave as if they were rationally inattentive. In particular, we find that our subjects modulate their attention according to the relative profitability of the state-contingent actions and strongly respond to incentives when the payoff differential among alternatives is substantial. We also document, however, some gaps between RI theory and experimental behavior that requires adjustments in the baseline model, especially with regards to perceptual distance between similar states difficult to discriminate and perceived information processing costs compared to utility gains. These behavioral aspects can also interact with the complexity of the economic environment. Our setting provides a suitable research framework to further explore these effects and understand their implications in complex situations.

Given the large variety of applications suitable for our framework and the predictive success of the rational inattention model, we believe our paper would provide a useful tool for future theoretical and experimental analysis applying rational inattention to different economic settings.

References

Bushong, B., Rabin, M. and Schwartzstein, J. (2020). A Model of Relative Thinking. *The Review of Economic Studies*, **88** (1), 162–191.

Caplin, A. and Dean, M. (2013). *Behavioral implications of rational inattention with shannon entropy*. Tech. rep., National Bureau of Economic Research.

— and — (2015). Revealed preference, rational inattention, and costly information acquisition. *American Economic Review*, **105** (7), 2183–2203.

—, — and Leahy, J. (2019). Rational inattention, optimal consideration sets, and stochastic choice. *The Review of Economic Studies*, **86** (3), 1061–1094.

—, — and — (2021). Rationally inattentive behavior: Characterizing and generalizing shannon entropy. *Journal of Political Economy. Forthcoming.*

— and Martin, D. (2014). A Testable Theory of Imperfect Perception. *The Economic Journal*, **125** (582), 184–202.

Cheremukhin, A., Popova, A. and Tutino, A. (2015). A theory of discrete choice with information costs. *Journal of Economic Behavior and Organization*, **113**, 34–50.

Csaba, D. (2018). *Attentional Complements*. Tech. rep.

Dean, M. and Neligh, N. (2019). Experimental test of rational inattention. *Mimeo, Columbia University.*

Dewan, A. and Neligh, N. (2020). Estimating information cost functions in models of rational inattention. *Journal of Economic Theory*, **187**.

Dohmen, T., Falk, A., Huffman, D., Sunde, U., Schupp, J. and Wagner, G. G. (2011). Individual risk attitudes: Measurement, determinants, and behavioral consequences. *Journal of the European Economic Association*, **9** (3), 522–550.

Duffy, J. and Puzzello, D. (2021). The friedman rule: Experimental evidence. *International Economic Review. Forthcoming.*

Hébert, B. M. and Woodford, M. (2020). Neighborhood-based information costs. *American Economic Review*, **111** (10), 3225–55.

Kőszegi, B. and Szeidl, A. (2013). A model of focusing in economic choice. *The Quarterly journal of economics*, **128** (1), 53–104.

Kryvtsov, O. and Petersen, L. (2021). Central bank communication that works: Lessons from lab experiments. *Journal of Monetary Economics*, **117** (C), 760–780.

Maćkowiak, B., Matějka, F. and Wiederholt, M. (2022). Rational inattention: a review. *Journal of Economic Literature. Forthcoming.*

Matějka, F. and McKay, A. (2015). Rational inattention to discrete choices: A new foundation for the multinomial logit model. *The American Economic Review*, **105** (1), 272–298.

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

— (2006). Rational inattention: A research agenda. *American Economic Review: Papers and Proceedings*, **96**, 158–163.

Appendix

A General Instructions for the Laboratory Experiment

Please make sure your phone and all other electronic devices are turned off and put away. If you have a question at any point, please raise your hand and someone will come to you. Please do not communicate with anyone else in this study or do anything to distract anyone else in this study.

You will face 160 decision periods. A single period selected at random will determine your payment (in addition to the $5 you are receiving for participating in this study).

Every decision period you will be shown a picture of 100 colored balls. There may only be two colors of balls: RED and BLUE. Or there may be four colors of balls: RED, BLUE, GREEN, and GRAY.

There will always be a total of 100 balls, but you are not told how many balls there are of any specific color.

Each period there will be a table at the top right of your screen that tells you how many balls there are of the most common color and how many balls there are of the 2nd most common color. If there are four colors, the table will also tell you how many balls there are of 3rd most common color and how many balls there are of the 4th most common color. Each color is equally likely to be the most common, the second most common, and so on.

Every period, you must click on one of the four colored buttons on the lower right portion of your screen. If the color you click is the most common color, you will earn the 1st place prize. If the color you click is the 2nd most common you will earn the 2nd place prize. When there are four colors, if you click on the 3rd most common color you will earn the 3rd place prize and if you click on the least common color you will earn the 4th place prize. All prizes are in $US.

The table with the number of balls and the prize can change every period.

Let's look at the example.

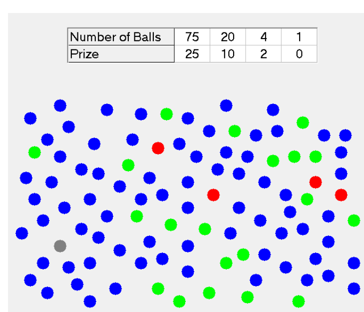

The table tells you there are 75 balls of the most common color and if you click that color you will earn $25. The table tells you there are 20 balls of the 2nd most common color and that color is worth $10. There are 4 of the 3rd most common color and that color is worth $2 and there is 1 of the least most common color and that color is worth $0. In this example, there are 75 BLUE balls, 20 GREEN balls, 4 RED balls, and 1 GRAY ball. So in this example you would earn

$25 for clicking ■ $10 for clicking ■ $2 for clicking ■ $0 for clicking ■

Please raise your hand if you have a question. If you are ready to begin the study you may press Start.