Heterogeneous Costs and the Decision to Learn^{*}

David Walker-Jones[†]

May 29, 2024

Abstract

A novel data enrichment demonstrates that experiment subjects are more likely to invest effort into learning about the value of options if simple choice parameters, like price, differ from previous choice problems. This increase in effort in 'unfamiliar' choice problems means that the behavior of many subjects violate even the most flexible model of costly learning if the cost for information is assumed to be constant across choice problems with the same prior beliefs. This observation motivates the introduction of heterogeneous decision makers into a standard and more restrictive (posterior separable) model of costly learning to better fit the data.

Keywords: Posterior Separable, Rational Institution, Experiment

^{*}I would like to thank Andrew Caplin, Mark Dean, Rahul Deb, Tommaso Denti, Henrique de Oliveira, Yoram Halevy, Johannes Hoelzemann, and Colin Stewart, for their helpful advice and support. Funding for the experiment was provided by the University of Toronto. The experiment protocol was approved by the University of Toronto Social Sciences, Humanities & Education REB (protocol number 39850). An earlier version of this paper was called "Optimal Choice Mistakes."

[†]University of Surrey, d.walker-jones@surrey.ac.uk.

1 Introduction

In our modern world decision makers (DMs) are increasingly inundated with information. Incorporating this information into choices requires the costly investment of time and effort. As a result of the cost of learning, DMs often do not acquire all the information that may be relevant to their decision.

The scarcity of attention that results from an over abundance of information motivates the study of rational inattention (Sims, 2003), a model in which DMs acquire information "as if" they solve an optimization problem, weighing the benefits of better choices against the costs of acquiring more detailed information. The work of Sims (2003) has inspired a rapidly growing literature that demonstrates the significance of costly learning in a number of economics' sub-fields (Maćkowiak, Matějka, & Wiederholt, 2023).

Understanding and accurately modelling the cost of information is important for answering standard economic questions for at least two reasons. First, costly learning is relevant to predicting what a DM would choose if parameters such as price change. If the information a DM acquires changes when the price of an option changes, knowing this is crucial for accurately predicting the probability of the option being selected at a new price. Second, understanding what information a DM acquires is significant for understanding the preferences of a DM. The revealed preference approach is predicated on the assumption that the option the DM selects is preferred to the options they¹ do not select. When partial information acquisition can result in the DM selecting an option that has a relatively low payoff, it is more difficult to recover preferences. To identify preferences, a good understanding of what information the DM is acquiring is needed.

Studying costly learning is difficult because it depends on inputs that are difficult to measure such as time, effort, and cognitive resources, and it is thus difficult to

¹I use 'they' instead of she or he. When I use 'they' as a singular noun I conjugate the verb, i.e. I use 'do' instead of 'does,' as recommended by the APA style guide (*Singular "They"*, n.d.).

quantify the relationship between these inputs and outputs such as the quality or accuracy of decisions. As a result, it can be hard to know what mathematical structure should be imposed onto the costs of information.

This paper provides characterizations of state dependent choice probabilities that can arise from various information cost functions in a natural setting. Importantly, I show that the testable implications of the models are not vacuous and vary depending on the assumptions on the information cost function. This paper is also the first to discuss and characterize the testable implications of a "random utility" version of the "standard posterior separable" model, which allows for heterogeneous beliefs and cost functions for information, and is useful for analysis of both individual level and aggregate data. Aggregating the behavior of heterogeneous DMs that each pay for information according to a posterior separable model predicts changes in demand that are inconsistent with a representative DM version of the posterior separable model, but are still more specific than the most general model of costly learning. This result may be surprising to some readers whose intuition is based on random utility models since, in such models, if the value distribution of DMs is not restricted to be in a particular parametric class then aggregating over heterogeneous DMs always produces behavior that is consistent with some representative DM.

This paper also introduces an experiment to test the heterogeneous model. The experimental design is novel as it allows me to uncover whether subjects indeed behave as if their information cost is random. The data enrichment shows that most subjects fluctuate back and forth between learning and not learning in a way that violates even the most flexible model of costly learning if it is assumed that the cost function for information is constant across choice problems in which subjects have the same prior belief. Fatigue is unsurprisingly shown to be an issue, but even allowing for fatigue cannot rationalize the data of many subjects as choice problems that are novel, in the sense that their parameters differ from previous choice problems, seem to increase subjects' willingness to invest effort into acquiring information. In the literature on rational inattention it is standard to assume that the cost function for information is "posterior separable" (Caplin & Dean, 2013), which is to say the cost of reaching different posterior beliefs is additively separable, as this assumption is conducive to analysis, has been provided with micro-foundation from dynamic learning models (e.g., Morris & Strack, 2019; Bloedel & Zhong, 2021; Hébert & Woodford, 2023), and axiomatic models of costly learning often produce posterior separable models (e.g., Mensch, 2018; Pomatto, Strack, & Tamuz, 2023). Denti (2022) provides an important characterization of the behavior that can be rationalized in general contexts by posterior separable learning costs, but further shows that subjects in an experiment conducted by Dean and Neligh (2023) frequently violate the standard and quite flexible posterior separable model of costly learning.

This paper explores a natural explanation for why the standard posterior separable model can fail and how it can be amended to account for observed behavioral patterns. The solution is motivated by addressing a challenge associated with studying costly learning; the necessitated aggregation of data.

Even estimating the outcome of costly learning is difficult because incomplete information acquisition results in behavior that is generally stochastic in nature; what a researcher wishes to study is the probabilities of different options being selected in different states of the world, what is referred to in the literature as "state dependant stochastic choice data" (Caplin & Dean, 2015; Caplin & Martin, 2015). Collecting this type of data requires aggregation of some kind. At the very least, aggregation of different decision outcomes from the same individual making choices repeatedly in similar but slightly varied choice problems is required to estimate such probabilities. The impact of aggregation, even though it is necessitated empirically, has not yet received much attention in the theory literature.

When the specifics of the experiment conducted by Dean and Neligh (2023) are considered, however, the aggregation of data seems to provide a natural explanation for the shortfalls of the posterior separable model. Consider a DM learning about which of two possible states of the world has been realized, call them state 1 and state 2, each of which is equally likely *a priori*. Suppose the DM also has three options to pick between: option 1, which is best if the learning of the DM indicates that there is a high probability that state 1 has been realized, option 2, which is best if the learning of the DM indicates that there is a high probability that state 2 has been realized, and option 3, which is only best if the probability of each of the two states is close to one half. Denti (2022) shows with that paper's Proposition 4 that, generically, the posterior separable model predicts that there are possible states of the world. The experiment of Dean and Neligh (2023), in contrast, indicates that one should expect a subject will likely select all three of the options in this paragraph's example with positive probabilities.

If a non-constant cost function for information is introduced, as is done in this paper, then experiment subjects selecting all three options from the example in the previous paragraph with positive probabilities begins to make sense. Suppose that initially a subject is putting effort into learning and they always convince themselves that either state 1 or 2 is likely enough that they select option 1 or 2. But then the repeated choices that are necessary for aggregating data to get state dependant stochastic choice data cause the subject to fatigue (which one can argue is analogous to an increase in the cost of information), and eventually the subject stops investing effort into learning and simply picks the option that is best at their prior belief; option 3.

Is the gap between predicted and observed behavior in this example a failure of the posterior separable model or an issue with assuming that the cost function for information is constant? This paper argues that it is the latter, though such an assumption is the standard in the field of rational inattention.

2 Model

Consider a DM choosing between two options, option X and option Y. Option Y is referred to as the **uncertain option** because there is some uncertainty about the payoff the DM gets from selecting it: The DM's payoff from option Y, $u(\omega) \in \mathbb{R}$, is determined by the **state** $\omega \in \Omega = \{\underline{\omega}, \overline{\omega}\}$, with $u(\underline{\omega}) < u(\overline{\omega})$. The DM knows the distribution $\mu \in \Delta(\Omega)$ over states, which is referred to as their **prior belief**.² Option X is referred to as the **safe option** because the DM knows they receive a payoff of $p \in \mathbb{R}$ from selecting it, and p is referred to as the **price** of selecting option Y because the DM must give up the opportunity to have a payoff of p to select option Y.

Given μ , the **behavior** at a finite set of prices $\mathcal{P} \subseteq \mathbb{R}$ (a finite set of values for option X), is a mapping $s : \mathcal{P} \to S \equiv [0, 1] \times [0, 1]$. Namely, at each price $p \in \mathcal{P}$ the researcher observes the **information outcome** $s(p) = (\underline{s}(p), \overline{s}(p)) =$ $(\Pr(Y|\underline{\omega}, p), \Pr(Y|\overline{\omega}, p))$, that describes the probability of the DM selecting option Y at each ω , denoted $\Pr(Y|\omega, p) \equiv 1 - \Pr(X|\omega, p)$. This paper's model focuses on outcomes because this is the observable behavior that is measured in a typical dataset.

The DM is said to **learn** at a price p if $\Pr(Y|\underline{\omega}, p) \neq \Pr(Y|\overline{\omega}, p)$ and $\mu(\overline{\omega}) \in (0, 1)$. In other words, the DM is learning whenever their probability of selecting option Y changes with the realization of ω , because this is indicative of the DM at least partially differentiating between $\underline{\omega}$ and $\overline{\omega}$.

2.1 Costly Learning

A cost function (for information) $C_{\mu} : S \to \mathbb{R}_+$ describes the minimal cost of achieving each potential information outcome.³ The cost function does not depend on price, which is to say this paper's model assumes that a price change does not impact the cost to the DM of achieving a given information outcome. The realized

²The belief μ assigns a strictly positive probability to each state unless stated otherwise.

³The cost function has a subscript μ because I do not impose that it is constant in μ , as would be the case if I instead, for instance, assumed a "uniformly" posterior separable cost function (Caplin, Dean, & Leahy, 2022).

cost $C_{\mu}(s(p))$ might change when p changes, but the function C_{μ} does not change if p changes.

If the DM can achieve an outcome without doing any learning, then it is natural to think the outcome should be costless. So, the DM should be able to select option Y, or option X, or randomize over these options, all without incurring any learning costs. **Assumption 1:** $\forall x \in [0, 1], C_{\mu}(x, x) = 0$. Further, it is natural to assume that the DM can randomize over information outcomes without cost. So, if an information outcome is a convex combination of other information outcomes, then the cost of the information outcome should be weakly lower than the appropriate convex combination of the costs of the other information outcomes, which implies that more information in a Blackwell (1951, 1953) sense is weakly more costly. **Assumption 2:** C_{μ} is a weakly convex function.⁴

Assumption 1 and Assumption 2 represent the minimal structure that is assumed of all cost functions for information in this paper. Given the belief of the DM, μ , the **cost function** for information outcomes is defined as a mapping from S onto the positive reals, $C_{\mu}: S \to \mathbb{R}_+$, which satisfies Assumption 1 and Assumption 2.

Definition: Given a prior belief μ and a set of prices \mathcal{P} , the behavior is **rationalized** by a costly learning model if there are payoffs for option Y, $u(\underline{\omega})$ and $u(\overline{\omega})$ with $u(\underline{\omega}) < u(\overline{\omega})$, and a cost function for information outcomes C_{μ} such that, $\forall p \in \mathcal{P}$,

$$s(p) = (\underline{s}(p), \,\overline{s}(p)) \in \underset{(\underline{s}, \,\overline{s}) \in S}{\operatorname{arg\,max}} \left(p + \underline{s}\mu(\underline{\omega})(u(\underline{\omega}) - p) + \overline{s}\mu(\overline{\omega})(u(\overline{\omega}) - p) - C_{\mu}(\underline{s}, \,\overline{s}) \right).$$

Given p and μ , the expected **gain in payoff** the DM receives when they choose an information outcome $s = (\underline{s}, \overline{s})$ instead of picking option X with certainty is defined by:

$$\underline{s}\mu(\underline{\omega})(u(\underline{\omega})-p)+\overline{s}\mu(\overline{\omega})(u(\overline{\omega})-p)-C_{\mu}(\underline{s},\,\overline{s}).$$

Maximizing the gain in payoff eases exposition and is equivalent to maximizing the

⁴By 'weakly convex' I mean convex but not necessarily strictly convex.

expected payoff as the difference between them is a simple re-normalization of option payoffs. When p increases the gain in payoff from selecting option Y decreases. As a result, when p increases, the gain in payoff from an information outcome $s = (\underline{s}, \overline{s}) \in$ S decreases in proportion to the **unconditional probability** of selecting option Y that the information outcome results in, which is $\underline{s}\mu(\underline{\omega}) + \overline{s}\mu(\overline{\omega})$. So, the gain in payoff from information outcomes that create lower unconditional probabilities of selecting option Y decrease less quickly when p increases. This suggests the DM should choose an information outcome with a lower unconditional probability of selecting option Y when price increases, as is shown by Theorem 1.

Theorem 1. Given a prior belief μ and a finite and non-empty set of prices \mathcal{P} , the behavior is rationalized by a costly learning model if and only if the following three properties are satisfied:

- (i) The DM is weakly more likely to select option Y when its payoff is high: $\forall p \in \mathcal{P} : \Pr(Y|\underline{\omega}, p) \leq \Pr(Y|\overline{\omega}, p).$
- (ii) The DM is weakly less likely to select option Y when price increases:

 $\Pr(\mathbf{Y}|p) \equiv \mu(\overline{\omega})\Pr(\mathbf{Y}|\overline{\omega}, p) + \mu(\underline{\omega})\Pr(\mathbf{Y}|\underline{\omega}, p)$ is weakly decreasing in p.

(iii) If there is a price at which the DM randomizes over selecting option X and option Y without learning, then they do not learn at any price, and select either option X or option Y with probability one at all other prices:

if $\exists p \in \mathcal{P}$: $\Pr(Y|\underline{\omega}, p) = \Pr(Y|\overline{\omega}, p) \in (0, 1)$, then $\forall \tilde{p} \in \mathcal{P} \setminus p$: $\Pr(Y|\tilde{p}) \in \{0, 1\}$. **Proof.** See Appendix 1.

If no structure is imposed on the learning costs of the DM beyond Assumption 1 and Assumption 2 then Theorem 1 describes the behavioral patterns that are consistent with this most general model of costly learning. Randomization is significant in property (iii) because it indicates that the DM could instead optimally pick X or Y, which, if they learned at some other price, could be used to create a violation of (ii).

If I further assume that the researcher knows the expected payoff from option Y, which is $\mu(\underline{\omega})u(\underline{\omega}) + \mu(\overline{\omega})u(\overline{\omega})$, then predictions can be made for if a price change impacts whether or not the DM will learn, as is formalized in Corollary 1. Many results in the rational inattention literature, for instance in the work of Caplin et al. (2022), assume that the utility function of the DM is known, and thus assuming that the expected payoff of Y is known is a relatively weak assumption, and in the experiment presented in this paper the expected payoff from option Y in terms of probability points is known.

Corollary 1. Given a prior belief μ , and payoffs for option Y, $u(\underline{\omega})$ and $u(\overline{\omega})$ with $u(\underline{\omega}) < u(\overline{\omega})$, suppose the behavior is rationalized by a costly learning model with cost function for information outcomes C_{μ} , and that there are two prices $p_1, p_2 \in \mathcal{P}$ with $p_1 < p_2$.

(i) If the prices are below the expected payoff from option Y, i.e. $p_1 < p_2 \leq \mu(\underline{\omega})u(\underline{\omega}) + \mu(\overline{\omega})u(\overline{\omega})$, then the DM learns at p_2 if they learn at p_1 :

 $\Pr(\mathbf{Y}|\underline{\omega}, p_1) \neq \Pr(\mathbf{Y}|\overline{\omega}, p_1) \Rightarrow \Pr(\mathbf{Y}|\underline{\omega}, p_2) \neq \Pr(\mathbf{Y}|\overline{\omega}, p_2).$

(ii) If the prices are above the expected payoff from option Y, i.e. $\mu(\underline{\omega})u(\underline{\omega}) + \mu(\overline{\omega})u(\overline{\omega}) \leq p_1 < p_2$, then the DM learns at p_1 if they learn at p_2 :

 $\Pr(Y|\underline{\omega}, p_2) \neq \Pr(Y|\overline{\omega}, p_2) \Rightarrow \Pr(Y|\underline{\omega}, p_1) \neq \Pr(Y|\overline{\omega}, p_1).$

Proof. Suppose the DM learns at p_1 . If $p_1 < p_2 \leq \mu(\underline{\omega})u(\underline{\omega}) + \mu(\overline{\omega})u(\overline{\omega})$, then when p increases, the gain in payoff from $s(p_1)$ strictly decreases, but not as quickly as the gain in payoff from not learning and selecting option Y, so the DM must learn at p_2 . Suppose instead the DM learns at p_2 . If $\mu(\underline{\omega})u(\underline{\omega}) + \mu(\overline{\omega})u(\overline{\omega}) \leq p_1 < p_2$, then when p decreases the gain in payoff from $s(p_2)$ increases while the gain in payoff from not learning and selecting X stays the same, so the DM must learn at p_1 .

If the researcher can observe if the DM is acquiring information at different prices, which is a binary variable, Corollary 1 provides a means of testing a very general model of costly learning without needing to estimate choice probabilities, a variable that is continuous and inherently difficult to estimate.

2.2 The Posterior Separable (PS) Learning Model

The standard in the field of rational inattention is to model the DM paying for information according to a posterior separable (PS) cost function (Caplin & Dean, 2013; Caplin et al., 2022). This means the DM selects what to learn by choosing how confident to be in their choices after they finish learning. So, they pick $Pr(\overline{\omega}|X, p)$ and $Pr(\overline{\omega}|Y, p)$, the probability that state $\overline{\omega}$ has been realized after selecting each of the two options.

To apply a PS model I need a function that measures how 'informed' different posteriors are. Denote this weakly convex function $c : [0, 1] \rightarrow \mathbb{R}$. When the DM learns they pay for the information based on the change in c it creates. Again, weak convexity of c ensures that more information (in a Blackwell (1951, 1953) sense) is weakly more costly. Given belief μ , if the DM learns at a price p and their behavior is rationalized by a costly learning model, then $\Pr(\overline{\omega}|X, p) < \mu(\overline{\omega}) < \Pr(\overline{\omega}|Y, p)$, and it is easy to compute the unconditional probabilities of the DM selecting the two options. Then, when the DM learns at a price p in the PS model they pay the following: $\Pr(Y|p)(c(\Pr(\overline{\omega}|Y, p)) - c(\mu(\overline{\omega}))) + \Pr(X|p)(c(\Pr(\overline{\omega}|X, p)) - c(\mu(\overline{\omega})))$,

for their information, which is the probability of them selecting Y times the 'increase in information' (as measured by c) that occurred before they selected Y, plus the chance of them selecting option X times the 'increase in information' (as measured by c) that occurred before they selected X. An example of a commonly used measure the 'informativeness' of a posterior is Shannon Entropy (Shannon, 1948), e.g. in the work of Matějka and McKay (2015).

Definition: Given a prior belief μ , the cost function for information outcomes C_{μ} is a **PS cost function** for information outcomes if there is **measure of informedness**,

a weakly convex $c: [0, 1] \to \mathbb{R}$, such that $\forall s \in S$ with $\underline{s} < \overline{s}$:

$$C_{\mu}(\underline{s}, \overline{s}) = (\overline{s}\mu(\overline{\omega}) + \underline{s}\mu(\underline{\omega})) \left(c \left(\frac{\overline{s}\mu(\overline{\omega})}{\overline{s}\mu(\overline{\omega}) + \underline{s}\mu(\underline{\omega})} \right) - c(\mu(\overline{\omega})) \right) + (1 - \overline{s}\mu(\overline{\omega}) - \underline{s}\mu(\underline{\omega})) \left(c \left(\frac{(1 - \overline{s})\mu(\overline{\omega})}{(1 - \overline{s})\mu(\overline{\omega}) + (1 - \underline{s})\mu(\underline{\omega})} \right) - c(\mu(\overline{\omega})) \right)$$

Definition: Given a prior belief μ and a set of prices \mathcal{P} , the behavior is **rationalized** by a **PS model** if there are payoffs for option Y, $u(\underline{\omega})$ and $u(\overline{\omega})$ with $u(\underline{\omega}) < u(\overline{\omega})$, and a PS cost function for information outcomes C_{μ} such that, $\forall p \in \mathcal{P}$,

$$s(p) = (\underline{s}(p), \,\overline{s}(p)) \in \underset{(\underline{s}, \,\overline{s}) \in S}{\operatorname{arg\,max}} \left(\underline{s}\mu(\underline{\omega})(u(\underline{\omega}) - p) + \overline{s}\mu(\overline{\omega})(u(\overline{\omega}) - p) - C_{\mu}(\underline{s}, \,\overline{s}) \right).$$

Theorem 2. Given a prior belief μ and a finite and non-empty set of prices \mathcal{P} , the behavior is rationalized by a PS model if and only if it is rationalized by a costly learning model and $\Pr(\overline{\omega}|X, p)$ and $\Pr(\overline{\omega}|Y, p)$ are both weakly increasing over the set of p where the DM learns.

Proof. See Appendix $1.^5$

Theorem 2 says that when p increases in a PS model the DM is more confident they have made the right decision when selecting option Y and less confident they made the right decision when selecting option X.

2.3 An Aggregate PS Model: Fatigue and Novelty

Even if the structure imposed by a particular class of costly learning models is compelling in a setting, datasets used for analysis do not typically focus on an individual DM. Typically datasets aggregate behavior over DMs with potentially heterogeneous (accurate) prior beliefs and costs for information outcomes. Even when

 $^{^5\}mathrm{Ambuehl}$ (2017) shows necessity of the conditions for a smooth and strictly convex c.

a dataset does feature observations from an individual, that individual might have varying private pieces of information (that change their prior belief) and a varying information cost function due to either something akin to fatigue or variation in the choice environment that is not observable to the researcher, e.g. variation in the number of store clerks on the floor when the individual makes the decision. This subsection characterizes an aggregate version of the PS model in Section 2.2, and shows that aggregating choice data from heterogeneous PS DMs (either different individuals or varying versions of the same individual) can result in behavior that cannot be rationalized with a PS model. This is in contrast with aggregation of the flexible costly learning model in Section 2.1, which, it is easy to show, does not change the characterization of what behavior can be rationalized.

Informally, behavior is rationalized by an aggregate PS model if there are different types of DM, each with their own probability of occuring, prior belief, and PS cost function for information outcomes, such that each type has rationalized behavior, and if behavior is averaged over the different types then the observed behavior is obtained. This is akin to the way DMs are modeled in Random Utility models.

Definition: Given a prior belief μ , the behavior is **rationalized by an aggregate PS model** if there are payoffs for option Y, $u(\underline{\omega})$ and $u(\overline{\omega})$ with $u(\underline{\omega}) < u(\overline{\omega})$,⁶ $T \in \mathbb{N}$ types of DM each of which has probability of occurring $\pi_t > 0$, a belief about the probability of $\overline{\omega}$ occurring $\mu_t(\overline{\omega}) \in [0, 1]$, and behavior s_t , which is $\Pr_t(Y|\omega, p)$ for each state and price and type, such that each type's behavior is rationalized by a PS model⁷ with payoffs for option Y $u(\underline{\omega})$ and $u(\overline{\omega})$, and:

(i) The probabilities of the different types sum to one, and their mean belief is the

⁶I assume that $u(\underline{\omega})$ and $u(\overline{\omega})$ are the same for each type as this is more challenging. Given the result in Theorem 3, allowing utility to differ across types would not substantially change the characterization of rationalized behavior as long as Y is preferred in $\overline{\omega}$, the exception being that property (iii) from Theorem 1 could be violated.

⁷Note that different types can have their behavior rationalized by different PS models, namely I do not impose that different types have the same PS cost function for information outcomes.

distribution over states observed by the researcher:

$$\sum_{t=1}^{T} \pi_t = 1, \quad \sum_{t=1}^{T} \mu_t(\overline{\omega})\pi_t = \mu(\overline{\omega}).$$

(ii) Given any pair of price p and state ω , the mean behavior is the behavior observed by the researcher:

$$\forall p \in \mathcal{P}, \forall \omega \in \Omega : \sum_{t=1}^{T} \Pr_t(\mathbf{Y}|\omega, p) \frac{\mu_t(\omega)\pi_t}{\mu(\omega)} = \Pr(\mathbf{Y}|\omega, p).$$

Theorem 3. Given a prior belief μ and a finite and non-empty set of prices \mathcal{P} , the behavior is rationalized by an aggregate PS model if and only if it is rationalized by a costly learning model and $\Pr(Y|\underline{\omega}, p)$ and $\Pr(Y|\overline{\omega}, p)$ are both weakly decreasing in p.

Proof. See Appendix 1.

Theorem 3 says that the behavior can always be rationalized with heterogenous DMs that all behave in line with a PS model if and only if when price increases the probability of selecting option Y weakly decreases in both states. Using Bayes' Rule, it is easy to show that such behavior can violate the predictions of Theorem 2. Further, one can visually see the difference in predicted behavior when PS DMs are aggregated by comparing figures 4 and 5 in the Online Appendix. As a result, aggregating over DMs that behave in line with a PS model can produce behavior that cannot be rationalized by a PS model. Further, aggregating over PS DMs produces more specific predictions than the general model of costly learning introduced in Section 2.1, as can be seen by comparing figures 3 and 5 in the Online Appendix.

3 The Experiment

In this paper's experiment, whether or not a subject chooses to learn in a decision problem is observed, which is not typical in the literature on rational inattention. This added detail is quite useful as the choice to acquire more information results in a binary outcome, and using it for analysis with Corollary 1 does not require the estimation of probability, a continuous variable.

In the experiment, which was completed by 243 undergraduate students from the University of Toronto, each subject faced 92 rounds of **decision problems** in which they chose between a safe option X and an uncertain option Y, and, if they desired, they could also choose to 'get more information' (GMI) about the payoff from option Y before selecting an option.

In each of these decision problems the payoff from option X is always a $p \in \{0.25, 0.5\}$, which is clearly displayed to the subject. In the first 40 and last 40 decision problems p = 0.25, and in the middle 12 decision problems (decision problems 41 through 52) p = 0.5. The payoff of option Y is a $u(\omega) \in \{0, 1\}$, which is not displayed.

The payoffs of the options are measured in terms of percentage points. Subjects pick options to try to gradually increase their chance of winning a monetary prize at the end of the experiment.⁸ For example, if a subject always selects option Y, and in 50 decision rounds it has a payoff of 1, and the rest of the time it has a payoff of 0, then they have a 50% chance of winning the prize when the experiment ends. The prize the subject can win in the draw is either \$20, \$25, or \$30, depending on the subject's treatment group.

In each decision problem the subject is presented with a big dot that is either red or green, as is depicted in Figure 1. The big dot induces the belief of the subject about the payoff from option Y. If the big dot is red there is a $\frac{1}{4}$ chance that option

⁸In this experiment subjects are 'paid' with percentage points in a decision problem so that within subject analysis can be conducted in an incentive compatible way. See the Online Appendix for a discussion of incentive compatibility.

Y is the better option $(\mu(\overline{\omega}) = \frac{1}{4})$, and if the big dot is green there is a $\frac{3}{4}$ chance that option Y is the better option $(\mu(\overline{\omega}) = \frac{3}{4})$. In each decision problem there is a $\frac{3}{4}$ chance that the big dot is red, and a $\frac{1}{4}$ chance that the big dot is green.

Figure 1: Option X, Option Y, or Get more information (GMI): Please respond to the following investment decision:

Round number: 9 of 100 Option X increases your probability	of winning the \$30 prize by 0.25 percentage points
Please choose an option:	
	~
Option X (increases your probabili	ty of winning the prize by 0.25 percentage points)
Option Y (increases your probabili	ty of winning the prize by 0.25 percentage points on average)
Get more information	
R	

If a subject chooses to 'get more information,' (GMI) 100 small dots appear, each of which is either red or green, as is depicted in FIgure 2. There are always 49 of one color of small dot, and 51 of the other color.⁹ If 51 of the small dots are red, then the payoff from option Y is 0, and if 51 of the small dots are green, then the payoff from option Y is 1. Figure 1 and Figure 2 display an example of a decision problem before and after the subject has chosen to GMI. In total, subjects only chose to GMI in 9,233 of 22,356 decision problems, with substantial variation in how many times different individuals chose to GMI.

Subjects participated on-line due to COVID-19.¹⁰ In addition to the decision

 $^{^9 {\}rm Subjects}$ are provided with this information in their training so that their uncertainty about the returns to effort are minimized.

¹⁰The experiment was programmed using oTree (Chen, Schonger, & Wickens, 2016).

Figure 2: After choosing 'Get more information': Please respond to the following investment decision:



problems discussed above, each subject was trained, completed a quiz, and went through eight 'rounds' of practice problems to familiarize themselves with the interface before the 92 rounds of decision problems. For a more detailed description of the experiment, see the Online Appendix.

3.1 Discussion of Experiment Design: Connection to Theory

For the models studied in Section 2 to be appropriate in the decision environment of the experiment, experiment subjects need to use the big dot to correctly determine the prior distribution over states in each round. In any experiment on rational inattention, and in particular experiments where there are changes in the desired "prior belief" of subjects, there is a danger that subjects do not properly internalize the "prior." I use big colorful dots to convey the prior so that the prior is as salient of a feature of the choice environment as possible. Further, as Figure 1 shows, in each round the numeric information about the prior contained in the colorful big dot, and its implication for the expected value of option Y, is repeated next to the choice options in the drop-down menu that subjects use to select alternatives, and only two relatively simple priors are used throughout the experiment. In addition, there are several quiz questions that all subjects face about the big dots (see, in particular, Figures 18 and 19, but also Figures 11 and 14, in the Online Appendix for examples), and thus subjects that did well on the quiz seem to be able to convert the big dots into a prior belief in the intended way. The robustness of data results to quiz performance is addressed in Table 1 below. Subjects are also reminded in ever round, above the button that takes them to the next screen, about the information contained in the big dots (see Figure 20 in the Online Appendix).

As is typical in experiments on rational inattention, it is also assumed that subjects know the value of the safe option. To aid the appropriateness of this assumption, there are only two points within the 92 decision problems at which the value of the safe option changes. Before each block of rounds with the same payoff for option X subjects are made aware of the payoff and the fact that it would stay the same for a certain number of rounds (see Figures 21 and 22 in the Online Appendix). Further, as can be seen in Figure 1, the payoff from option X is displayed in two places in each round.

In the theory models it is said a DM learns if their probability of selecting option Y differs across the states of the world. In the experiment it is instead said that a subject learns if they choose to "get more information" (GMI). This is done because in the theory model it is assumed that choice probabilities are perfectly observed, whereas in any real dataset (such as the one generated in this paper's experiment) only an estimate of choice probabilities can be obtained, and frequently the estimate relies on relatively few observation. Choice probabilities not being statistically significantly different is thus not a satisfactory indicator of whether or not learning occurred. Reaction times can also be used as a proxy for the decision to learn, but this seems to be a nosier indicator of the decision to learn as there is nothing stopping subjects from taking short "breaks" during an experiment to check the time, stretch, rub their eyes, or look at their cellphones. I do, however, use reaction times to supplement the analysis in Section 3.3, and consider the possibility that subjects sometimes made mistakes when selecting if they wanted to GMI.

3.2 Aggregate Data

Because subjects sometimes do and sometimes do not learn (choose to GMI) at at each pair of p and μ I observe in the experiment, data that is aggregated over subjects rejects the predictions of Corollary 1 (see the next subsection for a more detailed discussion), and thus the most flexible model of costly learning this paper studies that only imposes Assumption 1 and Assumption 2, if the cost function for information outcomes is assumed to be constant across rounds with the same prior belief. This is not surprising since it is natural to think that subjects might have changes in their cost function for information over the course of the experiment due to factors such as fatigue. Though, as is discussed in the next subsection, incorporating fatigue into a model of costly learning is not sufficient for rationalizing the information decisions of many subjects.

If the data on the decision to GMI is ignored and only the state and price dependant choice probabilities for option Y are considered then, fixing either prior belief, aggregate demand changes in line with the predictions of both the representative DM version of the PS model and the heterogeneous DM version (Theorem 2 and Theorem 3 respectively) as is shown by Figure 6 in the Online Appendix. Theorem 3 is still an important result, however, because it demonstrates that (i): aggregating over heterogeneous PS DMs can produce behvior that rejects the representative DM version of the PS model, and (ii): even if heterogeneous DMs are introduced the PS model still produces narrower predictions than the most general costly learning model introduced in Section 2.1 (see Theorem 1) and rules out that there is a state of the world where the chance of the risky option being selected increases when its price does, which one can argue is normatively appealing.

3.3 Within Subject Heterogeneity: Fatigue and Novelty

Corollary 1 implies that, if the cost function for information outcomes is constant across rounds with the same prior (as is assumed throughout this paragraph), if $\mu(\overline{\omega}) = \frac{1}{4}$ a subject has more incentive to learn if the price is p = 0.25 compared to p = 0.5. Corollary 1 thus implies that, in rounds with $\mu(\overline{\omega}) = \frac{1}{4}$, if a subject sometimes chooses GMI in rounds with p = 0.5 then they should choose GMI in all rounds with p = 0.25. While if $\mu(\overline{\omega}) = \frac{3}{4}$, Corollary 1 says a subject has more incentive to learn if the price is p = 0.5. Corollary 1 thus implies that, in rounds with $\mu(\overline{\omega}) = \frac{3}{4}$, if a subject sometimes chooses GMI in rounds with p = 0.25 then they should choose GMI in all rounds with p = 0.5. Table 1 shows that many subjects thus violate the predictions of the most general model of costly learning that only imposes Assumption 1 and Assumption 2, and that this conclusion is robust to concerns about lack of understanding as it is true even for subjects that did quite well on the quiz (rows 2, 3, and 4, of Table 1).

To help ensure that the violations of Corollary 1 are not being generated by subjects either accidentally selecting or not selecting to GMI, Table 1 also conducts some analysis that allows for some "mistakes" and additionally controls for quiz performance (rows 5 through 8). Pervasive violations of Corollary 1 are still found when some choice are explained away as "mistakes." This is not surprising because the experiment is designed so that these types of mistakes are unlikely. Subjects pick if they do or do not want to GMI from a drop-down menu (see Figure 1), which then displays their choice, and they must then subsequently hit a "Next" button to finalize their decision that is quite some distance from the drop-down menu, below both the big dot for the round and a paragraph that reminds them of the meaning of the big dot (see Figure 20 in the Online Appendix).

One natural explanation for why the predictions of Corollary 1, and the very

(violate the costly learning model that only imposes Assumptions 1 and 2)				
Row: 1	If analysis ignores fatigue and "mistakes"	56%		
2	And is restricted to subjects that got $5/10$ or more on the quiz	55% (122/223)		
3	And is restricted to subjects that got $7/10$ or more on the quiz	54% (94/173)		
4	And is restricted to subjects that got $9/10$ or more on the quiz	49% (49/99)		
5	If analysis allows for some "mistakes"	45%		
6	And is restricted to subjects that got $5/10$ or more on the quiz	46% (102/223)		
7	And is restricted to subjects that got $7/10$ or more on the quiz	$46\% \ (80/173)$		
8	And is restricted to subjects that got $9/10$ or more on the quiz	46% (46/99)		
9	If analysis allows for fatigue	37%		
10	And is restricted to subjects that got $5/10$ or more on the quiz	$36\% \ (81/223)$		
11	And is restricted to subjects that got $7/10$ or more on the quiz	$35\% \ (60/173)$		
12	And is restricted to subjects that got $9/10$ or more on the quiz	31% (31/99)		

Table 1: Percent of subjects that violate Corollary 1

Row 1 indicates the percent of subjects that, when $\mu(\overline{\omega}) = \frac{1}{4}$, sometimes choose to acquire information (choose to GMI) if p = 0.5 but do not always choose GMI if p = 0.25, or, when $\mu(\overline{\omega}) = \frac{3}{4}$, sometimes choose GMI if p = 0.25 but do not always choose GMI if p = 0.5. Rows 2, 3, and 4, conduct the same exercise as row 1 but restrict attention to subjects that got quiz scores of at least 5, 7, and 9, out of 10 respectively. Row 5 conducts the same exercise as row 1 but a subject is only considered to have not GMI at a pairing of μ and p if there are at least two instances of them not choosing GMI at said pairing, and they are only considered to have GMI at a pairing of μ and p if there is at least one round in which they chose GMI at said pairing and then took at least 10 seconds to make a decision. Rows 6, 7, and 8, conduct the same exercise as row 5 but restrict attention to subjects that got quiz scores of at least 5, 7, and 9, out of 10 respectively. Row 9 indicates the percent of subjects that, in rounds with $\mu(\overline{\omega}) = \frac{1}{4}$, chose not to GMI for p = 0.25 and then subsequently (in a later round) chose to GMI for p = 0.5, or, in rounds with $\mu(\overline{\omega}) = \frac{3}{4}$, chose not to GMI for p = 0.5 and then subsequently chose to GMI for p = 0.25. Rows 10, 11, and 12, conduct the same exercise as row 9 but restrict attention to subjects that got quiz scores of at least 5, 7, and 9, out of 10 respectively.

general model of costly learning it studies, are violated by so many subjects is that subjects may fatigue over the course of the experiment and as a result choosing to GMI may be less appealing as the experiment progresses. This can be modelled by allowing C_{μ} to increase as a subject progresses through the decision problems. Let C_{μ}^{r} denote the cost function of a subject when their belief is μ and they are facing their r^{th} decision problem. A subject's cost function for information outcomes is said to be **consistent with fatigue** if, given belief μ , and integer $t > 0, \forall s \in S$ with $\underline{s} \leq \overline{s}$: $C^r_{\mu}(\underline{s}, \overline{s}) \leq C^{r+t}_{\mu}(\underline{s}, \overline{s})$. As is also demonstrated in Table 1, allowing for fatigue substantially reduces the percent of subjects that violate the very flexible model of costly learning introduced in Section 2.1, but many subjects still violate the predictions of Corollary 1, and this result holds even in the subsets of subjects that did best on the quiz.

It seems that fatigue is not sufficient for rationalizing the behavior of subjects, at least in part, because subjects are more likely to learn in rounds that are novel in the sense that they feature a price and or prior belief that differs from previous rounds. Table 2 reports the results of 3 different Logit regressions on the probability of subjects choosing to GMI in rounds.

	no on pro	Journa	or onn
Constant	-1.645	-1.676	-1.747
	(0.050)	(0.052)	(0.054)
	· /	· · · ·	
Round number	-0.026	-0.025	-0.025
	(0.001)	(0.001)	(0.001)
	(0001)	(0.001)	(0.001)
$ n-\mu(\overline{\mu}) $	-2.957	-2.964	-3.033
$ P \mu(\omega) $	(0.258)	(0.258)	(0.258)
	(0.250)	(0.200)	(0.258)
Number of other rounds	0.072	0.072	0.072
Number of other rounds	0.073	(0.073)	0.073
the subject chose GMI in	(0.001)	(0.001)	(0.001)
Dummy for $\mu(\overline{\omega}) = \frac{5}{4}$	0.737	0.698	0.649
	(0.124)	(0.126)	(0.126)
Decision problem different	0.153	0.148	0.129
than 1 round before	(0.047)	(0.047)	(0.047)
Decision problem different		0.096	0.072
than 2 rounds before		(0.046)	(0.047)
than 2 rounds before		(0.040)	(0.047)
Desigion problem different			0.224
Decision problem unterent			0.224
than 3 rounds before			(0.046)

Table 2: Logit regressions on probability of GMI

If the model fitted in the second last column in Table 2 is used to predict the probability of a subject that has gone through half of the decision problems and chose

A decision problem is said to be different than the decision problem a certain number of rounds before if the price and or the belief differ across the rounds.

to GMI in 26 of their other decision problems (26 is the median number of times to GMI) choosing GMI in a round with $p = \mu(\overline{\omega}) = 0.25$, then the decision problem being the same as the previous two rounds (same price and prior as the previous two rounds) results in a 23.6% chance of them choosing GMI whereas the decision problem being different from the previous two rounds (different price and or prior from the previous two rounds) results in a 28.2% chance of them choosing GMI. If, instead, the model fitted in the last column in Table 2 is used to predict the probability of a subject that has gone through half of the decision problems and chose to GMI in 26 of their other decision problems choosing GMI in a round with $p = \mu(\overline{\omega}) = 0.25$, then the decision problem being the same as the previous three rounds results in a 22.5% chance of them choosing GMI whereas the decision problem being different from the previous three rounds results in a 30.8% chance of them choosing GMI. Proportionally, both of these changes are large. This is surprising given how little variation there is between different decision problems in the experiment. Further, the coefficients on the round number in Table 2 is one indication of the fatigue that subjects experience during the experiment.

The impacts of fatigue and the "novelty" of a decision problem that are demonstrated in Table 2 suggest that even if behavior is aggregated across the decisions of a single individual and not across multiple individuals, allowing for the cost function for information to vary across decision problems, as is done in Section 2.3, seems like a natural modelling decision.

4 Literature Review

Caplin and Dean (2015) and Lipnowski and Ravid (2023) are also interested in determining if choice data can be rationalized with inattention. Caplin and Dean (2015) provide two testable conditions that, given a belief and a utility function, are satisfied if and only if the data can be rationalized with a costly learning model. So, given behavior s and payoffs for option Y $u(\underline{\omega})$ and $u(\overline{\omega})$, their result allows one to test whether or not the behavior is rationalized by a costly learning model. Lipnowski and Ravid (2023) complimentarily show that even if the cost of information is well understood, predicting behavior is essentially impossible if the utility function is unknown, but that structure can be imposed across choice problems. The results from my paper also study a setting where the utility function is unknown, and provide conditions that characterize the set of observed behaviors that can be rationalized by a given costly learning model when appropriate payoffs are selected in the special case of a price change, but, further, introduces heterogeneous learning costs.

There are a number of other papers that also experimentally test the implications of models of costly learning. The experiment in this paper is inspired by the work of Dean and Neligh (2023), but in their paper they do not have a change in parameters that is equivalent to the change in price that is the primary focus of this paper, and do not observe the outcome of the subjects' decision to 'get more information.' Dewan and Neligh (2020) study a different set of costly learning tasks experimentally, but again they do not observe the decision to 'get more information,' and do not have a change in parameters that is equivalent to a change in price. When these papers change option values they primarily do so in a multiplicative fashion, i.e. they do something analogous to multiplying p and $u(\omega)$ by a constant so incentive to learn is unambiguously higher or unambiguously lower. Understanding how inattention changes when price changes is important because changes in price are so commonly observed in the real world, whereas changes in belief or a multiplicative change in payoffs are more difficult to observe. Ambuehl (2017) and Ambuehl, Ockenfels, and Stewart (2022) experimentally and theoretically explore environments with changes in parameters that are equivalent to changes in price, but they do not observe the decision to 'get more information.'

There is a growing literature that demonstrates the importance of inattention and choice mistakes in standard economic settings, even in ones where it seems acquiring information should be costless. Chetty, Looney, and Kroft (2009), for instance, show that including sales tax in posted prices reduces demand, presumably because consumers were underestimating sales tax when making purchasing decisions, and Taubinsky and Rees-Jones (2018) actually show that in such situations accounting for the heterogeneity of subjects' under-reaction to taxes is crucial for accurately estimating the efficiency loss from taxation. Papers in recent years have demonstrated the significance of inattention in a wide variety of fields such as finance (Huberman, 2001), labor search (Acharya & Wee, 2020), trade (Dasgupta & Mondria, 2018), and voting behavior (Shue & Luttmer, 2009).

5 Conclusion

This paper uses a novel data enrichment to show that experiment subjects are more likely to invest effort into learning about the value of options if simple choice parameters, like price, differ from previous choice problems. This increase in effort in 'unfamiliar' choice problems and fatigue mean that the behavior of many subjects violate even the most flexible model of costly learning, one that only assumes that not learning is costless and that randomizing over learning strategies has a cost equal to the average of the learning strategies' costs, if the cost for information is assumed to be constant across choice problems with the same prior beliefs. This observation motivates the introduction of heterogeneous decision makers into a standard and more restrictive (posterior separable) model of costly learning to better fit the data.

It is shown that the introduction of heterogeneous decision makers into a posterior separable model can rationalize choice behavior that cannot be rationalized by a representative decision maker version of the posterior separable model. However, introducing heterogeneous decision makers into a posterior separable model still provides more precise predictions than the most general model of costly learning that is studied. Further, it is argued that variation in the cost function for information, which there is evidence of in the experiment, explains violations of the predictions of the posterior separable model that have been identified by other papers. This is true because fatigue over the course of an experiment could cause subjects to stop learning and simply pick an option that is best at their prior belief, which can create behavioral patterns that contradict the posterior separable model even though the posterior separable model has a number of compelling foundations and is the standard in the field of rational inattention.

Appendix 1: Proof of Theorems

Proof of Theorem 1. I begin with necessity of each point. (i): if $\exists p \in \mathcal{P}$ with $\Pr(Y|\underline{\omega}, p) > \Pr(Y|\overline{\omega}, p)$, then behavior is not rationalized because the DM would have achieved a strictly higher gain in payoff at p by choosing either (depending on the value of $u(\underline{\omega})$ relative to p) $s = (\Pr(Y|\overline{\omega}, p), \Pr(Y|\overline{\omega}, p))$ or s = (1, 1)).

(ii) : Suppose not, so $p_1 < p_2$ and: $\Pr(Y|p_1) < \Pr(Y|p_2)$. But I have a contradiction again since if $s(p_1)$ is optimal at p_1 then it provides a weakly higher gain in payoff at p_1 than $s(p_2)$ does, and, when price increases to p_2 the decrease in gain in payoff from $s(p_1)$ is strictly smaller than the decrease in gain in payoff from $s(p_2)$: $(p_2 - p_1)\Pr(Y|p_1) < (p_2 - p_1)\Pr(Y|p_2)$.

(iii): Again, I proceed with a proof by contradiction. Suppose $\exists p_1 \in \mathcal{P}$ such that: $\Pr(Y|\underline{\omega}, p_1) = \Pr(Y|\overline{\omega}, p_1) \in (0, 1)$, and $\exists p_2 \in \mathcal{P} \setminus p_1$ such that $\Pr(Y|p_2) \in$ (0, 1). Notice that optimality of the DM's behavior at p_1 implies s = (0, 0) and s = (1, 1) are both optimal when price is p_1 since all information outcomes s = (x, x)have zero cost for $x \in [0, 1]$. If $p_2 < p_1$, then I have a contradiction because the DM could optimally pick s = (1, 1) at p_1 , so $\Pr(Y|p_1) = 1$, which combined with (ii) implies the DM is not behaving optimally at p_2 . If $p_2 > p_1$, then I have a contradiction because the DM could optimally pick s = (0, 0) at p_1 , so $\Pr(Y|p_1) = 0$, which combined with (ii) implies the DM is not behaving optimally at p_2 .

To show (i), (ii), and (iii) are together sufficient I assume they are all satisfied and order the prices in $\mathcal{P} = \{p_1, p_2, \ldots, p_n\}$, so $p_1 < p_2 < \cdots < p_n$. Further, I can assume that there is a price in \mathcal{P} where the DM learns, because if not I can easily rationalize the behavior by defining C_{μ} for $\underline{s} < \overline{s}$ using a steep enough plane through the diagonal of S (points (x, x) with $x \in [0, 1]$). (iii) thus tells me that there is not behavior at a price $p_i \in \mathcal{P}$ with $s(p_i) = (x, x)$ and $x \in (0, 1)$. It is further without loss to assume that $s(p_1) = (1, 1)$ and $s(p_n) = (0, 0)$, since if this is not the case I can add prices to \mathcal{P} , and data to the behavior, so that this is the case, and then rationalizing the richer dataset rationalizes the original dataset. There is then a highest $p \in \mathcal{P}$ such that $\Pr(Y|p) = 1$, denote it p_h , and a lowest $p \in \mathcal{P}$ such that $\Pr(Y|p) = 0$, denote it p_l . Notice $p_h < p_l$. Pick a mean $m = p_h$, which is a reference point for conducting mean preserving spreads on $u(\underline{\omega})$ and $u(\overline{\omega})$ and is shifted at the end of the argument. Then pick $u(\underline{\omega}) < p_1$ and $u(\overline{\omega}) > p_n$ such that $\mu(\underline{\omega})u(\underline{\omega}) + \mu(\overline{\omega})u(\overline{\omega}) = m$. Remember, for all $x \in [0, 1] : C_{\mu}(x, x) = 0$. I now begin recursively assigning costs to the other information outcomes. Define:

$$C_{\mu}(s(p_{l-1})) = \mu(\underline{\omega}) \Big(\underline{s}(p_{l-1})(u(\underline{\omega}) - p_l) \Big) + \mu(\overline{\omega}) \Big(\overline{s}(p_{l-1})(u(\overline{\omega}) - p_l) \Big)$$
$$= \underline{s}(p_{l-1})(m - p_l) + \mu(\overline{\omega}) \Big(\overline{s}(p_{l-1}) - \underline{s}(p_{l-1}) \Big) (u(\overline{\omega}) - p_l).$$

This means the DM is indifferent between $s(p_{l-1})$ and $s(p_l)$ when price is p_l , and thus strictly prefers $s(p_{l-1})$ to $s(p_l)$ when price is p_{l-1} because when price decreases the gain in payoff from $s(p_{l-1})$ strictly increases since there is a strictly positive probability of the DM selecting option Y when they choose $s(p_{l-1})$ (by construction) while the gain in payoff from $s(p_l)$ is zero (by construction). If this $C_{\mu}(s(p_{l-1}))$ is strictly positive I continue, and if it is not I do a mean preserving spread on $u(\underline{\omega})$ and $u(\overline{\omega})$ so it is. Next, if l - 2 > h, I let:

$$C_{\mu}(s(p_{l-2})) = \underline{s}(p_{l-2})(m - p_{l-1}) + \mu(\overline{\omega}) \Big(\overline{s}(p_{l-2}) - \underline{s}(p_{l-2}) \Big) (u(\overline{\omega}) - p_{l-1}) \\ -\underline{s}(p_{l-1})(m - p_{l-1}) - \mu(\overline{\omega}) \Big(\overline{s}(p_{l-1}) - \underline{s}(p_{l-1}) \Big) (u(\overline{\omega}) - p_{l-1}) + C_{\mu}(s(p_{l-1})) \Big)$$

This means the DM is indifferent between $s(p_{l-2})$ and $s(p_{l-1})$ when price is p_{l-1} , and thus weakly prefers $s(p_{l-2})$ to $s(p_{l-1})$ when price is p_{l-2} . If this $C_{\mu}(s(p_{l-2}))$ is strictly positive I continue, if it is not I do a mean preserving spread on $u(\underline{\omega})$ and $u(\overline{\omega})$ (updating $C_{\mu}(s(p_{l-1}))$) accordingly) so $C_{\mu}(s(p_{l-2}))$ is strictly positive, which works since the value of $-\mu(\overline{\omega}) \left(\overline{s}(p_{l-1}) - \underline{s}(p_{l-1})\right) (u(\overline{\omega}) - p_{l-1}) + C_{\mu}(s(p_{l-1}))$ does not change. I continue in this fashion until I have set $C_{\mu}(s(p_{h+1}))$. If I keep the mean m the same (equal to p_h), then the DM strictly prefers $s(p_{h+1})$ to $s(p_h)$ when price is p_{h+1} , since they prefer $s(p_{h+1})$ to $s(p_l)$, which they strictly prefer to $s(p_h)$. I now increase $u(\underline{\omega})$ and $u(\overline{\omega})$ by the same amount so that the mean m increases, and all $C_{\mu}(s(p_i)) > 0$ so that the equations I used to define costs are still satisfied, until the DM is indifferent between $s(p_{h+1})$ to $s(p_h)$ when price is p_{h+1} . As a result, the DM strictly prefers $s(p_h)$ to $s(p_{h+1})$ when price is p_h since there is a strictly higher unconditional chance the DM selects option Y when they pick $s(p_h)$ compared to $s(p_{h+1})$ by construction.

At each price p_i I assigned a cost to the information outcome so that the DM is indifferent at p_i between $s(p_i)$, and $s(p_{i-1})$. This implies, out of the set of information outcomes I observe in the behavior, the DM is selecting their strategy optimally at each price, since as price decreases the gain in payoff from a strategy increase by the unconditional probability of choosing option Y, and as price decreases the unconditional probability of selecting option Y increases.

Now, I assign an arbitrarily high value to $C_{\mu}(0, 1)$ (so that (0, 1) is strictly worse than s(p) for all $p \in \mathcal{P}$), if I have not already assigned a value to it, and then define C_{μ} on $s \in S$ such that $\underline{s} < \overline{s}$ to be the the maximal convex function that is equal to the $C_{\mu}(s(p))$ I assigned for all $p \in \mathcal{P}$.

Why is this possible? Is it instead possible I assigned a $C_{\mu}(s(p_i))$ that was strictly above the relevant convex combination of other $C_{\mu}(s(p_j))$'s I assigned? This is not possible, as I can show with a quick proof by contradiction. Assume there is a set of prices $\tilde{\mathcal{P}} = \{p_m, \ldots, p_k\} \subseteq \mathcal{P}$, and a price $p_i \in \mathcal{P}$ such that there are positive weights α_j that sum to one such that $\alpha_m s(p_m) + \cdots + \alpha_k s(p_k) = s(p_i)$, but at the same time $\alpha_m C_{\mu}(s(p_m)) + \cdots + \alpha_k C_{\mu}(s(p_k)) < C_{\mu}(s(p_i))$. Since $\underline{s}(u(\underline{\omega}) - p)\mu(\underline{\omega}) + \overline{s}(u(\overline{\omega}) - p)\mu(\overline{\omega})$ is linear in the probabilities \underline{s} and \overline{s} , this implies when price is p_i the DM strictly prefers randomizing over their strategies from the $\tilde{\mathcal{P}} = \{p_m, \ldots, p_k\}$ prices compared to selecting $s(p_i)$, but this implies there is a $p_j \in \tilde{\mathcal{P}} = \{p_m, \ldots, p_k\}$ such that the DM strictly prefers $s(p_j)$ to $s(p_i)$ at p_i , which was ruled out by my recursive definition for the costs of information outcomes.

Similarly, at each p_i the information outcome of the DM is optimal given choice from the entire set of S since at each s with $\underline{s} \leq \overline{s}$ the cost is a convex combination of costs from information outcomes used for prices in \mathcal{P} and the corner (0, 1), and if the DM would prefer to switch to a different s at p_i , then they strictly prefer randomizing over the set of information outcomes used to generate the cost of s, and again one of the information outcomes used to generate the cost at s would then have been strictly preferred at p_i , and my recursive definition for the cost of information outcomes (and choice of arbitrarily high $C_{\mu}(0, 1)$) rules this out.

Proof of Theorem 2. I assume there is a $p \in \mathcal{P}$ such that $\underline{s}(p) < \overline{s}(p)$, otherwise the necessary conditions are trivially established, and sufficiency is easy to establish by making learning costly enough (since belief is fixed).

When the DM pays for information according to a measure of informedness (weakly convex c), the way for them to maximize their expected gain in payoff is to maximize the weighted average over option specific net gains in utility $V(Y|p, \cdot)$ and $V(X|\cdot)$, defined in the next paragraph. Each option specific net gain in utility takes into account the expected gain in payoff of the relevant option, X or Y, given the DM's posterior, and the cost of the posterior reached by the DM when they choose said option, where the posterior can be described by the resultant probability of $\overline{\omega}$ being realized, $\Pr(\overline{\omega}|Y, p)$ or $\Pr(\overline{\omega}|X, p)$ respectively.

If the DM does no learning they choose X or Y depending on the expected gain in payoff from selecting option Y, $m - p = \mu(\underline{\omega})u(\underline{\omega}) + \mu(\overline{\omega})u(\overline{\omega}) - p$, and which is larger as a result, $V(Y|p, \mu(\overline{\omega})) = m - p$ or $V(X|\mu(\overline{\omega})) = 0$. If the DM does some learning, then when they choose X their gain in payoff is $V(X|Pr(\overline{\omega}|X, p)) =$ $-c(Pr(\overline{\omega}|X, p)) + c(\mu(\overline{\omega}))$, and when they select Y their gain in payoff is $V(Y|p, Pr(\overline{\omega}|Y, p)) = Pr(\overline{\omega}|Y, p)(u(\overline{\omega}) - p) + (1 - Pr(\overline{\omega}|Y, p))(u(\underline{\omega}) - p)$ $-c(Pr(\overline{\omega}|Y, p)) + c(\mu(\overline{\omega}))$. Both $V(Y|p, \cdot)$ and $V(X|\cdot)$ are weakly concave functions since c is weakly convex. The DM is maximizing the weighted average of the two. Notice that in any optimal solution, if the DM is learning, $Pr(\overline{\omega}|X, p) < \mu(\overline{\omega}) <$ $Pr(\overline{\omega}|Y, p)$. As a result, to find the optimal solution of the DM I must find the cord from the top of $V(X|\cdot)$ on the left of $\mu(\overline{\omega})$ to the top of $V(Y|p, \cdot)$ on the right, so that I have a weakly concave closure.

When p increases, either the DM stops learning, which means the concave closure at $\mu(\overline{\omega})$ is $V(X|\mu(\overline{\omega}))$, or the DM continues to learn, in which case, $V(Y|p, \cdot)$ shifts downward at every point by the change in price, and the point where the cord that creates the weakly concave closure hits $V(Y|p, \cdot)$ and $V(X|\cdot)$ must then both weakly move to the right, which means $\Pr(\overline{\omega}|X, p)$ and $\Pr(\overline{\omega}|Y, p)$ are both weakly increasing, which establishes necessity.

Next I show sufficiency by constructing a weakly convex c function that generates the observed behavior. It is without loss to assume $\mathcal{P} = \{p_1, p_2, \ldots, p_n\}$, with $p_1 < \cdots < p_n$, and $\Pr(Y|\Omega, \mathcal{P})$ are such that $n \ge 4$ (in the example graphs, n = 6), with $\Pr(Y|p_1) = 1$, $\Pr(Y|p_n) = 0$, and $\Pr(Y|p_i) \in (0, 1)$ for $p_i \in \mathcal{P} \setminus \{p_1, p_n\}$, since if this is not the case I can generate such a dataset by adding more prices with behavior that satisfy the conditions since rationalizing this richer dataset rationalizes the original dataset. I draw the graphs step by step and the different components are going to help me pick $u(\underline{\omega})$ and $u(\overline{\omega})$, and tell me a suitable c. The horizontal axis goes from zero to one, I call the horizontal coordinate z since X is already being used. The vertical axis may take positive and negative values, I call this coordinate height, whether it be positive or negative, since Y is being used.

First (graph a): I draw a line segment from $z = \Pr(\overline{\omega}|X, p_{n-1})$ and a positive height, to z=1 and a negative height, so that when $z = \mu(\overline{\omega})$, the height is 0, and so that the height at $\Pr(\overline{\omega}|Y, p_{n-1})$ is -w such that $p_{n-1} - p_2 - w > 0$. The part of the segment between $z = \Pr(\overline{\omega}|\mathbf{X}, p_{n-1})$ and a $z = \mu(\overline{\omega})$ I make red, the segment between $z = \Pr(\overline{\omega}|Y, p_{n-2})$ and z = 1 I make blue, and the rest I erase. Second (graph b): I take the blue segment and increase the height by $p_{n-1} - p_{n-2}$, then I draw a new line segment from the blue line at $z = \Pr(\overline{\omega}|Y, p_{n-2})$, through the red segment at $z = \Pr(\overline{\omega}|X, p_{n-1})$ and continue on to $z = \Pr(\overline{\omega}|X, p_{n-2})$. The part of the new segment between $z = \Pr(\overline{\omega}|\mathbf{X}, p_{n-1})$ and $z = \Pr(\overline{\omega}|\mathbf{X}, p_{n-2})$ I make red, the part of the new segment between $z = \Pr(\overline{\omega}|Y, p_{n-2})$ and $z = \Pr(\overline{\omega}|Y, p_{n-3})$ I make blue, and the rest I erase. Third (graph c): I take the blue segment and increase the height by $p_{n-2} - p_{n-3}$, then I draw a new line segment from the blue line at $z = \Pr(\overline{\omega}|\mathbf{Y}, p_{n-3})$, through the red segment at $z = \Pr(\overline{\omega}|\mathbf{X}, p_{n-2})$ and continue on to $z = \Pr(\overline{\omega}|\mathbf{X}, p_{n-3})$. The part of the new segment between $z = \Pr(\overline{\omega}|\mathbf{X}, p_{n-2})$ and $z = \Pr(\overline{\omega}|\mathbf{X}, p_{n-3})$ I make red, the part of the new segment between $z = \Pr(\overline{\omega}|\mathbf{Y}, p_{n-3})$ and $z = \Pr(\overline{\omega}|Y, p_{n-4})$ I make blue, and the rest I erase.¹¹ Eventually (graph d), after continuing in the above fashion, I take the blue segment and increase the height by $p_3 - p_2$, then I draw a new line segment from the blue line at $z = \Pr(\overline{\omega}|\mathbf{Y}, p_2)$, through the red segment at $z = \Pr(\overline{\omega}|X, p_3)$ and continue on to z = 0. The part of the new segment between $z = \Pr(\overline{\omega}|X, p_3)$ and z = 0 I make red, the part of the new segment between $z = \mu(\overline{\omega})$ and $z = \Pr(\overline{\omega}|Y, p_2)$ I make blue, and the rest I erase.

The red line segments I take to be $-c(z) + c(\mu(\overline{\omega}))$ for $z \in [0, \mu(\overline{\omega})]$. Next, let b(z) denote the height of the blue segment for $z \in [\mu(\overline{\omega}), 1]$ (graph d). $b(\mu(\overline{\omega})) > 0$

¹¹In the graphs $\Pr(\overline{\omega}|\mathbf{X}, p_{n-2}) = \Pr(\overline{\omega}|\mathbf{X}, p_{n-3})$, which is meant to provide some insight into what happens when $\Pr(\overline{\omega}|\mathbf{X}, p)$ or $\Pr(\overline{\omega}|\mathbf{Y}, p)$ are only weakly decreasing in p.



because either the slope of the blue segments is negative (where it is defined), in which case $b(\mu(\overline{\omega}))$ is more than $b(\Pr(\overline{\omega}|Y, p_{n-1}))$, which is strictly positive based on how I chose -w at the beginning, or the slope of the blue segments is positive or zero immediately to the right of $z = \mu(\overline{\omega})$, in which case $b(\mu(\overline{\omega}))$ is greater or equal to the height of the red segments at $\Pr(\overline{\omega}|X, p_{n-1})$, which is strictly positive by construction. Further, $b(\mu(\overline{\omega})) - p_{n-1} + p_2 < 0$, again based on how I picked w. Next I pick the mean quality to be $m = p_2 + b(\mu(\overline{\omega})) \in (p_2, p_{n-1})$ so that when p = mthe DM is indifferent between choosing X without learning and choosing Y without learning. Next, I pick $u(\underline{\omega})$ and $u(\overline{\omega})$ so that $\mu(\underline{\omega})u(\underline{\omega}) + \mu(\overline{\omega})u(\overline{\omega}) = m$, and so that (to help ensure convexity of c):

$$u(\overline{\omega}) - u(\underline{\omega}) + \frac{c(\mu(\overline{\omega})) - c(\Pr(\overline{\omega}|\mathbf{X}, p_{n-1}))}{\mu(\overline{\omega}) - \Pr(\overline{\omega}|\mathbf{X}, p_{n-1})} \geq \frac{b(\Pr(\overline{\omega}|\mathbf{Y}, p_2)) - b(\mu(\overline{\omega}))}{\Pr(\overline{\omega}|\mathbf{Y}, p_2) - \mu(\overline{\omega})}$$

Next, for $z \in [\mu(\overline{\omega}), 1]$ I let $-c(z) + c(\mu(\overline{\omega})) = b(z) - b(\mu(\overline{\omega})) - (z - \mu(\overline{\omega}))(u(\overline{\omega}) - u(\underline{\omega}))$. Finally, I fix $c(\mu(\overline{\omega}))$ so that $\min_{z \in [0, 1]} c(z) = 0$.

Proof of Theorem 3. I begin with a helpful lemma.

Lemma 1. Given prices \underline{p} and $\overline{p} > \underline{p}$, payoffs for option Y $u(\underline{\omega})$ and $u(\overline{\omega})$ such that



 $u(\underline{\omega}) \leq \underline{p}$ and $u(\overline{\omega}) \geq \overline{p}$, and a type t with a prior belief such that $\mu_t(\overline{\omega}) \in (0, 1)$ and $\mu_t(\underline{\omega}) \equiv 1 - \mu_t(\overline{\omega})$, I can create a measure of informedness $c_t : [0, 1] \to \mathbb{R}_+$ such that DMs of type t are indifferent between doing no learning ($\underline{s} = \overline{s}$) and perfectly observing the payoff from option Y ($\underline{s} = 0, \overline{s} = 1$) when the price is $p \in \{\underline{p}, \overline{p}\}$, strictly prefers perfectly observing the payoff from option Y over all other learning strategies (any other probabilities of selecting Y in each state) when the price is $p \in (\underline{p}, \overline{p})$, and strictly prefers doing no learning over all other learning strategies when $p \notin [\underline{p}, \overline{p}]$, if the payoffs from option Y are a mean preserving spread of the prices: $\mu_t(\underline{\omega})u(\underline{\omega}) + \mu_t(\overline{\omega})u(\overline{\omega}) = \mu_t(\underline{\omega})\underline{p} + \mu_t(\overline{\omega})\overline{p}$.

Proof of Lemma 1. In the fashion of the proof of Theorem 2, I construct the functions V(X|z) in red on the left of $\mu_t(\overline{\omega})$ and V(Y|p, z) in blue on the right of $\mu_t(\overline{\omega})$ (graphs below). When $p = \mu_t(\underline{\omega})u(\underline{\omega}) + \mu_t(\overline{\omega})u(\overline{\omega}) = \mathbb{E}_t[u(\omega)]$, these functions must both be equal to zero. So, I am going to draw V(X|z) and $V(Y|\mathbb{E}_t[u(\omega)], z)$ as line segments on either side of $\mu_t(\overline{\omega})$, so that $V(X|\mu_t(\overline{\omega})) = V(Y|\mathbb{E}_t[u(\omega)], \mu_t(\overline{\omega})) = 0$ (graph e), and so that they satisfy two other properties. (i): First, when the blue line segment is shifted up by $\mathbb{E}_t[u(\omega)] - \underline{p}$, it must be that if I were to extend the blue line segment so that it reaches z = 0, it hits the red line segment at z = 0 (graph f), so

that when $p = \underline{p}$ the DM is indifferent between learning nothing and everything, and at all lower prices they strictly prefer learning nothing. (ii): Second, it must be that when the blue line segment is shifted down by $\overline{p} - \mathbb{E}_t[u(\omega)]$, it must be that if I were to extend the red line segment so that it reaches z = 1, it hits the blue line segment at z = 1 (graph g), so that when $p = \overline{p}$ the DM is indifferent between learning nothing and everything, and at all higher prices the DM strictly prefers learning nothing. Together, these properties imply that the slope of the blue line segment is strictly greater than the slope of the red line segment, and the DM strictly prefers learning everything over all other learning strategies when $p \in (p, \overline{p})$.



How do I find line segments that satisfy the two properties? I select the height of the red segment at zero (-b in graph e) and the blue segment at one (w in graph e) so that:

(i):
$$\frac{w}{\mu_t(\underline{\omega})} = b + w + (\mathbb{E}_t[u(\omega)] - \underline{p}), \text{ (ii)}: \frac{b}{\mu_t(\overline{\omega})} = b + w - (\overline{p} - \mathbb{E}_t[u(\omega)])$$
$$\Leftrightarrow b = \frac{\mu_t(\overline{\omega})w}{\mu_t(\underline{\omega})} - (\mathbb{E}_t[u(\omega)] - \underline{p}), b = \frac{\mu_t(\overline{\omega})w}{\mu_t(\underline{\omega})} - \frac{\mu_t(\overline{\omega})}{\mu_t(\underline{\omega})}(\overline{p} - \mathbb{E}_t[u(\omega)]).$$
But:
$$\mathbb{E}_t[u(\omega)] = \mu_t(\underline{\omega})\underline{p} + \mu_t(\overline{\omega})\overline{p} \Rightarrow -(\mathbb{E}_t[u(\omega)] - \underline{p}) = -\frac{\mu_t(\overline{\omega})}{\mu_t(\underline{\omega})}(\overline{p} - \mathbb{E}_t[u(\omega)]),$$
so:
$$b = \frac{\mu_t(\overline{\omega})w}{\mu_t(\underline{\omega})} - (\mathbb{E}_t[u(\omega)] - \underline{p}) \Leftrightarrow b = \frac{\mu_t(\overline{\omega})w}{\mu_t(\underline{\omega})} - \frac{\mu_t(\overline{\omega})}{\mu_t(\underline{\omega})}(\overline{p} - \mathbb{E}_t[u(\omega)]).$$

Thus, given any b, I simply make $w = \frac{\mu_t(\underline{\omega})}{\mu_t(\overline{\omega})}b + (\overline{p} - \mathbb{E}_t[u(\omega)]).$

For $z \in [0, \mu_t(\overline{\omega})]$, I let $-c(z) + c(\mu_t(\overline{\omega})) = V(X|z)$. For $z \in [\mu_t(\overline{\omega}), 1]$, I let $-c(z) + c(\mu_t(\overline{\omega})) = V(Y|\mathbb{E}_t[u(\omega)], z) - (z - \mu_t(\overline{\omega}))(u(\overline{\omega}) - u(\underline{\omega}))$. Finally, fix $c(\mu(\overline{\omega}))$ so that $\min_{z \in [0,1]} c(z) = 0$. It is easy to show that, given how I picked w based on b, this procedure produces a weakly convex c.

I now return to the **proof of Theorem 3.** I begin by showing that for any type t it must be that $\Pr_t(Y|\underline{\omega}, p)$ and $\Pr_t(Y|\overline{\omega}, p)$ are weakly decreasing in p. If $\Pr_t(Y|\underline{\omega}, p) = \Pr_t(Y|\overline{\omega}, p) = 1$ clearly both need to decrease. If $\Pr_t(Y|\underline{\omega}, p) <$ $\Pr_t(Y|\overline{\omega}, p) \leq 1$, then if $\Pr_t(Y|\underline{\omega}, p)$ increases, Theorem 1 requires $\Pr_t(Y|\overline{\omega}, p)$ decreases, and Bayes' Rule tells us:

$$\Pr_t(\overline{\omega}|\mathbf{Y}, p) = \frac{\Pr_t(\mathbf{Y}|\overline{\omega}, p)\mu_t(\overline{\omega})}{\Pr_t(\mathbf{Y}|\underline{\omega}, p)\mu_t(\underline{\omega}) + \Pr_t(\mathbf{Y}|\overline{\omega}, p)\mu_t(\overline{\omega})} = \frac{1}{1 + \frac{\Pr_t(\mathbf{Y}|\underline{\omega}, p)\mu_t(\underline{\omega})}{\Pr_t(\mathbf{Y}|\overline{\omega}, p)\mu_t(\overline{\omega})}}$$

so $\Pr_t(\overline{\omega}|\mathbf{Y}, p)$ decreases, which violates Theorem 2. If $\Pr_t(\mathbf{Y}|\underline{\omega}, p) \leq \Pr_t(\mathbf{Y}|\overline{\omega}, p) < 1$, and $\Pr_t(\mathbf{Y}|\overline{\omega}, p)$ increases, then Theorem 1 tells me $\Pr_t(\mathbf{Y}|\underline{\omega}, p)$ decreases, and Bayes' Rule tells us:

$$\Pr_{t}(\overline{\omega}|\mathbf{X}, p) = \frac{\Pr_{t}(\mathbf{X}|\overline{\omega}, p)\mu_{t}(\overline{\omega})}{\Pr_{t}(\mathbf{X}|\underline{\omega}, p)\mu_{t}(\underline{\omega}) + \Pr_{t}(\mathbf{X}|\overline{\omega}, p)\mu_{t}(\overline{\omega})} = \frac{1}{1 + \frac{\Pr_{t}(\mathbf{X}|\underline{\omega}, p)\mu_{t}(\underline{\omega})}{\Pr_{t}(\mathbf{X}|\overline{\omega}, p)\mu_{t}(\overline{\omega})}}$$

so $\Pr_t(\overline{\omega}|\mathbf{X}, p)$ decreases, which violates Theorem 2. As a result $\Pr(\mathbf{Y}|\underline{\omega}, p)$ and $\Pr(\mathbf{Y}|\overline{\omega}, p)$ both being weakly decreasing is necessary for $\Pr(\mathbf{Y}|\Omega, \mathcal{P})$ to be rationalized by an aggregate PS model.

Sufficiency is the challenge. I assume there is a $p \in \mathcal{P}$ such that $\underline{s}(p) < \overline{s}(p)$, otherwise sufficiency is easy to establish by making learning costly enough (since belief is fixed). It is without loss to assume $\mathcal{P} = \{p_1, p_2, \ldots, p_n\}$, with $p_1 < \cdots < p_n$. Notice that it is also without loss to assume that between any two adjacent prices p_i and p_{i+1} , if $\Pr(Y|p)$ decreases, then only one of $\Pr(Y|\underline{\omega}, p)$ and $\Pr(Y|\overline{\omega}, p)$ decreases, since otherwise I can always enrich $\Pr(Y|\Omega, \mathcal{P})$ in a way so that behavior is still rationalized by a costly learning model, and so that this is true. I can then rationalize the richer dataset, which in turn rationalizes the less rich dataset. Similarly, I can assume without loss that $Pr(Y|p_1) = 1$ and $Pr(Y|p_n) = 0$, and that there are at least two pairs of prices between which $Pr(Y|\overline{\omega}, p)$ strictly decreases. Notice that the change between p_1 and p_2 is then a reduction in $Pr(Y|\underline{\omega}, p)$, while the change between p_{n-1} and p_n is then a reduction in $Pr(Y|\overline{\omega}, p)$.

I am now going to start pairing reductions in $\Pr(Y|\underline{\omega}, p)\mu(\underline{\omega})$ with reductions in $\Pr(Y|\overline{\omega}, p)\mu(\overline{\omega})$. The goal is to rationalize each pair with a type using Lemma 1, but I have to be careful about how I construct the pairs, and I start off with an initial pairing that intentionally fails in an eventually fixable way. Initially (henceforth initial pairing), pair reductions $\delta_{\underline{\omega}}$ in $\Pr(Y|\underline{\omega}, p)\mu(\underline{\omega})$ between two adjacent prices, with reductions $\delta_{\overline{\omega}}$ in $\Pr(Y|\overline{\omega}, p)\mu(\overline{\omega})$ between higher adjacent prices that are not p_{n-1} and p_n (none of the reduction in $\Pr(Y|\overline{\omega}, p)\mu(\overline{\omega})$ between p_{n-1} and p_n is paired in the initial pairing), so that the following four properties are satisfied. First, all reduction in $\Pr(Y|\underline{\omega}, p)$ is paired, all reduction in $\Pr(Y|\overline{\omega}, p)$ except that between p_{n-1} and p_n is paired. Second, each pairing has $\frac{\delta_{\underline{\omega}}}{\mu(\underline{\omega})} > \frac{\delta_{\overline{\omega}}}{\mu(\overline{\omega})}$. Third, the pair that has the strictly lowest $\delta_{\underline{\omega}}/(\delta_{\underline{\omega}} + \delta_{\overline{\omega}})$ is a pairing (denoted $(\delta_{\underline{\omega}}, \delta_{\overline{\omega}})$) between the lowest pair of prices that have a reduction in $\Pr(Y|\overline{\omega}, p)\mu(\overline{\omega})$, call them p_j, p_{j+1} . Fourth, if I pick $u(\overline{\omega}) = p_n$, and pick $u(\underline{\omega}) < p_1$ so that:

$$\frac{\tilde{\delta_{\omega}}u(\underline{\omega}) + \tilde{\delta_{\omega}}u(\overline{\omega})}{\tilde{\delta_{\omega}} + \tilde{\delta_{\omega}}} = \frac{\tilde{\delta_{\omega}}p_1 + \tilde{\delta_{\omega}}p_j}{\tilde{\delta_{\omega}} + \tilde{\delta_{\omega}}},\tag{1}$$

$$\frac{(\mu(\underline{\omega}) - \tilde{\delta_{\underline{\omega}}})u(\underline{\omega}) + (\mu(\overline{\omega}) - \tilde{\delta_{\overline{\omega}}})u(\overline{\omega})}{(\mu(\underline{\omega}) - \tilde{\delta_{\underline{\omega}}}) + (\mu(\overline{\omega}) - \tilde{\delta_{\overline{\omega}}})} < \frac{(\mu(\underline{\omega}) - \tilde{\delta_{\underline{\omega}}})p_2 + (\mu(\overline{\omega}) - \tilde{\delta_{\overline{\omega}}})p_{j+1}}{(\mu(\underline{\omega}) - \tilde{\delta_{\underline{\omega}}}) + (\mu(\overline{\omega}) - \tilde{\delta_{\overline{\omega}}})}.$$
 (2)

Why is it possible to satisfy all four properties simultaneously? The second property is possible because of the first property, and ensures that a reduction in $\Pr(Y|\underline{\omega}, p)$ is always paired with a smaller reduction in $\Pr(Y|\overline{\omega}, p)$. Then, when I pick the pair with the strictly lowest $\delta_{\underline{\omega}}/(\delta_{\underline{\omega}} + \delta_{\overline{\omega}})$ in the third property, I can make sure the reduction in $\Pr(Y|\underline{\omega}, p)$ is paired with an arbitrarily close in size reduction in $\Pr(Y|\overline{\omega}, p)$, which makes the two values:

$$\frac{\tilde{\delta_{\omega}}u(\underline{\omega}) + \tilde{\delta_{\overline{\omega}}}u(\overline{\omega})}{\tilde{\delta_{\underline{\omega}}} + \tilde{\delta_{\overline{\omega}}}} \text{ and } \frac{(\mu(\underline{\omega}) - \tilde{\delta_{\underline{\omega}}})u(\underline{\omega}) + (\mu(\overline{\omega}) - \tilde{\delta_{\overline{\omega}}})u(\overline{\omega})}{(\mu(\underline{\omega}) - \tilde{\delta_{\underline{\omega}}}) + (\mu(\overline{\omega}) - \tilde{\delta_{\overline{\omega}}})}$$

arbitrarily close together, which ensures the fourth property can be satisfied since satisfying (1) and continuity then implies (2).

The behavior of the $(\tilde{\delta}_{\underline{\omega}}, \tilde{\delta}_{\overline{\omega}})$ pairing can be rationalized by a single type according to Lemma 1, and this continues to be true as long as I spread $u(\underline{\omega})$ and $u(\overline{\omega})$ away from each other so (1) is satisfied. Further, the second property then implies such spreads that maintain the equality in (1) increase the mean $m = \mu(\underline{\omega})u(\underline{\omega}) + \mu(\overline{\omega})u(\overline{\omega})$. Then, looking at how the ratios change when the weights change for reductions $\delta_{\underline{\omega}}$ between p_i and p_{i+1} and $\delta_{\overline{\omega}}$ between p_k and p_{k+1} with i + 1 > 1 and k + 1 > j:

$$\frac{\partial \frac{\delta_{\underline{\omega}} u(\underline{\omega}) + \delta_{\overline{\omega}} u(\overline{\omega})}{\delta_{\underline{\omega}} + \delta_{\overline{\omega}}}}{\partial \delta_{\underline{\omega}}} = \frac{\delta_{\overline{\omega}} (u(\underline{\omega}) - u(\overline{\omega}))}{(\delta_{\underline{\omega}} + \delta_{\overline{\omega}})^2} < \frac{\delta_{\overline{\omega}} (p_{i+1} - p_{k+1})}{(\delta_{\underline{\omega}} + \delta_{\overline{\omega}})^2} = \frac{\partial \frac{\delta_{\underline{\omega}} p_{i+1} + \delta_{\overline{\omega}} p_{k+1}}{\delta_{\underline{\omega}} + \delta_{\overline{\omega}}}}{\partial \delta_{\underline{\omega}}}, \quad (3)$$
so $\frac{\delta_{\underline{\omega}} u(\underline{\omega}) + \delta_{\overline{\omega}} u(\overline{\omega})}{\delta_{\underline{\omega}} + \delta_{\overline{\omega}}} < \frac{\delta_{\underline{\omega}} p_{i+1} + \delta_{\overline{\omega}} p_{k+1}}{\delta_{\underline{\omega}} + \delta_{\overline{\omega}}}.$

What I do next is I reduce each $\delta_{\underline{\omega}}$ to $\delta_{\underline{\omega}}$ (leaving the excess unpaired), so:

$$\frac{\hat{\delta_{\underline{\omega}}}u(\underline{\omega}) + \delta_{\overline{\omega}}u(\overline{\omega})}{\hat{\delta_{\underline{\omega}}} + \delta_{\overline{\omega}}} = \frac{\hat{\delta_{\underline{\omega}}}p_{i+1} + \delta_{\overline{\omega}}p_{k+1}}{\hat{\delta_{\underline{\omega}}} + \delta_{\overline{\omega}}} \ge \frac{\hat{\delta_{\underline{\omega}}}p_2 + \delta_{\overline{\omega}}p_{j+1}}{\hat{\delta_{\underline{\omega}}} + \delta_{\overline{\omega}}} > \frac{\hat{\delta_{\underline{\omega}}}p_1 + \delta_{\overline{\omega}}p_j}{\hat{\delta_{\underline{\omega}}} + \delta_{\overline{\omega}}}.$$
 (4)

Then (3) tells me there is a unique $\hat{\delta_{\omega}}$ that satisfies the equality in (4). Further, the inequalities from (4) and (1) and (2) imply that for each resultant pairing:

$$\frac{\hat{\delta_{\underline{\omega}}}}{\hat{\delta_{\underline{\omega}}} + \delta_{\overline{\omega}}} < \frac{\tilde{\delta_{\underline{\omega}}}}{\tilde{\delta_{\underline{\omega}}} + \tilde{\delta_{\overline{\omega}}}}, \text{ and } \frac{\hat{\delta_{\underline{\omega}}}}{\hat{\delta_{\underline{\omega}}} + \delta_{\overline{\omega}}} < \frac{\mu(\underline{\omega}) - \tilde{\delta_{\underline{\omega}}}}{(\mu(\underline{\omega}) - \tilde{\delta_{\underline{\omega}}}) + (\mu(\overline{\omega}) - \tilde{\delta_{\overline{\omega}}})}.$$
(5)

So (4), (5) and (2) say:
$$\frac{\hat{\delta_{\omega}}u(\underline{\omega}) + \delta_{\overline{\omega}}u(\overline{\omega})}{\hat{\delta_{\omega}} + \delta_{\overline{\omega}}} = \frac{\hat{\delta_{\omega}}p_{i+1} + \delta_{\overline{\omega}}p_{k+1}}{\hat{\delta_{\omega}} + \delta_{\overline{\omega}}} \ge \frac{\hat{\delta_{\omega}}p_2 + \delta_{\overline{\omega}}p_{j+1}}{\hat{\delta_{\omega}} + \delta_{\overline{\omega}}}$$
$$> \frac{(\mu(\underline{\omega}) - \tilde{\delta_{\omega}})p_2 + (\mu(\overline{\omega}) - \tilde{\delta_{\omega}})p_{j+1}}{(\mu(\underline{\omega}) - \tilde{\delta_{\omega}}) + (\mu(\overline{\omega}) - \tilde{\delta_{\omega}})u(\underline{\omega})} + (\mu(\overline{\omega}) - \tilde{\delta_{\omega}})u(\overline{\omega})}{(\mu(\underline{\omega}) - \tilde{\delta_{\omega}}) + (\mu(\overline{\omega}) - \tilde{\delta_{\omega}})},$$

which means that each pair $(\hat{\delta_{\omega}}, \delta_{\overline{\omega}})$, not including $(\tilde{\delta_{\omega}}, \tilde{\delta_{\overline{\omega}}})$, has a higher mean than the mean that is left after removing only $(\tilde{\delta_{\omega}}, \tilde{\delta_{\overline{\omega}}})$, so the pairings with $\hat{\delta_{\omega}}$'s have reduced the mean of what is left over.

But, I need to match the reductions in $\Pr(Y|\underline{\omega}, p)\mu(\underline{\omega})$ that are now unpaired with the reduction in $\Pr(Y|\overline{\omega}, p)\mu(\overline{\omega})$ between p_{n-1} and p_n . So, beginning with the lowest pair of prices with unmatched change and working my way up I take all of the unmatched reduction in $\Pr(Y|\underline{\omega}, p)\mu(\underline{\omega})$ between p_m and p_{m+1} , denoted $\overline{\delta_{\omega}}$, and match it with enough reduction in $\Pr(Y|\overline{\omega}, p)\mu(\overline{\omega})$ from between p_{n-1} and p_n , denoted $\overline{\delta_{\omega}}$, so that:

$$\frac{\overline{\delta_{\underline{\omega}}}u(\underline{\omega}) + \overline{\delta_{\overline{\omega}}}u(\overline{\omega})}{\overline{\delta_{\underline{\omega}}} + \overline{\delta_{\overline{\omega}}}} = \frac{\overline{\delta_{\underline{\omega}}}p_{m+1} + \overline{\delta_{\overline{\omega}}}p_n}{\overline{\delta_{\underline{\omega}}} + \overline{\delta_{\overline{\omega}}}},$$

$$\frac{\partial \frac{\delta_{\underline{\omega}}u(\underline{\omega}) + \delta_{\overline{\omega}}u(\overline{\omega})}{\delta_{\underline{\omega}} + \overline{\delta_{\overline{\omega}}}}}{\overline{\delta_{\underline{\omega}}} + \overline{\delta_{\overline{\omega}}}} = \frac{\delta_{\underline{\omega}}(u(\overline{\omega}) - u(\underline{\omega}))}{(\overline{\delta_{\underline{\omega}}} + \overline{\delta_{\overline{\omega}}})^2} > \frac{\delta_{\underline{\omega}}(p_n - p_{m+1})}{(\overline{\delta_{\underline{\omega}}} + \overline{\delta_{\overline{\omega}}})^2} = \frac{\partial \frac{\delta_{\underline{\omega}}p_{m+1} + \overline{\delta_{\overline{\omega}}}p_n}{\overline{\delta_{\underline{\omega}}} + \overline{\delta_{\overline{\omega}}}}, \quad (6)$$

so there is a unique $\overline{\delta_{\overline{\omega}}}$ for each $\overline{\delta_{\underline{\omega}}}$. But,

$$\frac{\bar{\delta_{\omega}}p_{m+1} + \bar{\delta_{\omega}}p_n}{\bar{\delta_{\omega}} + \bar{\delta_{\omega}}} > \frac{\bar{\delta_{\omega}}p_2 + \bar{\delta_{\omega}}p_{j+1}}{\bar{\delta_{\omega}} + \bar{\delta_{\omega}}} > \frac{(\mu(\underline{\omega}) - \tilde{\delta_{\omega}})p_2 + (\mu(\overline{\omega}) - \tilde{\delta_{\omega}})p_{j+1}}{(\mu(\underline{\omega}) - \tilde{\delta_{\omega}}) + (\mu(\overline{\omega}) - \tilde{\delta_{\omega}})} > \frac{(\mu(\underline{\omega}) - \tilde{\delta_{\omega}})u(\underline{\omega}) + (\mu(\overline{\omega}) - \tilde{\delta_{\omega}})u(\overline{\omega})}{(\mu(\underline{\omega}) - \tilde{\delta_{\omega}}) + (\mu(\overline{\omega}) - \tilde{\delta_{\omega}})},$$

where the second inequality is true due to (3) or (6), and third is due to (2), which in a fashion similar to how (4) implies (5), implies:

$$\frac{\bar{\delta_{\underline{\omega}}}}{\bar{\delta_{\underline{\omega}}} + \bar{\delta_{\overline{\omega}}}} < \frac{\tilde{\delta_{\underline{\omega}}}}{\tilde{\delta_{\underline{\omega}}} + \tilde{\delta_{\overline{\omega}}}}, \text{ and } \frac{\bar{\delta_{\underline{\omega}}}}{\bar{\delta_{\underline{\omega}}} + \bar{\delta_{\overline{\omega}}}} < \frac{(\mu(\underline{\omega}) - \tilde{\delta_{\underline{\omega}}})}{(\mu(\underline{\omega}) - \tilde{\delta_{\underline{\omega}}}) + (\mu(\overline{\omega}) - \tilde{\delta_{\overline{\omega}}})}.$$
(7)

This all means I am doomed to failure since the mean required by each remaining pair is strictly higher than the mean of the unmatched reductions I have left and I run out of reduction in $\Pr(Y|\overline{\omega}, p)\mu(\overline{\omega})$ from between p_{n-1} and p_n before I have paired all the unpaired reduction in $\Pr(Y|\underline{\omega}, p)\mu(\underline{\omega})$.

How do I do better? I return to the initial pairing, but increase $u(\overline{\omega})$ and decrease $u(\underline{\omega})$ so that (1) is satisfied, which increases the mean of what is left over after the $(\delta_{\underline{\omega}}, \delta_{\overline{\omega}})$ pairing (eventually (2) is violated as a result, but I do not need (2) to be satisfied by the final solution, and its eventual violation makes this strategy work). Then, (5) tells me that the $\delta_{\underline{\omega}}$ I had previously picked to satisfy the equality in (4) were too low, so this iteration I reduce each $\delta_{\underline{\omega}}$ less (but still reduce given how I picked $(\delta_{\underline{\omega}}, \delta_{\overline{\omega}})$ to satisfy the second property), which means I have less unmatched reduction in $\Pr(Y|\underline{\omega}, p)\mu(\underline{\omega})$ to pair after producing the $\delta_{\underline{\omega}}$'s. Further, for any amount of unmatched $\Pr(Y|\underline{\omega}, p)\mu(\underline{\omega})$, the previous amount of reduction in $\Pr(Y|\overline{\omega}, p)\mu(\overline{\omega})$ from between p_{n-1} and p_n I would have matched it with is too large given (7), since I spread $u(\underline{\omega})$ and $u(\overline{\omega})$ but (1) is satisfied, so the unmatched $\Pr(Y|\overline{\omega}, p)\mu(\overline{\omega})$ goes farther, and eventually this strategy works.

References

- Acharya, S., & Wee, S. L. (2020). Rational inattention in hiring decisions. American Economic Journal: Macroeconomics, 12(1), 1–40.
- Allais, M. (1953). Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'école américaine. Econometrica: Journal of the Econometric Society, 503–546.
- Ambuehl, S. (2017). An offer you can't refuse? incentives change how we inform ourselves and what we believe. CESifo Working Papers (6296).
- Ambuehl, S., Ockenfels, A., & Stewart, C. (2022). Who opts in? composition effects and disappointment from participation payments. The Review of Economics

and Statistics, 1-45.

- Azrieli, Y., Chambers, C. P., & Healy, P. J. (2018). Incentives in experiments: A theoretical analysis. *Journal of Political Economy*, 126(4), 1472–1503.
- Azrieli, Y., Chambers, C. P., & Healy, P. J. (2020). Incentives in experiments with objective lotteries. *Experimental Economics*, 23(1), 1–29.
- Blackwell, D. (1951). Comparison of experiments. In Proceedings of the second berkeley symposium on mathematical statistics and probability (Vol. 2, pp. 93– 103).
- Blackwell, D. (1953). Equivalent comparisons of experiments. The annals of mathematical statistics, 265–272.
- Bloedel, A. W., & Zhong, W. (2021). The cost of optimally-acquired information. Unpublished Manuscript, November.
- Caplin, A., & Dean, M. (2013). Behavioral implications of rational inattention with shannon entropy (Tech. Rep.). National Bureau of Economic Research.
- Caplin, A., & Dean, M. (2015). Revealed preference, rational inattention, and costly information acquisition. American Economic Review, 105(7), 2183–2203.
- Caplin, A., Dean, M., & Leahy, J. (2022). Rationally inattentive behavior: Characterizing and generalizing shannon entropy. *Journal of Political Economy*, 130(6), 1676–1715.
- Caplin, A., & Martin, D. (2015). A testable theory of imperfect perception. The Economic Journal, 125(582), 184–202.
- Chen, D. L., Schonger, M., & Wickens, C. (2016). otree—an open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, 9, 88–97.
- Chetty, R., Looney, A., & Kroft, K. (2009). Salience and taxation: Theory and evidence. American economic review, 99(4), 1145–77.
- Dasgupta, K., & Mondria, J. (2018). Inattentive importers. Journal of International Economics, 112, 150–165.

- Dean, M., & Neligh, N. (2023). Experimental tests of rational inattention. Journal of Political Economy, 131(12), 3415–3461.
- Denti, T. (2022). Posterior separable cost of information. American Economic Review, 112(10), 3215–3259.
- Dewan, A., & Neligh, N. (2020). Estimating information cost functions in models of rational inattention. Journal of Economic Theory, 187(105011).
- Greiner, B. (2015). Subject pool recruitment procedures: organizing experiments with orsee. Journal of the Economic Science Association, 1(1), 114–125.
- Hébert, B., & Woodford, M. (2023). Rational inattention when decisions take time. Journal of Economic Theory, 208, 105612.
- Huberman, G. (2001). Familiarity breeds investment. The Review of Financial Studies, 14(3), 659–680.
- Lipnowski, E., & Ravid, D. (2023). Predicting choice from information costs. arXiv preprint arXiv:2205.10434.
- Maćkowiak, B., Matějka, F., & Wiederholt, M. (2023). Rational inattention: A review. Journal of Economic Literature, 61(1), 226–273.
- Matějka, F., & McKay, A. (2015). Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review*, 105(1), 272–98.
- Mensch, J. (2018). Cardinal representations of information. Available at SSRN 3148954.
- Morris, S., & Strack, P. (2019). The wald problem and the relation of sequential sampling and ex-ante information costs.
- Pomatto, L., Strack, P., & Tamuz, O. (2023). The cost of information: The case of constant marginal costs. American Economic Review, 113(5), 1360–1393.
- Shannon, C. E. (1948). A mathematical theory of communication. The Bell System Technical Journal, 27(3), 379–423.

Shue, K., & Luttmer, E. F. (2009). Who misvotes? the effect of differential cognition

costs on election outcomes. American Economic Journal: Economic Policy, 1(1), 229–57.

- Sims, C. A. (2003). Implications of rational inattention. Journal of monetary Economics, 50(3), 665–690.
- Singular "they". (n.d.). Retrieved 2020-11-17, from https://apastyle.apa.org/ style-grammar-guidelines/grammar/singular-they
- Taubinsky, D., & Rees-Jones, A. (2018). Attention variation and welfare: theory and evidence from a tax salience experiment. *The Review of Economic Studies*, 85(4), 2462–2496.

Walker-Jones, D. (2024, May). Heterogeneous costs and the decision to learn.

Online Appendix

There is some repetition of Section 3 from Walker-Jones (2024) to make things more clear, but the explanation of the experiment in this online appendix is not complete, and is meant to complement the description in Section 3 of Walker-Jones (2024).

Participants were recruited for the experiment using ORSEE (Greiner, 2015) from the Toronto Experimental Economics Laboratory recruitment pool. Subjects signed up ahead of time for a particular day, either the 21st, 22nd, 23rd, 28th, 29th, or 30th of September 2020.¹² Subjects were told ahead of time they would be sent a link at 8 AM EDT, and would have until 8 PM EDT to finish the experiment.¹³ In total 270 undergraduate students from the University of Toronto attempted to partake in the experiment, and 243 of them completed the experiment.¹⁴ I refer to the 243 students that completed the experiment as the 'subjects.'

Each subject began by consenting to participate in the experiment. They then saw a welcome screen that explained that they were about to be trained and quizzed. They were told: "Please read the following instructions carefully. Understanding what is going on will help you earn more money. You will first be trained, and then you will be quizzed, to make sure you understand how everything works. There will be 10 questions in the quiz. You must answer each question correctly before you can move on to the next question. When you complete the quiz you will earn the \$5

 $^{^{12}\}mathrm{I}$ had 24 subjects on the 21st, 39 on the 22nd, 48 on the 23rd, 43 on the 28th, 46 on the 29th, and 70 on the 30th.

¹³On the 22nd the session was extended until 9 PM EDT because students experienced crashes. Then on September 28th the Economic Department's servers were unexpectedly down until almost 10 AM EDT, so students were given until 10 PM EDT to complete the experiment.

¹⁴There were 315 undergraduate students that signed up to receive a link for the experiment but 45 of them did not use the link. This proportion is in line with the "turnout" rates of other experiments run with the recruitment pool. A handful of students experienced one of two errors on September 22nd. One seemed to be caused by a server problem, while the other was caused by a missing image file. Four of the students that experienced errors did not want to re-start. Two students started the experiment with less than 20 minutes left before their session finished and could not finish as a result. This left 264 students, but some of them chose to drop out, as is explained later, so 243 ended up completing the experiment.

training payment. In addition, each quiz question you answer correctly on the first try earns you another \$0.5 training bonus payment, so you can earn up to \$10 just by doing well during the training." I wanted to incentivize students to pay attention during the training so that they understood the environment and their choices would be indicative of their preferences. All payments were in Canadian dollars.

The training consisted of three pages, which are figures 7, 8, and 9, below. In the second training page subjects were required to select 'get more information' (GMI) so that they knew what happened if they did. If they did not, then they got an error message asking them to do so. The 10 quiz questions can be seen at the end of this appendix in figures 10 through 19. If a subject got a quiz question wrong the order of the answers in the drop down menu randomly permuted so that subjects could not just try them in order. The quiz verified that subjects understood the environment, but also gave me another opportunity to train them. In the data I can see how many times each subject attempted each question. I also know how long subjects spent on each page. Among subjects the average time spent on the welcome page, training pages, and quiz pages was 17.9 minutes, and the average quiz score was 7.6/10. So, subjects in general put a lot of effort into understanding the environment, and were quite successful in doing so.

After the quiz subjects faced 100 rounds of 'investment decisions' (the last 92 of which are the decision problems I study in the body of the paper). In each round the subject selects one of two options, option X or option Y. In all 100 rounds the DM selects from their options in a drop down menu, and the order of the options is randomly generated in each round.

The subject earned probability points in each round. Subjects gradually increased their collection of probability points so as to have a higher chance of winning the draw at the end of the 100 rounds. If they won the draw they would receive a monetary prize of either \$20, \$25, \$30. The prize they would receive if they won was determined by the subject's treatment group, and was displayed to the subject throughout the training pages and rounds.¹⁵ In a round the smallest increase the subject can get in their probability of winning the prize is zero percentage points, and the largest increase is one percentage point. In each round option X is worth an amount of percentage points strictly between 0 and 1, and option Y is worth 0 or 1 percentage points.

I paid subjects in probability points so I could attempt within subject analysis in an incentive compatible way. The ideal when it comes to incentive compatibility is that in each of the 100 rounds the subject chooses what they would have chosen if they had only had to make a decision in that one round (Azrieli, Chambers, & Healy, 2018, 2020). In the setting of my experiment if I pay the subject money in each round then their marginal value for money could decrease as the rounds progress (wealth effects). Further, it would be difficult to separate this change in preferences from fatigue. The standard solution is to use a "random incentive system" (RIS), and pay the DM based on their decision in that round (Allais, 1953). In this paper's experiment setting this strategy is not incentive compatible. The DM needs to consider the cost of their learning in the setting studied in this paper, and they cannot defer their learning until they know which round is selected. If I increase the number of rounds they see, and keep the monetary payments the same, their accuracy should go down. The benefit (in expectation) from making a good selection in a round decreases while their cost of learning stays the same. Paying subjects in probability points is mathematically equivalent to RIS if subjects reduce compound lotteries and know each round is selected with equal chance.

If subjects do not accurately internalize an objective lottery over rounds, then paying in probability points might be better for eliciting preferences from a pedagogical perspective. This strategy is certainly not infallible, however. The value of a one percent increase in the probability of winning the prize might depend on the subject's current probability of winning the prize. I think this is of particular concern in two

 $^{^{15}{\}rm Because}$ the subjects participated on-line, they were transfered the money electronically within 24 hours of their session ending.

cases, when the subject's probability of winning the prize is moving away from 0%, and when it is moving to 100%. I think a probability point should have essentially the same value in two different rounds if neither of these two cases are occurring. This means that if I can move a subject chance of winning away from 0% and 100% I could reduce certainty effects, while certainty effects would be persistent with RIS.

The first 8 rounds of investment decisions familiarize subjects with the experiment interface and allow me to move a subject's chance of winning away from 0% and 100%. In the first 8 rounds subjects could not 'get more information' and they had to make a decision based on the payoff from the safe option X and the big dot. If the subject chooses the safe option in any of the first 8 rounds then they know they cannot achieve a 100 percent chance of winning the prize, and certainty effects are mitigated at least partially. If the subject chooses X in any of the first 8 rounds their probability of winning the prize is strictly above zero, and I do not need to worry about them making a decision later purely because it means they guarantee a chance of winning the prize. Further, the first 8 rounds test the understanding of the DMs and if they are reducing compound lotteries in the rounds.

In the first 8 rounds each subject saw the same sequence of big dots (which form beliefs) and payoffs for the safe option X (prices). The big dots were green, green, red, red, red, red, red, and red, and the payoffs for option X were 0.8, 0.7, 0.2, 0.5, 0.4, 0.3, 0.24, and 0.26.

In each decision problem subjects had the option to stop doing decision problems. They were told that if they decided to 'exit' they would maintain their current probability of winning the prize, and immediately find out if they had won or not, forfeiting the chance to further increase their chance of winning. If a DM fatigues, and they no longer think it is worth their time to select the better option, but they cannot stop without loosing any probability of winning they had accrued, then they might rush through the experiment, selecting options not because they are indicative of the DM's preferences, but because the DM is trying to finish as fast as they can. So, if a subject prefers to stop doing decision problems, I let them, so my data would be less noisy and more indicative of preferences.

I had 21 subjects decide not to finish the experiment, which left me with a sample of 243 subjects. Less than ten percent of students dropping out is not overly concerning since they were participating over the internet and could have had any number of distractions arise.

It was hoped that red and green would have intuitive value to the subjects. The downside with red and green is that about one in twenty people have red/green color blindness. To combat this issue, each dot displayed has a black letter in its middle, either R of G, so that subjects can still differentiate between red and green dots even if they are color blind. See Figure 2 for an example. I did not receive any complaints about the dots being hard to differentiate.

Before any subjects participated in the experiment ten different "treatments" of image files were generated ($100 \times 10 = 1000$ rounds of images, each consisting of a big dot image and an image of 100 small dots). This means that 1000 times a big dot color was drawn. The chance of the big dot being green was $\frac{1}{4}$ and the chance of the big dot being red was $\frac{3}{4}$. After the big dot color for the round was determined the composition of the small dots was drawn. In each instance there was a $\frac{3}{4}$ chance that 51 of the small dots would match the color of the big dot and a $\frac{1}{4}$ that 51 of the small dots would not match the color of the big dot. Either way, the order of the small dots was randomly generated given the drawn proportion. After this process was completed, the first eight rounds of images were used for all ten treatments, but other than that the images were not altered. This means, for instance, that if a subject was assigned to image "treatment" nine, their first eight rounds of images were the first eight rounds of images generated, and their last 92 rounds of images were the 809th through 900th rounds of images generated.



Figure 3 depicts S and the implications of Theorem 1 (Walker-Jones, 2024). The red line is the information outcomes that result in the same probability of selecting option Y given either realization of the state (these are the information outcomes Assumption 1 (Walker-Jones, 2024) requires to have a cost of zero). In Theorem 1 (Walker-Jones, 2024), part (i) requires that each s(p) be on or above the red line. The blue line, defined by $\underline{s}\mu(\underline{\omega}) + \overline{s}\mu(\overline{\omega}) = \underline{s}(p_1)\mu(\underline{\omega}) + \overline{s}(p_1)\mu(\overline{\omega})$, is the collection of points that create the same unconditional probability of selecting option Y. Part (ii) requires that $s(p_2)$ be on or below the blue line. Part (iii) requires that s(p) not be in the interior of S and on the red line unless for all other $\tilde{p} \in \mathcal{P}$ either $s(\tilde{p}) = (0, 0)$ or $s(\tilde{p}) = (1, 1)$.



Figure 4 depicts S and the implications of Theorem 2 (Walker-Jones, 2024). The blue lines are the lines through $s(p_1)$ from (0, 0) and (1, 1). When the DM is learning and price increases, $\Pr(Y|\underline{\omega}, p)$ and $\Pr(Y|\overline{\omega}, p)$ must both weakly decrease, but the proportional reduction in $\Pr(Y|\underline{\omega}, p)$ must be weakly larger than the proportional reduction in $\Pr(Y|\overline{\omega}, p)$, so $s(p, \mu)$ must remain weakly above the blue line from (0, 0) through s(p). Further, $\Pr(X|\underline{\omega}, p)$ and $\Pr(X|\overline{\omega}, p)$ must both weakly *increase*, but the proportional increase in $\Pr(X|\underline{\omega}, p)$ must be weakly smaller than the proportional increase in $\Pr(X|\underline{\omega}, p)$ must be weakly smaller than the proportional increase in $\Pr(X|\underline{\omega}, p)$, so $s(p, \mu)$ must remain weakly below the blue line from (1, 1) through s(p) (assuming I am not starting from $s(p, \mu) = (1, 1)$).

Figure 5: Implications of Theorem 3 (aggregate PS model, $p_1 < p_2$)



Figure 5 depicts S and the implications of Theorem 3 (Walker-Jones, 2024), which requires that $\Pr(Y|\underline{\omega}, p)$ and $\Pr(Y|\overline{\omega}, p)$ are weakly decreasing in p. In Figure 5, this means that $s(p_2)$ is below and to the left of the old information outcome $s(p_1)$, weakly between the two blue lines but above the red line.



Figure 6 shows the aggregated behavior of subjects in the experiment from Walker-Jones (2024) with rounds with $\mu(\overline{\omega}) = \frac{1}{4}$ on the left side and rounds with $\mu(\overline{\omega}) = \frac{3}{4}$ on the right side. The changes in aggregate observed behavior are completely in line with the predictions of the PS model (Theorem 2 (Walker-Jones, 2024)). When $\mu(\overline{\omega}) = \frac{1}{4}$: $\Pr(Y|\underline{\omega}, p = 0.25) = 0.182$, $\Pr(Y|\underline{\omega}, p = 0.5) = 0.098$, $\Pr(Y|\overline{\omega}, p = 0.25) = 0.513$, and $\Pr(Y|\overline{\omega}, p = 0.5) = 0.391$. When $\mu(\overline{\omega}) = \frac{3}{4}$: $\Pr(Y|\underline{\omega}, p = 0.25) = 0.645$, $\Pr(Y|\underline{\omega}, p = 0.5) = 0.511$, $\Pr(Y|\overline{\omega}, p = 0.25) = 0.943$, and $\Pr(Y|\overline{\omega}, p = 0.5) = 0.893$.

Figure 7:

Training (1 of 3):

The experiment consists of 100 rounds of investment decisions, which does not include the training round (Round 0). Each of the 100 rounds provides you with the opportunity to increase your probability of winning a prize of \$25 at the end of the experiment. This prize is in addition to whatever you earn during the training process.

You will not be able to increase your probability of winning the prize by more than 1 percentage point in any one round. Ideally, you would thus like to increase your probability of winning the prize by 1 percentage point in each round, so that you win the prize with a 100% probability when the experiment is over, but it is extremely unlikely that this will be possible.

In each investment decision you get to choose between two investment options, option X and option Y. Option X always increases your probability of winning the prize by a displayed positive amount, which is always less than 1 percentage point.

Option Y, in contrast, either does not increase your probability of winning the prize at all, or it increases your probability of winning the prize by 1 percentage point.

To help provide you with some information about the chance that option Y increases your probability of winning the prize, in each round there is one large dot that is either red with a black R in the middle of it (example below on right), or green with a black G in the middle of it (example below on left).



When the large dot is red, this is a bad sign for the value of option Y, and there is only a 1 out of 4 chance (25% chance) that option Y increases your probability of winning the prize by 1 percentage point, and a 3 out of 4 chance (75% chance) that option Y leaves your probability of winning the prize unchanged. This means that when the large dot is red, on average option Y increases your probability of winning the prize by 0.25 percentage points.

When the large dot is green, this is a good sign for the value of option Y, and there is a 3 out of 4 chance (75% chance) that option Y increases your probability of winning the prize by 1 percentage point, and only a 1 out of 4 chance (25% chance) that option Y leaves your probability of winning the prize unchanged. This means that when the large dot is green, on average option Y increases your probability of winning the prize by 0.75 percentage points.

In each round there is a 3 out of 4 chance (75% chance) the big dot is red, and a 1 out of 4 chance (25% chance) the big dot is green. In the first 8 rounds (rounds 1 through 8) the large dot is the only source of information available to you.



Figure 8:

Training (2 of 3):

In the last 92 rounds (rounds 9 through 100), you can get more information if you choose to. When you choose to 'Get more information,' you are shown 100 small dots, each of which is red with a black R in the middle, or green with a black G in the middle. Together, the small dots tell you for sure whether or not option Y increases your probability of winning the \$25 prize by 1 percentage point in the round you are in.

In each round that you can 'Get more information,' there are either 49 small red dots and 51 small green dots, or 51 small red dots and 49 small green dots.

If there are 49 small red dots and 51 small green dots, then option Y increases your probability of winning the prize by 1 percentage point.

If there are 51 small red dots and 49 small green dots, then option Y does not increase your probability of winning the prize.

Under the dotted line below is an example of a round where you can choose to get more information. Notice that the round number is displayed, and the amount that option X increases your probability of winning the prize is displayed.

The large dot is green, so you know there is a 3 out of 4 chance (75% chance) option Y increases your probability of winning the prize by 1 percentage point, but you cannot know for sure unless you choose to 'Get more information.'

Please select 'Get more information' so you see what happens when you do.

Round number: 0 of 100

Option X increases your probability of winning the \$25 prize by 0.5 percentage points

Please choose an option:



Remember:

If the large dot is red, then there is a 3 out of 4 chance (75% chance) option Y increase your probability of winning the prize by 0 percentage points, and a 1 out of 4 chance (25% chance) it increases your probability of winning the prize by 1 percentage point. This means that when the large dot is red, on average option Y increases your probability of winning the prize by 0.25 percentage points.

If the large dot is green, then there is a 3 out of 4 chance (75% chance) option Y increases your probability of winning the prize by 1 percentage point, and a 1 out of 4 chance (25% chance) it increases your probability of winning the prize by 0 percentage points. This means that when the large dot is green, on average option Y increases your probability of winning the prize by 0.75 percentage points.

If you choose to 'Get more information,' 100 small dots will appear that you can use to determine how much option Y increases your probability of winning the prize.



Figure 9:

Training (3 of 3):

Under the dotted line below is an example of what you would see after requesting to 'Get more information' in a round. Selecting this option has made the 100 small dots for this round appear.

Try counting the number of small red dots and or the number of small green dots below. You should find there are 49 red dots and 51 green dots, which means option Y would increase your probability of winning the prize by 1 percentage point. So in this round, option Y increases your probability of winning the prize by more than option X. This training round (Round number 0) does not count towards your eventual probability of winning the prize.

After you have finished the quiz, if you want to quit the experiment you will be given the option to do so at the bottom of the screen in each round. When you choose to exit, you still have the chance to win the \$25 prize, you do not lose any probability you have acquired of winning the prize, but you do lose the chance to further increase your probability of winning the prize.

Round number: 0 of 100

Option X increases your probability of winning the \$25 prize by 0.5 percentage points

Please choose an option:



Remember:

In each round, option Y always increases your probability of winning the prize by either 0 or 1 percentage points.

In each round, there are always either 49 small red dots and 51 small green dots, or 51 small red dots and 49 small green dots. If there are 49 small red dots and 51 small green dots, then option Y increases your probability of winning the prize by 1 percentage point.

If there are 51 small red dots and 49 small green dots, then option Y increases your probability of winning the prize by 0 percentage points.

Next

Figure 10:

Quiz time! (Question 1 of 10)

It is now time to do the quiz. Remember, to move on to the next question you must first answer the current question correctly. Completing the quiz earns you the \$5 training payment, and in addition to the training payment, you can earn another \$0.5 for each question you get correct on your first attempt to answer it.

Which of the following is correct?



Figure 11: Quiz time! (Question 2 of 10)

If you see the big dot below in a round, which of the following is NOT correct?



Figure 12: Quiz time! (Question 3 of 10)

In the last 92 rounds, if you see the the big dot below, if you choose to 'Get more information,' how many small green dots with black Gs will be displayed?



Figure 13: Quiz time! (Question 4 of 10)

In the image below there are 51 small green dots. If you see such an image in a round, what happens when you select option Y?



Figure 14: Quiz time! (Question 5 of 10)

If the big dot below is visible in a round, which of the following is correct?



Figure 15: Quiz time! (Question 6 of 10)

In the image below there are 51 small red dots. If you see such an image in a round, what happens when you select option Y?



Figure 16: Quiz time! (Question 7 of 10)

Count the number of small green dots in the image below. If you see such an image in a round, what happens when you select option Y?



Figure 17: Quiz time! (Question 8 of 10)

When you choose to 'Get more information' in a round, which of the following is correct?



Figure 18: Quiz time! (Question 9 of 10)

If the large dot below is the only information available, then which of the following is correct?



Figure 19: Quiz time! (Question 10 of 10)

If the large dot below is the only information available, then which of the following is correct?



Figure 20:

Remember:

If the large dot is red, then there is a 3 out of 4 chance (75% chance) option Y increases your probability of winning the prize by 0 percentage points, and a 1 out of 4 chance (25% chance) it increases your probability of winning the prize by 1 percentage point. This means that when the large dot is red, on average option Y increases your probability of winning the prize by 0.25 percentage points.

If the large dot is green, then there is a 3 out of 4 chance (75% chance) option Y increases your probability of winning the prize by 1 percentage point, and a 1 out of 4 chance (25% chance) it increases your probability of winning the prize by 0 percentage points. This means that when the large dot is green, on average option Y increases your probability of winning the prize by 0.75 percentage points.

If you choose to 'Get more information' 100 small dots will appear that you can use to determine how much option Y increases your probability of winning the prize.

Next

Figure 21: For your information: for the next 40 rounds option X will increase your probability of winning the prize by 0.25 percentage points

For your information: of the 92 rounds remaining, 12 of the rounds have an option X that increases your probability of winning the prize by 0.5 percentage points, and 80 of the rounds have an option X that increases your probability of winning the prize by 0.25 percentage points.

Remember: in all remaining rounds you can choose to 'Get more information' if you want to.

Next

Figure 22:

For your information: for the next 12 rounds option X will increase your probability of winning the prize by 0.5 percentage points

For your information: of the 52 rounds remaining, 12 of the rounds have an option X that increases your probability of winning the prize by 0.5 percentage points, and 40 of the rounds have an option X that increases your probability of winning the prize by 0.25 percentage points.

Remember: in all remaining rounds you can choose to 'Get more information' if you want to.

