Decomposing Inattention

Kevin Laughren[*]

November 12, 2023

Abstract

First collect some information, then choose an action. In such contexts, apparent mistakes may in fact be optimal if information acquisition is costly. But when can an observer conclude that a mistake was not optimal? I investigate the causes of mistakes in a novel online experiment with an intent-to-collect information stage. I test the axioms of the costly information acquisition model of Caplin and Dean (2015) and estimate within-subject bounds on attention costs. Optimal inattention explains most observed mistakes, but same-day dynamic inconsistency also appears to be a cause of mistakes.

[*]Thank you to Henry Schneider and David Freeman for supervision and financial support while I was writing and conducting research, respectively. Thank you to Mark Dean for helpful suggestions on an earlier draft. Thank you to conference participants at the 2022 Econometric Society North America conference, the 2022 Canadian Economics Association conference, the 2021 Economic Science Association Global conference, Simon Fraser University seminar participants, and my peers from iFree 2020 for helpful comments. All errors and omissions are my own. Research conducted while at Simon Fraser University Department of Economics, currently with Smith School of Business, Queen's University. Email: kevinlaughren@gmail.com

1 Introduction

Gathering information can be costly, yet it helps inform decision-makers (DMs) about available options. If the cost of gathering information exceeds its benefit, it is optimal for a DM to make choices under incomplete information. For example, overpaying is optimal when the expected cost of locating a lower price exceeds the potential savings.

The theoretical literature on attention describes the process by which DMs gather information and make choices. Two central papers in this literature are Sims (2003) and Caplin and Dean (2015), which introduce generalized models of rational inattention. In these models, a DM chooses whether to collect or not collect information based on the relative costs and benefits of applying attention to collect information. While the attention choice is optimal ex ante, it may lead to "mistakes" in the actual choice ex post. In line with the literature, I use "mistake" to denote an action that does not maximize a DM's expected payoff under complete information, acknowledging that the ex post mistake may be ex ante optimal under incomplete information.

The objective of this paper is to identify the various causes of mistakes, classifying them as ex ante optimal or suboptimal in the process. I conduct an online experiment in which I divide the choice process into three stages: a participant's initial intent-to-collect information, their actual information collection, and their application of that information to choose an action. Analyzing where in the choice process mistakes occur enables me to classify them as ex ante optimal or suboptimal, and determine the causes for each.

This paper contributes to the literature on costly attention by demonstrating the benefit of observing intent-to-collect information and the role of dynamic inconsistency in generating mistakes. In an online experiment, I separate the typical theoretical information collection stage into distinct intent-to-collect and actual collection stages, providing more granular choice process than data than was available in earlier experiments. The experiment allows participants to deviate from theoretically predicted behavior, allowing me to observe suboptimal mistakes at each of the three stages of the choice process. These mistakes can appear in the:

1. Intent-to-collect stage: Costly inconsistencies can occur across choices, for example

suboptimal inattention occurs when a participant expresses an intent-to-not-collect for one prize value while also expressing an intent-to-collect for all smaller prize values.[1]

2. Collection stage: *dynamic inconsistency* occurs when a participant collects less information than their expressed intent-to-collect.

3. Prediction stage: *ineffective attention* occurs when a participant collects enough information to identify the optimal action but fails to choose it.

Although a single participant mistake can always be attributed to optimal inattention by assuming sufficiently high attention cost, a set of choices made by a participant implies bounds on their attention cost. These bounds can be used to rule out optimal inattention by testing whether a participant's choices align with the axioms of the costly attention model Caplin and Dean (2015), which are necessary and sufficient conditions for a set of observed choices to be generated by optimal decision-making under generalized attention costs and information benefits. These conditions are further explained in Section 2.

The experiment consists of 10 paid rounds of a state-prediction task, with each round featuring two states that are equally probable ex ante, {R = red circle, B = blue triangle}. If a participant correctly predicts the state, they win a monetary prize. Participants can randomize their prediction immediately and incur no time cost, or they can incur time cost to collect information that will reveal the state. The experiment provides information that is structured as a Poisson process: participants who choose to collect information must pay attention by clicking a button every 10 seconds until the state is eventually revealed. Participants can stop collecting information at any time. A participant who collects some information but stops before revealing the state incurs unnecessary time costs, and such behaviour is dynamically inconsistent if the marginal cost of collecting information is constant across time. In Section 5, I relax the constant time-cost assumption and confirm dynamic inconsistency as a much more convincing explanation for early stopping than convex time costs.

[1] *Suboptimal attention* also occurs at this stage but is not a focus of this paper. Suboptimal attention occurs, for example, when a participant expresses intent-to-collect for one prize value while also expressing intent-to-not-collect for all larger prize values. Suboptimal attention cannot cause a mistake, but it may precede a mistake made in a later stage.

Participants make an information-collection choice for 20 prize values in the first stage of each round. In each round, one of the 20 prize values is randomly selected as the choice-that-counts, and a participant's first stage choice for that prize value is implemented. With these choices I estimate bounds on a participant's attention costs and use these bounds to classify information collection choices as ex ante optimal or suboptimal. I identify optimal inattention as the cause of 73% of observed mistakes.

The second stage is the actual information collection. Every 10 seconds, a prompt is randomly drawn from a finite set of prompts and the state associated with that prompt is displayed. One prompt reveals the state, and all other prompts are useless and do not affect posterior beliefs. Participants can collect this prompt-state information in a table by clicking a button, but failing to click before the next prompt is displayed prevents the information from being collected. Participants can stop collecting information at any time and use this table of collected information to inform their prediction. Dynamic inconsistency occurs when a participant stops collecting information before drawing the relevant prompt despite previously stating an intention to collect information for that prize value. I identify dynamic inconsistency as the cause of 9% of mistakes.

In the third stage, participants can at any time choose to randomize their prediction of the state or use the information collected in the second stage to predict the state. Ineffective attention occurs when a participant has gathered the relevant information in the second stage to make the correct prediction but either chooses to randomize the prediction or makes the wrong prediction.[2] I identify ineffective attention as the cause of 13% of mistakes.

There are two variations of the third stage: the *automatic treatment* and the *manual treatment*. All participants complete five rounds of each treatment, enabling estimation of within-subject treatment effects. In the automatic treatment, participants do not make any predictions of their own and instead the computer chooses optimally on their behalf using the information a participant collected in the second stage table. In the manual treatment, participants are required to make a prediction. Observing a participant's intent-to-collect information across the two treatments allows me to estimate their attention cost as a sum of

[2]A less costly form of ineffective attention occurs when a participant has not collected information that reveals the state but chooses to predict R (or B), forgoing 1pp of expected value by not randomizing.

time cost (which affects both treatments) and cognitive cost of collecting information (which affects only the manual treatment). I find that participants choose to collect information at similar rates across the automatic and manual treatments, suggesting cognitive cost is small relative to time cost. Adjusting for beliefs about mistakes, participants prefer collecting information in the manual treatment.

To summarize the experimental results, participant choice data is highly consistent with the axioms of Caplin and Dean (2015). The model imposes 2300 testable conditions on a participant's choices, and 41% of participants make choices that satisfy all 2300 tests. In fact, for the median participant fewer than 2% of their choices are inconsistent with the axioms, implying that participants make attention choices as if they have a known attention cost function. Moreover, 98 of 101 participants choices are significantly closer to the costly information representation of Caplin and Dean (2015) than randomly generated choices would be. Optimal inattention, suboptimal inattention, dynamic inconsistency, and ineffective attention account for 73%, 5%, 9%, and 13% of mistakes, respectively. Dynamic inconsistency has not been frequently discussed in the costly attention literature, but this experiment makes clear that it should be accounted for even in simple contexts and even if expected attention costs are in the range of one to thirty minutes.

Related literature

Many economic models of decision-making beginning with Sims (2003) focus on the optimal choice of information collection and subsequent action for a DM with a bounded or costly ability to obtain information. Sims' DM makes an information choice modeled as a reduction in Shannon entropy, and this modeling approach has been expanded by Caplin *et al.* (2022) and used by Matejka and McKay (2014) who derive rational inattention microeconomic foundations for the multinomial logit as a model of discrete choice probabilities.[3] Modeling information as Shannon entropy is limited in its generality, and theorists have expanded the literature to consider broader types of attention choices including random attention (Cattaneo *et al.*, 2020; Aguiar *et al.*, 2023), consideration sets (Manzini and Mariotti, 2007, 2014), and

[3]This contrasts McFadden (1973) who had originally founded the multinomial logit model as resulting from i.i.d. unobserved errors on consumer utility.

sparse attention (Gabaix, 2014; Caplin *et al.*, 2018; Enke, 2020). I am agnostic to the specific formulation of attention costs so I build testable conditions using the axiomatic approach of Caplin and Dean (2015), whose general attention cost function $K()$ can incorporate all of these attention cost formulations.[4]

This experiment was inspired by the approach of Dean and Neligh (2023), a working paper of which was available prior to this experiment's design. Dean and Neligh also asked participants to discern a binary state using red and blue elements. My experiment expands on their design because I instead directly measure the intent-to-collect information in a separate first stage. Like the resume study by Bartoš *et al.* (2016), I rely on button clicks to directly measure attention in the second stage, which is an inexpensive alternative to eye-tracking and other costly physiological measures of attention. The separation of the first two stages allows me to identify dynamic inconsistency as a source of mistakes, and my data provides a clear separation between inattention and ineffective attention.

In a lab experiment with similar goals to this paper but a different attention task (stating the number of dots on a screen, knowing there are between 38 and 42 inclusive), Dewan and Neligh (2020) find "evidence on continuity and convexity of costs is mixed", as they "cannot reject that 29 of 42 responsive lab subjects (69.0%) have discontinuities in their response functions", noting "observing a discontinuity in the performance function indicates a violation of convexity". I also find little evidence of convex time cost using choice time stamps and intent-to-collect choices, but I did not set out to test convexity ex ante and so withhold those exploratory analyses for Discussion Section 5.

The paper proceeds as follows. In Section 2 I outline a model of costly attention and derive testable conditions the model of Caplin and Dean (2015) imposes on experimental data. In Section 3 I describe the three-stage experiment design, and in Section 4 I report the results. In Section 5 I discuss two alternative ex post explanations – convex time costs and a preference for choice agency – and conclude.

[4] See also Caplin and Martin (2015) for a Bayesian Expected Utility approach to an NIAS-driven model. This branch of the theoretical literature has continued with Hébert and Woodford (2021, 2023) and Caplin *et al.* (2022), theories published after the design of this experiment that are worthy of experimental economists' attention.

2 A model of costly attention

Consider a DM who has the option to gather information at a cost before making a choice. In many of the aforementioned theories of costly attention this problem is approached as a two-step process: first, the DM chooses an information structure (defined below) at a cost; second, they select the action with the highest expected benefit based on the posterior beliefs formed by the information structure. In this section I apply the notation and approach of Caplin and Dean (2015) to a binary state-prediction task with discrete attention choices and derive testable conditions from their model.

Model primitives

There are two equally-likely states of the world $\omega \in \Omega = \{R, B\}$. I model the choices of a who DM will win a monetary prize P if they correctly predict the state, and this DM is assumed to have a utility function over prizes $u(P) \in \mathbb{R}$ with preferences that permit an expected utility representation. The DM has a correct prior belief $\mu(R) = \mu(B) = \frac{1}{2}$. The DM updates their belief to posterior $\gamma(R)$ using an information structure π that maps from the true state to a distribution of posteriors about the state: $\pi : \Omega \to \Delta\Omega$.

The set of information structures available to the DM at any time is restricted to be binary: $\Pi = \{\pi_0, \pi_L\}$, where π_0 represents the uninformative information structure that returns the prior belief immediately, and π_L represents an information structure that reveals the state at a random time by randomly sampling from a Poisson process with parameter λ_L. This Poisson process is characterized as an opportunity to sample (iid with replacement every 10 seconds) from the set of all prompts of length L, one of which reveals the state and the remaining $2^L - 1$ of which do not affect the posterior.

There are two treatments $T \in \{automatic, manual\}$ that restrict how the DM enters their prediction: the automatic treatment restricts the DMs prediction to that which is implied by their posterior belief (eliminating the possibility of some types of mistakes), but the manual treatment allows for prediction mistakes.

A decision problem A is defined by a prize value P, a prompt length L, and a treatment T. The set of possible decision problems is $\mathcal{F} = \mathcal{P} \times \mathcal{L} \times \mathcal{T}$. The gross benefit of an information structure G maps a decision problem and an information structure to a real value: $G : \mathcal{F} \times \Pi \to \mathbb{R}$. The DM has some expected attention cost $K : \Pi \to \mathbb{R}$, which is unobserved but assumed to be known by the DM. I assume that $K(\pi)$ is weakly increasing in L, as the value of L determines the expected time to observe the state.

After choosing an information structure the DM has three available actions: $a \in \mathcal{A} = \{red, blue, guess\}$. The action-state pair (a, ω) produce a realized utility payoff $a(\omega) \in \mathbb{R}$ with utility value $u(a(\omega))$. The nature of the decision problem implies $u(red(R)) = u(blue(B)) = u(P)$ and $u(red(B)) = u(blue(R)) = u(0)$. A DM who chooses $a = guess$ wins prize P with a probability of 51% regardless of the true state, so $u(guess(\omega)) = u(P)$ with probability 51% and $u(guess(\omega)) = u(0)$ with probability 49% $\forall \omega \in \Omega$. The 51% prize probability implies choosing $a = guess$ is the unique payoff-maximizing action for an expected-utility-maximizing DM whose belief has not changed from the prior.

Costly information representation

The model implies that in each decision problem $A = (P, L, T)$, a DM chooses $\pi \in \{\pi_0, \pi_L\}$ to maximize the ex ante expected net benefit $E\big[G(A, \pi) - K(\pi)\big]$, then forms a posterior $\gamma(R) \in \{0, \frac{1}{2}, 1\}$ and chooses $a \in \mathcal{A} = \{red, blue, guess\}$ to maximize expected utility. The theoretical literature describes the combination of decision problems and choice data as state-dependent stochastic choice (SDSC) data, which assumes that a theorist observes the probability a DM chooses action a in decision problem A when the state is ω for a collection of decision problems $D \subset \mathcal{F}$.

The primary theoretical contribution of Caplin and Dean (2015) is their derivation of two testable conditions that are necessary and sufficient for SDSC data to have a costly information representation. First, each final action choice must be optimal considering the posterior, a condition known as No Improving Action Switches (NIAS, defined in Appendix A). Second, a DM's choices to gather information cannot be restructured across problems to yield greater expected information benefits at no additional attention cost, a condition known as No Improving Attention Cycles (NIAC, defined in Appendix A). If a DM's SDSC data have a costly information representation, the DM is behaving as if they have a known attention cost function and optimally choose information and actions given this function. Those with costly information representations have their choices modeled precisely with the two-stage limited attention models. Any mistake by a DM with a costly information representation must be optimal inattention.

A SDSC data set has a *costly information representation* if the DM:

(i) makes ex ante optimal choice of attention structure π given expected costs and benefits: $\pi \in \arg\max_{\pi' \in \Pi}\{G(A, \pi') - K(\pi')\}$, and

(ii) chooses an action a that is optimal given the posterior formed with observed information: $\sum_{\omega \in \{R,B\}} \gamma(\omega) u(a(\omega)) \geq \sum_{\omega \in \{R,B\}} \gamma(\omega) u(a'(\omega))$ for all $a' \in \mathcal{A}$.

3 Experimental design

Here I introduce a three-stage experiment to collect incentive-compatible SDSC data that permits within-subject tests of the costly information representation outlined in Section 2. The experiment decomposes the step of choosing an information structure into two separate stages: intent-to-collect information and actual information collection. The final stage of the experiment is action selection, which is standard in this literature.

In each round of the experiment participants view a sequence of binary elements, called a *prompt*, and choose whether to gather information before predicting the binary element that follows, called the *state*. The experiment frames this as a chance to predict the binary element (either a red circle or blue triangle) following the prompt, with participants receiving a monetary reward for each correct prediction.

Participants can complete the experiment at their own pace over a three-day period. This design choice provides several advantages in terms of evaluating participants' choices and assessing the value of time saved. By allowing participants to complete the experiment at their convenience, it is expected that they will choose times when their opportunity cost is relatively low, suggesting that the bounds I estimate can be interpreted as lower bounds. Had this experiment been conducted in a lab, the relative benefit of saving time by choosing $a = guess$ might arguably be zero if participants must wait for others to finish before receiving payment and exiting the lab. In this remote setup the perceived value of time saved might be higher as participants can allocate their time savings without lab constraints.

This experiment design enables the identification of four causes of mistakes: optimal inattention, suboptimal inattention, dynamic inconsistency, and ineffective attention. In the first stage, participants may choose not to collect information – this choice is considered *optimal inattention* if it is rational given the attention costs implied by other attention choices, and *suboptimal inattention* if the choice needs to be dropped for the remaining data to satisfy a costly information representation. In the second stage, participants might stop gathering information before acquiring any information that affects beliefs, and I classify this as *dynamic inconsistency*. In the final stage, participants may fail to make an optimal choice despite having previously received a prompt that perfectly reveals the state, which I classify as *ineffective attention*. Figure 1 depicts the three stages from a participant's perspective and my interpretation of mistakes at each stage.

First stage: intent-to-collect information

A participant who chooses the inattentive information structure π_0 is assigned a random prediction immediately that has a 51% chance of winning the prize. Alternatively a participant may choose the attentive information structure π_L to randomly draw prompts of length L and record the states that follow each draw before making a prediction. The gross benefit of choosing π_L instead of π_0 is an increase in the probability of receiving a prize from 51% to 100%. Participants make this binary intent-to-collect choice for each of the 20 different prize levels $\mathcal{P} = \{\$0.25, \$0.50, ..., \$4.75, \$5.00\}$ in each of ten paid rounds with no feedback (i.e., 200 choices).

The expected time to reveal the state is $10(2^L)$ seconds with $L \in \mathcal{L} = \{3, 4, 5, 6, 7\}$. A participant faces each prompt length L twice, once in the manual treatment and once in the automatic treatment. In the manual treatment only, participants report what they believe is their personal probability of ineffective attention ($q(\pi) \in \mathbb{R}$) using an incentive compatible binarized scoring rule Karni (2009); Holt and Smith (2009),[5] and in Section 4 I use these beliefs to provided an estimate of cognitive cost that controls for expected mistakes.

[5]See Schotter and Trevino (2014) for a review of experimental belief-elicitation mechanisms.

Second stage: information collection

In the second stage, the experiment software draws a random prompt with replacement every 10 seconds and reveals the associated state. A participant can collect this information in a data table but must be attentive to do so. The participant has 10 seconds to add this observation to a table using a button click. If they miss this 10 second window, the observation can no longer be recorded. The data table is organized to mirror the prediction task, presenting every possible prompt and counting the frequency of observed red and blue states associated with each prompt. Within a player-round it is impossible to observe both red and blue states for a single prompt so each row of the frequency table has one cell that is fixed at zero observations.[6] There is no time limit on collecting information, but because the state is revealed once the prompt of interest has been observed once there is no benefit to collecting more information.

The prompt length L increases from three to seven elements across rounds, doubling the expected time to reveal the true state with each additional element. If a participant chooses to collect information they are shown a randomly drawn prompt every 10 seconds, with the subsequent binary element revealed. These prompts are randomly selected with replacement from all binary sequences of length L. If a randomly drawn prompt matches the round's relevant prompt, the state is revealed. There is a $\frac{1}{2^L}$ probability the new prompt will update the posterior to $\gamma(R) \in \{0, 1\}$, and a $\frac{2^L-1}{2^L}$ probability the new information leaves the posterior unchanged at $\gamma(R) = \mu(R) = \frac{1}{2}$. Thus, the 'time to receive the prompt revealing the state' follows a Poisson process with parameter $\lambda_L = \frac{1}{10(2^L)}$, with $L \in \{3, 4, 5, 6, 7\}$ implying $\lambda_L \in \{\frac{1}{80}, \frac{1}{160}, \frac{1}{320}, \frac{1}{640}, \frac{1}{1280}\}$. It is not feasible for an optimizing participant to use the provided information in unanticipated ways that improve their gross benefits, as Zhong (2022) proved that a Poisson process is the optimal information structure for a DM in this environment.

When a participant fails to add an observation to the table with a click, there is a discrepancy between the information that has been presented and the information that has been collected. If a missed observation would have revealed the state, the 'true' posterior differs from the 'observed' posterior that is formed using only the recorded observations in

[6]See the third panel of Figure 1 for an example.

the data table. In such cases, I classify participants who could have collected information that revealed the state but failed to do so as demonstrating ineffective attention.

Third stage: translating information to action

The i.i.d. nature of the Poisson information flow restricts the posterior space to three possible beliefs: $\gamma(R) \in \{0, \frac{1}{2}, 1\}$, and $\gamma(R) = \frac{1}{2}$ is the only Bayesian belief until the relevant observation has been drawn because the prior belief is $\mu(R) = \frac{1}{2}$ and all other prompts have zero information about the state. A DM chooses an action $a \in \{red, blue, guess\}$ to maximize $E\Big[G(A, \pi|\gamma)\Big]$. Since there are only three plausible posteriors and three actions, it is easy to check if a participant chose the optimal action given the available information.

Implementation

The experiment was conducted entirely online over two sessions beginning 7 April 2021 and 25 May 2021. Participants were recruited through the Simon Fraser University Experimental Economics Lab portal. The experiment took place over three days, starting with a 30-minute live video conference introduction that covered informed consent, instructions, testing, and questions. After the introduction, participants could complete the experiment at any time of their choosing. Participants had access to a short version of the instructions on every screen, and could revisit the complete instructions at any time.

Out of the 104 participants who attended the introduction and received the show-up fee, 101 completed the experiment. Participants in the first session faced prompt lengths that increased by round, and participants in the second session faced prompt lengths that randomly varied by round. As there were no significant order effects between the two groups, I pool their results throughout.

On average, participants earned a variable pay of \$19.91, spending 34 minutes on information collection. A participant who spent no time collecting information would be expected to earn a variable pay of \$6.70 through correct randomized guesses. This means that participants earned an average of \$23.31 per hour while collecting information. Including the show-up fee of \$7 and comprehension quiz pay averaging \$2.30, the average total pay for participants was \$29.21.[7]

I programmed the experiment in Python using oTree (Chen *et al.*, 2016) as the underlying structure to capture data and distribute unique participant links for online participation. I programmed a custom JavaScript interface for the interactive information collection stage.

Data

Participants completed 10 paid rounds, five in each of the automatic and manual treatments, and they do not learn the outcome or payoff of any round until all rounds are complete to limit learning and wealth effects as explanations for changes in behavior across rounds. The 10 paid rounds followed a paid comprehension quiz and two practice rounds with feedback. The median participant scored 5/6 on the comprehension quiz, and answers were provided once the quiz was scored.

Each participant i chooses their choice of information structure $\pi^i \in \{\pi_0, \pi_L\}$ in each decision problem A. There are 10 paid rounds, each round has 20 prize levels requiring an attention choice. One of the 20 choices is randomly selected and paid in each round.

In the automatic treatment, the computer interprets the information collected by a participant and makes the optimal prediction $a^i \in \{red, blue\}$ as soon as the state is revealed. In the manual treatment, participants face the same set of decision problems but must interpret the data and make the choice from $a^i \in \{red, blue, guess\}$ themselves. In the manual treatment only, a participant provides their belief they will make an error in the prediction stage if they are collecting information using π_L, defined as $q^i_L \in [0, \frac{1}{2}]$.

Let $I^i_{PLT} \in \{0, 1\}$ be the indicator variable equal to 1 if and only if participant i chose to pay attention when the prize was P with information structure π_L in treatment $T \in \{automatic, manual\}$. I suppress the subscripts for T and i whenever the context is a single treatment or participant. I refer to a participant's full set of binary information collection choices as their *attention allocation.*

[7] It is difficult to estimate an hourly wage because players were permitted to leave the experiment and return at any time during the three-day window, and some players may have taken long breaks between rounds.

Testable conditions

Observable choices differ from theorists' assumptions about SDSC data in meaningful ways. The utility function $u()$ is not known or directly observed through choice data, and any finite number of choices implies only coarse observation of the choice probabilities in a given decision problem. However, experimental data can be richer than a SDSC in some dimensions because I can directly observe the choice of information structure π in each decision problem and the exact information received. These differences in observables require the costly information representation axioms to be modified for this context. I modify the NIAC and NIAS conditions of Caplin and Dean (2015) into the Optimal Action (Oa) and Optimal Information ($O\pi$) conditions below to account for these differences between SDSC data and choice data from an experiment.

Optimal action (Oa) is a property of each decision problem, and requires that the observed actions are optimal given the posterior beliefs of the DM. In this experiment, there is a unique optimal action for any given posterior. It is necessary that:

$$\begin{aligned} a = guess &\iff \gamma(R) = \frac{1}{2} \\ a = red &\iff \gamma(R) = 1 \\ a = blue &\iff \gamma(R) = 0 \end{aligned} \tag{Oa}$$

If condition (Oa) is not met, a DM has demonstrated ineffective attention. The condition (Oa) is analogous to Caplin and Dean (2015)'s NIAS condition.

Optimal information ($O\pi$) is a property of the entire SDSC data set generated by a DM and requires that there could be no gross payoff improvement by reassigning the chosen information structures across decision problems. The condition ($O\pi$) is analogous Caplin and Dean (2015)'s NIAC condition. In a binary information choice environment, the optimal information condition ($O\pi$) can be checked through a series of pairs of choices. Each pair of choices that have the same prize P or prompt length L offer an opportunity for an information allocation mistake:

$$I_{PLT} \leq I_{P*LT} \quad \forall P, P* \in \mathcal{P} \quad s.t. \quad P < P*; \quad \forall L \in \mathcal{L} \tag{$O\pi1$}$$

$$I_{PLT} \geq I_{PL*T} \quad \forall L, L* \in \mathcal{L} \quad s.t. \quad L < L*; \quad \forall P \in \mathcal{P} \tag{$O\pi 2$}$$

To satisfy ($O\pi 1$), if a DM collected information when the prompt length was L for prize P in treatment T, they must also collect information when the prompt length is L for any higher prizes in that treatment. To satisfy ($O\pi 2$), if a DM collected information when the prompt length was L for prize P in treatment T, then they must also collect information when the prize is P in rounds with shorter prompt lengths $L < L*$, because they have a lower expected time cost.

A DM's choice data has a costly information representation if it satisfies all the conditions jointly - (Oa), ($O\pi 1$), and ($O\pi 2$). By checking the (Oa) and ($O\pi$) conditions in the experimental data, we can determine whether choices can be rationalized by a limited attention model with a known attention cost function. If both conditions hold, we can classify mistakes and gain insights into the nature of the attention costs and the decision-making process of the participants.

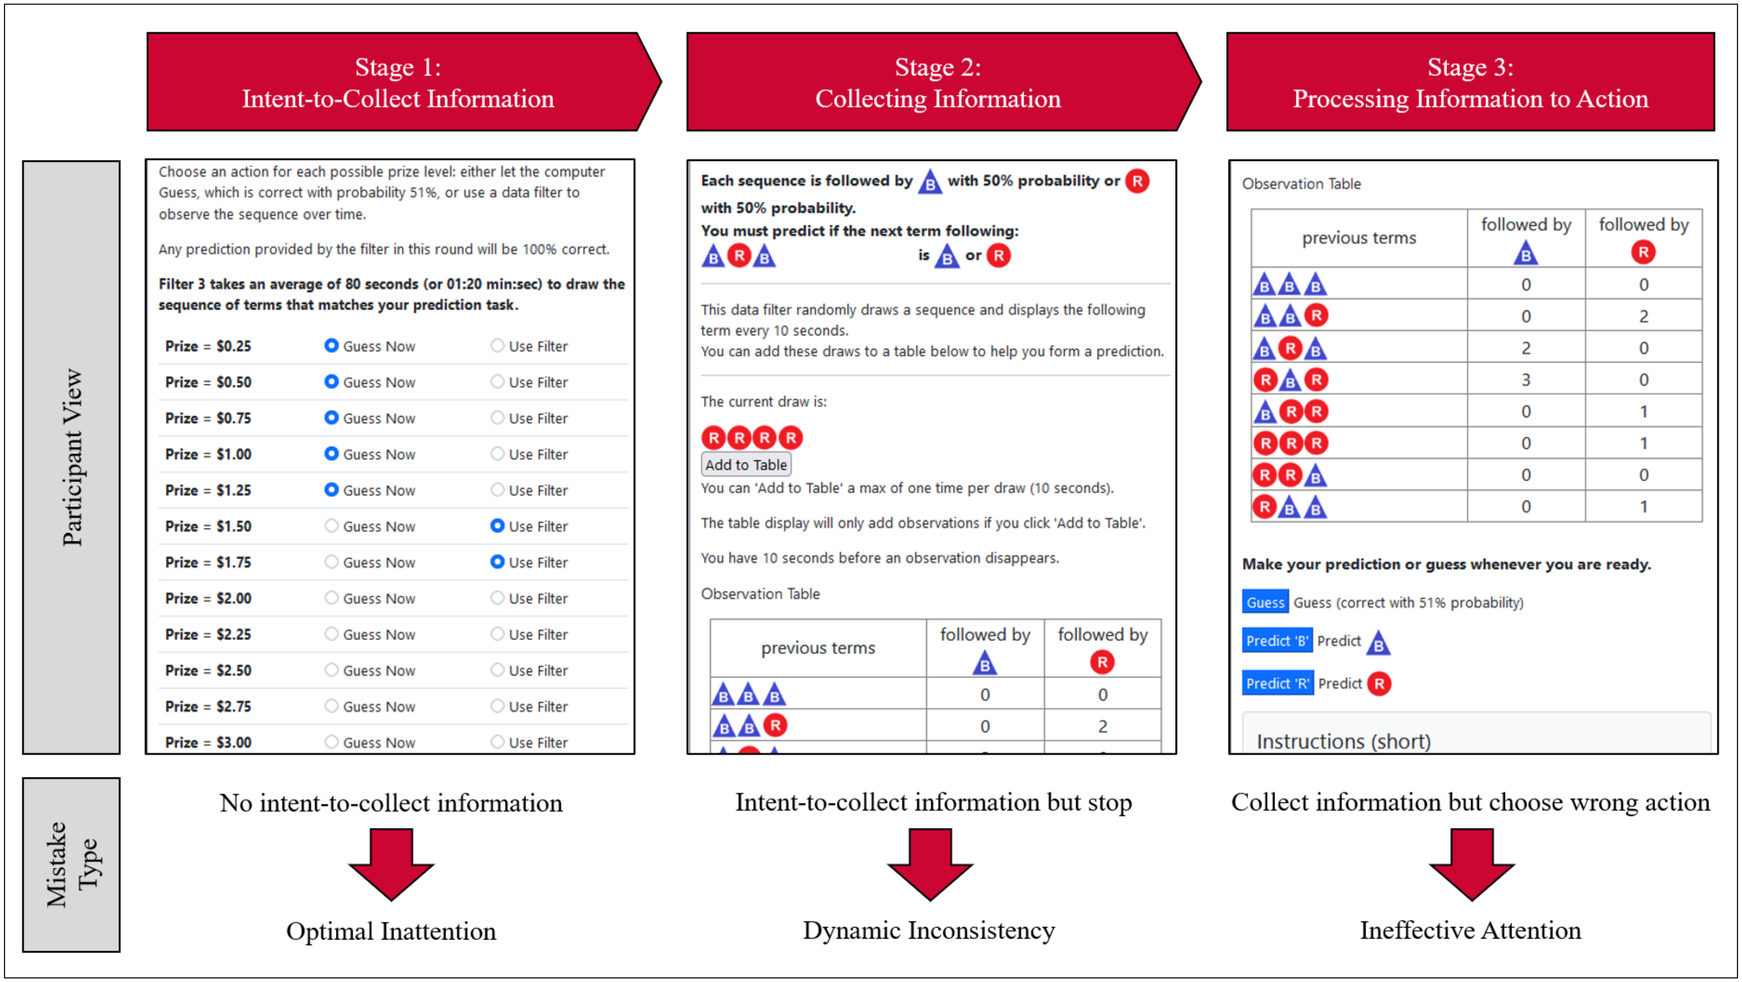

Figure 1: Timeline of Experiment

Notes: This three-stage choice process experiment separates the intent-to-collect information from the actual collection and use of information. The stage in which a mistake occurs reveals the cause of the mistake. The choice in the first row of Stage 1 to "Guess Now" for a \$0.25 prize is an example of optimal inattention (given this participant also chose to "Guess Now" for \$0.50, \$0.75, \$1.00, and \$1.25). If a participant chose to "Guess Now" for the choice that counts in Stage 1, they proceed immediately to Stage 1 of the next round. If a participant chose to collect information by selecting "Use Filter" for the choice that counts in Stage 1, they proceed to Stage 2. Dynamic inconsistency occurs when a participant stops collecting information in Stage 2 before revealing the true state. An example of ineffective attention is when a participant collects information that reveals the true state but fails to make the correct prediction in Stage 3.

4 Results

The attention allocation is my primary tool for analysis, but participants also provide a prediction $a \in \{red, blue, guess\}$ and a decision time t in each round. The experiment records several other timestamps, including when it was first possible to form an informed prediction, when the player actually clicked to collect prediction-relevant information, and the posterior at the player's decision time $\gamma_t(R) \in \{0, \frac{1}{2}, 1\}$.

Figure 2 displays the distribution of choices across all rounds. Note that 532 of 1010 observations would be classified as a mistake under complete information, 381 of these 532 were Stage 1 choices to guess rather than collect information and incur time costs.

Result 1: Participants behave as if they have a known attention cost function. The median participant requires only 2% of their choices to be dropped to satisfy all axiomatic tests, and 41% of participants satisfy all axiomatic tests without dropping any data.

A participant's choices can be rationalized by a costly information representation if they satisfy equations (Oa), ($O\pi 1$), and ($O\pi 2$):

- To satisfy equation (Oa), participants must choose action $a = guess$ (random guess) until the moment the state is revealed, and after that time, they must only make correct predictions.
- To satisfy equation ($O\pi 1$), when participants face prompt length L in treatment T and choose to collect information with prize P, they must also collect information for any greater prize values.
- To satisfy equation ($O\pi 2$), when participants face prize P in treatment T and choose to collect information with prompt length L, they must also collect information for any shorter prompt lengths when the prize is P in treatment T.

Optimal information

A costly information representation implies the optimal information conditions ($O\pi1$) and ($O\pi2$) are satisfied for all possible comparisons in a participant's choice data. Under a null hypothesis that participants choose the binary I_{PLT} randomly and independently with probability $\frac{1}{2}$, a failure of ($O\pi1$) will be generated with probability $\frac{1}{4}$, i.e., when $I_{PLT} = 1$ and $I_{P*LT} = 0$.

There are $\frac{|\mathcal{P}|^2-|\mathcal{P}|}{2}$ opportunities to fail equation ($O\pi1$) for any L, and $\frac{|\mathcal{L}|^2-|\mathcal{L}|}{2}$ opportunities to fail equation ($O\pi2$) for any P. With $|\mathcal{P}| = 20$, $|\mathcal{L}| = 5$, and two treatments, this leaves 1,900 opportunities to fail equation ($O\pi1$) and 400 opportunities to fail equation ($O\pi2$) per participant.[8]

The results show that participants' attention allocation is generally consistent with the conditions outlined in equations ($O\pi1$) and ($O\pi2$), suggesting that the model is suitable for the data in this experiment:

- Participants fail only 1.3% of the tests of equation ($O\pi1$), with 71% of participants failing none of the 1,900 conditions imposed on the data. This condition requires a single switch in the first stage choices from not collecting information to collecting information as the prize P increases *within* a single decision page.

- Participants fail 7.3% of the tests of equation ($O\pi2$), with 58% of participants failing none of the 400 conditions imposed on the data. This condition requires a single switch from collecting information to not collecting information as the prompt length L increases *across* decision pages.

For only three of 101 participants there is a failure to reject a hypothesis of random binomial data generation in favour of a one-tailed alternative hypothesis of fewer rejections, implying that 98 of 101 participants generate choice datasets that are significantly closer to a costly information representation than randomly generated data.[9]

I construct the predictive success measure inspired by Selten, treating each participant's choice data as one observation. This measure takes the proportion of participants whose

[8]190 conditions per prompt length L * 5 values of L * 2 treatments; 10 conditions per price P * 20 values of P * 2 treatments.

[9]The one-tailed test is conducted at a 0.025 significance level.

choice data perfectly satisfy ($O\pi 1$) and ($O\pi 2$) and subtracts the probability a randomly generated set of choice data would satisfy the same conditions Demuynck and Hjertstrand (2019). Overall, 41% of participants satisfy all 2,300 information conditions versus a near zero probability of a random dataset satisfying those conditions. The Selten score of 41% being positive suggests a costly information model has predictive success for the attention allocation data.

Optimal action

The experiment records the exact time when there is sufficient information to make a correct prediction ${t_i}^*$, as well as a participant's decision time t^i and action a_i. Since the random information collecting process is independent across participants, ${t_i}^*$ varies by participant.

For a costly information representation to hold, any observed action must be optimal given the posterior at the time of the action. If a participant makes a choice at a time $t < {t_i}^*$, then they must choose $a_i = guess$ because it yields a prize with a 51% probability, which is higher than the 50% probability obtained with either $a_i = red$ or $a_i = blue$. Once $t > {t_i}^*$, the posterior is either $\gamma(R) = 1$ or $\gamma(R) = 0$, and for these cases (Oa) requires $a_i = red$ and $a_i = blue$, respectively. An (Oa) violation occurs if a participant chooses to make any specific prediction before the state is revealed, and an (Oa) violation occurs if a participant randomizes or makes an incorrect prediction after the state is revealed.

There are two information states that could be used to form a posterior: one that uses all the drawn data (the true information state), and another that only uses the observations that the player clicked to store in a data table. There is only one marginal violation of (Oa) in the entire sample when using the true information state instead of the clicked observations, suggesting that participants were unlikely to use their memory as a replacement for the observation table. I use the true information state throughout.

Regardless of whether an action is chosen before or after the state has been revealed, only one of the three available actions $a \in \{red, blue, guess\}$ is optimal, so if action choices are generated randomly the average participant would fail $\frac{2}{3}$ of the 10 (Oa) conditions. The median participant failed zero of the 10 (Oa) conditions, and the mean number of conditions failed was 0.77, indicating that all failures were concentrated in a minority of participants.

Sixty-two of 101 participants (61%) satisfied all (Oa) conditions. If action choices are generated randomly the probability of one participant satisfying all (Oa) conditions is $(\frac{1}{3})^{10} \approx 0$. In fact 24% of participants passed all $(O\pi)$ and (Oa) conditions. Therefore, the overall Selten score is 24%.

Goodness-of-fit

There is a substantial literature on measuring the variation in observed choice datasets from an axiomatic standard (see Demuynck and Hjertstrand (2019) for a review of measures and recent advances in their computation). One common measure of a dataset is the Houtman-Maks Index (HMI), which measures the maximal proportion of choices that can collectively satisfy every condition of $(O\pi1)$ and $(O\pi2)$ (or analogously drops the fewest observations to have a consistent dataset). In this experiment, I calculate an HMI specific to the attention allocation and another one using all choices, which includes the 10 conditions on paid actions imposed by (Oa).

The median HMI(attention) among participants is 0.99, and the median HMI(all) is 0.98, which indicates that participant choices are highly consistent with a costly information representation. This consistency suggests that the few observations that fail one of the costly information axioms may be interpreted as mistakes. For comparison, in a sample of uniform randomly generated datasets the median HMI(attention) is 0.63.

After computing the HMI for each participant's attention allocation, there is a set of observations that if omitted, leaves a dataset that is fully consistent with $(O\pi)$.[10] If a dropped observation was a choice to collect information ($I_{PLT} = 1$) it is recorded as suboptimal attention, and if a dropped observation was a choice to be inattentive ($I_{PLT} = 0$) it is recorded as suboptimal inattention. Some participants have multiple subsets of the same size that satisfy the costly information representation axioms, one subset that involves dropping a 0 and one that involves dropping a 1. These ambiguous cases are not included in the counts of classified mistakes.

The top row of Figure 3 shows that of the actions categorized as mistakes in the 10 paid

[10]This omitted set is one interpretation of the set of attention allocation violations a participant made ex ante, taking the maximal consistent dataset as the true preference.

rounds, 36 of 532 can be described as errors at the first stage. In these cases a participant made a choice deemed suboptimal relative to the preponderance of their other intent-to-collect choices. Of these 36 attention allocation errors, 24 were cases of suboptimal inattention and 12 were cases of suboptimal attention.

Hereafter when discussing the proportion of mistake causes, I use data from the five manual rounds displayed in Figures 4 and 5. The automatic treatment restricts a participant from incorrectly choosing $a = blue$ or $a = red$, and thus some causes of mistakes are impossible in the five automatic rounds. Only one cause of ineffective attention is possible in the automatic treatment: a participant can fail to click a relevant observation in the 10-second window, then stop collecting information and randomize before that prompt is drawn again.[11] This is distinct from dynamic inconsistency, where a participant starts the information collecting process but stops and randomizes before the relevant observation is ever drawn.

Result 2: Seventy-three percent of mistakes are caused by optimal inattention, but dynamic inconsistency and ineffective attention are also important, causing 13% and 9% of mistakes, respectively.

This experiment was designed to track the information state and the implied posterior at all times, which allows the differentiation between different causes of mistakes. While a costly information representation can rationalize many mistakes as optimal ex ante, it cannot rationalize mistakes resulting from dynamic inconsistency and ineffective attention, which both occur after the choice to collect information.

Any randomized guess or incorrect prediction is characterized as a mistake in this experiment, as either fails to maximize the gross payoff. However, when attention is costly, many ex ante choices to randomize are optimal inattention as they optimally trade off the costs and benefits of attention. If an $a = guess$ mistake is not an (Oa) violation, then a participant's action was the best choice given the information state. If the relevant first stage choice for that round was $I_{PLT} = 0$, and this observation is not among the subset dropped in the process of estimating the HMI, it is classified as being caused by optimal inattention.

[11] This fail-to-click ineffective attention occurred 5 times in the 5 automatic rounds.

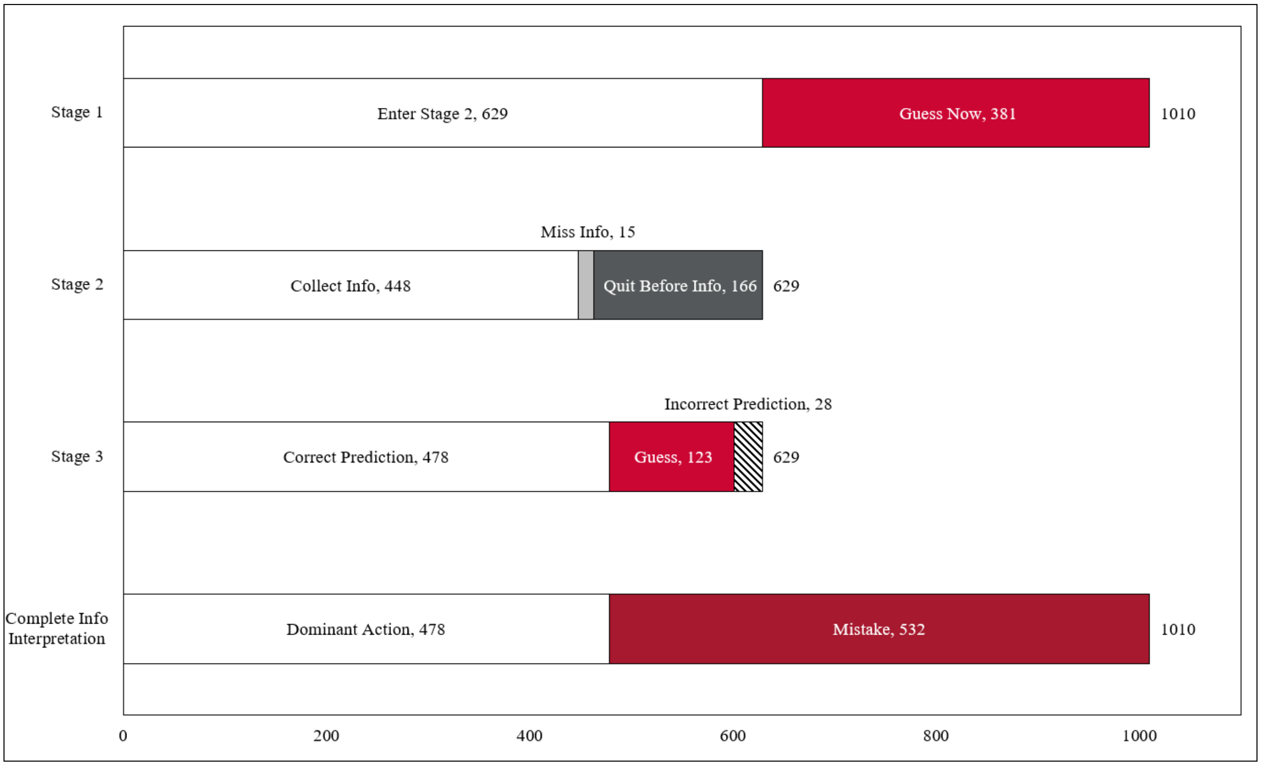

Figure 2: All Choices By Stage

Notes: Choices from 101 participants who complete 10 paid rounds each. Over half (532/1010) of choices would be interpreted as a mistake if complete information is assumed.

However, if observations of $I_{PLT} = 0$ needed to be dropped to form a costly information representation, they are interpreted as suboptimal inattention, as the majority of comparable attention allocation choices imply that participant's benefits of collecting information exceeded their costs. Similarly, if the relevant first stage choice for that round was $I_{PLT} = 1$, that participant's choice can be classified as suboptimal attention or optimal attention based on whether it was included in the largest dataset fully consistent with $(O\pi)$.

If a mistake is also an (Oa) violation, it implies the action was not optimal given the information state. A mistake where a participant failed to interpret the posterior into the optimal action is classified as ineffective attention. There are two types of ineffective attention: one when the state was fully revealed and a participant randomized or made an incorrect prediction, and the other when a participant made an incorrect prediction before receiving information that would change their prior.

The remaining cause of mistake has a participant choose to collect information but then choose to randomize in the second stage before any information is revealed. This is a form

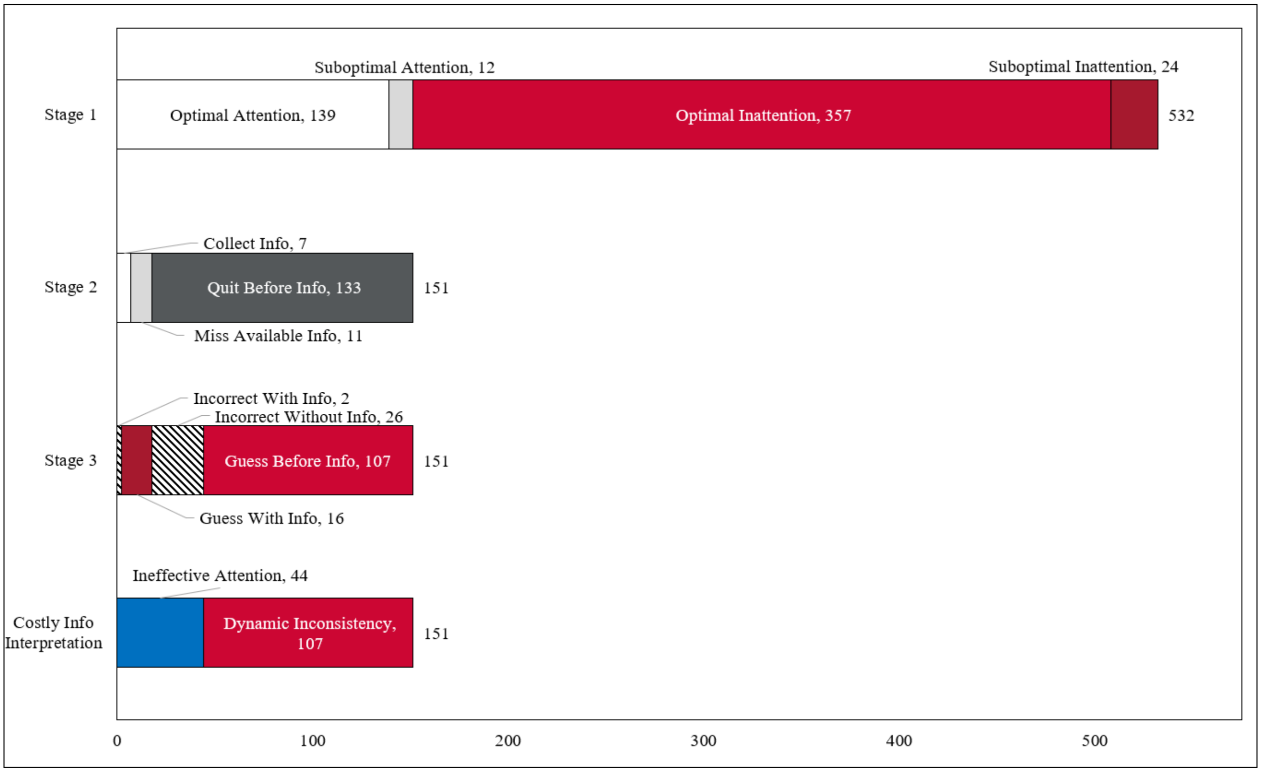

Figure 3: Classifying Mistakes: All Choices

Notes: Of the 532 mistakes across all rounds, two-thirds (357/532) are classified as optimal inattention and 5% (24/532) are classified as suboptimal inattention. Of the remaining 151 mistakes, 107 were cases of participants choosing to enter Stage 2 but then choosing to guess before information arrived, interpreted as dynamic inconsistency. The remaining 44 mistakes were cases of participants who made an incorrect prediction or who guessed when they had sufficient info to make a prediction, interpreted as ineffective attention.

of dynamic inconsistency but is not an (Oa) error. These participant's previous self wanted to collect information, but when they reached the information collection stage they stopped collecting early. This experiment was not designed to test whether a participant faced an unanticipated shock to time costs during the second stage, or if they were previously naive about their future self, or if there was some other cause of their dynamic inconsistency. But it is a distinct cause of mistakes worth separating from ineffective attention.

Figure 5 shows that 199 of 272 mistakes (73%) in the manual treatment can be explained as optimal inattention. Ineffective attention accounts for 36 of 272 mistakes (13%), these are observations where a participant chose to collect information but then failed to choose the best action given the information received. Another 23 of 272 mistakes (9%) can be explained as dynamic inconsistency, these are observations where a participant chose to collect information

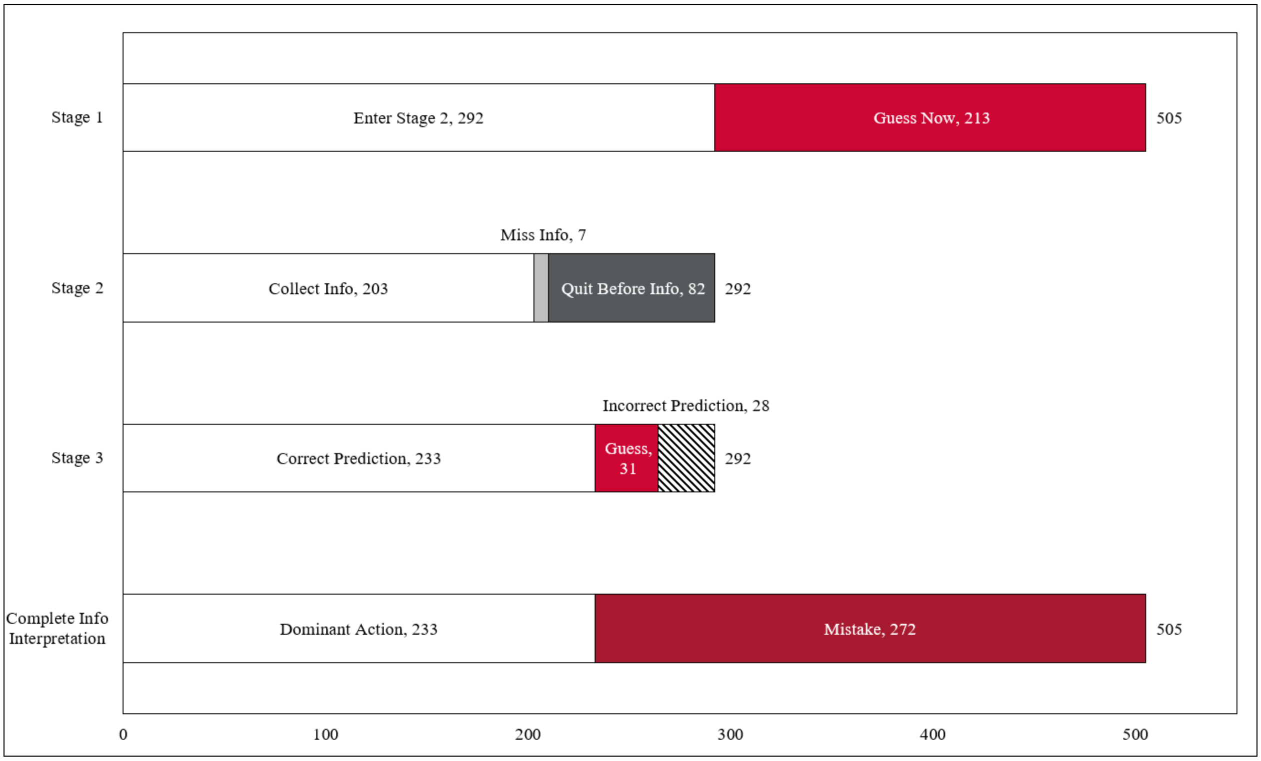

Figure 4: Choices By Stage: Manual Treatment Rounds

Notes: Choices from 101 participants who complete 5 paid manual rounds each. Over half (272/505) of choices would be interpreted as a mistake if complete information is assumed.

in the first stage and later chose to randomize with $a = guess$ after incurring some time cost but having collected no relevant information. Only 14 of 272 mistakes (5%) are driven by a violation of $(O\pi)$ in the attention allocation, these are observations where a participant made several choices to collect information for lower prizes and one difficult-to-explain choice to randomize for a larger prize, interpreted as suboptimal inattention.

Result 3: Participants make low-cost mistakes. Observed mistakes are only 30% as costly as mistakes would be for randomly generated choices.

While any non-optimal action is an (Oa) violation, their expected costs take three different values, measured as a change in prize probability. If a participant chooses $a = blue$ or $a = red$ when the posterior $\gamma(R) = \frac{1}{2}$, her probability of winning a prize decreases from 51% to 50%, for an expected utility loss of $0.51u(P) - 0.50u(P) = 0.01u(P)$, denoted 1pp. If a participant chooses $a = guess$ when $\gamma(R) = 1$ or $\gamma(R) = 0$, her probability of winning a prize decreases from 100% to 51%, for a loss of 49pp. Finally, if a participant chooses

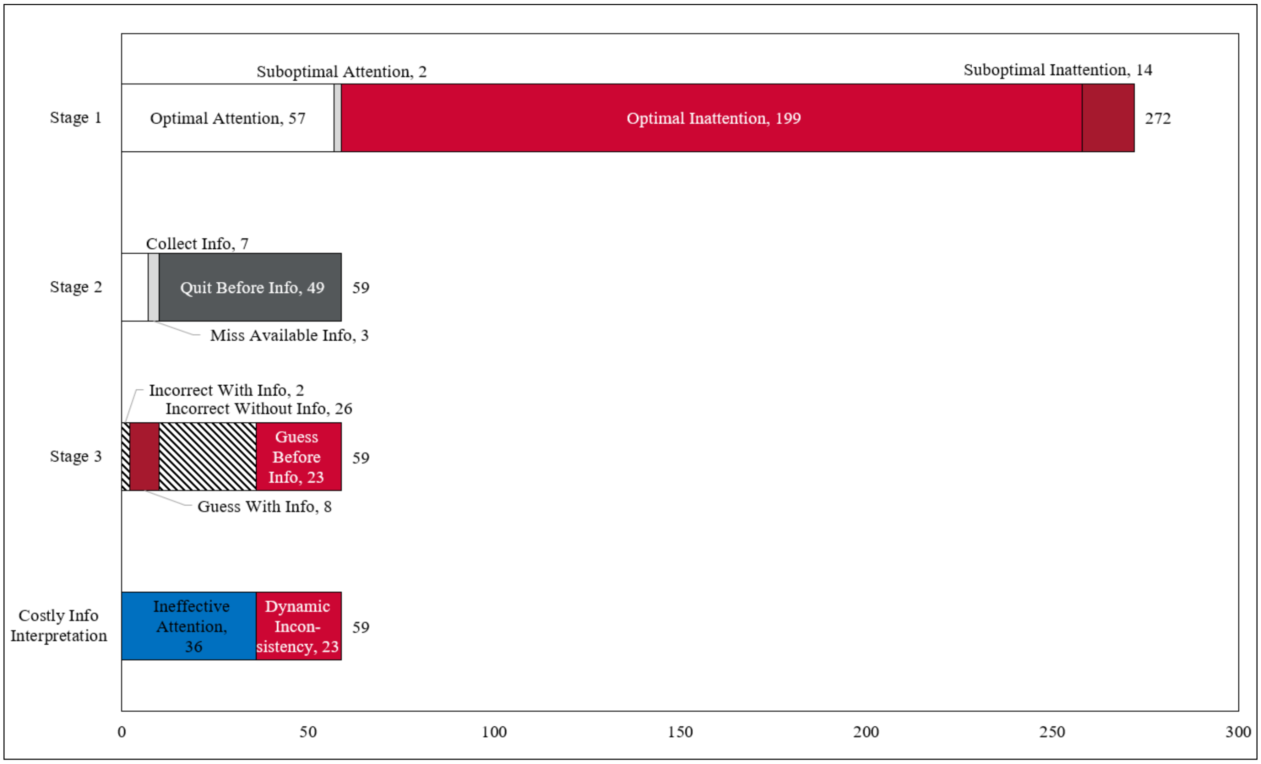

Figure 5: Classifying Mistakes: Manual Treatment Rounds

Notes: Of the 272 mistakes within manual treatment rounds, 73% (199/272) are classified as optimal inattention and 5% (14/272) are classified as suboptimal inattention. Of the remaining 59 mistakes, 23 were cases of participants choosing to enter Stage 2 but then choosing to guess before information arrived, interpreted as dynamic inconsistency. The remaining 36 mistakes were cases of participants who made an incorrect prediction or who guessed when they had sufficient info to make a correct prediction, interpreted as ineffective attention.

$a = red$ when $\gamma(R) = 0$ or $a = blue$ when $\gamma(R) = 1$, her probability decreases from 100% to 0%, for a loss of 100pp. Experiment participants make less costly errors than a set of errors generated by a random dataset. Mistakes costing 1pp account for 75% of participant (Oa) violations, 22% of (Oa) violations are the 49pp variety, and only 3% of (Oa) violations are 100pp mistakes. The mean participant (Oa) violation costs 15pp, versus an average cost of 50pp from a randomly generated dataset. Assuming linear prize utility, this implies that the average participant violation of (Oa) costs a minimum of \$0.39 in expectation while a randomly generated violation costs more than three times as much, with a minimum of \$1.31.[12]

[12]This minimum is based on the unconditional expectation of prize value, but a player in the attention stage knows that none of the prize values for which they chose to randomize were chosen for payment, so the expected prize conditional on reaching the attention stage is higher for participants who pay less attention

It is difficult to assign a monetary value to the cost of a mistake caused by dynamic inconsistency because rather than failing to maximize the probability of a prize, a participant failed to minimize their attention costs for a given prize probability. These participants enter the information collecting stage and incur at least 10 seconds of time cost but ultimately choose to randomize, which they could have done in the first stage with no time cost incurred.

Result 4: The majority of attention cost is time cost. Estimated cognitive costs are near zero. Belief-adjusted cognitive costs are negative, implying participants have a preference to manually make a choice. I interpret a participant's reaction to attention costs under an assumption that expected attention cost $K(\pi)$ can be separated into time cost $T(\pi)$ and cognitive cost $C(\pi)$, so that $K(\pi) = T(\pi) + C(\pi)$. The automatic treatment suppresses the cognitive requirement as much as possible as a participant needs only to click a button every 10 seconds and the computer will eventually make a correct prediction and win a prize with certainty. I normalize whatever cognitive requirement remains in the automatic treatment to $C(\pi) = 0$, and interpret the reaction to expected time across rounds.

In the automatic treatment, for any prize P and prompt length L, a participant with a costly information representation will choose to randomize ($\pi = \pi_0$ and $a = guess$) when:

$$\frac{1}{2}u(P) - K(\pi_L) < 0 \tag{1}$$

Likewise, a participant with a costly information representation will choose to make an informed prediction ($\pi = \pi_L$ and $a \in \{red, blue\}$) when:

$$\frac{1}{2}u(P) - K(\pi_L) > 0 \tag{2}$$

In the automatic treatment time cost $T(\pi)$ can be estimated by observing the switching point between inattention and information collection as the prize level (P) increases or the prompt length (L) decreases. In the manual treatment, we can estimate the overall attention cost ($K(\pi)$) by analyzing the same switching point. By assuming $K_i(\pi) = T_i(\pi) + C_i(\pi)$,

at lower prize values.

I calculate the implied within-subject cognitive cost as the difference between the attention cost in the manual treatment and the time cost in the automatic treatment. I normalize $K(\pi_0) = 0$ so the cost of being inattentive is zero.

To better contextualize time costs I convert them into dollar values by assuming that the utility function is linear over the small prize range considered in the study.[13]

The results show that the majority of attention cost is time cost, with average cognitive costs near zero. Estimated time costs in the first stage are negatively correlated with time spent collecting information in the second stage, suggesting that participants are aware of and responsive to time costs ($\rho = -0.45$, $p < 0.001$). Additionally, belief-adjusted cognitive costs are negative, implying a preference to manually make a choice rather than relying on the automatic treatment.

A costly information representation implies a cutoff strategy over prizes for any given prompt length and treatment; the cutoff prize value where a participant begins to collect information can be estimated by the sum of all first stage information choices for a given round's attention cost L. This proxy maintains every participant's intended probability of reaching the information collecting stage without selectively dropping any observation, and identifies the cutoff perfectly for those participant rounds when no observations fail ($O\pi$). If $\sum_p I_{PLT} = n$, this proxy is interpreted as a participant being willing to collect information for the n largest prizes and not willing for the $(20 - n)$ smallest prizes. This allows for upper and lower bound identification of attention costs. In utility terms, it implies $\frac{1}{2}u(p_{20-n}) \leq T(\pi_L) \leq \frac{1}{2}u(p_{21-n})\ \forall n \geq 1$, and $\frac{1}{2}u(p_{20}) \leq T(\pi_L)$ if $n = 0$. Assuming $u(P) = P$ over the prize values, I can estimate a dollar value of time cost from a participant's first stage choices in the automatic treatment, and calculate the implied reservation wage.

Figure 6 displays the aggregated estimates of attention and time cost. Average time cost is positive and increasing in expected time. Average cognitive cost is near zero, implying indifference between making ones own prediction in the manual treatment relative to letting the computer predict in the automatic treatment. Participants enter Stage 2 equally often in manual and automatic rounds, implying time is the primary driver of choice. However there is

[13] Rabin (2000) showed that linearity over small prizes is not a substantially stronger assumption than an expected utility representation.

substantial heterogeneity in cognitive cost, with many participants displaying preference for manual predictions over automatic, resulting in a negative-valued cognitive cost. A negative-valued cognitive cost may be interpreted as a preference for choice agency despite the weakly greater probability of mistake.

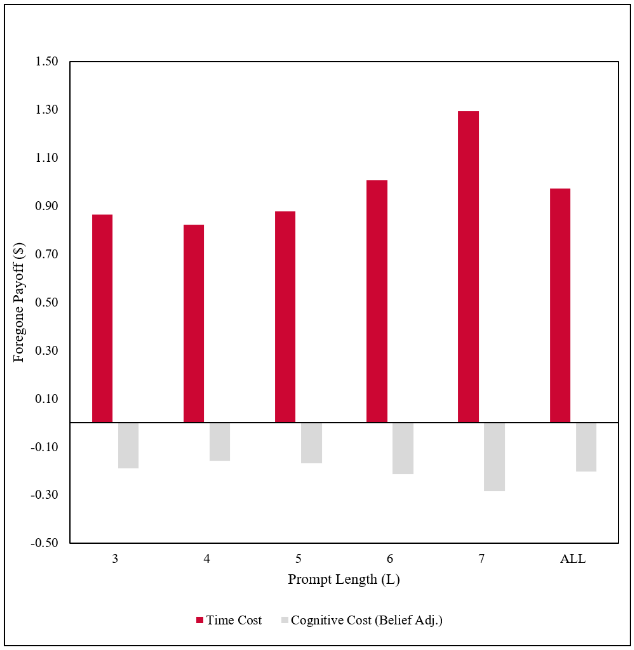

Figure 6: Average Time Costs and Cognitive Costs

Notes: Expected time to reveal the state doubles for each unit increase in prompt length. After adjusting for beliefs about the probability of mistake, the average cognitive cost is negative implying that participants are willing to collect information with smaller expected benefits in the manual treatment than the automatic treatment. All of the following tests were conducted with a T-test at 95% significance level: There are no significant differences in time cost between prompt lengths $L = 3$, $L = 4$, and $L = 5$. Time cost for $L = 6$ is significantly greater than time cost for $L = 5$ and $L = 4$, but I fail to reject the hypothesis of zero difference between $L = 6$ and $L = 3$. Time cost for prompt length $L = 7$ is significantly greater than time cost for all other prompt lengths.

Because there is a non-zero chance of choosing an action that violates (Oa) in the manual treatment, the gross benefit of paying attention in the manual treatment is weakly less than in the automatic treatment, and this reduces the prize probability conditional on choosing to pay attention. In the manual treatment participants report a belief of their own error rate in an incentive compatible mechanism. The naive and adjusted cognitive costs are both displayed in Figure 6. Cognitive costs are small in all cases, zero in the naive case and slightly negative in the belief-adjusted cased. This implies that participants prefer the manual rounds after adjusting the gross benefit to account for the probability that they could make a mistake. This result could be interpreted a preference for choice agency, i.e., a preference to complete the puzzle over mindlessly clicking while waiting for a computer to do so, which I discuss further in Section 5. An equally plausible explanation for this pattern of choices is that participants are indifferent between the two treatments and fail to include their mistake probability in their first stage attention choice. In either case, it is the time cost rather than the cognitive costs acting as the key driver of attention allocation choices.

5 Discussion

5.1 Convex time costs

My interpretation of some mistakes as dynamic inconsistency is founded on an assumption of linear time costs. The experiment by Brown *et al.* (2011) calls into question the validity of this assumption, as they demonstrated that real-time search in a laboratory showed the same "falling reservation wage phenomenon", which had been previously documented in empirical and experimental labor search contexts (Kasper, 1967; Schotter and Braunstein, 1981).Their finding suggests that participants experience convex time cost. Convex time costs could rationalize both a declining reservation wage and the experiment behavior I labeled as dynamic inconsistency. Here I explore the evidence in the experimental data for convex time costs, with special attention on the observations I categorized as dynamic inconsistency (DI) in the main analysis.

Brown *et al.* (2011) conduct a reservation wage experiment where new wage offers arrive at a random time interval determined by a Poisson process. When a new wage offer arrives participants must state a reservation wage before the value of the new wage offer is revealed. Brown, Flinn, and Schotter use treatments with a mix of actual waiting and monetary search costs. Their Wait-No Cost treatment is similar to the experiment in this paper because both have participants wait an uncertain amount of time before receiving information that could improve payoffs via a Poisson process.[14]

In the experiment presented in this paper all time costs are innate instead of incentivized

[14]Their experiment uses the Poisson parameter $\lambda \in [0.05, 0.5]$, whereas my experiment parameter values of L and 10 seconds per draw imply $\lambda_L \in \{\frac{1}{80}, \frac{1}{160}, \frac{1}{320}, \frac{1}{640}, \frac{1}{1280}\}$, a slower Poisson process.

by monetarily discounting a participant payouts as implemented by the Cost treatments of Brown *et al.* (2011). It is reasonable to assume innate time costs are greater in my experiment than in the Wait-No Cost treatment of Brown *et al.* (2011) because my experiment was conducted outside the lab with freedom for participants to stop and return later, so participants face a meaningful opportunity cost in each moment they are collecting information. The lab participants in Brown *et al.* (2011) had to wait in the lab quietly for all participants to complete all searches before receiving payment so opportunity costs are arguably close to zero in the early parts of their experiment environment. Brown *et al.* (2011) found reservation wages declined more quickly for their Wait-Cost treatment than their No Wait-Cost treatment, so their participants did display a response to innate time costs.

Recall the general model of Caplin and Dean (2015) with attention cost $K(\pi)$ and gross attention benefit $G(A, \pi)$ for decision problem A and information structure π. I now assume attention cost to be convex in elapsed time t so that $K_t(t, \pi) > 0$ and $K_{tt}(t, \pi) > 0$. Under these assumptions a DM faces a dynamic (rather than static) problem. At every t, the DM evaluates the expected marginal gross attention benefit of continuing to collect information relative to the marginal attention cost, implying the existence of some $\bar{t} \in \{0, 10, 20, ...\}$ at which it is no longer acceptable to continue collecting information if the state has not yet been revealed (see Appendix C for a complete derivation). For simplicity I use (DI) to denote the observations where a participant stopped the information collection stage before revealing the state, while acknowledging they may not actually represent dynamic inconsistency if time costs are convex.

Assuming experiment participants have a costly information representation where costs are convex in time, two testable comparative statics for the experiment data are:

Comparative statics of convex time costs

1. $mean(t_i|\pi_L, DI) > mean(t_i|\pi_L, non\text{-}DI)$
2. $rate(DI|\pi_L, Round\ r + x) > rate(DI|\pi_L, Round\ r) \quad x \geq 1$

Notes: Conditioning on π_L controls for difficulty and expected time in the information collection stage.

Comparative Static 1 would be apparent in data where stopping information collection (and being classified as DI) was a result of unlucky information collection causing convex time costs to exceed some $\bar{t}$. Comparative Static 1 would be apparent regardless of whether participants are narrow bracketers of time cost across rounds Ellis and Freeman (2020). Figure 7 shows a pattern of DI observations stopping significantly earlier than non-DI observations for all round types. Pooling across rounds, I test a null hypothesis that $mean(t_i|\pi_L, DI) = mean(t_i|\pi_L, non\text{-}DI)$ and find DI and non-DI have significantly different mean with a 2-tailed t-statistic value = 2.61. This is strong evidence against convex time costs because the DI observations quit significantly faster than those who collect enough information to reveal the state, which is the opposite of what is predicted by convex time cost.

Comparative Static 2 would be apparent in data where participants have convex time costs and quit earlier in later rounds as total experiment time accumulates. For accumulated time to affect decisions requires that participants are not narrow bracketers across rounds. Figure 8 shows the rate of DI conditional on entering the information collection stage, and sorts observations by prompt length and instance. The second instance of a prompt length occurs five rounds after the first instance. The data permit a test of the null hypothesis that $rate(DI|\pi_L, Round\ r+x) > rate(DI|\pi_L, Round\ r) \quad x = 5$, because identical prompt lengths were always five rounds apart. There is not a significant difference in the rate of DI in the earlier versus later rounds. Pooling across all 1st instances of a prompt length and all 2nd instances of a prompt length, a Fisher exact test fails to reject the null hypothesis of equal means with a p-value of 0.398. This suggests that convex time costs are not a meaningful driver of quitting the information collection stage.

There is no evidence to believe convex time costs played a significant role in behavior in this experiment context, but convex time costs seem intuitive in many contexts and should continue to be considered especially in more strenuous environments than this online experiment with flexible timing.

5.2 Choice agency

An unanticipated pattern revealed itself in the experiment data: participants quit the information collection stage early (DI) disproportionately more in the automatic treatment than in the manual treatment. For every prompt length $L \in \{3, 4, 5, 6, 7\}$ there were

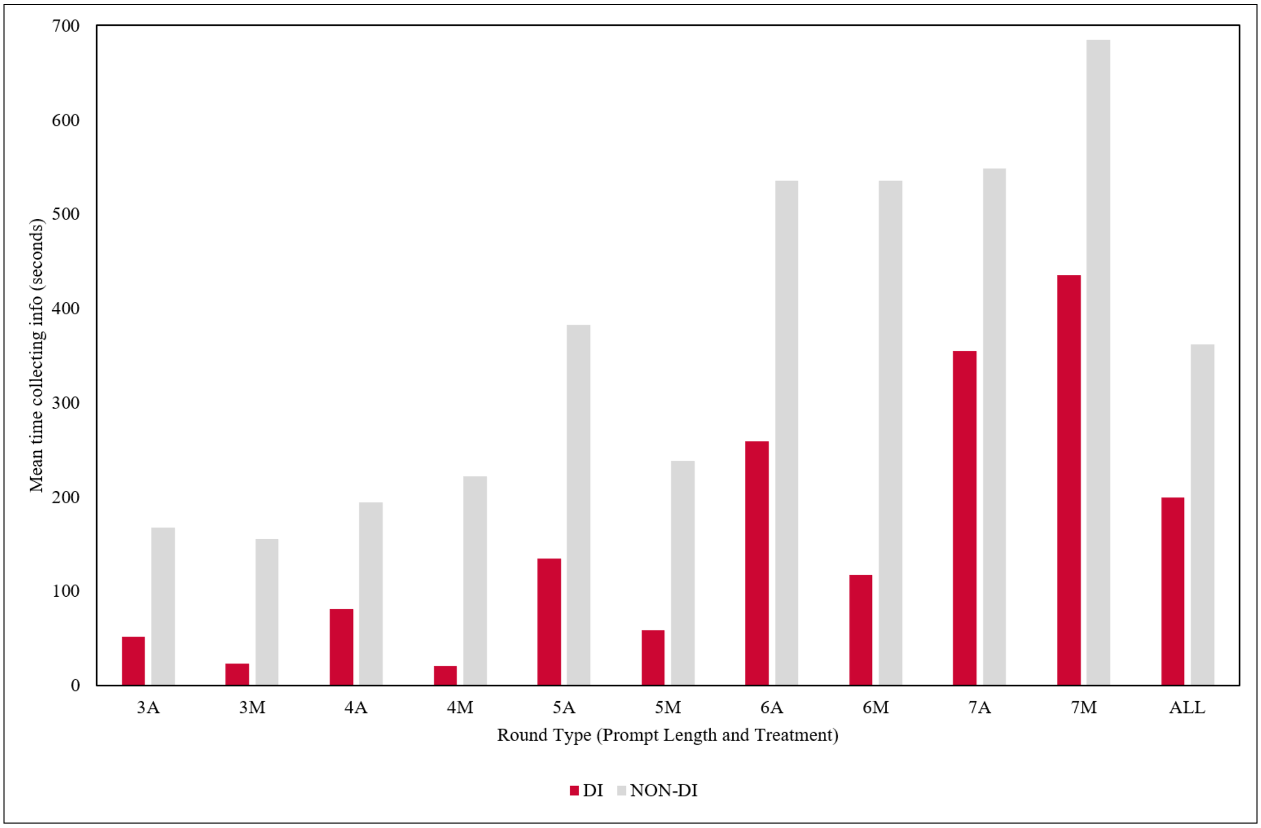

Figure 7: Average Time by Round Type

Notes: Conditional on entering the information collection stage, dynamic inconsistency (DI) observations stop earlier than other observations across all round types. A t-test of null hypothesis that $mean(\bar{t}|\pi_L, DI) = mean(\bar{t}|\pi_L, non\text{-}DI)$ is rejected with a t-value of 2.61, indicating that DI observations stopped significantly earlier than other observations and convex time costs are not likely a driver of DI.

15 or more participants exhibiting DI in the automatic treatment and 8 or fewer participants exhibiting DI in the manual treatment. Figure 9 shows the rate of DI conditional on entering the information collection stage, and sorts observations by prompt length and automatic versus manual treatments. Fisher's exact test rejects a null hypothesis that $rate(DI|\pi_L, automatic) = rate(DI|\pi_L, manual)$ with a p-value < 0.001. Once entering the information collection stage, participants quit much more quickly in the automatic treatment, suggesting they are more willing to collect information if there is an expectation that they use it themselves instead of having a computer automatically make the best prediction.

Only the manual treatment permits a participant to enter an incorrect prediction, so the automatic treatment has weakly greater payoffs than the manual treatment in expectation. Yet participants stop collecting before receiving valuable information twice as frequently in the automatic treatment relative to the manual treatment. It seems participants were more

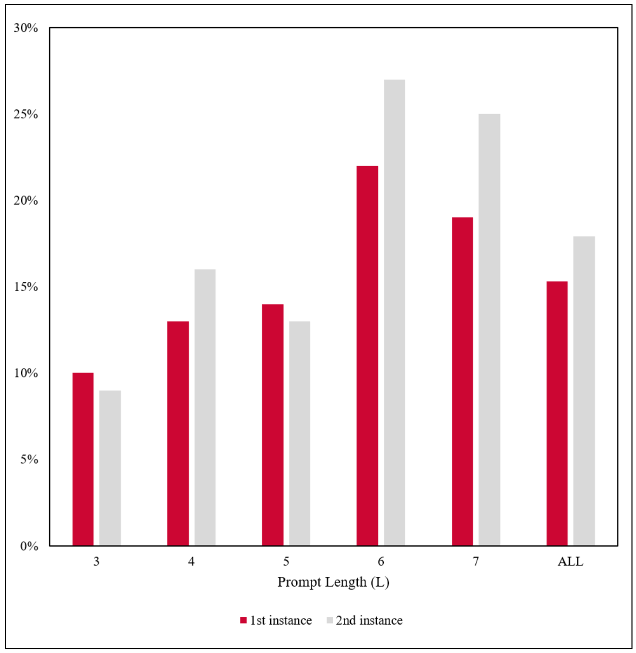

Figure 8: Rate of Dynamic Inconsistency by Prompt Length and Instance
Notes: Observations are grouped by a participant's first or second instance of facing prompt length L. Conditional on entering the information collection stage, the rate of dynamic inconsistency is not significantly different in later rounds of the experiment when controlling for prompt length. A Fisher exact test fails to reject the null hypothesis that $rate(DI|L, \text{1st instance}) = rate(DI|L, \text{2nd instance})$ with a p-value = 0.398.

willing to incur time and attention costs if they were required to make a choice after going through the click-every-ten-seconds gauntlet, even though that choice created the possibility of a payoff-reducing mistake.

Previous research has documented a preference for choice agency in both directions: there are situations where individuals would pay not to choose (Sunstein (2014) gives the examples of medical care and retirement plans), and there are situations where individuals would pay to control their choice (Owens *et al.*, 2014; Freundt *et al.*, 2023). This paper provides a new case of participants paying to control their own choice. Participants in this experiment consistently incur more attention costs in the manual treatment for weakly lower attention benefits, implying a utility increase in the manual treatment relative to the automatic treatment.

Conclusion

I introduce a novel experimental design to isolate different causes of mistakes when attention is costly. I test the extent to which choice data in this experiment satisfy the costly infor-

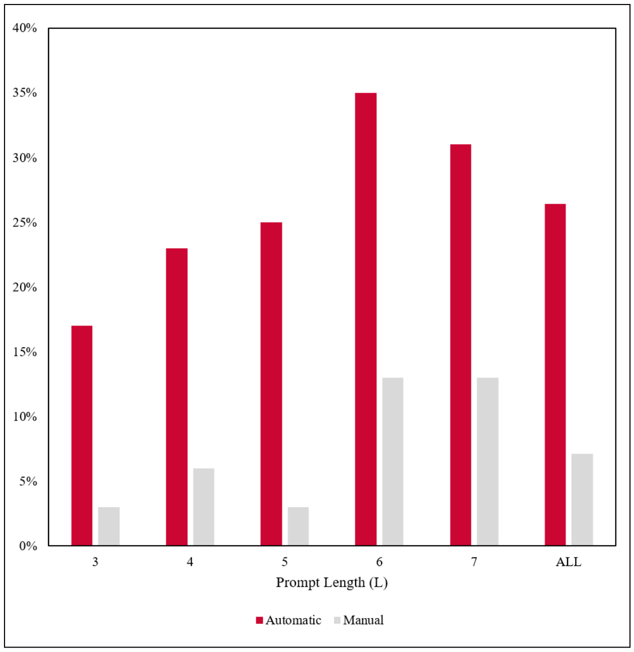

Figure 9: Rate of dynamic inconsistency by round type

Notes: Conditional on entering the information collection stage, the rate of dynamic inconsistency is greater in automatic rounds than in manual rounds across all prompt lengths. A Fisher exact test of the hypothesis that $rate(DI|Automatic) = rate(DI|Manual)$ is rejected with a p-value < 0.001, indicating significantly more dynamic inconsistency in the automatic treatment.

mation representation axioms of Caplin and Dean (2015), and classify choices that would be viewed as mistakes under full information. The results show that 73% of mistakes are due to optimal inattention, 13% are caused by ineffective attention, 9% by dynamic inconsistency, and 5% by suboptimal inattention. These findings suggest that optimal inattention can explain many mistakes, but even in a simple environment a substantial number of mistakes are caused by errors at the information collecting and processing stages. There is no evidence of convex time costs driving participants to quit the information collection stage early, but I did find evidence that the automatic treatment led to more stopping the information collection stage, and this suggests a possible role for choice agency in behavioral modeling. Researchers studying choice domains with significant information requirements should consider modeling or controlling for dynamic inconsistency, as it is an evident cause of mistakes.

References

Aguiar, V. H., Boccardi, M. J., Kashaev, N. and Kim, J. (2023). Random utility and limited consideration. *Quantitative Economics*, **14** (1), 71–116.

Bartoš, V., Bauer, M., Chytilová, J. and Matějka, F. (2016). Attention discrimination: Theory and field experiments with monitoring information acquisition. *American Economic Review*, **106** (6), 1437–1475.

Brown, M., Flinn, C. J. and Schotter, A. (2011). Real-time search in the laboratory and the market. *American Economic Review*, **101** (2), 948–74.

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

—, — and Leahy, J. (2018). Rational inattention, optimal consideration sets, and stochastic choice. *Review of Economic Studies*, **86** (3), 1061–1094.

—, — and — (2022). Rationally inattentive behavior: Characterizing and generalizing shannon entropy. *Journal of Political Economy*, **130** (6), 1676–1715.

— and Martin, D. (2015). A testable theory of imperfect perception. *Economic Journal*, **125** (582), 184–202.

Cattaneo, M. D., Ma, X., Masatlioglu, Y. and Suleymanov, E. (2020). A random attention model. *Journal of Political Economy*, **128** (7), 2796–2836.

Chen, D. L., Schonger, M. and Wickens, C. (2016). otree—an open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, **9**, 88–97.

Dean, M. and Neligh, N. L. (2023). Experimental tests of rational inattention. *Journal of Political Economy*.

Demuynck, T. and Hjertstrand, P. (2019). Samuelson's approach to revealed preference theory: Some recent advances. In R. A. Cord, R. G. Anderson and W. A. Barnett (eds.), *Paul Samuelson Master of Modern Economics*, *Ch.9*, Springer.

Dewan, A. and Neligh, N. (2020). Estimating information cost functions in models of rational inattention. *Journal of Economic Theory*, **187**, 105011.

Ellis, A. and Freeman, D. J. (2020). Revealing choice bracketing. *arXiv preprint arXiv:2006.14869*.

Enke, B. (2020). What you see is all there is. *Quarterly Journal of Economics*, **135** (3), 1363–1398.

Freundt, J., Herz, H. and Kopp, L. (2023). Intrinsic preferences for autonomy. *CESifo Working Paper*.

Gabaix, X. (2014). A sparsity-based model of bounded rationality *. *Quarterly Journal of Economics*, **129** (4), 1661.

Holt, C. A. and Smith, A. M. (2009). An update on bayesian updating. *Journal of Economic Behavior & Organization*, **69** (2), 125–134.

Hébert, B. and Woodford, M. (2021). Neighborhood-based information costs. *American Economic Review*, **111** (10), 3225–55.

— and — (2023). Rational inattention when decisions take time. *Journal of Economic Theory*, **208**, 105612.

Karni, E. (2009). A mechanism for eliciting probabilities. *Econometrica*, **77** (2), 603–606.

Kasper, H. (1967). The asking price of labor and the duration of unemployment. *Review of Economics and Statistics*, pp. 165–172.

Manzini, P. and Mariotti, M. (2007). Sequentially rationalizable choice. *American Economic Review*, **97** (5), 1824–1839.

— and — (2014). Stochastic choice and consideration sets. *Econometrica*, **82** (3), 1153–1176.

Matejka, F. and McKay, A. (2014). Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review*, **105** (1), 272–298.

McFadden, D. (1973). Conditional logit analysis of qualitative choice behavior. *Frontiers in Econometrics*.

Owens, D., Grossman, Z. and Fackler, R. (2014). The control premium: A preference for payoff autonomy. *American Economic Journal: Microeconomics*, **6** (4), 138–61.

Rabin, M. (2000). Risk aversion and expected-utility theory: A calibration theorem econometrica. *September*, **68**, 5.

Schotter, A. and Braunstein, Y. M. (1981). Economic search: an experimental study. *Economic inquiry*, **19** (1), 1–25.

— and Trevino, I. (2014). Belief elicitation in the laboratory. *Annu. Rev. Econ.*, **6** (1), 103–128.

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

Sunstein, C. R. (2014). Choosing not to choose. *Duke Law Journal*, **64**, 1.

Zhong, W. (2022). Optimal dynamic information acquisition. *Econometrica*, **90** (4), 1537–1582.

Appendices

Appendix A Deriving model axioms

This section derives the axiomatic conditions on experimental data - (Oa), $(O\pi 1)$, and $(O\pi 2)$ - using the model and notation of Caplin and Dean (2015).

Each DM action a is a mapping from the state space Ω to the prize space X. Formally, $F = X^{\Omega}$ is the grand set of actions and $\mathcal{F} \equiv \{A \subset F| \ |A| < \infty\}$ is the grand set of decision problems.

In each decision problem A the DM chooses an information structure $\pi : \Omega \rightarrow \Delta(\Gamma)$ which is a stochastic mapping from objective states of the world to subjective signals. Subjective signals are modeled using posterior beliefs $\gamma \in \Gamma = \Delta(\Omega)$. Let $\Gamma(\pi)$ denote the set of possible posteriors when using information structure π. For a given information structure π, $\pi(\gamma|\omega)$ is the probability of posterior belief γ conditional on the state being ω. Denote $\mu(\omega)$ as the prior belief the state is ω, and denote $\gamma(\omega)$ as the posterior belief.

Denote the gross payoff of using information structure π in decision problem A as $G(A, \pi)$. In this experiment context with three actions and two states, this gross benefit is defined as:

$$G(A, \pi) = \sum_{\gamma \in \Gamma(\pi)} \left[\sum_{\omega \in \{R,B\}} \mu(\omega)\pi(\gamma|\omega) \right] \left[\max_{a \in \{red, blue, guess\}} \sum_{\omega \in \{R,B\}} \gamma(\omega)u(a(\omega)) \right] \tag{3}$$

The first term in brackets captures the prior probability of posterior γ given the information structure chosen, and the second term in brackets captures the gross benefit of choosing the optimal action given the posterior.

Caplin and Dean (2015) suppose that a DM has an information cost function $K()$, an attention function π_A which captures the DM's choice of information structure in a given decision problem A, and a choice function $C_A : \Gamma(\pi_A) \rightarrow \Delta(A)$ which maps posterior beliefs to action probabilities.

A state-dependent stochastic choice dataset (D, P) is a collection of decision problems $D \subset F$ and a related set of state-dependent stochastic choice functions $P = \{P_A\}_{A \in D}$ where $P_A : \Omega \rightarrow \Delta(A)$. Denote $P_A(a|\omega)$ as the probability the DM chooses action a conditional on state ω in decision problem A, and $\hat{P}_A(a|\omega)$ its experimental analogue.

Caplin and Dean (2015) state that a state-dependent stochastic choice dataset (D, P) has a *costly information representation* if there exists $K(), \pi_A(), C_A()$ such that for all $A \in D$:

1. Information is optimal: $\pi_A \in \arg\max_{\pi \in \Pi} \{G(A, \pi) - K(\pi)\}$
2. Choices are optimal: Given $C_A(a|\gamma) > 0$:

   $\sum_{\omega \in \Omega} \gamma(\omega)u(a(\omega)) \geq \sum_{\omega \in \Omega} \gamma(\omega)u(b(\omega))$ for all $b \in A$
3. The data is matched: $P_A(a|\omega) = \sum_{\gamma \in \Gamma(\pi_A)} \pi_A(\gamma|\omega)C_A(a|\gamma)$

Caplin and Dean (2015) show a state-dependent stochastic choice dataset (D,P) has a costly information representation if and only if it satisfies the testable criteria *No Improving Attention Cycles* (NIAC) and *No Improving Action Switches* (NIAS). NIAC is a property of the entire dataset (D, P) and requires that there could be no gross payoff improvement by reassigning the chosen information structures across decision problems. NIAS is a property of each decision problem, and requires that the observed actions chosen are optimal given the revealed posterior beliefs of the DM.

NIAS violations

NIAS is a property of each decision problem, and requires that the observed actions chosen are optimal given the revealed posterior beliefs of the DM. Given $\mu \in \Gamma, A \in D, P_A \in P$, and $a \in Supp(P_A)$, the revealed posterior $\bar{\gamma}_A^a \in \Gamma$ is defined by:

$$\bar{\gamma}_A^a(\omega) \equiv Pr(\omega | a \text{ chosen from } A)$$
$$\frac{\mu(\omega)P_A(a|\omega)}{\sum_{\nu \in \Omega} \mu(\nu)P_A(a|\nu)}$$

One formulation of the NIAS condition is given by the equation labeled (3) in Caplin and Dean (2015): For any $A \in D, a \in Supp(P_A)$, and $b \in A$:

$$\sum_{\omega \in \Omega} \bar{\gamma}_A^a(\omega)u(a(\omega)) \geq \sum_{\omega \in \Omega} \bar{\gamma}_A^a(\omega)u(b(\omega)) \tag{4}$$

A violation of Condition (4) in an experimental dataset implies that the DM failed to choose the optimal action given their posterior beliefs γ in some decision problem A. One example of such a NIAS violation is when a posterior belief is strong enough that the optimal action is a prediction (e.g., if $\gamma(R) = 1$, the optimal $a = red$) and the DM is observed to choose one of the suboptimal actions $a = blue$ or $a = guess$. I interpret such a violation of NIAS as *ineffective attention*; the choice of π may have been ex ante optimal, but the DMs use of the information was not effective.

NIAC violations

NIAC is a property of the entire dataset and requires that there could be no gross payoff improvement by reassigning the chosen information structures across decision problems.

Caplin and Dean (2015) define the NIAC as follows: Given μ and $u : X \to \mathbb{R}$, dataset (D, P) satisfies NIAC if, for any set of decision problems $A^1, A^2, ..., A^J \in D$ with $A^J = A^1$:

$$\sum_{j=1}^{J-1} G(A^j, \bar{\pi}_{A^j}) \geq \sum_{j=1}^{J-1} G(A^j, \bar{\pi}_{A^{j+1}}) \tag{5}$$

A violation of the NIAC condition (5) in an experimental dataset would suggest that the DM failed to allocate attention efficiently across decision problems.

Costly information representation: translating to this experiment context

In the simple context of this experiment the second term in (3) in brackets is equal to $u(P)$ when the state is known ($\gamma(R) \in \{0, 1\}$) and is equal to $\frac{1}{2}u(P)$ when the state is not known ($\gamma(R) = \frac{1}{2}$). The space of posteriors is restricted by the two information structures (π_0, π_L in each decision problem: $\Gamma(\pi_0) = \{\frac{1}{2}\}$, $\Gamma(\pi_L) = \{0, \frac{1}{2}, 1\}$ with $\pi_L(\gamma = \frac{1}{2}|\omega) \to 0$ as the information sample gets large. If a participant enters the attention stage with an intention of waiting for a signal that changes the posterior away from the prior $\mu(R) = \frac{1}{2}$, the gross attention benefit of any decision problem A is drastically simplified to:

$$
\begin{aligned}
G(A, \pi_L) &= u(P) \text{ if } \pi = \pi_L \\
G(A, \pi_0) &= \frac{1}{2}u(P) \text{ if } \pi = \pi_0
\end{aligned}
$$

By considering π_0 (randomized guessing) as the default action, the gross *marginal* benefit of paying attention in this context is:

$$
G^*(A, \pi_L) = \frac{1}{2}u(P) \tag{6}
$$

A DM has a *costly information representation* (Caplin and Dean, 2015) if they both (i) make ex ante optimal choices of attention structure given expected costs and benefits ($\pi_A \in \arg\max_{\pi\in\Pi}\{G(A, \pi) - K(\pi)\}$), and (ii) their chosen action a is optimal given their observed information ($\sum_{\omega\in\{R,B\}} \gamma(\omega)u(a(\omega)) \geq \sum_{\omega\in\{R,B\}} \gamma(\omega)u(b(\omega))$ for all $b \in A$). If a participant behaves in this way, any 'mistake' must be case of optimal inattention, because (i) rules out suboptimal inattention and (ii) rules out ineffective attention. These conditions are difficult to test directly because the attention cost $K()$ is unobserved. However in this experiment context these conditions can be simplified to those used in the main text:

$$
\begin{aligned}
a = \textit{guess} &\iff \gamma(R) = \frac{1}{2} \\
a = \textit{red} &\iff \gamma(R) = 1 \\
a = \textit{blue} &\iff \gamma(R) = 0
\end{aligned} \tag{Oa}
$$

$$
I_{PLT} \leq I_{P*LT} \quad \forall P, P* \in \mathcal{P} \quad s.t \quad P < P*; \quad \forall L \in \mathcal{L} \quad \forall T \in \{\textit{automatic}, \textit{manual}\} \tag{$O\pi1$}
$$

$$
I_{PLT} \geq I_{PL*T} \quad \forall L, L* \in \mathcal{L} \quad s.t \quad L < L*; \quad \forall T \in \{\textit{automatic}, \textit{manual}\} \tag{$O\pi2$}
$$

Appendix B Automatic treatment results

Figure 10: Choices By Stage: Automatic Treatment Rounds

Notes: Choices from 101 participants who each complete 5 paid rounds of the automatic treatment. Over half (260/505) of choices would be interpreted as a mistake if complete information is assumed. An incorrect prediction in Stage 3 is not possible in the automatic treatment.

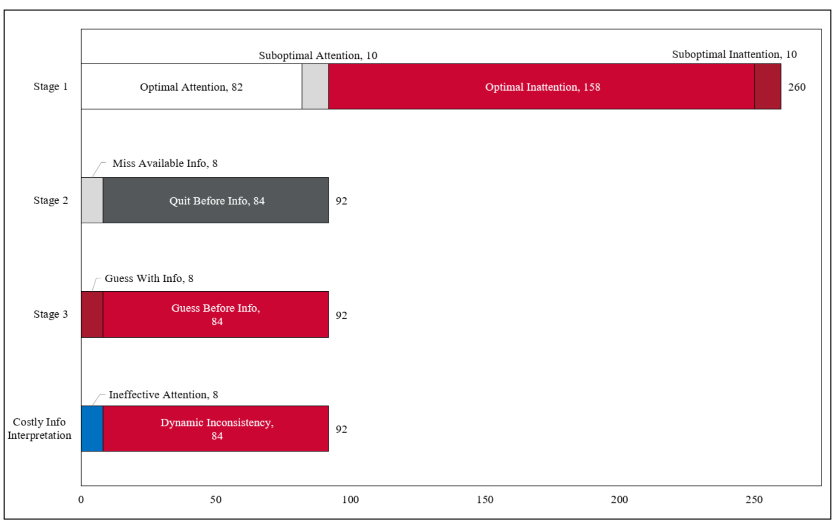

Figure 11: Classifying Mistakes: Automatic Treatment Rounds

Notes: Of the 260 mistakes within automatic treatment rounds, 61% (158/260) are classified as optimal inattention. Of the remaining 92 mistakes, 84 were cases of participants choosing to enter Stage 2 but then choosing to guess before information arrived, interpreted as dynamic inconsistency. The remaining 8 mistakes were cases of participants who missed information that would allow the computer to make a correct prediction by failing to click in the 10-second window, interpreted as ineffective attention.

Appendix C Convex time cost comparative statics

Recall the model of Caplin and Dean (2015) with attention cost $K(\pi)$ and gross attention benefit $G(A, \pi)$ for decision problem A and information structure π. I now assume attention cost to be convex in elapsed time t so that $K_t(t, \pi) > 0$ and $K_{tt}(t, \pi) > 0$. Under these assumptions a DM faces a dynamic (rather than static) problem. At every t, the DM evaluates the expected marginal gross attention benefit of continuing to collect information relative to the marginal attention cost. First, consider whether the DM wants to continue for a full 10 seconds after each of $t \in \{0, 10, 20, ...\}$ (as continuing for nine or fewer seconds has a gross attention benefit of zero). The DM continues to collect information if and only if:

$$\begin{aligned}
&EU(continue) &&\geq EU(stop)\\
&\frac{1}{L}\bigg(u(P) - K(t+10, \pi)\bigg) + \frac{L-1}{L}\bigg(\frac{1}{2}u(P) - K(t+10, \pi)\bigg) &&\geq \frac{1}{2}u(P) - K(t, \pi)\\
&\frac{1}{L}\bigg(\frac{1}{2}u(P)\bigg) &&\geq K(t+10, \pi) - K(t, \pi) \quad (7)
\end{aligned}$$

The marginal benefit versus marginal cost condition in Inequality 7 captures the $\frac{1}{L}$ probability that the prompt arriving in 10 seconds reveals the state to improve the prize probability by $\frac{1}{2}$ versus the marginal attention cost that will be incurred in the next 10 seconds.

Note that a DM who has collected information for $t' \in \{1, 2, ..., 9\}$ seconds since the previous prompt at $t \in \{0, 10, 20, ...\}$ chooses to continue if and only if $\frac{1}{L}[\frac{1}{2}u(P)] \geq K(t + 10 - t', \pi) - K(t + t', \pi)$. Compared to Inequality 7, the marginal benefit of the next bit of information is unchanged and the marginal cost of that information is lower at $t + t'$ than at t, so it follows that any DM who chooses to continue at $t \in \{0, 10, 20, ...\}$ also chooses to continue at $t + t'$ for $t' \in \{1, 2, ..., 9\}$. This simplifies the dynamic problem to a choice for each $t \in \{0, 10, 20, ...\}$.

The left-hand side of Inequality 7 is fixed[15] and the right-hand side is increasing in t, implying the existence of some $\bar{t} \in \{0, 10, 20, ...\}$ at which it is no longer acceptable to continue collecting information: $K(\bar{t} + 10, \pi) - K(\bar{t}, \pi) \geq \frac{1}{L}\bigg(\frac{1}{2}u(P)\bigg) \geq K(\bar{t}, \pi) - K(\bar{t} - 10, \pi)$.

Because the amount of time to reveal the state is random, participants incur heterogenous time costs. Costs $K()$ are convex in t, so if this convexity is a key driver of behavior I would expect that a participant who stopped collecting information before revealing the state (classified as DI) did so because they got unlucky draws and breached their $\bar{t}$. Thus, convex time costs imply that $mean(\bar{t}|\pi_L, DI) = mean(\bar{t}|\pi_L, non\text{-}DI)$, which I call Comparative Static 1 and test in Section 5.

If the variable over which $K()$ is convex is *total experiment time* and not *round time*, then I would expect that participants spend less time collecting information in later rounds as more participants breach their $\bar{t}$. Thus, time costs that are convex over total experiment time imply $rate(DI|\pi_L, Round\ r{+}x) > rate(DI|\pi_L, Round\ r) \quad \forall x \geq 1$, which I call Comparative Static 2 and test in Section 5.

[15] until the state is revealed, then the benefit of continuation is zero.

Appendix D Experiment instructions

Participants sign on to a live Zoom meeting to complete an ID check and consent form. Participants click through the screens below on their own devices as the experimenter reads the instructions aloud, takes questions, and offers tech support. The Zoom meeting ends after the instructions and participants have until 23:59 two days later to complete the experiment. Participants can exit the experiment at any time and return with saved progress using a unique link generated by oTree (Chen *et al.*, 2016).

Instructions

This experiment is conducted using your own internet-connected device. You earn a participation fee of $7.00 CAD today.

You can sign in and complete the experiment at any time from 00:01 2021-May-25 to 23:59 2021-May-27. If you exit the browser, you can sign back in with the same link and your progress is saved.

It is possible to earn between $0.00 CAD and $53.00 CAD in additional prizes based on your decisions and an element of chance.

Any additional prizes you earn during this time will be paid by email transfer along with your $7.00 participation fee on 2021-May-29.

Primary Task

To earn additional prizes in this experiment, you must correctly predict whether the next term in a sequence is B = Blue Triangle B or R = Red Circle R.

In each round you will be given one specific sequence - for example: R B B - and asked whether B or R comes next.

Your choices are to (1) Guess, (2) Predict B, or (3) Predict R.

For any sequence you are given, the next term is randomly selected to be R with 50% probability or B with 50% probability. This randomization occurs independently for each participant, round, and sequence.

If you 'Guess', the computer will choose on your behalf and will be correct with 51% probability, so it can be better to choose Guess than predict randomly.

As an alternative to a Guess, you will be given an opportunity to collect information about the sequence using a tool called a **data filter**.

Every 10 seconds, a data filter randomly draws a sequence and displays the single term which followed, just like your prediction task.

Example

You will win a monetary prize if you can predict whether the term following R R R is B or R.

You choose to spend time using data filter 3 instead of guessing immediately.

Every 10 seconds, data filter 3 randomly draws one of the possible sequences made of 3 B or R terms - such as (R B R), (R R B), (R B B), etc. - and displays the randomly selected R or B term that follows.

Eventually, data filter 3 will draw the same 3 term sequence as you are given in your prediction problem (R R R), and display whether the next term is B or R.

Below is a screenshot of data filter 3 recording a draw of R R R followed by B:

('type M' will be explained on the next page)

Round 8: Data Filter 3 type M

Filter Type M: You must make a Prediction or 'Guess'

Each sequence is followed by B with 50% probability or R with 50% probability.
You must predict if the next term in the sequence following: R R R
is B or R

This data filter randomly draws a sequence and displays the following term every 10 seconds.
You can add these draws to a table below to help you form a prediction.

The current draw is:

R R R B

Add to Table

You can 'Add to Table' a max of one time per draw (10 seconds).

The table display will only add observations if you click 'Add to Table'.

You have 10 seconds before an observation disappears.

Observation Table

previous terms	followed by B	followed by R
B B B	0	3
B B R	0	4
B R B	1	0
R B R	0	0
B R R	0	0
R R R	1	0
R R B	2	0
R B B	1	0

Make your prediction or guess whenever you are ready.

Guess Guess (correct with 51% probability)

Predict 'B' Predict B

Predict 'R' Predict R

Every 10 seconds filter 3 selects one of the sequences of length 3 and displays the term that follows. You can store this information in the observation table by clicking 'Add to Table'.

The first row of the observation table shows that B B B has been followed by B zero (0) times, and has been followed by R three (3) times.

The second row of the observation table shows that B B R has been followed by B zero (0) times, and has been followed by R four (4) times.
In this example round, you are being paid to predict what follows R R R, so you will be most interested in the row of the table which says that R R R, has been followed by B one (1) time, and has been followed by R zero (0) times.

Instructions

This experiment has 12 total rounds: Two practice rounds, which are not paid, followed by ten paid rounds.

There is one practice round with filter type A and one practice round with filter type M.

You then play a set of five paid rounds with filter type A and a set of five paid rounds with filter type M. The order of these two sets is random for each participant.

Filter Types A and M

Filter type A (A for 'Automatic') has the computer make a correct prediction on your behalf as soon as you collect enough data.

Filter type M (M for 'Manual') has you enter your own prediction, you can use the filter for as long as you wish before making a prediction.

Example

See Filter type A and Filter type M below.

Filter type A does not have a button for 'Predict R' or 'Predict B'.

In both filters you can click 'Add to Table' to add observations to the table.

In both filters you can 'Guess' at any time to proceed to the next round.

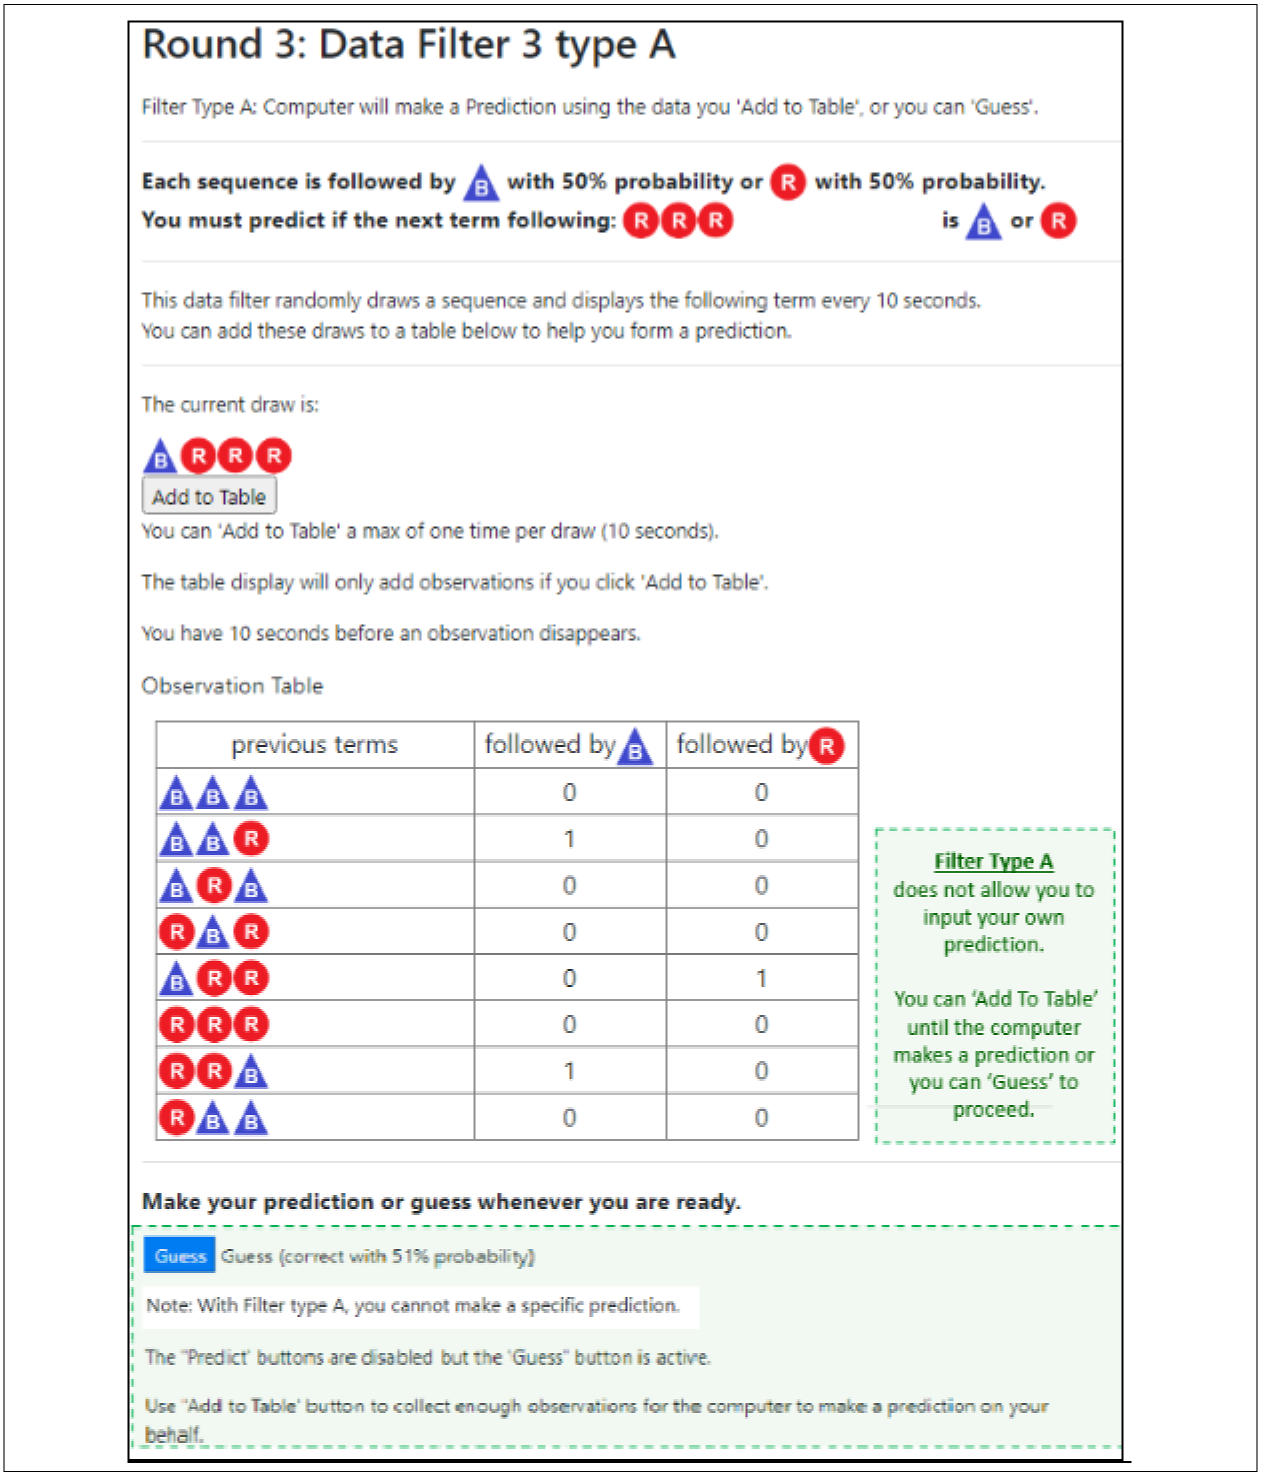
Round 3: Data Filter 3 type A
Filter Type A: Computer will make a Prediction using the data you 'Add to Table', or you can 'Guess'.
Each sequence is followed by B with 50% probability or R with 50% probability.
You must predict if the next term following: R R R is B or R
This data filter randomly draws a sequence and displays the following term every 10 seconds.
You can add these draws to a table below to help you form a prediction.
The current draw is:
B R R R
Add to Table
You can 'Add to Table' a max of one time per draw (10 seconds).
The table display will only add observations if you click 'Add to Table'.
You have 10 seconds before an observation disappears.
Observation Table
previous terms
followed by B
followed by R
B B B 0 0
B B R 1 0
B R B 0 0
R B R 0 0
B R R 0 1
R R R 0 0
R R B 1 0
R B B 0 0
Filter Type A
does not allow you to input your own prediction.
You can 'Add To Table' until the computer makes a prediction or you can 'Guess' to proceed.
Make your prediction or guess whenever you are ready.
Guess
Guess (correct with 51% probability)
Note: With Filter type A, you cannot make a specific prediction.
The "Predict" buttons are disabled but the "Guess" button is active.
Use "Add to Table" button to collect enough observations for the computer to make a prediction on your behalf.

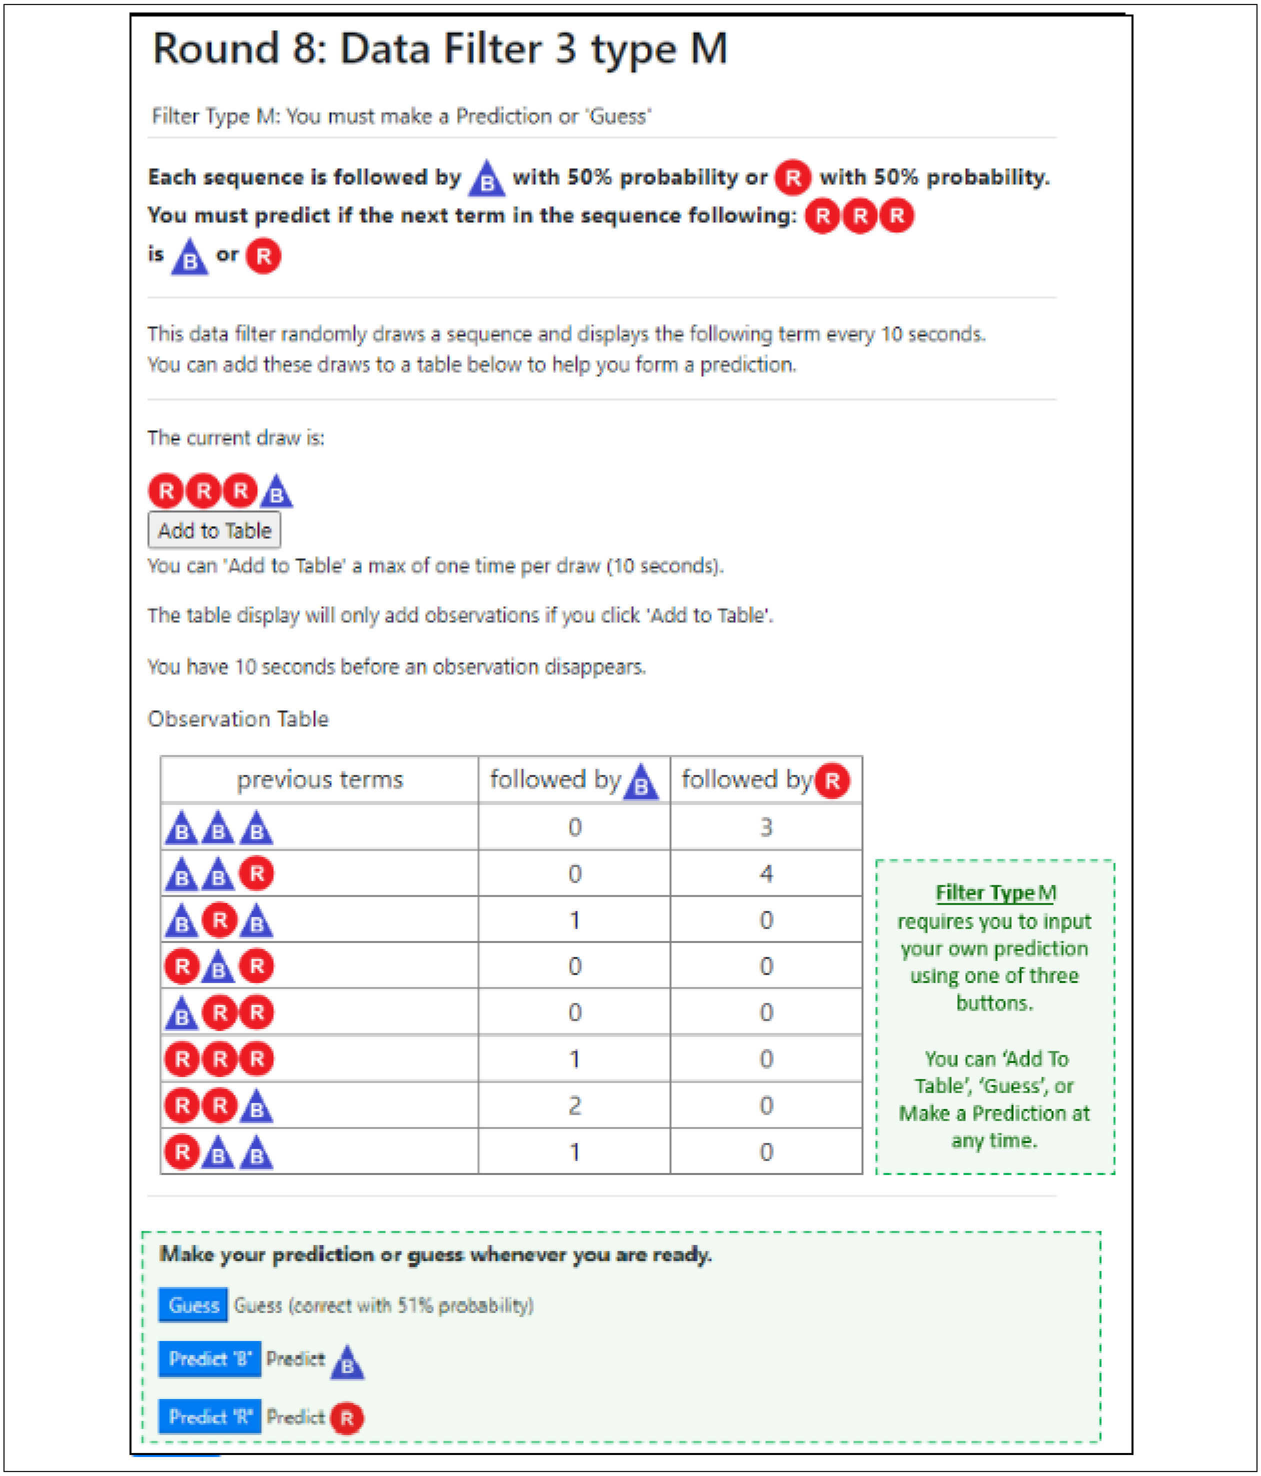
Round 8: Data Filter 3 type M
Filter Type M: You must make a Prediction or 'Guess'
Each sequence is followed by B with 50% probability or R with 50% probability.
You must predict if the next term in the sequence following: R R R
is B or R
This data filter randomly draws a sequence and displays the following term every 10 seconds.
You can add these draws to a table below to help you form a prediction.
The current draw is:
R R R B
Add to Table
You can 'Add to Table' a max of one time per draw (10 seconds).
The table display will only add observations if you click 'Add to Table'.
You have 10 seconds before an observation disappears.
Observation Table
previous terms | followed by B | followed by R
B B B | 0 | 3
B B R | 0 | 4
B R B | 1 | 0
R B R | 0 | 0
B R R | 0 | 0
R R R | 1 | 0
R R B | 2 | 0
R B B | 1 | 0
Filter Type M
requires you to input your own prediction using one of three buttons.
You can 'Add To Table', 'Guess', or Make a Prediction at any time.
Make your prediction or guess whenever you are ready.
Guess
Guess (correct with 51% probability)
Predict 'B'
Predict B
Predict 'R'
Predict R

Instructions

Filter A: Automatic Prediction

Filter type A will alert you as soon as you have added enough observations to the table for the computer to make a correct prediction.

See Filter 3 type A before and after the prediction alert in the two screenshots below.

Round 3: Data Filter 3 type A

Filter Type A: Computer will make a Prediction using the data you 'Add to Table', or you can 'Guess'.

Each sequence is followed by B with 50% probability or R with 50% probability.
You must predict if the next term following: R R R is B or R

This data filter randomly draws a sequence and displays the following term every 10 seconds.
You can add these draws to a table below to help you form a prediction.

The current draw is:

B R R R

Add to Table

You can 'Add to Table' a max of one time per draw (10 seconds).

The table display will only add observations if you click 'Add to Table'.

You have 10 seconds before an observation disappears.

Observation Table

previous terms	followed by B	followed by R
B B B	0	0
B B R	1	0
B R B	0	0
R B R	0	0
B R R	0	1
R R R	0	0
R R B	1	0
R B B	0	0

Make your prediction or guess whenever you are ready.

Guess Guess (correct with 51% probability)

Note: With Filter type A, you cannot make a specific prediction.

The 'Predict' buttons are disabled but the 'Guess' button is active.

Use 'Add to Table' button to collect enough observations for the computer to make a prediction on your behalf.

Note the alert generated when the observation table adds R R R, followed by R:

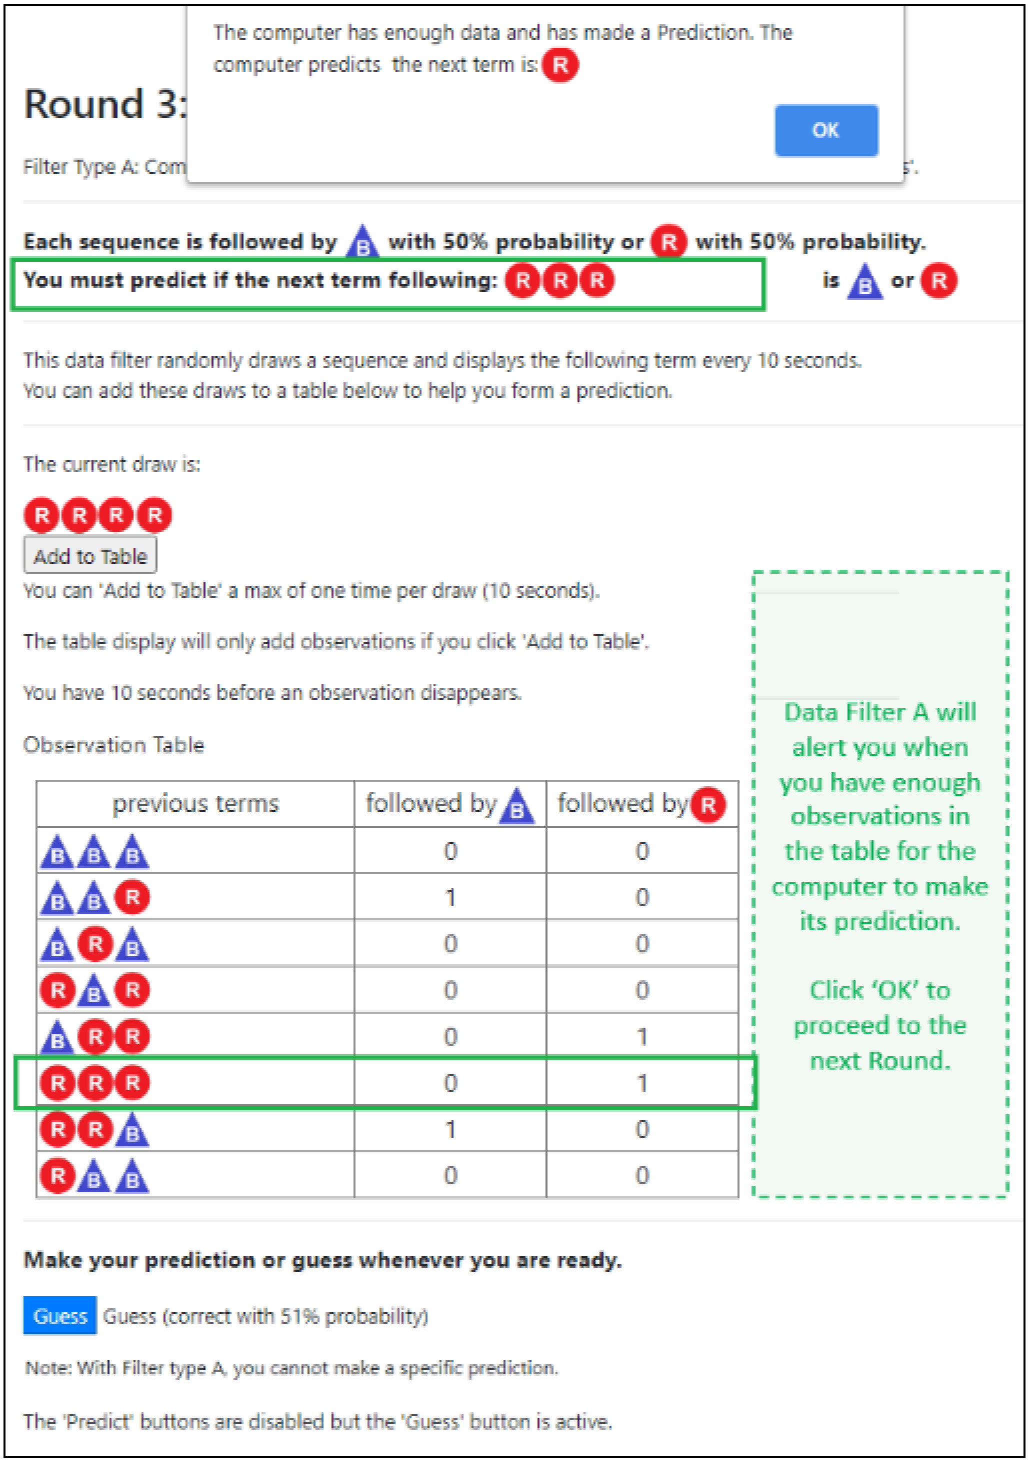

The computer has enough data and has made a Prediction. The computer predicts the next term is: R

OK

Round 3:

Filter Type A: Com ... s'.

Each sequence is followed by B with 50% probability or R with 50% probability.

You must predict if the next term following: R R R **is B or R**

This data filter randomly draws a sequence and displays the following term every 10 seconds.
You can add these draws to a table below to help you form a prediction.

The current draw is:

R R R R

Add to Table

You can 'Add to Table' a max of one time per draw (10 seconds).

The table display will only add observations if you click 'Add to Table'.

You have 10 seconds before an observation disappears.

Observation Table

previous terms	followed by B	followed by R
B B B	0	0
B B R	1	0
B R B	0	0
R B R	0	0
B R R	0	1
R R R	0	1
R R B	1	0
R B B	0	0

Data Filter A will alert you when you have enough observations in the table for the computer to make its prediction.

Click 'OK' to proceed to the next Round.

Make your prediction or guess whenever you are ready.

Guess Guess (correct with 51% probability)

Note: With Filter type A, you cannot make a specific prediction.

The 'Predict' buttons are disabled but the 'Guess' button is active.

Time-Prize Tradeoff

Filters require time to make predictions, but Guessing takes no time.

In each round you will be told the expected time it will take for a filter to reach a prediction, and indicate whether you want to 'Guess Now' or 'Use Filter' for prize increments ranging between $0.25 and $5.00.

The expected time required changes in each round and is published at the top of the decision list for that round.

You must decide ahead of time what you would do for each prize level.

After you submit your choice for each prize, one prize will be randomly selected as the 'choice that counts', and that is the prize value you play for in that round.

Below is a screenshot of the list of choices. Note the average time in this round of 80 seconds to draw the sequence of terms that matches your prediction task:

Round 3: Data Filter 3 type A Choices

Choose an action for each possible prize level: either let the computer Guess, which is correct with probability 51%, or use a data filter to observe the sequence over time.

Any prediction provided by the filter in this round will be 100% correct.

Filter 3 takes an average of 80 seconds (or 01:20 min:sec) to draw the sequence of terms that matches your prediction task.

Prize = $0.25	○ Guess Now	○ Use Filter
Prize = $0.50	○ Guess Now	○ Use Filter
Prize = $0.75	○ Guess Now	○ Use Filter
Prize = $1.00	○ Guess Now	○ Use Filter
Prize = $1.25	○ Guess Now	○ Use Filter
Prize = $1.50	○ Guess Now	○ Use Filter
Prize = $1.75	○ Guess Now	○ Use Filter
Prize = $2.00	○ Guess Now	○ Use Filter
Prize = $2.25	○ Guess Now	○ Use Filter
Prize = $2.50	○ Guess Now	○ Use Filter
Prize = $2.75	○ Guess Now	○ Use Filter
Prize = $3.00	○ Guess Now	○ Use Filter
Prize = $3.25	○ Guess Now	○ Use Filter
Prize = $3.50	○ Guess Now	○ Use Filter
Prize = $3.75	○ Guess Now	○ Use Filter

The filter available to you will change by round.

The average time a filter requires doubles with each increase in Filter number. So Filter 4 requires twice as much time as Filter 3, and Filter 5 requires twice as much time as Filter 4 (on average).

Quiz

Following the two practice rounds, there will be a 6 question multiple choice quiz to ensure you understand the experiment. Each quiz question you answer correctly on the first try is worth an additional $0.50 CAD. Questions answered incorrectly on the first try are not paid but you are encouraged to review the answers and short instructions prior to the paid rounds.

Total Payment

You face a total possible payoff between $7 and $60 based on your choices and the randomly selected prizes in each list.

You earn $7 for participating regardless of the outcome of the experiment.

You can earn up to $3 in the Quiz that follows the Practice Round ($0.50 per correct answer).

In addition to the participation fee and quiz payment, there are 10 paid rounds, 5 with Filter type A and 5 with Filter type M.

For the 5 rounds with Filter type A, the minimum payment is $0, maximum payment is $25, and the expected payment when always choosing Guess is $6.70.

For the 5 rounds with Filter type M, the minimum payment is also $0, and the maximum payment is also $25.

When you use Filter M, you will be asked for your belief about your personal probability of making a mistake with the filter. Your probability of winning a prize with Filter type M is affected by your mistakes and your belief. This is explained in more detail when you first use Filter type M.

Payment Method

Your payment will be electronically transferred to the email you used to sign up for the experiment, on 2021-May-29. Contact econexp@sfu.ca if you have any issues receiving payment.

When you click 'Next', the practice rounds will begin.

Next