Consumer Inferences from Product Rankings: The Role of Beliefs in Search Behavior*

Jessica Fong[1], Olivia R. Natan[2], and Ranmit Pantle[3]

[1]University of Michigan

[2]University of California, Berkeley

[3]Northwestern University

February 9, 2025

Abstract

In online markets, consumers tend to search and purchase prominently positioned products. We develop an experimental paradigm to distinguish between two mechanisms driving this behavior: position-specific search costs and beliefs about expected returns to search. Using incentivized lab experiments, we find that both mechanisms exist, and short-term randomization of rankings alone does not separate the two mechanisms. Failing to account for beliefs leads to biased estimates of search costs and incorrect consumer welfare predictions under alternative recommendation systems, such as platform self-preferencing. We discuss how to extend our findings to real world search settings and solutions for estimating unbiased search costs in such environments.

JEL: D83, C93, L81, D84

*We are grateful for comments from Anocha Aribarg, Chu Ivy Dang, Andrey Fradkin, Rafael Greminger, Sam Hirshman, Aaron Kaye, Bowen Luo, Brett Gordon, Przemek Jeziorski, Ilya Morozov, Sarah Moshary, Yeşim Orhun, Gautam Rao, Heiner Schumacher, Katja Seim, Raluca Ursu, Caio Waisman, Anna Tuchman, and audiences and seminar participants at Dartmouth Tuck, University of Michigan Dearborn, CU Boulder, Wharton, the CODE 2024 conference, CESifo Economics of Digitization, and Marketing Science 2024 conference. Contact information: Jessica Fong (jyfong@umich.edu), Olivia R. Natan (olivia.natan@berkeley.edu), Ranmit Pantle (ranmit.pantle@kellogg.northwestern.edu).

1 Introduction

Due to the vast array of products available online, recommendation systems have become increasingly influential in digital markets, shaping both consumer choice and competition among firms (Armstrong et al., 2009). These systems, such as product ranking algorithms, play an important role in determining which products consumers see and in what order. Although recommendation systems can help consumers discover products more easily, they can also be used to manipulate consumer choice in ways that can be harmful to the consumer. For example, consumers are more likely to search and consequently purchase products that appear towards the top of the search results (De los Santos and Koulayev, 2017; Ursu, 2018; Compiani et al., 2023). In response to this behavior, platforms may place more profitable products in these positions (Lam, 2023; Peitz, 2023). The potential for consumer harm due to platform manipulation raises important questions about how recommendation systems impact consumer behavior.

To evaluate how recommendation systems impact consumer behavior and welfare, we need to understand *why* consumers are more likely to search recommended products. In particular, does searching recommended products take less effort (i.e., lower search costs), or do consumers believe recommended products are "better" (i.e., have higher expectations about the returns to search)? Distinguishing between these two mechanisms is challenging because they can produce the same patterns in search behavior, and consumers' expectations are typically not collected. However, making this distinction is essential for accurately predicting welfare under alternative recommendation systems (e.g., ranking algorithms) because expectations about the returns to search adjust as consumers learn the new algorithm, but search costs (e.g., effort) remain the same regardless of the algorithm.

In this paper, we show that (1) both mechanisms can exist and (2) failing to account for consumer beliefs about their returns to search for recommended products leads to biased estimates of search costs, which in turn leads to incorrect welfare predictions under alternative recommendation systems. We demonstrate this in the context of position effects—the idea that the placement of products impacts user behavior. We focus on this context for several reasons. First, many papers document that people are more likely to search more prominently ranked products, such as those at the top of a list (e.g., Ursu, 2018; De los San-

tos and Koulayev, 2017; Compiani et al., 2023). This behavior is most commonly attributed to position-specific search costs, which represent the effort to scroll farther down the page.[1] However, consumers may believe that more prominently ranked products are higher surplus because platforms typically prioritize better products in their rankings (e.g., Kaye, 2024). For example, Expedia's default rankings are ordered by "Recommended." Such beliefs would also induce less frequent search of less prominently-ranked products. Consequently, not accounting for this mechanism results in overestimating the search costs for these products. Second, quantifying position-specific search costs has important policy implications about platform power. In particular, policy-makers are concerned about the ability of digital platforms to manipulate consumer choice through their ranking algorithms. If position effects are cost-driven, prominently ranking lower-quality products reduces consumer surplus by increasing search costs or resulting in choosing worse products. If driven by beliefs, consumers learn to avoid positions with lower expected returns to search. Therefore, the extent of the platforms' influence on choice and consumer surplus depends on the true magnitude of position-specific search costs.

Position-specific costs and position-specific beliefs about the payoffs from search typically cannot be separately identified, as differences in beliefs and differences in costs both shift search behavior in the same way. This is true even when product rankings are randomized. Position-specific search costs are often measured using short-term experiments in which treatments are undisclosed. Because participants do not know that the products are randomly ordered, they may still hold beliefs formed before the experiment—that the top-ranked products have higher expected returns to search.[2] As a result, any position effects exhibited during the experiment may still be driven by beliefs in addition to search costs, despite randomization. To overcome this challenge, we design an incentive-compatible experiment that shuts off the beliefs mechanism, which allows us to measure search costs. In the experiment, participants search for a product in a stylized setting similar to that of Weitzman (1979). The product has two vertical attributes, one observed costlessly and the other observed only after a costly search. Participants perform this search task repeatedly,

[1] Exceptions include Nocke and Rey (2023) and Athey and Ellison (2011) attribute position effects to beliefs—position correlates with the product's utility that is observed prior to search.

[2] For the remainder of this paper, we refer to "beliefs" as beliefs about the expected returns to search.

so we can detect any learning over time about the ranking algorithms.

The order in which products are displayed and whether participants are informed about the ranking algorithm depend on the participant's experimental condition. In one condition, which we refer to as Random-Informed, products are ordered randomly, and consumers are explicitly informed about this random ordering. This condition allows us to identify position-specific search costs. Because we inform participants that products are ordered randomly, they should not hold position-specific beliefs. Therefore, assuming that participants understand that products are ordered randomly, any position effects in this condition are attributed to position-specific search costs. In another condition, Random, the products are sorted randomly but participants are not informed. In this condition, participants may face position-specific costs, as in Random-Informed, but also may have position-specific prior beliefs (i.e., participants may come into the experiment with pre-existing beliefs about returns to search). In the third condition, Strong, products are generally sorted in decreasing order of the hidden attribute, such that more prominently positioned products have greater returns to search, on average. As in Random, participants are not informed about the ranking algorithm. In this condition, any position effects may be due to position-specific costs, position-specific prior beliefs, and learned position-specific beliefs. Identifying and extrapolating the search costs estimated from Random-Informed allows us to identify these beliefs in Random and Strong. In other words, any residual position effects in Random and Strong, after accounting for the position-specific search costs estimated from Random Informed, can be attributed to beliefs.

Using data generated by the experiment, we estimate a sequential search model to quantify search costs. We assume that participants search in descending order of a reservation utility index, stopping when an already-searched option is better than the best unsearched option (Weitzman, 1979). We use a Bayesian framework, similar to Morozov (2023), to account for heterogeneity in both baseline search costs and in position-specific search costs across people. The model estimates highlight beliefs as a major component of position effects. The estimates also show evidence of learning about the informativeness of product rank in Strong; consumers exhibit increasing position effects and converge towards the payoff-maximizing search order across search tasks.

To better understand how the beliefs can confound search costs in the field, we conduct another study that is designed to replicate the design of typical search ranking experiments. In this study, we introduce changes to the ranking algorithm mid-experiment and do not disclose these changes to participants. We consider two types of changes. The first change involves the ranking algorithm starting as *Strong* (i.e., generally decreasing in the hidden attribute) in the initial tasks and then switching to *Random*. This setup replicates typical search experiments where the ranking algorithm changes to *Random* without the consumers' knowledge. In such experiments, search costs are estimated from the *Random* tasks, but participants may still hold position-specific beliefs learned from the *Strong* tasks. The second change simulates a specific form of platform or retailer steering: the algorithm switches from *Strong* to one that moves products with the lowest value to the consumer in the top two ranks. This arrangement, which we refer to as the *Pref* algorithm, reflects the concern that ranking algorithms might prioritize products that are better for the platform or firm but worse for the consumer, i.e., self-preferencing.

We find that participants adjust their search over several periods when facing a new, undisclosed ranking algorithm. The larger position effects observed under the *Strong* algorithm (compared to Random-Informed) persist even after the algorithm switches from *Strong* to *Random*. This persistence suggests that their beliefs formed during the *Strong* algorithm periods carry over to the *Random* algorithm periods. Consequently, the conventional approach of introducing randomization without disclosure can misattribute position effects generated by beliefs to costs.

Separating cost-driven position effects and belief-driven position effects is a critical component in evaluating the impact of product position on consumer surplus. To generate welfare predictions, we estimate a modified version of our search model using data from Study 2. Using the descriptive patterns from search over time, we adapt our model of search to allow for Bayesian learning about position-specific payoffs in addition to position-specific costs. We use the model estimates to demonstrate how failing to account for the potential evolution of beliefs leads to wrong welfare conclusions. In particular, using search costs estimated from short-term experiments with incorrect consumer beliefs predicts a 3.3% loss in consumer surplus when the platform switches to the *Pref* algorithm, while search costs

estimated from a model that accounts for learning predicts a 0.84% loss in consumer surplus. That is, the model that excludes beliefs overestimates the loss in surplus by a factor of four.

In order to separate beliefs and search costs in practice, experiments should be run for long enough for consumers to fully update their beliefs to the new algorithm, and only the search behavior from after beliefs have fully updated should be used to estimate search costs. However, this may require the experiment to run for a long time, which may not be feasible. We also consider other methods to address this problem, such as assuming correct preexperiment beliefs about returns to search in model estimation (e.g., Kaye, 2024). This method assumes that consumers hold correct beliefs about the informativeness of product ranking *prior to* the experiment that orders products randomly. We find that if consumers are slow to learn, as in our setting, assuming correct beliefs can lead to incorrect welfare predictions in the short run and even to the opposite direction of search costs. Therefore, assuming correct pre-experiment beliefs is best suited to settings where consumers are very experienced with the search environment and the product market and for short-term experiments where learning about the experimental ranking algorithm may be limited. Another solution to account for beliefs in field data is to directly survey consumers on their beliefs and account for these in the model.

In addition to documenting the role of consumer beliefs in improving the measurement of economic primitives in search settings, our paper can shape future work evaluating alternative recommendation algorithms. Given the prevalence of ordered information, platforms and other intermediaries that can rank items hold power over firms supplying products and services. In evaluating the welfare effects of information design, we are often interested in the degree of steering by platforms (i.e., manipulation of such rankings to benefit the intermediary) and the related costs to consumers and firms. However, the core approach of this research is to estimate the position or position effects in the current setting/equilibrium (versus a benchmark of randomly ordered products), and to assume such position effects would hold counterfactually if alternative rankings were supplied. In reality, consumers are not immediately informed about changes in ranking algorithms. This paper extends our understanding of how beliefs impact welfare and its measurement, including the transition

to a new equilibrium in response to a change in platform policy.

Our paper also contributes more broadly to the importance of accounting for beliefs when experimentally estimating model primitives, of which estimating search costs is just one example. Measuring advertising effectiveness is another setting where this problem can arise. When using short-term field experiments to evaluate the impact of promotional advertising, researchers may find that consumers' responses are influenced by their preexisting beliefs about the brand or product. Thus, consumers' behaviors during the experiment depend on their beliefs formed pre-experiment. If consumers believe that an advertisement is part of a temporary promotion rather than a permanent change, their behavior during the experiment may not accurately predict their behavior for when the firm permanently increases advertising levels. Pricing is another example. If consumers think a discount is a one-time event rather than a new pricing strategy, the estimated price elasticity using the experimental data may not reflect their price elasticity in a new equilibrium. The solutions presented in this paper—informing consumers of the policy change so that beliefs are updated accordingly, soliciting pre-experimental beliefs (Aridor et al., 2024), or assuming rational expectations for experienced consumers—apply to these settings as well.

Relevant Literature and Contribution This paper contributes to four related areas of literature.

First, our paper relates to existing work on consumer search (De los Santos et al., 2012; Moraga-González et al., 2023; Dinerstein et al., 2018). We focus on position effects in product search and advertising. Position effects are well documented across many domains: users click more on alternatives near the top of the list in a paid keyword search (e.g., Jeziorski and Segal (2015), Ghose and Yang (2009), Agarwal et al. (2011), Jeziorski and Moorthy (2018), and Yao and Mela (2011)) as well as in product search (e.g., De los Santos and Koulayev (2017), Chen and Yao (2017), Ursu (2018), Compiani et al. (2023), and Donnelly et al. (2024)). In fact, these effects are so commonplace that work in these areas accounts for position effects wherever possible (Ursu et al., 2023), sometimes even assuming the existence of the effect ex ante (Choi and Mela, 2019; Lam, 2023; Derakhshan et al., 2022). Our work builds on these findings by taking a closer look at whether these position

effects result from equilibrium inferences about firm behavior by consumers or whether they reflect fundamentals of information processing.

Second, we relate to the empirical literature on consumer search behavior with learning (Rothschild, 1978; Koulayev, 2013; De Los Santos et al., 2017). Much of this literature focuses on learning about prices (e.g., Wu et al., 2022). Jindal and Aribarg (2021) demonstrate that consumers have heterogeneous prior beliefs about prices, which are important to take into account when estimating individual search costs. Other papers model learning about other attributes, such as match values (Ursu et al., 2020), and the distribution of product quality (Hu et al., 2019). Hodgson and Lewis (2023) consider whether consumers extrapolate beliefs from previously searched options to unsearched options that have similar attributes, and Gardete and Hunter (2024) focus on learning about within-product attributes. We build on this literature by incorporating learning about the structure of the search environment (i.e., the informativeness of product rankings).

Third, this paper relates to a growing body of work in economics and quantitative marketing that studies consumer search in controlled laboratory settings. Papers in this area investigate a wide variety of search behavior, including the role of revisiting products during a search (Dang et al., 2022), social influence and product arrangement (Jameei Osgouei et al., 2023), how advertising impacts the search funnel (Morozov and Tuchman, 2024), and heuristic approximations of optimal information acquisition (Gabaix et al., 2006). These works also study undirected search behaviors, including the role of contextual information (Karle et al., 2023) and learning about the distribution of product values (Casner, 2021). We contribute to this literature by creating an experiment design that can distinguish between position-specific beliefs and position-specific costs.

Finally, we relate to work incorporating beliefs into structural models (see DellaVigna (2018) for a review).[3] These papers generally consider new markets or first-time consumers, scenarios in which agent behavior can be observed "from the beginning." For example, Doraszelski et al. (2018) and Huang et al. (2022) study firm pricing following deregulation in the UK energy and Washington state liquor markets, respectively. Firm behavior eventually converges to an equilibrium, although the speed of adjustment can be heterogeneous across

[3]The standard approach to beliefs is to impose rational expectations about market conditions (e.g., Aguirregabiria, 2021, Bresnahan and Reiss, 1991, Aguirregabiria and Mira, 2007, among many others).

firms (e.g., Goldfarb and Xiao, 2011). For consumers, models of learning about brand value (see Ching et al. (2013) for a summary) demonstrate the important role of dynamic beliefs in explaining several features of consumer behavior, including inertia, and imply incomplete information long after consumers enter a market. In education markets, households may hold inaccurate beliefs about school quality, which leads to too little search (Agte et al., 2024). Together, these studies suggest that the rational expectations assumption may not hold in settings where a stable equilibrium has yet to be reached. We add to these works by explicitly studying the (confounding) role of beliefs when estimating structural primitives using short-term experiments.

The paper proceeds as follows. In Section 2, we first describe a motivating model of directed consumer search and highlight the difficulty of separately identifying position-specific search costs and position-specific beliefs. Next, in Section 3, we describe the design and results of Study 1, outlining our structural model and discussing estimation results. In Section 4, we describe the design and results of Study 2, outline the adaptations we make to the structural model, and describe our estimation results. We discuss counterfactuals in Section 5. Section 6 highlights how our experimental setting differs from that of real world search, and how to generalize our findings to such environments. Section 7 concludes.

2 The Role of Beliefs in Search Models

To illustrate how beliefs about recommendation quality can be confounded with model primitives in our setting, we first describe the consumer search model in detail. We consider a directed sequential search framework similar to that of Weitzman (1979) and Kim et al. (2010). Consumers face a finite, ordered set of product alternatives $j \in 1, \dots J$, such as one page of search results on a retail website or platform.[4] The utility that a consumer receives from choosing product j is

$$u_j = a_j + b_j, \tag{1}$$

[4]We abstract from pagination for simplicity, which may be interpreted as consumers exclusively studying the first page of results. See Greminger (2022) for an approach for handling scrolling and pagination in search.

where a_j and b_j are vertical attributes which are identically valued by all consumers.[5]. We assume that the distributions of a_j and b_j are independent and consumers are aware of this. The attribute a_j is observed prior to search (e.g., a_j is visible without clicking on the product's page) and thus is costless to observe. The attribute b_j is observed only if the consumer clicks on the product, incurring a cost c_j.[6] The distribution of b_j is known to consumers: $b_j \sim N(\mu_j, \sigma_b^2)$.

The extension to the case where a_j and b_j are not independent is straightforward. For instance, a_j and b_j having a correlation coefficient of ρ is equivalent to a search context with transformed variables $\tilde{a}_j$ and $\tilde{b}_j$, where $\tilde{a}_j = a_j + E[b_j|a_j]$ and $\tilde{b}_j = b_j - E[b_j|a_j]$. By definition, $\tilde{a}_j$ and $\tilde{b}_j$ are uncorrelated. The variance of $\tilde{b}_j$ is ${\sigma_b}^2(1-\rho^2)$ which is less than ${\sigma_b}^2$, since correlation reduces the uncertainty resolved by search. A setting with correlated attributes is therefore equivalent to a setting with uncorrelated attributes but with a lower variance of the hidden attribute. We revisit this extension in Section 6.

The consumer's reservation utility is defined as the value at which the consumer is indifferent between the expected benefit of searching another product and choosing among the already-searched products (i.e., the maximum of already-searched utilities). These reservation utilities, z_j, are the solutions to

$$c_j = \mathcal{B}(z_j) = \int_{z_j}^{\infty} (u_j - z_j) f(u_j) du_j. \tag{2}$$

Because the post-search utility term b_j has a normal distribution, we can use the solution derived in Kim et al. (2010). The reservation utility is

$$z_j = a_j + \mu_j + \zeta(\frac{c_j}{\sigma_b})\sigma_b, \tag{3}$$

where ζ is a scalar function.[7]

[5]We abstract from heterogeneity in valuation and horizontal differentiation to parallel our experimental paradigm. We discuss how to extend our model and experimental results to horizontal preferences over differentiated products in Section 6.

[6]Consistent with literature, we define a product as having been "searched" if the consumer clicks on it to view the hidden attribute (b_j). Products which are not searched are referred to as "unsearched products."

[7]In particular, "the function $\zeta(x)$ solves the following implicit equation $x = (1-\Phi(\zeta))(\lambda(\zeta)-\zeta)$ where Φ is the cumulative standard normal distribution and λ is the standard normal hazard rate, $\phi(\zeta)/(1-\Phi(\zeta))$, in which ϕ is the standard normal probability distribution function." (Kim et al., 2010, p. 1011).

Under these assumptions, optimal search and choice behavior follows three rules:

1. *Selection rule:* Products are searched (among remaining unsearched products) in descending order of reservation utility. The next product searched (if any) is always that with the highest reservation utility among the remaining unsearched options.

2. *Stopping rule:* Consumers stop searching when the highest utility found so far is greater than the maximum of the reservation utilities for all remaining unsearched products.

3. *Choice rule:* After ending their search, consumers choose the option with the maximum utility among the searched set.

Thus far, we have only reiterated the basic model outlined in Weitzman (1979) and specified using a normally-distributed error term in Kim et al. (2010), but we have not yet discussed how position effects can arise. We now describe the potential mechanisms through which consumers tend to search products towards the top of the page. For the rest of this paper, we refer to "more prominently ranked" or "top-ranked" products to mean products positioned toward the top of the page.

One simple explanation for position effects is that more prominently ranked products may have higher values of a_j. However, this does not fully explain why consumers are more likely to search more prominently ranked products. Similar to previous studies (e.g., Ursu, 2018), we later show that position effects exist even when a_j is not correlated with rank. Therefore, we focus on two alternative mechanisms.

First, the search costs c_j, which are model primitives, can increase with rank $r(j)$, such that $\frac{\partial c_j}{\partial r(j)} > 0$. Higher search costs for less prominently ranked products can be attributed to higher effort costs (e.g., scrolling) or attention. More specifically, these increased costs may reflect the difficulty of navigating through ordered products, or they may reflect the higher attention costs to search products lower down the list.

Second, consumers might expect more prominently ranked products to yield greater post-search utility b_j. This belief can arise due to the design of many e-commerce platforms, in which products are sorted based on popularity or other product attributes that are observed only after search. For example, Kaye (2024) documents that online travel

agency Expedia's default algorithm positions listings with better hidden attributes (location desirability) towards the top of the page. As a result, consumers' beliefs about the returns to searching top-ranked items might be higher.

This mechanism implies $\frac{\partial \mu_j}{\partial r(j)} < 0$.

In summary, rank enters the reservation utilities in two places in the model: in costs ($c_{r(j)}$) and in beliefs about the expected value of b_j, the hidden attribute ($\mu_{r(j)}, \sigma_{b,r(j)}$). A third alternative mechanism—that the variance of the returns to search varies by rank—is outside the scope of this paper.[8] Taking these mechanisms into account, the reservation utility for product j is then

$$z_j = a_j + \mu_{r(j)} + \zeta(\frac{c_{r(j)}}{\sigma_{b,r(j)}})\sigma_{b,r(j)}. \tag{4}$$

Both mechanisms have similar effects on search behavior. Reservation utilities will be lower for less prominently ranked products; products closer to the bottom of the page (i.e., larger $r(j)$) have higher search costs, which in turn decreases $\zeta(\cdot)$, and they have a lower expected b_j, $\mu_{r(j)}$. This implies that consumers are less likely to search less prominently ranked products first (thus impacting search order) and more likely to stop searching earlier, as less prominently ranked products will compare less favorably to already-searched products (thus impacting search depth). Consequently, one cannot separately identify costs from beliefs solely by considering the presence of position effects in observational data.

Although the two mechanisms can produce similar search behavior, they yield different predictions of behavior under counterfactual ranking algorithms. If position effects are due to position-specific costs, the search cost for a product in rank r is fixed regardless of the ranking algorithm because search costs are considered to be policy invariant. However, under the "beliefs" mechanism, a change in the ranking algorithm eventually affects consumers' beliefs about the correlation between $r(j)$ and b_j. As a result, the two mechanisms have different predictions of position effects under alternative ranking algorithms.

To illustrate, suppose that we observe position effects when products are sorted from best to worst. If we attribute the position effects solely to search costs, a reversal in ranking order to worst-to-best would still lead consumers to search higher-ranked (but

[8] We consider robustness to different beliefs about post-search variance in our estimates in Section 3.4.

worse) products first. Conversely, if we attribute the effects to beliefs, and if consumers learn from experience, we would expect consumers to become more likely to search the bottom-ranked products first over time. These two predictions have different implications for consumer welfare. If position effects are primarily cost-driven, rankings that prominently feature lower-quality products will substantially decrease welfare, as consumers either incur higher total search costs or end up with worse products. If position effects are due to beliefs and beliefs adjust, consumers will learn to avoid positions with lower quality products. In other words, consumers are more susceptible to platform steering if position effects are primarily driven by search costs. Understanding to what extent each mechanism influences position effects is important both for welfare evaluation and for managerial purposes, where accurate prediction is the goal.

One solution is to assume rational expectations about the post-search utility (Kaye, 2024). If the researcher is willing to impose rational expectations, such that consumers know the expected $\mu_{r(j)}$, then estimation with observational data alone can identify the effect of position on costs. However, this assumption may be too strong in settings where ranking algorithms change frequently relative to search frequency, or where algorithms or post-search attributes are hard to evaluate. For example, if quality is difficult to assess, consumers may be unable to learn the relationship between rank and post-search product utility.

For our purposes, we choose to abstract away from horizontal differentiation to make it easier to identify the two mechanisms of interests - beliefs and costs. If we included horizontal differentiation in the model, we would also need to estimate heterogeneous preferences over attributes across consumers. This would add noise to our ability to separate beliefs from costs and raise the amount of data required to detect differences.

In the next sections, we describe a series of controlled experiments designed to achieve several objectives. First, these experiments enable us to separately measure the roles of beliefs and costs in consumers' responses to search rankings, allowing us to quantify each mechanism's contribution to position effects. Second, they document the speed with which consumers form beliefs about hidden attributes, both when the ranking algorithm is constant and when it changes. Third, the data from the experiments are used to estimate a structural

model of search behavior. With these estimates, we can then simulate several alternative ranking algorithms and evaluate their impact on consumer welfare. In particular, we can compare welfare effects when the roles of costs and beliefs are accurately attributed versus when position effects are mistakenly attributed solely to costs.

3 Study 1: Separating Costs and Beliefs

The objective of the first study is to separately identify the effects of beliefs and costs in producing position effects in product search behavior. We design a simple search environment that (1) abstracts from unobserved preference heterogeneity and (2) experimentally removes the role of beliefs about rankings in one condition.[9]

3.1 Design

Participants are asked to select one product from a vertical list of ten products. Each product j has two vertical attributes, "bonus A" (a_j) and "bonus B" (b_j). Bonus A is visible to participants before searching, and bonus B is observed only after the participant "searches" the product by clicking on it. These bonuses are independent, with $a_j \sim TN(20, \sigma = 4)$ and $b_j \sim TN(40, \sigma = 12)$, which are truncated normal distributions bounded below by zero. Assuming that these bonuses are independent allows for more straightforward identification of position-specific search costs and simplifies the optimal search problem for participants.[10] Participants are informed about the mean, standard deviation and 95% confidence interval of b_j and that the bonuses are drawn independently (i.e., bonus A is not informative of bonus B).[11] For each product they search, they incur a search cost of one point, and they earn $a+b$ points from the product they choose minus their search costs. Revisits are costless. There is no outside option, so all participants must search at least one product in order to

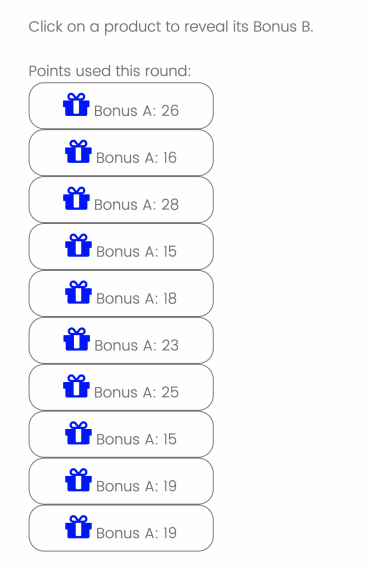
[9] Study 1 is pre-registered. The pre-registration is available at https://aspredicted.org/SWN_MK9.

[10] While this assumption simplifies identification, it deviates from real-world search settings where visible, and hidden and observed attributes are correlated (e.g., price and match value). We discuss how to generalize our results to settings with correlated attributes in Section 6.

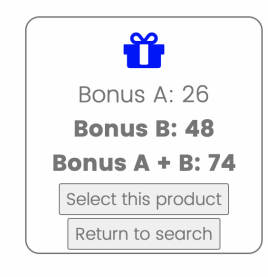
[11] We use the following language in the instructions to participants: "Bonus B has an average value of 40 points, and a standard deviation of 12. Put otherwise, 95% of products have their Bonus B between 16 and 64" and "The value of Bonus A is not informative of the value of Bonus B. In other words, Bonus A and B are not correlated. Products with high Bonus A can have low Bonus B, and vice versa." Full instructions are included in Appendix J.1.

complete the task. Figure 1 illustrates what the task looks like for the participants. Not all products may be visible without scrolling.[12] Therefore, position-specific search costs may be in part due to effort costs, such as scrolling or the effort to click on a product farther down on the page. They may also be attributed to the general tendency to read top down. For example, using eye tracking, Guan and Cutrell (2007) found that individuals are more likely to look at the top-positioned items first and rarely look at items towards the bottom of the page. So, position-specific search costs also capture position-specific differences in attention.

Figure 1: Illustration of Search Task



(a) Search task



(b) Bonus B reveal

Notes: Panel (a) shows an example of the list of products that a participant sees. Panel (b) shows an example of what appears once the participant clicks on a product.

Participants first complete a practice round, which does not affect their bonus payment, to become familiar with the interface. To ensure that they understand the instructions and task structure, the participants then answer three comprehension checks. The first two comprehension checks test for understanding of how bonus payments are determined.[13] If

[12] Whether all products are visible depends on the participant's screen size.

[13] The first comprehension check asks, "Suppose you clicked on product X once, and product Y twice, and you selected product Y, which has a bonus A value of 20 and a bonus B value of 40. What is your final payoff (bonus points after click costs) for this round?" The second asks a similar question but with different

the participant answers incorrectly, the instructions are shown again, and the participant is given another chance. If they answer both comprehension checks incorrectly twice, they are disqualified from the study. The third comprehension check tests for whether participants understand that bonus A and bonus B are uncorrelated. We revisit the comprehension checks in model estimation (Section 3.3).

Participants then complete the incentivized search task ten times. The bonuses A and B for each task are redrawn and are independent across tasks. After completing the ten tasks, participants receive a bonus payment, in addition to the base pay, of the total points they earn on all tasks. For every 400 points, they receive \$1 of bonus pay.[14] Participants also may earn an additional 40 bonus points (\$0.10) at the end of the survey for accurately answering a question that elicits their beliefs about bonus B.[15] At the end of the survey, we collect participants' self-reported age, gender, and online shopping frequency (e.g., never, less than once a week, etc). We ask for online shopping frequency because we hypothesize that those who shop online more may have more exposure to ordered product lists, and thus may have stronger prior beliefs about rank.

Participants were recruited on Prolific.com among US participants and were required to use a desktop device. Participants were randomized into one of three conditions, which differ only in their ranking algorithm and disclosure of the ranking algorithm: Strong, Random, and Random-Informed.[16] In the Strong condition, products are ordered in each task so that bonus B is strongly correlated with the product's rank. Products closer to the top of the page have higher bonus Bs.[17] In this condition, participants are not informed of how rankings are generated. In the Random condition, products are randomly ordered in each task, and as in Strong, participants are not informed of the ranking algorithm. In the Random-Informed condition, products are randomly ordered in each task, and participants are told that products have been randomly shuffled.[18]

bonus values.

[14]These values were chosen to calibrate bonus payments to be between \$1.50 and \$2.50.

[15]We ask participants about two products that have the same bonus A but in positions 1 and 3 in the rankings. We elicit their belief about the probability that the product in position 1 has a higher bonus B than the product in position 3. The correct answer depends on the participant's ranking algorithm.

[16]We report randomization and manipulation checks in Appendix I.2.1.

[17]To generate these imperfectly informative or "noisy" rankings, we draw an additional, independent noise term $\sim N(0, 5)$ and add it to product bonus B values, then sort in descending order of the perturbed value.

[18]This study was preregistered and run with a coding error in a fourth condition, where the ranking

Why include these conditions? The Random-Informed condition is designed to isolate the rank-specific search costs. Because participants in this condition are informed that products are randomly ordered, any tendency to search more prominently positioned products is driven by search costs rather than by beliefs. Therefore, the Random-Informed condition allows us to estimate search costs. Given these estimated search costs from the Random-Informed condition, we use variation in behavior in the Random and Strong conditions to isolate beliefs about rank-specific payoffs. Any residual position effects after accounting for rank-specific search costs in the Random and Strong conditions can be attributed to beliefs.

We include the Random condition for two reasons. First, comparing Random to Random-Informed allows us to detect whether participants enter the experiment with equilibrium beliefs from other search settings. In other words, do participants have position-specific prior beliefs? Second, we can ensure that any additional position effects observed in Strong, relative to those observed in Random-Informed, are due to the ordering of products and not due to participants not knowing the algorithm.

3.2 Experimental Results

After excluding participants who repeatedly failed comprehension checks and those who completed the survey too quickly, our final dataset contains 961 participants.[19] Across all search tasks where participants searched at least two products, participants chose the product with the largest bonus in 93% of tasks.

Table 1 reports the summary statistics for each condition. In all three conditions, the median number of products searched (clicked on) is 1, while the mean ranges from 2.1 to 2.2 products.[20] Participants search the most in the Random condition and the least in the Strong condition, but these differences are not statistically significant. However, the total bonus earned significantly varies across conditions. Participants in Strong earn significantly more than participants in the other conditions.

algorithm is supposed to switch from Strong to Random after five tasks. We dropped this condition from our analysis in Study 1 and re-ran this condition in Study 2.

[19] As noted in our pre-registration, we omit participants who completed the survey in less than the first percentile, which is 3.3 minutes.

[20] The search intensity in our setting is similar to that of other experimental search papers. For example, Morozov and Tuchman (2024) find that the average participant searches 1.9 products.

Table 1: Summary Statistics across Conditions (Study 1)

	Condition Mean (SD)		
	Random	Random Informed	Strong
Num Searches	2.22 (2.04)	2.14 (1.77)	2.11 (1.82)
Bonus Earned	66.41 (11.67)	67.18 (10.94)	68.42 (11.4)
Click Max(a) First	0.41 (0.49)	0.44 (0.5)	0.41 (0.49)
Time Spent per Task	13.32 (16.71)	13.4 (17.99)	13.88 (19.43)
N(Tasks)	3300	3170	3140
N(Participants)	330	317	314

Notes: Table reports mean and standard deviation by condition for each variable. Num Searches is the number of products searched in a task. Bonus Earned is the net bonus earned in a task (bonus of selected product net of nominal search costs). Click Max(a) First is a dummy variable which is equal to one if the participant clicked on the product with the highest bonus A value first in a task. Time Spent per Task is the seconds spent on a task from page load until final click. The prior statistics summarize across participants and tasks within a condition.

Before presenting the empirics for the search model, we report model-free evidence of position-specific search costs, beliefs, and learning. First, we find strong evidence of position effects in all conditions, consistent with what is documented in many other field and lab settings. Figure 2 plots the probability that participants search each item by rank across conditions. All conditions exhibit position effects: top-ranked items are more likely to be searched.[21] Note that these position effects are not driven by differences in pre-search payoffs because bonus A is independent of rank. In the Strong condition, moving a product one position down the page, on average, reduces the probability of a consumer searching it by the same amount as offering 0.55 fewer points in bonus A.[22] The average effect size is smaller in the Random and Random-Informed conditions: one rank is "worth" 0.3 bonus A points. In all conditions, position effects are most pronounced for the top five products, evidenced by the steeper slope, and flatten out for the bottom five. In the Strong condition, position effects are especially strong for the first product. In this condition, 33% of participants click on the first product. This suggests that our model should be flexible enough to capture the

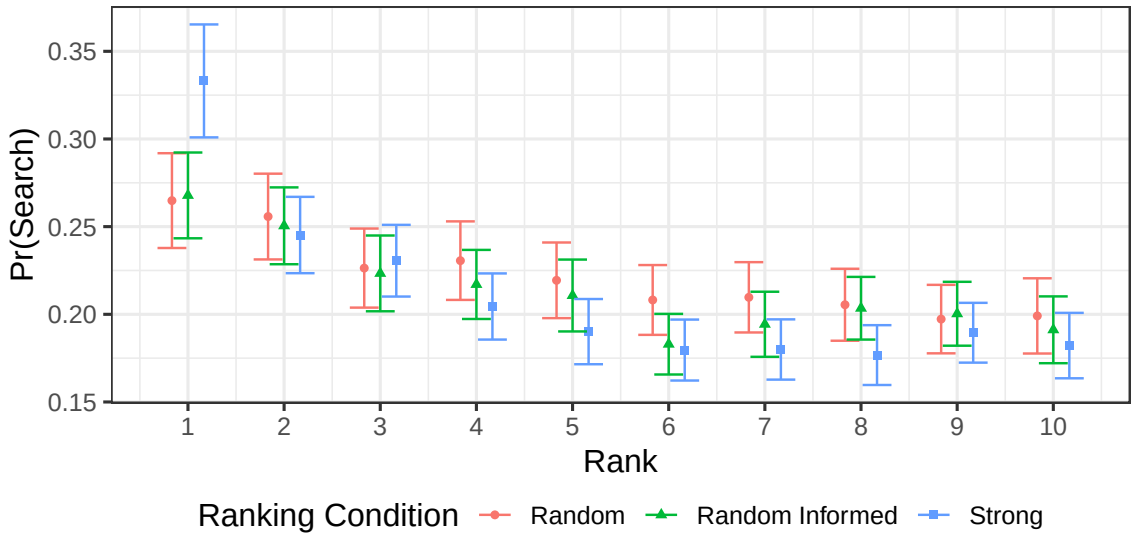
[21] Sixty-two percent of sequences with more than one product search involve out-of-order searching (returning to click on a higher-ranked product after first clicking on another option). This invalidates the simplifying assumption often made for tractability that search order is exclusively top-to-bottom. (e.g., (Choi and Mela, 2019; Lam, 2023)).

[22] These magnitudes are obtained from the results in Table I.1 in the appendix, in which we regress an indicator variable for whether the participant searches product j on an interaction between the treatment condition and product j's rank r.

increased search propensity for the first product in particular.

Figure 2 also provides evidence that both position-specific search costs and beliefs exist. The position effects we observe in the Random-Informed condition provide evidence for position-specific search costs. Recall that in this condition, we inform participants that the products are randomly ordered, essentially "shutting off" beliefs. The difference in position effects between the Strong and Random-Informed conditions provides evidence for position-specific beliefs about the returns to search. Furthermore, the lack of differences in position effects across conditions in the first search task, as shown in Figure 2 in the appendix, suggests that participants in the Strong condition are learning through experience that the first product offers the greatest returns to search. In addition, we find no evidence of prior position-specific beliefs, as there are no differences in position effects between the Random and Random-Informed conditions, whether averaged over all tasks or for the first task alone.[23]

Figure 2: Search Probability by Rank and Condition (Study 1)

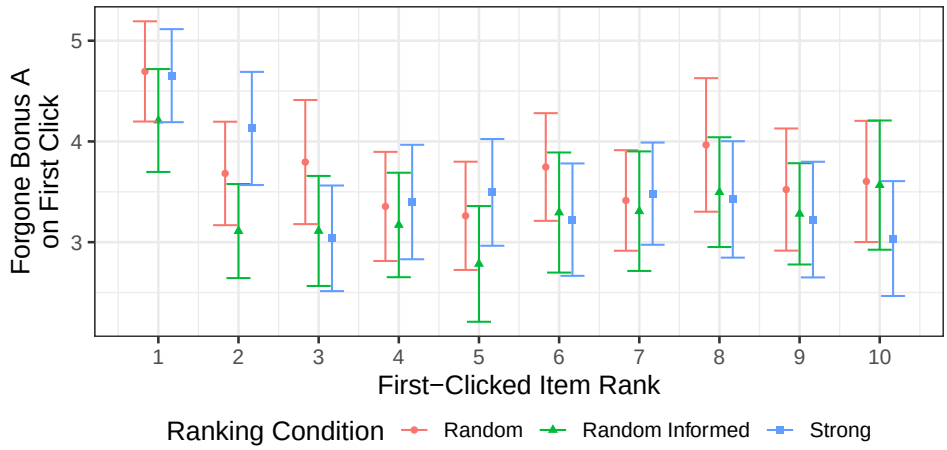


Notes: Figure shows the probability each item is searched by rank and condition across participants and tasks. 95% confidence intervals are based on standard errors clustered at the participant level.

Second, participants in all conditions forgo higher certain payoffs (bonus As) from other

[23]We also do not find heterogeneity in search behavior based on whether the participant is a high or low frequency online shopper. This suggests that frequent online shoppers do not have different prior position-specific beliefs. We present this evidence in Appendix C.

Figure 3: Forgone Bonus A by Rank of First-Clicked Item (Study 1)



Notes: This figure plots the average forgone bonus A of the first-clicked item by the rank of the first-clicked item. Error bars represent 95% confidence intervals, which are based on standard errors clustered at the participant level.

products to search products that are ranked towards the top of the page. How much bonus A they forgo helps identify the magnitude of position effects. Recall that the bonus As of top-ranked items are not higher than those for the bottom-ranked items, and the participant's total bonus for the search task depends on bonus A + bonus B. Therefore, because position does not provide information about bonus B in Random-Informed, participants should prioritize searching products with higher bonus As first to maximize their certain payoff, if there are no position-specific costs. That is, according to the model in Weitzman (1979), searching the product with the highest bonus A first is the optimal strategy for participants in Random-Informed if search costs do not vary by position. Figure 3 displays the average forgone bonus A of the first searched product by the product's rank and by condition. The forgone bonus A is defined as the maximum bonus A of all products in the task minus the bonus A of the searched product. Figure 3 shows that participants in all conditions forgo more points when starting their search at the top-ranked product. This implies that participants click on the top-ranked products at the expense of products that may have a higher bonus A but are towards the bottom of the page. Additionally, on average, participants forgo more than three bonus A points, even when they begin their search with a less prominently ranked item. This suggests that participants may be inattentive, which

we should account for in our model.

Table 2: Learning to Find Higher Bonus Bs (Study 1)

Dependent Variable:	b_B of First Searched
Model:	(1)
Variables	
Search Task	0.0146
	(0.0710)
Search Task × Condition = Random	0.1143
	(0.0989)
Search Task × Condition = Strong	0.2824***
	(0.1052)
Fixed-effects	
Participant (961)	Yes
Fit statistics	
Observations	9,610
R^2	0.13796
Within R^2	0.00226

Clustered (Participant) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: This table reports the OLS estimates of Equation 5. The dependent variable is the bonus B of the product that the participant searches first in each task. Thus, each observation is at the participant-task level. The baseline condition is Random-Informed.

Finally, we show that the larger bonuses earned by participants in the Strong condition are due to learning during the experiment (as opposed to position-specific prior beliefs). To demonstrate this, we consider the bonus B of the product that participants search first in a given search task. If the participant learns which positions yield higher bonus Bs, they should be more likely to search those positions first. Therefore, we estimate the following linear regression:

$$b_{it} = \alpha_i + \beta_1 SearchTask_{it} + \beta_2 (SearchTask_{it} \times Condition_i) + \epsilon_{it}. \tag{5}$$

The dependent variable b_{it} is the bonus B of the product that i searches first in task t. The variable $SearchTask_{it}$ denotes the task number, which is from one to ten; $Condition_i$ is participant i's experimental condition. We also include participant fixed effects α_i.[24]

[24]Note that this specification does not include the main effect of $Condition_i$ because it is collinear with

Table 2 reports the OLS estimates of Equation 5. The coefficients of the main effect of $Task$ and its interaction with $Condition = Random$ are not statistically significant. These results suggest that, as expected, participants in Random-Informed and Random do not learn to find better bonus Bs with experience. However, the statistically significant interaction with Strong indicates that participants in the Strong condition do learn to identify better bonus Bs through experience.

We expect that our results on learning potentially serve as an upper bound versus non-lab settings, where post-search utility requires further evaluation or even consumption experience. In field settings, attributes can be more difficult for consumers to evaluate and are rarely presented as univariate numerical payoffs. We return to these differences in Section 6.

3.3 Model and Estimation

We estimate a version of the directed sequential search model described in Section 2 using a Bayesian estimator similar to that Morozov (2023). As described in Morozov (2023), this estimator is more numerically stable and can better accommodate heterogeneity compared to frequentist methods. In our model, heterogeneity in position effects based on individual-level attributes plays an important role, as we describe below.

To simplify estimation, we re-parameterize participant i's reservation utility of product j in task t (described in Equation 3) as

$$z_{itj} = a_{itj} + E[b] + \delta_{itj}, \tag{6}$$

where a_{itj} is the pre-search, product-specific component of utility, and $E[b]$ is the average payoff from search, which is 40 points. Following Morozov (2023), instead of directly modeling search costs, we include the *search propensity*, which we denote by δ_{itj}. Search propensities capture rank-specific differences in the likelihood the consumer searches j, regardless of whether this is due to beliefs about bonus B or search costs[25].

participant fixed effects.

[25]We assume that participants are myopic across tasks - that is, they search in each task to maximize payoffs for that particular task and do not actively search more in earlier tasks to learn about position-specific distributions. Figure I.2 in the appendix shows that the number of products searched does not vary

The consumer i's propensity to search product j is specified as

$$\delta_{itj} = \beta_i^0 + \beta_i^r \cdot (r_{itj} - 1) + \epsilon_{itj}. \tag{7}$$

The term ϵ_{itj} represents idiosyncratic differences in search propensities that are uncorrelated with position and cannot be rationalized by bonus A, such as differences in attention costs (e.g., seeing the third product first). Realizations of ϵ can explain why participants in the Random-Informed condition search products that have a lower bonus A, even in positions that are more costly to search on average. Assuming that ϵ_{itj} is drawn from a normal distribution with mean zero and variance σ_ϵ^2, we can augment the data by drawing δ_{itj} directly from $\delta_{itj} \sim N(\beta_i^0 + \beta_i^r \cdot (r_{itj} - 1), \sigma_\epsilon^2)$.

We refer to the intercept, β_i^0, as the baseline search cost. This term can be interpreted as i's propensity to search the first-ranked product, which also includes i's search costs regardless of the product's rank, such as the one point search cost imposed by the study and the opportunity cost of time reading the search results. Holding all else constant, a larger β_i^0 increases the total number of searches the participant performs.

The term β_i^r is the effect of *rank* on search propensity relative to the first-ranked product. A more negative value of β_i^r would imply that user i is more likely to search products displayed at the top of the rankings.

Recall that in the Random-Informed condition, we attribute variation in search propensities with respect to rank to rank-specific search costs. Residual differences in the rank-specific search propensity (after accounting for the search cost estimated from Random-Informed) in the other conditions can then be interpreted as differences in beliefs. Therefore, it is important to measure the heterogeneity in position effects by the participant's treatment condition.

Because each participant completes only ten tasks, estimating a separate β for each participant would yield very noisy estimates. We therefore assume that β_i, the vector of individual coefficients, is drawn from a shared distribution using a hierarchical model and

significantly by search task, which supports this assumption. Note that like Weitzman (1979), our model ignores product revisits (i.e., consumers viewing a product multiple times in the same search task). Only a small fraction (4.56%) of search tasks in our data involved revisits.

depend on participant characteristics X_i, as follows:

$$\beta_i \sim N(\Delta X_i, V_\beta). \tag{8}$$

The matrix X_i includes (a) an intercept term, (b) an indicator for the Strong condition, (c) an indicator for the Random condition, (d) an indicator for whether user i passed all comprehension checks on the first attempt ($Comp_i$), and (e) an indicator for whether i's self-reported online shopping frequency is greater than once a week ($ShopFreq_i$). We allow β_i to vary by $Comp_i$ and $ShopFreq_i$ because participants that passed the comprehension checks on the first attempt exhibit stronger position effects and search more than those who do not, and we hypothesize that participants who shop online more frequently may have position-specific prior beliefs from their extensive experience with online product rankings. We report the reduced form effects on search by these attributes in Section C in the appendix.

The vector Δ contains the coefficients that shift the mean value of each β element by demographics and treatment condition, and V_β is the covariance matrix of β_i across individuals. The coefficient β_i^r is given by

$$\beta_i^r = \Delta_0^r + \Delta_S^r \mathbb{1}\{Cond_i = S\} + \Delta_R^r \mathbb{1}\{Cond_i = R\} + \Delta_{comp}^r Comp_i + \Delta_{shop}^r ShopFreq_i + \eta_i^r. \tag{9}$$

The intercept term Δ_0^r represents differences in search propensities for participants in Random-Informed; these should be attributed to search costs. The coefficients Δ_S^r and Δ_R^r represent any differences in search propensities by rank above and beyond those induced by search costs in the Strong and Random conditions, respectively. These differences are induced by beliefs.

Similarly, the baseline search cost β_i^0 varies by the same characteristics:

$$\beta_i^0 = \Delta_0^0 + \Delta_S^0 \mathbb{1}\{Cond_i = S\} + \Delta_R^0 \mathbb{1}\{Cond_i = R\} + \Delta_{comp}^0 Comp_i + \Delta_{shop}^0 ShopFreq_i + \eta_i^0. \tag{10}$$

Notice that in Equation 7, rank impacts search propensity linearly. However, position effects can manifest by especially increasing the participant's propensity to click on the first listing, as shown in Section 3.2. Therefore, we estimate an alternative specification

that captures the search propensity of the top-ranked product separately from the position effects across other positions:

$$\delta_{itj} = \beta_i^0 + \beta_i^1 \cdot \mathbb{1}\{r_{itj} = 1\} + \beta_i^r \cdot \Big(\mathbb{1}\{r_{itj} \neq 1\} \cdot (r_{itj} - 2)\Big) + \epsilon_{itj}. \tag{11}$$

In this specification, a positive β_i^1 indicates stronger position effects (i.e., that the participant is more likely to click on the first product relative to products lower on the page).

We estimate this model using a Gibbs sampler, which cycles through sampling the conditional posteriors of different parameters. Algorithm 1 provides an overview of the sampler, and Appendix B outlines the estimation steps in detail, describes the priors, which are non-informative, for all parameters, and shows the performance on simulated data.

Algorithm 1 Gibbs Sampler Overview

1: **for** $m = 1, \ldots, M$ **do**
2: Update reservation utilities $z^{(m)}|\beta^{(m-1)}, \sigma_\epsilon^{(m-1)}$
3: Update search propensity coefficients $\beta^{(m)}|z^{(m)}, \sigma_\epsilon^{(m-1)}$
4: Update search propensity coefficient shifters $\Delta^{(m)}, V_\beta|\beta^{(m-1)}$
5: Update error variance $\sigma_\epsilon^{(m)}|z^{(m)}, \beta^{(m-1)}$
6: **end for**

3.4 Results

The estimated parameters are presented in Table 3. The first column of parameter estimates (titled "Baseline") reports the estimates from the specification in Equation 7, and the second column (titled "Separate Position 1") reports the estimates from the specification in Equation 11. In both specifications, consumers show position effects in Random-Informed, which are due to search costs: $\Delta_0^r = -0.24$ in the "baseline" specification and $\Delta_0^r = -0.22$ in the "Separate Position 1" specification. This confirms that users find it more costly to search less prominently ranked products. In the Strong condition, there is an additional linear rank effect on search propensities due to beliefs, $\Delta_S^r = -0.19$. This is almost of the same magnitude as the rank effect due to costs. Search models that do not differentiate between the two would therefore significantly overestimate search cost differences across positions.

Recall that comparing the estimates between Random and Random-Informed informs us of whether participants have position-specific prior beliefs. The coefficients Δ^r_R and Δ^1_R are small in magnitude, and their 95% confidence intervals contain zero. This reaffirms the conclusions from Figure 2—that participants do not have position-specific prior beliefs.

The estimates in the "Separate Position 1" specification show that much of the rank effect arises from the increased propensity to search the first product on the page. The positive estimate of Δ^1_S implies that participants in the Strong condition are more likely to click the first product, relative to participants in Random and Random-Informed. This increased propensity relative to Random-Informed and Random is due to beliefs that are formed during the study.

The model estimates regarding heterogeneity by participant attributes reflect those reported in Section 3.2. Participants who passed all comprehension checks in the first attempt have a higher search propensity ($\Delta^0_{comp} = 19$ and $\Delta^0_{comp} = 18$ in both specifications) than those who did not. One possible explanation is that participants with a low opportunity cost of time (and hence low search costs) spend more time reading the instructions and are therefore more likely to pass the comprehension checks on the first attempt. These lower search costs are reflected in a higher search propensity. These participants also had stronger position effects ($\Delta^1_{comp} = -0.46$ in the first column, and $\Delta^1_{comp} = -0.36$ in the second). Participants with a higher online shopping frequency also had higher position effects ($\Delta^r_{shop} = -0.12$) and are more likely to search the first product ($\Delta^1_{shop} = 1$), although not significantly.

In order to compute search costs, we invert the augmented draws of δ: $c_{itj} = \sigma_b \zeta^{-1}(\frac{\delta_{itj}}{\sigma_b})$. These estimates imply that consumers face search costs for products in rank one of an average of \$0.05 and a median of \$0.03.[26] Figure 4 plots the distribution of implied dollar search costs for Random-Informed participants for the first item. These costs combine the effort of processing information and the physical effort of clicking and reading. Note that the dollar amounts are low, but the median time spent per task is nine seconds.

As previously mentioned, another way that position can influence search behavior is through beliefs about the variance of bonus b at each position. We had previously ab-

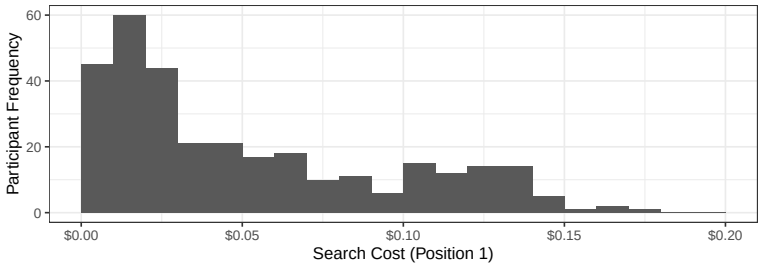
[26] These values are obtained from the estimates in the "Separate Position 1" specification.

Table 3: Search Model Estimation Results (Study 1)

		Estimate (SD)	
Variable	Coefficient	Baseline	Separate Position 1
Rank Coefficient β^r			
Intercept	Δ_0^r	-0.24 (0.1)**	-0.23 (0.096)**
Strong	Δ_S^r	-0.19 (0.12)	-0.063 (0.11)
Random	Δ_R^r	-0.096 (0.12)	-0.12 (0.11)
Comp	Δ_{comp}^r	-0.46 (0.095)***	-0.36 (0.091)***
ShopFreq	Δ_{shop}^r	-0.12 (0.1)	-0.066 (0.094)
Pos 1 Coefficient β^1			
Intercept	Δ_0^1	-	-0.36 (0.7)
Strong	Δ_S^1	-	1.9 (0.8)***
Random	Δ_R^1	-	-0.28 (0.78)
Comp	Δ_{comp}^1	-	2.7 (0.61)***
ShopFreq	Δ_{shop}^1	-	1 (0.66)
Baseline Search Propensity β^0			
Intercept	Δ_0^0	-22 (1.6)***	-22 (1.6)***
Strong	Δ_S^0	2.3 (2)	1.4 (1.8)
Random	Δ_R^0	0.2 (2)	0.35 (1.8)
Comp	Δ_{comp}^0	19 (1.5)***	18 (1.4)***
ShopFreq	Δ_{shop}^0	-1.7 (1.7)	-2.1 (1.6)
Heterogeneity			
Rank Coefficient	$\sqrt{V}_{\beta^r}$	1.3 (0.11)	1.1 (0.12)
Pos1 Coefficient	$\sqrt{V}_{\beta^1}$	-	6.8 (9.8)
Mean Search Propensity	$\sqrt{V}_{\beta^0}$	24 (32)	22 (28)
Reservation Utility	σ_ϵ	15 (4.6)	15 (6.1)

Notes: We report the statistical significance only for the search propensities and their shifters. Signif. Codes: $***: 0.01, **: 0.05, *: 0.1$. Posterior means and standard deviations are based on the thinned chain which drops the first 5,000 draws of the chain and keeps every tenth draw thereafter.

Figure 4: Distribution of Search Costs for Product 1 across Consumers (Study 1)



Notes: Figure plots the frequency distribution of the marginal search cost for the first position product across consumers.

stracted from this mechanism by assuming that participants do not learn about position-specific variances. However, because rank conveys information about bonus b in Strong, it mechanically reduces the variance of bonus b for each position. A potential concern is that the search costs we estimate are biased because we do not account for learning about this reduction in variance in Strong. To address this, we check the sensitivity of our search cost estimates by re-calculating search costs under the assumption that participants have full knowledge of bonus b's rank-specific variance in Strong. We find that the estimates of search costs are not very sensitive to the assumed variance of bonus B. In Study 1, the estimated search cost for position 1 was 5.11 cents (4.92–5.33) when assuming variance of bonus B = 144. If we instead assume a variance of 46 (equal to the variance of bonus B in the Strong condition), the estimate remains similar at 4.94 cents (4.71–5.16). Therefore, abstracting away from learning about variance of bonus b does not significantly impact our results.

How much of the position effect we find is due to costs? To measure this, we consider two metrics. First, what fraction of the difference in search rates from position one to position ten in Strong can be explained by costs? At the posterior mean, 35% of this gap is due to differences in costs, and the remaining 65% is attributed to differences in beliefs. Second, we document that the difference in search propensities due to beliefs in Strong between ranks one and two is larger than the entire difference due to costs across ten positions. These metrics both show that beliefs play a much larger role than search costs in producing position effects when rankings are informative of products' post-search attributes.

So far, the estimates we discussed are averaged over all tasks. However, recall that we find model-free evidence of participants learning through their search experience—position effects in the first task do not differ significantly between Random-Informed and Strong in the first search task (Figure I.3 in the appendix) but are statistically different on average across all tasks. In Appendix D, we estimate the search model and allow position effects to vary for tasks 1–5 and 6–10 and find that position effects in Strong are indeed larger in tasks 6–10 than tasks 1–5.

In summary, the search model estimates in Study 1 highlight three main results. First, participants' search propensities vary significantly with rank. This is consistent with prior work, which suggests that our experimental setting can replicate features of field search settings. Second, these position effects are a combination of both costs and beliefs, and beliefs are a major component of position effects in the Strong condition, with 65% of the magnitude of position effects coming from beliefs. Third, participants adjust their search behavior over time in a manner consistent with learning about the underlying algorithm from repeated searching and updating their beliefs to be closer to the ranking algorithm.

4 Study 2: Consumer Responses to Algorithm Changes

The results from Study 1 imply that position effects are not only static, unchangeable costs—rather, they arise in part from consumers' beliefs, which may evolve in equilibrium. In Study 2, we seek to understand the implications of these equilibrium effects on search behavior when there are changes in the ranking algorithm. We focus on changes in the ranking algorithm for several reasons. First, changes in ranking algorithms are frequently analyzed in counterfactuals. A recent example is Donnelly et al. (2024), who compare consumer welfare under several different personalized ranking algorithms. Understanding the speed and degree of learning is important not only for determining the accuracy of predictions for managerial purposes but also for welfare analysis. Analyzing practices like steering, where certain products are placed at the top of rankings, requires considering counterfactual rankings (Lee and Musolff, 2023; Lam, 2023). If consumers can learn the ranking algorithm, they are less likely to be misled or manipulated into specific search and

demand patterns. Furthermore, platform steering has a smaller impact on behavior the faster consumers learn the new algorithm. Conversely, the greater the influence of search costs on position effects, the more significant the impact of platform steering on consumers. Careful measurement of both belief evolution and rank-specific costs is critical for evaluating behavior under alternative platform designs.

Second, this study allows us to test our approach against typical paradigms for estimating rank-specific search costs. A common approach to measuring search costs is to randomize product rankings. The algorithm shifts from the platform's default ranking to a random ranking without disclosing this to consumers. These types of experiments can lead to incorrect estimates of the search costs. We include a condition that mimics this paradigm, Strong-to-Random, which we describe below, that allows us to compare how counterfactual ranking algorithms impact consumer surplus when position effects are correctly attributed to both beliefs and search costs relative to when they are misattributed to search costs only.

4.1 Design

The design of this study is very similar to that of Study 1. Individuals search through a vertical list of products with two additive, independently drawn bonus components. The products are generated from the same data generating process. However, in Study 2, participants complete fifteen tasks, rather than ten. This allows more opportunities for participants to learn the algorithm.[27] We also change the number of products per search task from ten to twenty. We do this because in Study 1, 10% of participants searched all ten products in at least one task, which consequently reduces the precision of our search cost estimates.[28] Increasing the number of products to twenty reduces the portion of the sample for which we cannot accurately estimate search costs.

Participants are randomized into one of six conditions, which are summarized in Table 4. Two of the conditions are Random-Informed and Strong; these are identical to those in Study 1, and the ranking algorithms remain constant across search tasks. The remaining

[27] In addition, rather than compensating participants for all tasks, we inform participants that 8 of the 15 tasks will be drawn at random and bonuses paid out based on those task performances. We also re-calibrate the point conversion to 500 points equals $1. We include participant instructions in Appendix J.2.

[28] More specifically, we can only estimate the lower bound of search costs for tasks where the participant searches all items.

Table 4: Study 2 Design

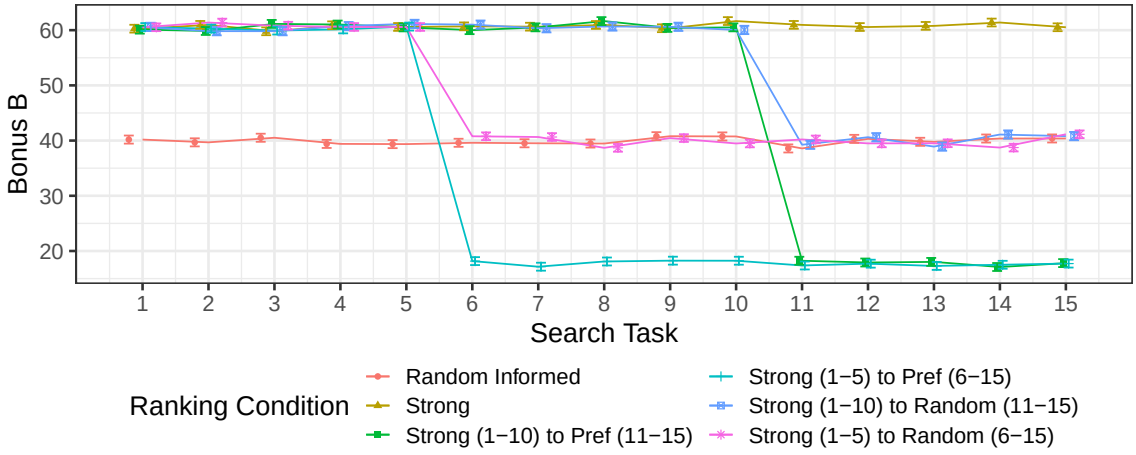
Condition	Ranking Algorithm		
	Tasks 1–5	Tasks 6–10	Tasks 11–15
Random-Informed	*Random*	*Random*	*Random*
Strong	*Strong*	*Strong*	*Strong*
Strong-to-Random5	*Strong*	*Random*	*Random*
Strong-to-Random10	*Strong*	*Strong*	*Random*
Strong-to-Pref5	*Strong*	*Pref*	*Pref*
Strong-to-Pref10	*Strong*	*Strong*	*Pref*

Notes: Ranking algorithms are defined as follows. *Random* sorts products randomly. *Strong* is defined as in Study 1, where products are sorted noisily in descending order of bonus B. *Pref* first sorts products as in Strong, but then puts the two lowest bonus B items in the first two positions (so that the third position has the highest expected value).

conditions include a change in the ranking algorithm after a set number of search tasks in the *Strong* algorithm. In one condition, Strong-to-Random5, the ranking algorithm is identical to the one in Strong for the first five tasks, and then switches to *Random* for the remaining ten tasks. We also have a condition Strong-to-Random10 in which the algorithm is *Strong* for the first ten tasks and *Random* for the last five. These two conditions mimic the typical experimental paradigm used to measure search costs, in which the algorithm switches from the platform's default to *Random* without disclosure to consumers. The remaining two conditions are Strong-to-Pref5 and Strong-to-Pref10, which are meant to represent a switch to platform steering. In these conditions, the first five and ten tasks, respectively, are sorted by the *Strong* algorithm. The algorithm then switches to the *Pref* algorithm, which is the following: we first sort the products using the *Strong* algorithm, then we move the two products with the lowest values of bonus B to the first and second position in the list. Recall that bonus A is independent of rank and bonus B, so the products sorted by *Strong* and *Pref* algorithms are indistinguishable to participants before search. Figure 5 illustrates how the bonus B of the first-ranked product differ by condition across search tasks. The only condition that informs participants of the ranking algorithm is Random-Informed. Also, mimicking most digital platforms' practice, we do not inform participants of the change in the ranking algorithm.

Like Study 1, Study 2 includes comprehension checks. The first comprehension check

Figure 5: Average Bonus B of First-Ranked Product by Condition (Study 2)



Notes: Error bars represent 95% confidence intervals, based on i.i.d. standard errors across participants.

tests their understanding of the bonus payment and the second tests for understanding that bonuses A and B are uncorrelated. Participants who incorrectly answer the second comprehension check twice are removed from the study.

In this study, we also directly elicit participants' prior and posterior beliefs about rank-specific payoffs. We display to participants the search results page, which contains only the bonus As of twenty ordered products, and ask them to predict the bonus Bs of the first-, tenth-, and twentieth-ranked products. We ask this question following comprehension checks (before the incentivized tasks) and again immediately after completing the fifteenth task.[29] In the pre-task belief elicitation, we do not incentivize participants' responses. However, in the post-task belief question, we reward participants with ten points for each of the three beliefs answered correctly for their condition.

4.2 Experimental Results

We collected data from 3,569 participants on Prolific. After removing participants who completed the survey too quickly (less than the first percentile, which is 4.2 minutes), our

[29]In a pilot study, we tested whether asking participants about their beliefs about bonus Bs before the search tasks influenced their subsequent search behavior. We did not find evidence of differences in search.

remaining sample includes 3,533 participants.[30]

We report the summary statistics in Table 5. Compared with Study 1, participants search less (1.92 products in Random-Informed in Study 2 compared to 2.14 in Random-Informed in Study 1). However, the mean bonus earned across the studies in Random-Informed are similar (67.18 and 67.16 in Studies 1 and 2, respectively).

Table 5: Summary Statistics across Conditions (Study 2)

	Condition Mean (SD)					
	RI	Str	Str to P (10)	Str to P (5)	Str to R (10)	Str to R (5)
Num Searches	1.92 (1.6)	1.78 (1.44)	1.86 (1.5)	1.97 (1.58)	1.92 (1.64)	1.99 (1.61)
Bonus Earned	67.16 (11.29)	69.97 (11.63)	68.95 (12.11)	68.98 (11.93)	68.66 (11.69)	68.1 (11.38)
Click Max(a) First	0.31 (0.46)	0.29 (0.45)	0.31 (0.46)	0.32 (0.47)	0.29 (0.46)	0.3 (0.46)
Time Spent per Task	13.41 (15.38)	12.02 (14.32)	12.92 (17.93)	13.57 (19.81)	12.51 (14.81)	13.53 (27.55)
N(Tasks)	8595	9195	8835	8910	8070	9390
N(Participants)	573	613	589	594	538	626

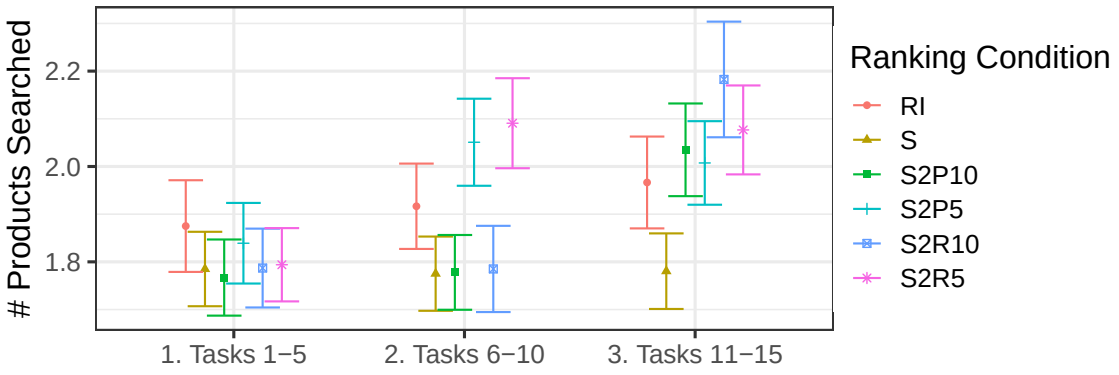
Notes: "Str" stands for Strong, "P" for Pref, "R" for Random, and "RI" for Random-Informed. Table reports mean and standard deviation by condition for each variable. Num Searches is the number of products searched in a task. Bonus Earned is the net bonus earned in a task (bonus of selected product net of nominal search costs). Click Max(a) First is a dummy variable which is equal to one if the participant searched the product with the highest ex-ante bonus value first on a task. Time Spent per Task is the seconds spent on a task from page load until final click. The prior statistics summarize across participants and tasks within a condition.

Because the ranking algorithm changes after five or ten search tasks in some conditions, we report results by each "task group" separately. Figure 6 plots the average number of products searched per task for tasks 1–5, 6–10, and 11–15. In tasks 6–10, the search depth in Random-Informed is significantly greater than Strong, Strong-to-Random10, and Strong-to-Pref10. However, search depth increases in Strong-to-Pref5 and Strong-to-Random5; recall that in these conditions, the product rankings in first five tasks are identical to those in Strong. This increase in search depth also occurs in tasks 11–15 for Strong-to-Random10 and Strong-to-Pref10, relative to Random-Informed. This pattern demonstrates that participants respond to the algorithm change; when the products they click on first yield a lower bonus, they continue searching.

Here, we confirm the same stylized facts we document in Study 1. Participants in all conditions display strong position effects. Figure 7 reports the probability that each item is

[30]This study, including this data cleaning process, is pre-registered on AsPredicted.org (#159069). The pre-registration is available at https://aspredicted.org/3RP_ZVB.

Figure 6: Search Depth by Rank by Condition (Study 2)



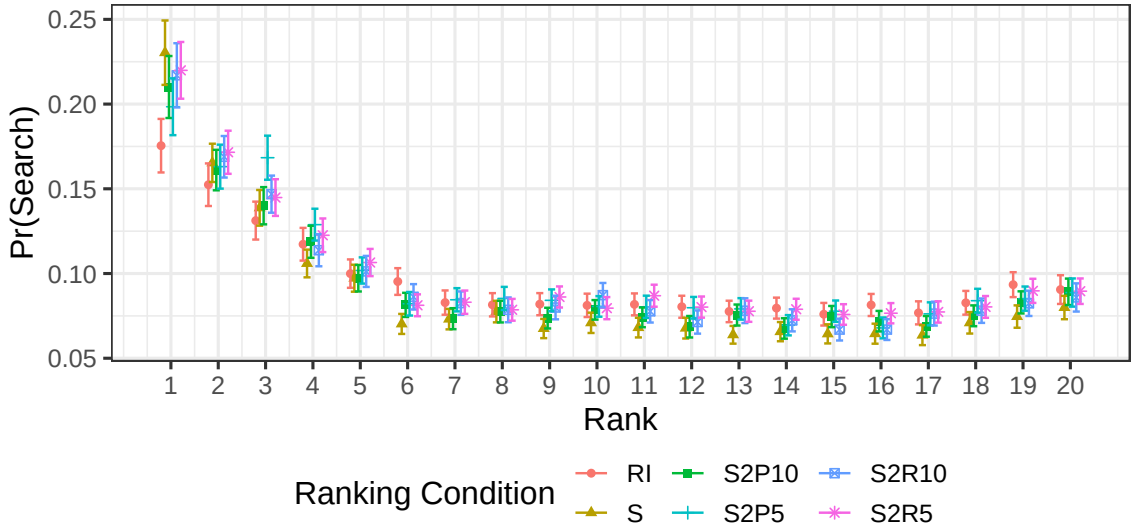
Notes: Error bars represent 95% confidence intervals, which are based on standard errors clustered at the participant level. "RI" refers to Random-Informed, "S" refers to Strong, "S2P10" and "S2P5" refer to Strong-to-Pref10 and Strong-to-Pref5, respectively, and "S2R10" and "S2R5" refer to Strong-to-Random10 and Strong-to-Random5, respectively.

clicked by rank across conditions. In Random-Informed, participants are more than twice as likely to search the first item as compared to the tenth item. The effect is more pronounced in all other conditions, where the probability of searching the first item is significantly higher than in Random-Informed. We confirm these results in regressions (Table 6), which shows that the average effect of rank is greater in all conditions relative to Random-Informed. In addition, we confirm that the use of undisclosed randomization (in Strong-to-Random) leads to incorrect beliefs in the short term. For example, using only the eleventh task, we estimate that participants in Strong-to-Random10 are just as likely to click the first ranked item as those in Strong, even though the products are randomly ordered. This probability of searching the first-ranked item in Strong-to-Random10 is 50% higher than in Random-Informed.

As in Study 1, we can use the forgone payoffs from bonus A to see the scale of position effects in cost terms. Figure 8 shows the bonus A left on the table with the first click by condition and position. For Random-Informed, participants forgo 1.7 additional bonus points (40% additional forgone bonus A) when starting their search with the first item versus the tenth.

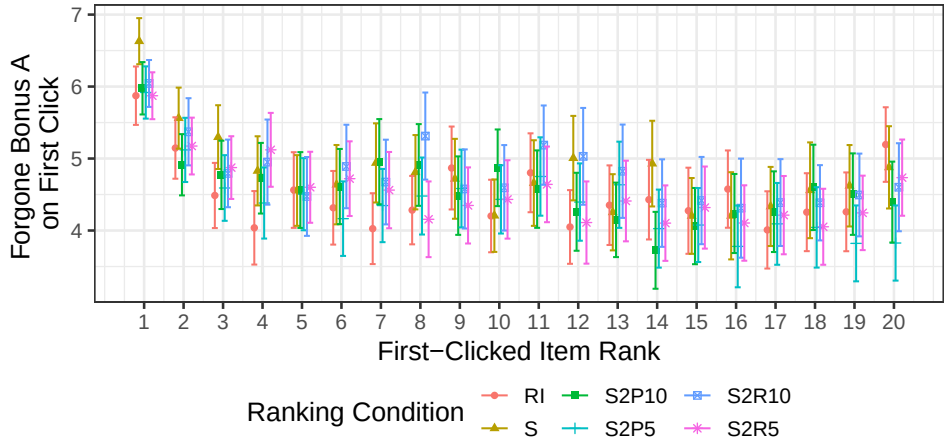
Recall that in this study, we directly elicit prior and posterior beliefs about expected

Figure 7: Search Probability by Rank by Condition (Study 2)



Notes: Error bars represent 95% confidence intervals, which are based on standard errors clustered at the participant level. "RI" refers to Random-Informed, "S" refers to Strong, "S2P10" and "S2P5" refer to Strong-to-Pref10 and Strong-to-Pref5, respectively, and "S2R10" and "S2R5" refer to Strong-toRandom10 and Strong-to-Random5, respectively.

Figure 8: Forgone Bonus A by Rank of First-Clicked Item (Study 2)



Notes: Error bars represent 95% confidence intervals, which are based on standard errors clustered at the participant level. Condition and rank-specific values average over all tasks (including changing ranking algorithms). "RI" refers to Random-Informed, "S" refers to Strong, "S2P10" and "S2P5" refer to Strong-to-Pref10 and Strong-to-Pref5, respectively, and "S2R10" and "S2R5" refer to Strong-toRandom10 and Strong-to-Random5, respectively.

Table 6: Effect of Rank on Search by Condition (Study 2)

Dependent Variable:	1(Product Searched)					
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
$r \times$ Condition = Random Informed	-0.0032***	-0.0032***	-0.0032***	-0.0032***	-0.0032***	-0.0032***
	(0.0003)	(0.0003)	(0.0003)	(0.0003)	(0.0003)	(0.0003)
$r \times$ Condition = Strong	-0.0049***	-0.0049***	-0.0049***	-0.0049***	-0.0049***	-0.0049***
	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)
$r \times$ Condition = Strong2Pref10	-0.0043***	-0.0042***	-0.0043***	-0.0042***	-0.0043***	-0.0042***
	(0.0003)	(0.0003)	(0.0003)	(0.0003)	(0.0003)	(0.0003)
$r \times$ Condition = Strong2Pref5	-0.0045***	-0.0044***	-0.0045***	-0.0044***	-0.0045***	-0.0044***
	(0.0003)	(0.0003)	(0.0003)	(0.0003)	(0.0003)	(0.0003)
$r \times$ Condition = Strong2R10	-0.0046***	-0.0046***	-0.0046***	-0.0046***	-0.0046***	-0.0046***
	(0.0003)	(0.0003)	(0.0003)	(0.0003)	(0.0003)	(0.0003)
$r \times$ Condition = Strong2R5	-0.0044***	-0.0044***	-0.0044***	-0.0044***	-0.0044***	-0.0044***
	(0.0003)	(0.0003)	(0.0003)	(0.0003)	(0.0003)	(0.0003)
a_j (demeaned)		0.0137***		0.0137***		0.0138***
		(0.0003)		(0.0003)		(0.0003)
Fixed-effects						
Condition (6)	Yes	Yes	Yes	Yes		
Search Task (15)			Yes	Yes	Yes	Yes
Participant (3,533)					Yes	Yes
Fit statistics						
Observations	1,059,900	1,059,900	1,059,900	1,059,900	1,059,900	1,059,900
R^2	0.00742	0.04259	0.00762	0.04278	0.03544	0.07080
Within R^2	0.00728	0.04245	0.00728	0.04245	0.00749	0.04388

Clustered (Participant) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes. Column 1 reports the OLS estimates of the following regression: $\mathbb{1}\{ProductSearched_{itj}\} = \alpha_{cond} + \beta \cdot r_{itj} + \delta(r_{itj} \cdot Cond_i) + \epsilon_{itj}$, where the dependent variable is an indicator for whether participant i searches product j in task t. The independent variables are fixed effects for i's experimental condition, the rank of product j, and the interaction of rank and the participant's experimental condition. In columns 2, 4, and 6, we also include product j's demeaned bonus A. We add search task fixed effects in columns 2–4 and participant fixed effects in columns 5–6.

returns to search. Participants report no significant differences in beliefs about rank-specific payoffs across conditions prior to the search tasks. In particular, none of the conditions differs from Random-Informed in the fraction of participants reporting that their beliefs about bonus B are identical across positions. Although we elicit beliefs in an incentivized manner after the search tasks, participants still do not report post-search beliefs that are consistent with the bonus B values they observe or the position effects they exhibit in their search behavior.[31] Figure I.4 in the appendix shows that all conditions are statistically indistinguishable in terms of the fraction of participants who report the same beliefs across positions. This conflicts with their observed search behavior, since we document very strong position effects in conditions that include the *Strong* algorithm. Because their reported posterior beliefs are not consistent with their behavior, we do not further use participants' stated post-search beliefs in model estimation.

4.2.1 Evidence of Learning

In this section, we present evidence that individuals learn the relationship between the product's position and its bonus B. We find evidence that participants use realized bonus Bs in position k to update their beliefs about the payoffs of position k in future tasks. However, we find no evidence of spillovers in learning across positions. That is, individuals do not seem to use the bonus B of the product in position k to update their beliefs about the bonus B of the products in positions $k+1$ or $k-1$. These patterns inform how we formally model learning in Section 4.3.

Payoffs increase across search tasks. The first piece of evidence of learning is that participants earn more as they gain more experience with the search tasks.[32] To demonstrate this, we estimate the following linear regression:

$$Earnings_{it} = \alpha_i + \alpha_{cond} + \beta_1 SearchTask_{it} + \beta_2 (SearchTask_{it} \times Condition_i) + \epsilon_{it}, \quad (12)$$

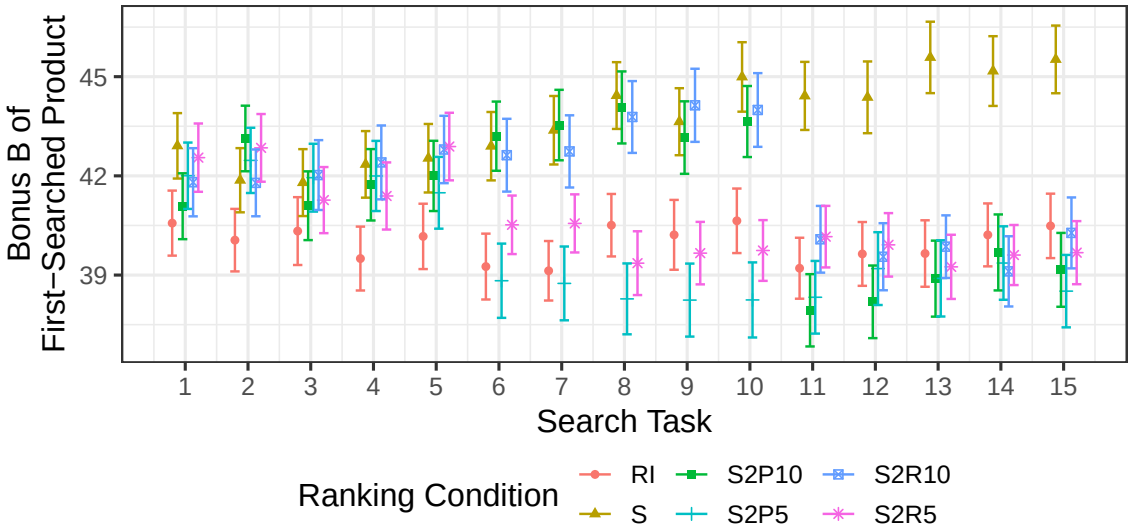
[31] There are several reasons why participants' observed behavior do not reflect their stated beliefs, including the incentive was not large enough to induce participants to put in effort to answer correct, or the question and response format were confusing.

[32] Recall that the earning for a task is the bonus A + bonus B of the selected product minus the number of products searched.

where α_i is the individual fixed effect, α_{cond} is the treatment condition fixed effect, and $SearchTask_{it}$ is the task number, which ranges from 1 to 20. We estimate this only for observations in conditions where the algorithm does not change or before the algorithm changes in conditions where it is supposed to.[33]

Column 1 of Table 7 reports the estimates of Equation 12. Participants in the Random-Informed condition earn 2.8 more points in the last task compared to the first, while those in the Strong condition earn 4.4 more points. This increase in earnings is primarily due to increases in the bonus B of the selected product, as shown by the similar coefficient magnitudes in column 2.[34] This suggests that, with more experience, participants become better at finding products with higher bonus B.

Figure 9: Bonus B of First Search by Task and Condition (Study 2)



Notes: Error bars represent 95% confidence intervals which are based on standard errors clustered at the participant level. "RI" refers to Random-Informed, "S" refers to Strong, "S2P10" and "S2P5" refer to Strong-to-Pref10 and Strong-to-Pref5, respectively, and "S2R10" and "S2R5" refer to Strong-to-Random10 and Strong-to-Random5, respectively.

[33]For example, the algorithm switches from *Strong* to *Random* after ten tasks in the condition Strong-to-Random10; we estimate Equation 12 using only tasks 1–10 from participants in the Strong-to-Random10 condition.

[34]Improvement in earnings in Random-Informed could be due to better stopping or continuation decisions. It is *not* due to selecting better bonus Bs because bonus Bs are randomly ordered. We confirm this in Figure 9, which plots the selected bonus B by condition and search task. The bonus B of the first-searched product per task does not increase with experience in Random-Informed.

Table 7: Change in Earned Bonuses across Tasks (Study 2)

Dependent Variables:	Earnings	Selected Bonus B
Model:	(1)	(2)
Variables		
Search Task	0.1448***	0.1218***
	(0.0271)	(0.0259)
Search Task × Condition = Strong	0.0890**	0.1121***
	(0.0378)	(0.0371)
Search Task × Condition = Strong2Pref10	0.1848***	0.2190***
	(0.0546)	(0.0528)
Search Task × Condition = Strong2Pref5	0.3101**	0.3570***
	(0.1382)	(0.1352)
Search Task × Condition = Strong2R10	0.1637***	0.1581***
	(0.0561)	(0.0542)
Search Task × Condition = Strong2R5	0.0523	0.0562
	(0.1374)	(0.1358)
Fixed-effects		
Condition (6)	Yes	Yes
Fit statistics		
Observations	35,160	35,160
R^2	0.01285	0.01573
Within R^2	0.00497	0.00497

Clustered (Participant) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes. This table reports the OLS estimates of Equation 12.

Learning across positions. How do participants learn to find better bonus Bs? Do participants learn that certain ranks yield higher bonus Bs, and does this learning spill over (Hodgson and Lewis, 2023) to adjacent positions? For example, if a participant clicks on the second-ranked product and uncovers a large bonus B, are they more likely to search the second-ranked product in the next search task? If learning spills over to nearby ranks, we might expect the participant to also be more likely to search the products in ranks one or three in the next task. Understanding whether spillovers exist allows us to distinguish learning about specific actions (searching specific ranks) from learning about the environment (searching products in adjacent ranks with predicted good payoffs) to model learning. To test for rank-specific learning, we estimate the following linear regression:

$$\begin{aligned}\mathbb{1}\{\text{Searched Product } r\}_{i,t+1} = \alpha_t + \alpha_r + \beta_1 \mathbb{1}\{\text{Searched Product } r\}_{i,t} + \\ \beta_2 \mathbb{1}\{\text{Searched Product } r\}_{i,t} \times b_{i,t,r} + \beta_3 b_{i,t,r} + \\ \beta_4 a_{i,t+1,r} + \beta_5 NSearched_{i,t+1} + \epsilon_{i,t+1} \quad (13)\end{aligned}$$

The dependent variable $\mathbb{1}\{\text{Searched Product } r\}_{i,t+1}$ is an indicator for whether participant i searches the product in rank r in the next task, $t+1$; $\mathbb{1}\{\text{Searched Product } r\}_{i,t}$ is an indicator for whether i searches the product in rank r in the current task, t. The term b_{itr} is the bonus B of the product in rank r in task t. We control for $a_{i,t+1,r}$, the bonus A of the r-ranked product in task $t+1$. This accounts for the possibility that individual i clicks on r in the next task simply because it happens to have a high bonus A. We also control for the number of products searched in $t+1$ ($NSearched_{i,t+1}$). This accounts for the possibility that i searches r simply because this individual searches many products. However, one may be concerned that $NSearched_{i,t+1}$ is influenced by actions at t. Therefore, we estimate another specification of the regression that includes only the first search in $t+1$, which allows us to drop $NSearched_{i,t+1}$. To measure whether learning spills over to nearby ranks, we also include alternative dependent variables: whether the individual searches the product $r+1$ and $r-1$.

We estimate this regression for participants in the Strong condition and for tasks where

the individual searches only one product at task t.[35] We impose the former restriction for two reasons. First, we do not expect to find position-specific learning in Random-Informed because participants are informed that all products are ordered randomly. Second, other conditions (e.g., Strong-to-Random5) have varying number of tasks with the strong algorithm, and the degree of learning can vary across tasks; focusing on the Strong condition allows us to keep the number of tasks constant. We impose the latter restriction of single-search tasks because it allows us to attribute the effect to the bonus B of product r, as opposed to the bonus Bs of other positions the participant searched.

Table 8 displays the OLS estimates of Equation 13. Column 1 reports the estimates for whether the individual searches the identically positioned product in $t+1$ as in t. The coefficients of $\mathbb{1}\{\text{Searched Product } r\}_{i,t}$ and its interaction with b_{itr} imply that conditional on searching a product, the individual is more likely to search the product in the same rank in the subsequent task if the searched product has a higher bonus B in the current task. Columns 2 and 3 report whether participants who discover that product r has a higher bonus B are more likely to search product $r+1$ or $r-1$ in the next task, respectively. The coefficients of the interactions in columns 2 and 3 are closer to zero and not statistically significant. Columns 4–6 repeat these regressions with an indicator for searching that item first as the dependent variable and drops the independent variable $NumSearchedt + 1$. Similarly, we find that the learning about one position does not spill over onto the adjacent products across tasks (columns 5 and 6, very small and n.s. coefficient on the interaction). Therefore, we find evidence of position-specific learning but no spillovers across positions. These results are consistent with individuals forming beliefs about the bonus B of product r based on past realizations of bonus B for product r, and such realizations do not impact beliefs about nearby products. These results motivate how we model learning.

4.3 Structural Model of Learning

The model in Study 2 differs from that in Study 1 in that it incorporates how individuals learn about how the expected bonus B varies by the product's position. This approach enables us to simulate behavior and calculate consumer surplus under counterfactual ranking

[35] Of all tasks in the Strong condition, 62% involve only one search.

Table 8: Learning Across Positions (Study 2)

Dependent Variables:	Search r	Search $r+1$	Search $r-1$	Search First r	Search First $r+1$	Search First $r-1$
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
1(r Searched in $t-1$)	-0.2105***	-0.0041	0.0083	-0.2027***	0.0218	0.0167
	(0.0260)	(0.0182)	(0.0171)	(0.0263)	(0.0170)	(0.0155)
1(r Searched in $t-1$) × Bonus B	0.0059***	0.0006	5.3×10^{-5}	0.0057***	-2.39×10^{-6}	-0.0003
	(0.0007)	(0.0004)	(0.0004)	(0.0007)	(0.0004)	(0.0004)
Bonus B	-0.0003**	9.89×10^{-5}	-8.01×10^{-5}	-0.0002*	0.0001	-4.82×10^{-5}
	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Next Bonus A	0.0099***	0.0098***	0.0100***	0.0074***	0.0072***	0.0075***
	(0.0005)	(0.0005)	(0.0005)	(0.0004)	(0.0004)	(0.0004)
Num Searched t	0.0495***	0.0496***	0.0490***			
	(0.0001)	(0.0004)	(0.0003)			
Fixed-effects						
Search Task (14)	Yes	Yes	Yes	Yes	Yes	Yes
r	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
# r	20	19	19	20	19	19
Observations	106,800	101,460	101,460	106,800	101,460	101,460
R^2	0.08758	0.07539	0.07983	0.03872	0.02299	0.03129
Within R^2	0.07135	0.06913	0.06277	0.02691	0.02024	0.01896

Clustered (Participant) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

algorithms.

An individual i's reservation utility for a product j in task t is

$$z_{itj} = a_{itj} + E[b_{itj}|r_{itj}] + \overbrace{\beta_i^0 + \beta_i^1 \mathbb{1}\{r_{itj} = 1\} + \beta_i^r\Big(\mathbb{1}\{r_{itj} \neq 1\} \cdot (r_{itj} - 2)\Big)}^{\delta_{itj}} + \epsilon_{itj}, \quad (14)$$

where δ_{itj} is the search propensity, and r_{itj} is the product's rank. This reservation utility is similar to Equation 11 with one exception: the expected bonus B varies by the product's rank.

We assume that i updates beliefs about the expected value of $b_{jt}|r_{tj}$ in a Bayesian fashion. Before engaging in the first task, i holds a prior belief about the expected bonus B for each product j. This prior varies with the product's rank. With a slight abuse of notation, we denote $b_{itj}|r_{itj}$ with b_{itr}. We assume that the prior is drawn from a normal distribution with mean b_{ir}^0 and variance σ_0^2. Note that σ_0^2 is neither product- nor task-specific and is assumed to be homogeneous across participants. We also assume that $b_{ir}^0 = 40$ because (1) participants were informed that mean of bonus B (across ranks) is 40, and (2) this is the modal self-reported prior belief in all conditions.

Once i searches product r in task $t = 1$, he observes b_{itr} and then updates the expected value of bonus B in rank r to be

$$\frac{\sigma_j^2 b_{ir}^0 + \sigma_0^2 b_{itr}}{\sigma_0^2 + \sigma_j^2}, \quad (15)$$

where $\sigma_j^2 = 144$ is the variance of bonus B. Recall that we inform participants of this variance in the study's instructions. Note that σ_j^2 does not vary by rank. This assumption warrants discussion because although this is true for the Random-Informed condition and participants are directly informed about σ_j^2, this assumption does not hold empirically for the other conditions. For example, in the Strong-to-Pref5 condition, the variance of bonus B in rank 1 is higher than that of rank 20 because the bonus B of rank 1 changes from the highest to the lowest after five search tasks. However, because participants are not informed of this change, we assume that σ_j^2 is constant across ranks to simplify estimation. Additionally, as reported in Section 3.4, not allowing individuals to learn about bonus b's variance does not significantly impact our estimates.

Given Equation 15, we can express participant i's expected value of bonus B for a product in rank r at task t as

$$\frac{1}{\sigma_0^{-2} + n_{itr}\sigma_j^{-2}} \left(\frac{b_{ij}^0}{\sigma_0^2} + \frac{n_{itr}\bar{b}_{itr}}{\sigma_j^2} \right), \tag{16}$$

where n_{itr} is the number of times that i has sampled the product in rank r between tasks 1 to $t-1$, and $\bar{b}_{itr}$ is the mean of the sampled bonus Bs in rank r in tasks 1 to $t-1$. We can further simplify Equation 16 to be expressed as a weight on the prior and the weight on the samples:

$$E[b_{itr}] = b_{ir}^0 \frac{m_{itr}}{1 + m_{itr}} + \bar{b}_{itr} \frac{1}{1 + m_{itr}}, \tag{17}$$

where $m_{itr} \equiv \frac{\sigma_j^2}{n_{itr}\sigma_0^2}$.

This specification of beliefs implies that we attribute differences in learning rates to differences in the variance of prior beliefs. In other words, if participants behave as if they completely update ($E[b_{itr}] = \bar{b}_{itr}$), then our model assumes that σ_0^2 is very large. Furthermore, since we assume that σ_0^2 is homogeneous, differences in beliefs across participants arise only due to different sampling of bonus Bs ($\bar{b}_{itr}$).

Returning to Equation 14, we assume that participants in the Random-Informed condition do not update their beliefs about bonus B. Thus, the reservation utility, when including all conditions, is

$$\begin{aligned} z_{itj} = a_{itj} &+ (b_{ij}^0 \frac{m_{itj}}{1 + m_{itj}} + \bar{b}_{itj} \frac{1}{1 + m_{itj}}) \cdot \mathbb{1}\{Cond_i \neq RI\} \\ &+ 40 \cdot \mathbb{1}\{Cond_i = RI\} + \overbrace{\beta_i^0 + \beta_i^1 \mathbb{1}\{r_{itj} = 1\} + \beta_i^r \Big(\mathbb{1}\{r_{itj} \neq 1\} \cdot (r_{itj} - 2) \Big) + \epsilon_{itj}}^{\delta_{itj}}. \end{aligned} \tag{18}$$

We further decompose β_i^r to the following:

$$\beta_i^r = \Delta_0^r + \Delta_{comp}^r Comp_i + \Delta_{shop}^r ShopFreq_i + \epsilon_i^r. \tag{19}$$

The expressions for β_i^0 and β_i^1 are similar to Equation 19. Note that there is a key difference between Equations 19 and 9 (Study 1). In Study 1, we included terms in the

expression for β_i^1 and β_i^r (Δ_S^1 and Δ_S^r respectively) to capture differences in search propensities by rank between Strong and Random-Informed. These differences were due to the beliefs that formed in the Strong condition. Since we explicitly model the evolution process of beliefs in Study 2, these beliefs are now captured by the $(b_{ij}^0 \frac{m_{itj}}{1+m_{itj}} + \bar{b}_{itj} \frac{1}{1+m_{itj}})$ term.

4.3.1 Estimation

The additional parameter to estimate in this model is the prior variance σ_0^2. To estimate our search model in the previous study, we used a Gibbs sampler. In particular, we were able to estimate search propensities using a Bayesian regression because the parameters β_i^0, β_i^1, and β_i^r enter the search propensity linearly (Equation 11). However, σ_0^2 does not enter the reservation value linearly in Equation 18. Rather, our specification results in a posterior density for the parameters that we can no longer directly sample from. Therefore, we modify our estimation method to include a Metropolis-Hastings step within the Gibbs sampler.[36] We verify we can recover the parameters on simulated data, which we describe in Appendix B.2.2.

4.3.2 Estimation Results

The results of this estimation are presented in Table 9. The estimated prior variance is 11, which reflects the initial uncertainty participants have about position effects (recall, the prior mean is $b = 40$ all ranks). Other parameters are similar to our Study 1 estimates: search costs are lower for the first position, position effects due to position-specific costs are smaller for positions 2 through 20, and participants who passed comprehension checks on the first attempt search more. The implied average search cost is \$0.06 for the first item and \$0.08 for the twentieth item.

We report how well the model fits the data using simulated search patterns in Appendix G. The model matches three key data patterns: (i) baseline position effects in Random-Informed, (ii) increasing position effects in Strong over time, and (iii) increased search in

[36] Specifically, we use a random walk proposal distribution: at the kth iteration, we draw ${\sigma_0^2}^{prop} = TN({\sigma_0^2}^{(k-1)}, h0)$. We then accept the proposal (${\sigma_0^2}^{(k)} = {\sigma_0^2}^{prop}$) with a probability proportional to the ratio of the posterior probabilities of the proposed and current values of σ_0^2. Otherwise, we reject the proposal (${\sigma_0^2}^{(k)} = {\sigma_0^2}^{(k-1)}$). The other steps of the Gibbs sampler remain the same as in Study 1.

Table 9: Estimation Results (Study 2)

Variable	Coefficient	Mean (SD)
Learning Parameter		
Prior Variance	σ_0^2	11 (0.47)
Rank Coefficient β^r		
Intercept	Δ_0^r	-0.27 (0.019)***
Comp FE	Δ_{comp}^r	0.025 (0.021)
ShopFreq FE	Δ_{shop}^r	-0.0098 (0.02)
Pos 1 Coefficient β^1		
Intercept	Δ_0^1	3.2 (0.37)***
Comp FE	Δ_{comp}^1	0.64 (0.39)
ShopFreq FE	Δ_{shop}^1	0.31 (0.37)
Baseline Search Propensity β^0		
Intercept	Δ_0^0	-28 (0.63)***
Comp FE	Δ_{comp}^0	9.4 (0.71)***
ShopFreq FE	Δ_{shop}^0	-1.5 (0.65)**
Heterogeneity		
Rank Coefficient	$\sqrt{V}_{\beta^r}$	0.51 (0.0092)
Pos1 Coefficient	$\sqrt{V}_{\beta^1}$	8.9 (3.4)
Mean Search Propensity	$\sqrt{V}_{\beta^0}$	18 (9.1)
Reservation Utility	σ_ϵ	16 (2.1)

Notes: We report the statistical significance only for the search propensities and their shifters. Signif. Codes: $***: 0.01, **: 0.05, *: 0.1$. Posterior means and standard deviations are based on the thinned chain which drops the first 1,000 draws of the chain and keeps every tenth draw thereafter.

position 3 in Strong-to-Pref conditions after the algorithm switches to *Pref*.

5 Impacts of Ranking Algorithms on Consumer Surplus

In this section, we use the model estimates from Study 2 to demonstrate that failing to account for beliefs leads to incorrect predictions of search behavior and biases welfare evaluation, even when rankings are randomized. We also propose some solutions to estimate search costs when it is not feasible to implement a Random-Informed experimental condition.

First, we compare behavior and welfare predictions with and without accounting for consumer beliefs. Because of growing concerns around the potential harms of platform steering, we focus on welfare implications of a scenario where the platform switches from a *Strong* ranking algorithm to the *Pref* algorithm, as in Study 2. Recall that the *Pref* algorithm places the two items with the least valuable Bonus Bs in the top two positions on the page, and the remaining products are arranged in noisily decreasing order of the hidden attribute.

We generate welfare predictions that account for consumer beliefs by estimating auxiliary models. We estimate the search model without beliefs and without learning on the Random-Informed data in Study 2, in which we attribute all position effects to search costs.[37] This provides us with the correct position-specific search costs. We then use these estimates to simulate search behavior for consumers over thirty tasks. In the first five tasks, the products are ordered by the *Strong* algorithm, and by the *Pref* algorithm in the remaining tasks. We allow for Bayesian updating of position-specific beliefs about bonus B.[38] Simulating search behavior provides us with each consumer's surplus per task, which is the total bonus earned (bonus A + B) minus the incurred search cost.

We compare this average surplus per task to that generated by an auxiliary model that does not account for consumer beliefs. The goal of this auxiliary model is to mimic inference drawn from typical empirical settings that randomize product ordering, such as the

[37] We use tasks 6 and 7 to ensure there is an apples-to-apples comparison with the data used to estimate the auxiliary model, which we describe next.

[38] We use the estimated prior variance in Table B.5.

commonly used Expedia field experiment data, where hotel listings were randomly ordered but users did not receive disclosure of the randomization (e.g., the data used in Ursu (2018), Compiani et al. (2023), and others) and on other platforms (Derakhshan et al., 2022). In such settings, the data generated from the random rankings are used to estimate search costs. To parallel this, we estimate the search model without position-specific beliefs using data generated from the Strong-to-Random5 condition immediately following the switch to (undisclosed) *Random* rankings.[39] In other words, we assume that all position effects in the *Random* portion of the Strong-to-Random5 condition are due to search costs.[40] We then simulate search behavior and calculate consumer surplus given these estimates. Specifically, we use the Strong-to-Random5 data from tasks 6 and 7 to estimate this auxiliary model. We describe the method for generating these surplus predictions in detail in Section H.1 in the appendix.

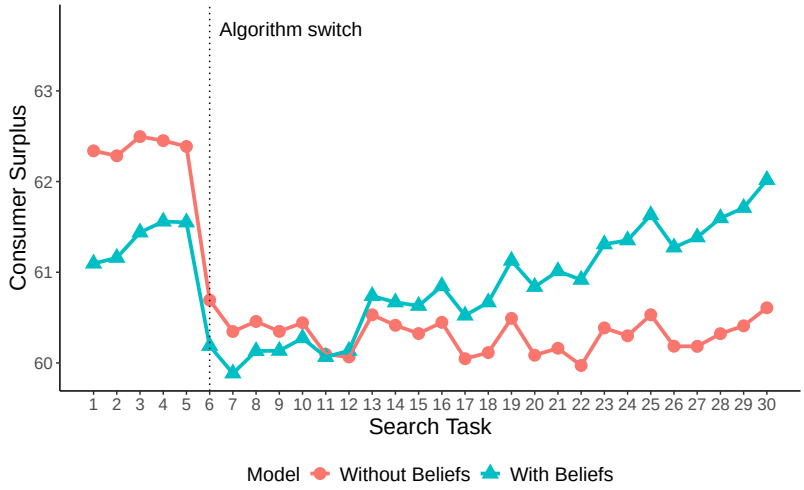
Figure 10 shows the average surplus per consumer for each task. It demonstrates that not accounting for beliefs (i.e., attributing all position effects to search costs in the presence of position-specific beliefs) overpredicts consumer surplus under the *Strong* algorithm. This occurs because under the *Strong* algorithm, consumers benefit from stronger position effects because products in top positions have higher hidden payoffs. The model without beliefs generates stronger position effects due to its (incorrect) steeper search costs with respect to position. Meanwhile, the model with beliefs has less steep (but correct) search costs, so consumers exhibit weaker position effects. Indeed, the difference in search costs between products 1 and 20, estimated from the model without beliefs is 1.8 times as large as that estimated from the model that accounts for beliefs.

Importantly, Figure 10 also shows that after the algorithm switches to the *Pref* ranking algorithm, the model without beliefs underpredicts surplus. That is, failing to account for consumer learning predicts greater consumer harm from this particular form of self-preferencing. Intuitively, the underprediction of surplus occurs because in the model without beliefs, (1) search costs with respect to position are overestimated, so that consumers

[39] We choose Strong-to-Random5 instead of Strong-to-Random10 to demonstrate that consumers form rank-specific beliefs in just a short exposure of five search tasks. Using Strong-to-Random10 would lead to stronger position effects due to beliefs, and thus, a larger error in estimated search costs and a larger misprediction of welfare.

[40] We report results for both auxiliary models in Table H.1 in the appendix.

Figure 10: Consumer Surplus with and without Accounting for Beliefs



Notes: Figure plots the simulated average consumer surplus per consumer for each search task. We calculate the average surplus (bonus chosen − incurred search costs) per consumer per task, given a draw of consumer parameters. We repeat this 200 times. The error bars represent the 95% confidence intervals, which are taken over these iterations. The ranking algorithm switches from *Strong* to *Pref* after task 5. "Without Beliefs" refers to the model estimated on the Strong-to-Random5 data after the switch to *Random* rankings. "With Beliefs" refers to the model estimated using the learning parameter in the main model and search costs recovered from Random-Informed.

are predicted to exhibit stronger position effects, and (2) they do not learn (i.e., form equilibrium beliefs) about the hidden attribute over time. Therefore, in the model with beliefs, consumer surplus increases over time due to learning, which leads to the consumer picking a product that has a higher total bonus while also reducing search costs.[41] In total, our model suggests that when accounting for beliefs, changing the algorithm from *Strong* to *Pref* reduces consumer surplus by 0.84% (95% CI: 0.2% - 1.48%) when comparing the average surplus per task for tasks 1–5 to 6–30.[42] However, the model without beliefs suggests that switching to the *Pref* algorithm reduces consumer surplus by 3.3% (95% CI: 2.52% - 4.17%), overstating the loss in surplus by four times relative to the model with beliefs. Furthermore, the model without beliefs underpredicts the level of consumer surplus by 2.3% at task 30.[43]

[41] This can be seen in Figure H.1 in the appendix, which plots the average of the selected product's bonus and total incurred search costs over search tasks.

[42] Note that this is an underestimate of the decrease in surplus because consumers have not completely learned the *Strong* algorithm by task 5, as evidenced by the increasing surplus per task for these tasks under the model with beliefs.

[43] This value is obtained by comparing the surplus at task 30 between the two models.

Discussion As the previous counterfactual shows, not accounting for beliefs can lead to incorrect search and surplus predictions. How can this issue be addressed in the field? One solution is to run the experiment for a long enough period of time for beliefs to adjust. However, this solution may be impractical. If the algorithm is subtle or differences in hidden attributes are difficult for consumers to discern, the experiment might need to run for an extended period of time to allow for consumers to fully update. For example, consumers might learn more quickly if the hidden attribute is simple, such as price, but less quickly if the hidden attribute is more difficult to evaluate, such as the quality of experience goods. Another challenge with this approach is that consumers learn at different rates. Those with high search costs take longer to learn the algorithm because they conduct fewer searches. In settings where consumers have heterogeneous search costs or high search costs, running the experiment long enough for beliefs to converge may not be feasible.

Another potential solution is to assume complete information (i.e., correct beliefs) in combination with a randomized experiment (Kaye, 2024). That is, one can estimate the search model using data from an experiment in which listings are randomly ordered and consumers are not informed, but assume that consumers know the correct rank-specific returns to search (i.e., the average hidden attribute value at each position) under the ranking algorithm before the switch. We test whether this is a viable solution using our data from Study 2. We estimate the model using data from the first two tasks after a switch from *Strong* to *Random* and assume that consumers have correct beliefs (as per the *Strong* algorithm) about the bonus B for a product in rank r. We can then compare the surplus predictions using these estimates to predictions from model estimates where search costs are separately measured (Study 2 estimates), in addition to the observed surplus from the experiment.

If we impose correct beliefs about the "old" firm policy, the estimated position effect of search costs has the *opposite* sign. That is, this assumption results in estimated search costs *decreasing* with position. This is not surprising: the assumption of correct beliefs implies complete knowledge about the algorithm, whereas we saw that learning is slow in Study 2. Position effects are therefore not as stark as they would be under this assumption (i.e., participants should be clicking on the top-ranked products much more if they were fully

aware of the underlying algorithm). The only way the model can rationalize this is by assuming that search costs decrease with position. As a result, when we simulate search under the *Pref* algorithm, these estimates predict that consumers are *better off* (+2.23% versus −0.84% under our model's estimates). This effect is driven by lower incurred search costs (−10.33% versus +4.25% under our model's estimates). We note that the "wrong" direction of the position effects resulting from search costs might be specific to our setting because the extent to which this assumption produces incorrect costs depends on the speed of learning. The faster that participants learn or the higher proportion of participants who had correct pre-randomization beliefs, the better the assumption of correct beliefs can recover search costs. Therefore, in field settings, this solution may be best suited for contexts where the correlation between hidden attributes and position is easily learned or when the sample includes consumers who have extensive experience with the search environment. We also note that the direction of the bias may vary depending on how recommendation algorithms prioritize hidden attributes in other settings. In our case, not accounting for beliefs biases search costs upward because consumers typically believe that preferred products are ranked more prominently. However, in settings where preferred products are typically ranked less prominently, failing to account for beliefs would bias search costs estimates downward. Alternatively, firms can use survey responses from consumers about their beliefs as model inputs.[44]

6 Comparisons to Real-World Search Settings

Our experimental design decisions were made to simplify the experiment and model estimation. However, the simplification creates some divergence from real-world product search settings in several key ways. Below, we outline these differences and discuss their implications.

Correlated attributes. First, we impose that a_j and b_j are uncorrelated. As previously mentioned, this choice simplifies identification and estimation but deviates from real-world

[44]This solution may not work in every setting. In our studies, consumers' self-reported beliefs were not always consistent with their actions. Recall that they reported no differences in expected bonus Bs by position post-search but exhibited stronger position effects than search costs can account for.

settings where pre- and post-search attributes are often correlated. As discussed in Section 2, the uncorrelated attributes setting can be easily adjusted to one where the pre- and post-search characteristics are correlated. In Appendix A, we present a supplemental experiment with an alternative data generating process that mimics real-world correlations. We find that participants engage in similar search behaviors, in terms of position effects and search intensity, as in the case with uncorrelated attributes.

Ranking algorithm. Second, our Strong algorithm only sorts based on b_j and not a_j. Since a_j is observed by participants, we do not believe that learning about whether (and to what extent) the ranking algorithm prioritizes a_j is of material consequence. On the other hand, sorting only by b_j allows for cleaner identification – any incremental position effects in Strong are *only* due to beliefs.

Horizontal preferences. Third, by using bonuses instead of product attributes, participants in our setting have vertical preferences over product attributes. This simplifies model estimation because we do not need to estimate participants' preferences. However, in real search settings, consumers may have horizontal preferences. This may produce several differences from our paradigm. It makes identification more difficult for the econometrician, since search costs and preference heterogeneity may be confounded without the panel style data we collect (Morozov et al., 2021). Also, horizontal preferences may also yield differentiated products, which are more costly to evaluate. As a result, consumers may learn more slowly because the signals about the ranking algorithm are more obscured. Additionally, with horizontal preferences, the self-preferencing algorithm in Study 2 and our counterfactuals may not harm all consumers as much as if preferences were purely vertical. Some consumers may prefer the products that are placed more prominently (i.e., those with lower bonus Bs) under a self-preferencing algorithm. While this might impact the magnitude of our consumer welfare estimates in the counterfactuals, the presence of horizontal preferences does not impact the main conclusion that consumers form beliefs about the correlation between prominence and product value and account for these beliefs in search. In terms of the model, if consumers' preferences are recovered, we can just as easily separate search costs

and beliefs by using individual-specific $a_{ij} = \beta_i^a a_j$ and $b_{ij} = \beta_i^b b_j$ in the model.[45]

Explicit marginal search costs. Fourth, we impose an explicit marginal search cost: each product costs one point to search. This differs from real search settings where consumers do not pay monetary costs to search each product. We impose the marginal search cost because participants in an online experiment setting may face different outside options than those who are conducting search in field settings. A source of this difference is that study participants are paid for their time, while consumers in the field are not. The differences in outside options impact their marginal search costs, as these search costs are relative to the normalized outside option (but should not create differences in position-specific search costs). As a robustness check, we implemented a small pilot study that is similar to Study 2 but with zero explicit marginal search costs (i.e., products are free to search). We find that this study generates search behavior that is less similar to field settings and to optimal search. Participants in Random-Informed are less likely to search the product with the highest bonus A first (which is the optimal search strategy). Additionally, participants search more products relative to field settings, and are more likely to search all products relative to Study 2.[46] In summary, having an explicit marginal search cost in the online experiment setting is important because it generates search behavior that more closely mirrors real-world patterns and predictions from Weitzman search model. We describe the details of this pilot study with zero marginal search in Appendix E.

Greater returns to search. Fifth, in both studies, the fifth and ninety-fifth percentiles of bonus B are 20 and 60, which equates to a difference of approximately \$0.10 USD. A concern is that the incentives to search are not very high, leading to behavior that is not consistent with the Weitzman search model. In other words, the monetary incentive may be too small to motivate participants to search in a way that maximizes their payoffs. Also, in real search settings, price dispersion is likely much greater than in our study. To address this, we first note that each search task takes each participant on average 13 seconds to

[45] The variables β_i^a and β_i^b represent individual-specific preferences towards bonuses A and B.

[46] Participants search on average between 4.5 to 7.4 products across condition, which is more than the 1.78-1.99 products searched in Study 2 and in other field settings (e.g. Morozov, 2023; Ursu, 2018). Participants searched all products in 12% of search tasks, as opposed to 0.1% in Study 2.

complete. This translates into an hourly rate of $27.69. So, while the incentives are low, the tasks are also conducted quickly. Additionally, the estimated search costs are in line with these incentives. In Study 1, the estimated search cost for a product in the first position is $0.05. This translates to an hourly cost of $13.80, which is a reasonable rate, especially relative to our study's average base payment of $9.30/hour. Furthermore, to ensure that our conclusions are not driven by low study incentives, we ran a follow-up study that "raises the stakes." This study is similar to Study 2, with a few exceptions: (1) we increased the payoffs from search by using a conversion rate of 200 points per dollar (instead of 500), (2) participants conduct ten search tasks (instead of fifteen) and only include the Strong and Random-Informed conditions (for cost considerations), (3) decreased the point value of the marginal search cost to 0.5 (instead of 1), and (4) payoffs were summed over all tasks (instead of eight out of fifteen). These changes translate into an increase in expected returns to search by ∼4x because a point is worth more than twice as much as before, and we compensate the participant for more tasks. We find that while raising the stakes increases search intensity (as expected), it does not impact position effects relative to those in Study 2. Therefore, we conclude that our findings from our main studies are not impacted by lower incentives to search. We describe the "higher stakes" study in detail in Appendix F.

Ease of Attribute Evaluation. Finally, in our experiment, the total product payoff is presented for easy processing by participants, unlike many product market settings where quality must be assessed through further attention or through experience. We expect if attributes are more difficult to evaluate, the learning speed we observe is *faster* than in many real search settings.

7 Conclusion

In this paper, we show that beliefs play an important role in customer responses to recommendation systems, such as product rankings. We provide a framework for separating belief-driven and cost-driven position effects. Our experimental findings demonstrate that consumer beliefs contribute significantly to position effects. If beliefs are overlooked, this can lead to incorrect search cost estimates and welfare implications of alternative recom-

mendation systems. Although the magnitude of this decomposition between beliefs and costs is specific to our experimental design and setting, we show that consumers indeed learn about the ranking algorithm, and this learning leads to biased search cost estimates.

The bias in search cost estimates exists even when listings are presented in a randomized order. Our findings highlight the shortcomings of the common approach of extrapolating results from short-term field experiments without accounting for consumer beliefs. Specifically, we demonstrate that short-term experiments may mis-estimate consumer surplus and demand under different ranking policies if they fail to consider the gradual adaptation of consumer beliefs to changes in the recommendation system.

An important implication of our findings in the product positioning context is that search costs are typically overestimated because platforms typically order from best to worst. Higher search costs can lead to the conclusion that platforms are able to steer consumers' choices to a large extent simply through the product ranking algorithm. This implies that policies such as self-preferencing (or any policy that more prominently displays products that are worse for consumers) greatly reduces consumer welfare: consumers will not adjust their search in response to learning about the algorithm, and they incur (incorrectly) high search costs for conducting the same search and purchase.

We also present practical solutions for addressing the conflation of beliefs and search costs in both experimental and field settings. Although we document this problem in the context of product rankings, the issue of out-of-equilibrium beliefs induced by short-term experiments can also arise when measuring other model primitives, such as advertising effects (Goli et al., 2024) or price elasticity (Anderson and Simester, 2004) where consumer beliefs also play a pivotal role. In the case of pricing, short-term experiments typically overstate the magnitude of price elasticities.

Another implication for evaluating recommendation systems is that belief adjustment can be slow. Imposing correct preexperiment beliefs fails to consider that consumers learn slowly, and learning may be incomplete for consumers who rarely participate in the market or have particularly high search costs. A solution is to elicit beliefs as part of data collection where possible: direct measurement provides a short-term, data-driven way to better isolate policy-invariant primitives like search costs (Manski, 2004). Our solution in this paper is

to model the evolution of beliefs directly, which not only accounts for this slow adjustment but also allows for simulating long-run counterfactual scenarios.

We emphasize that accounting for beliefs is especially important when the platform's ranking algorithm is not transparent to consumers (e.g., Amazon's "Featured" algorithm or Expedia's "Recommended" algorithm). However, there are search environments where position may be less informative of the product's hidden attribute, such as when consumers are informed about the algorithm. Examples include price-ascending algorithms for undifferentiated products or algorithms that sort based on recency. In such cases, accounting for beliefs may be less important.

References

Agarwal, Ashish, Kartik Hosanagar, and Michael D Smith, "Location, location, location: An analysis of profitability of position in online advertising markets," *Journal of Marketing Research*, 2011, *48* (6), 1057–1073.

Agte, Patrick, Claudia Allende, Adam Kapor, Christopher Neilson, and Fernando Ochoa, "Search and Biased Beliefs in Education Markets," 2024. Working Paper.

Aguirregabiria, Victor, "Identification of firms' beliefs in structural models of market competition," *Canadian Journal of Economics/Revue canadienne d'économique*, 2021, *54* (1), 5–33.

_ **and Pedro Mira**, "Sequential estimation of dynamic discrete games," *Econometrica*, 2007, *75* (1), 1–53.

Anderson, Eric T and Duncan I Simester, "Long-run effects of promotion depth on new versus established customers: three field studies," *Marketing Science*, 2004, *23* (1), 4–20.

Aridor, Guy, Duarte Goncalves, Daniel Kluver, Ruoyan Kong, and Joseph Konstan, "The Informational Role of Online Recommendations: Evidence from a Field Experiment," 2024.

Armstrong, Mark, John Vickers, and Jidong Zhou, "Prominence and consumer search," *The RAND Journal of Economics*, 2009, *40* (2), 209–233.

Athey, Susan and Glenn Ellison, "Position auctions with consumer search," *The Quarterly Journal of Economics*, 2011, *126* (3), 1213–1270.

Bresnahan, Timothy F and Peter C Reiss, "Entry and competition in concentrated markets," *Journal of political economy*, 1991, *99* (5), 977–1009.

Casner, Ben, "Learning while shopping: an experimental investigation into the effect of learning on consumer search," *Experimental Economics*, 2021, *24* (1), 238–273.

Chen, Yuxin and Song Yao, "Sequential search with refinement: Model and application with click-stream data," *Management Science*, 2017, *63* (12), 4345–4365.

Ching, Andrew T, Tülin Erdem, and Michael P Keane, "Learning models: An assessment of progress, challenges, and new developments," *Marketing Science*, 2013, *32* (6), 913–938.

Choi, Hana and Carl F Mela, "Monetizing online marketplaces," *Marketing Science*, 2019, *38* (6), 948–972.

Compiani, Giovanni, Gregory Lewis, Sida Peng, and Peichun Wang, "Online Search and Optimal Product Rankings: An Empirical Framework," *Marketing Science*, 2023.

Dang, Chu Ivy, Raluca Ursu, and Pradeep K Chintagunta, "Repeated Product Searches and Choice Elimination: Evidence from a Lab Study," 2022. Working Paper.

De los Santos, Babur, Ali Hortaçsu, and Matthijs R Wildenbeest, "Testing models of consumer search using data on web browsing and purchasing behavior," *American economic review*, 2012, *102* (6), 2955–2980.

De Los Santos, Babur, Ali Hortaçsu, and Matthijs R Wildenbeest, "Search with learning for differentiated products: Evidence from e-commerce," *Journal of Business & Economic Statistics*, 2017, *35* (4), 626–641.

De los Santos, Babur and Sergei Koulayev, "Optimizing click-through in online rankings with endogenous search refinement," *Marketing Science*, 2017, *36* (4), 542–564.

DellaVigna, Stefano, "Structural behavioral economics," in "Handbook of Behavioral Economics: Applications and Foundations 1," Vol. 1, Elsevier, 2018, pp. 613–723.

Derakhshan, Mahsa, Negin Golrezaei, Vahideh Manshadi, and Vahab Mirrokni, "Product ranking on online platforms," *Management Science*, 2022, *68* (6), 4024–4041.

Dinerstein, Michael, Liran Einav, Jonathan Levin, and Neel Sundaresan, "Consumer price search and platform design in internet commerce," *American Economic Review*, 2018, *108* (7), 1820–1859.

Donnelly, Robert, Ayush Kanodia, and Ilya Morozov, "Welfare effects of personalized rankings," *Marketing Science*, 2024, *43* (1), 92–113.

Doraszelski, Ulrich, Gregory Lewis, and Ariel Pakes, "Just starting out: Learning and equilibrium in a new market," *American Economic Review*, 2018, *108* (3), 565–615.

Gabaix, Xavier, David Laibson, Guillermo Moloche, and Stephen Weinberg, "Costly information acquisition: Experimental analysis of a boundedly rational model," *American Economic Review*, 2006, *96* (4), 1043–1068.

Gardete, Pedro M and Megan Hunter, "Multiattribute search: Empirical evidence and information design," *Marketing Science*, 2024.

Ghose, Anindya and Sha Yang, "An empirical analysis of search engine advertising: Sponsored search in electronic markets," *Management Science*, 2009, *55* (10), 1605–1622.

Goldfarb, Avi and Mo Xiao, "Who thinks about the competition? Managerial ability and strategic entry in US local telephone markets," *American Economic Review*, 2011, *101* (7), 3130–3161.

Goli, Ali, David Reiley, and Hongkai Zhang, "Personalizing Ad Load to Optimize Subscription and Ad Revenues: Product Strategies Constructed from Experiments on Pandora," 2024. Working Paper.

Greminger, Rafael P, "Optimal search and discovery," *Management Science*, 2022, *68* (5), 3904–3924.

Guan, Zhiwei and Edward Cutrell, "An eye tracking study of the effect of target rank on web search," in "Proceedings of the SIGCHI conference on Human factors in computing systems" 2007, pp. 417–420.

Hodgson, Charles and Gregory Lewis, "You can lead a horse to water: Spatial learning and path dependence in consumer search," Technical Report, National Bureau of Economic Research 2023.

Hu, Mantian, Chu Dang, and Pradeep K Chintagunta, "Search and learning at a daily deals website," *Marketing Science*, 2019, *38* (4), 609–642.

Huang, Yufeng, Paul B Ellickson, and Mitchell J Lovett, "Learning to set prices," *Journal of Marketing Research*, 2022, *59* (2), 411–434.

Jeziorski, Przemyslaw and Ilya Segal, "What makes them click: Empirical analysis of consumer demand for search advertising," *American Economic Journal: Microeconomics*, 2015, *7* (3), 24–53.

_ **and Sridhar Moorthy**, "Advertiser prominence effects in search advertising," *Management Science*, 2018, *64* (3), 1365–1383.

Jindal, Pranav and Anocha Aribarg, "The importance of price beliefs in consumer search," *Journal of Marketing Research*, 2021, *58* (2), 321–342.

Karle, Heiko, Florian Kerzenmacher, Heiner Schumacher, and Frank Verboven, "Search Costs and Context Effects," 2023. Working Paper.

Kaye, Aaron, "The Personalization Paradox: Welfare Effects of Personalized Recommendations in Two-Sided Digital Markets," 2024. Working Paper.

Kim, Jun B, Paulo Albuquerque, and Bart J Bronnenberg, "Online demand under limited consumer search," *Marketing Science*, 2010, *29* (6), 1001–1023.

Koulayev, Sergei, "Search with dirichlet priors: estimation and implications for consumer demand," *Journal of Business & Economic Statistics*, 2013, *31* (2), 226–239.

Lam, H Tai, "Platform search design and market power," 2023. Working Paper.

Lee, Kwok Hao and Leon Musolff, "Entry Into Two-Sided Markets Shaped By Platform-Guided Search," 2023. Working Paper.

Manski, Charles F, "Measuring expectations," *Econometrica*, 2004, *72* (5), 1329–1376.

Moraga-González, José Luis, Zsolt Sándor, and Matthijs R Wildenbeest, "Consumer search and prices in the automobile market," *The Review of Economic Studies*, 2023, *90* (3), 1394–1440.

Morozov, Ilya, "Measuring benefits from new products in markets with information frictions," *Management Science*, 2023.

_ **and Anna Tuchman**, "Where does advertising content lead you? We created a bookstore to find out," *Marketing Science*, 2024.

_ **, Stephan Seiler, Xiaojing Dong, and Liwen Hou**, "Estimation of preference heterogeneity in markets with costly search," *Marketing Science*, 2021, *40* (5), 871–899.

Nocke, Volker and Patrick Rey, "Consumer search, steering and choice overload," *Journal of Political Economy*, 2023.

Osgouei, Ata Jameei, Andrew T Ching, Brian T Ratchford, and Shervin Shahrokhi Tehrani, "Estimating Position and Social Influence Effects in Online Search: An Empirical Generalized Weitzman Model," *Available at SSRN 4545610*, 2023.

Peitz, Martin, "How to apply the self-preferencing prohibition in the DMA," *Journal of European Competition Law & Practice*, 2023, *14* (5), 310–315.

Rothschild, Michael, "Searching for the lowest price when the distribution of prices is unknown," in "Uncertainty in Economics," Elsevier, 1978, pp. 425–454.

Ursu, Raluca M, "The power of rankings: Quantifying the effect of rankings on online consumer search and purchase decisions," *Marketing Science*, 2018, *37* (4), 530–552.

_ **, Qingliang Wang, and Pradeep K Chintagunta**, "Search duration," *Marketing Science*, 2020, *39* (5), 849–871.

Ursu, Raluca, Stephan Seiler, and Elisabeth Honka, "The sequential search model: A framework for empirical research," 2023. Working Paper.

Weitzman, Martin, "Optimal search for the best alternative," *Econometrica*, 1979, *78* (8).

Wu, Xiaosong, Matthew S Lewis, and Frank A Wolak, "Search with learning in the retail gasoline market," *The RAND Journal of Economics*, 2022.

Yao, Song and Carl F Mela, "A dynamic model of sponsored search advertising," *Marketing Science*, 2011, *30* (3), 447–468.

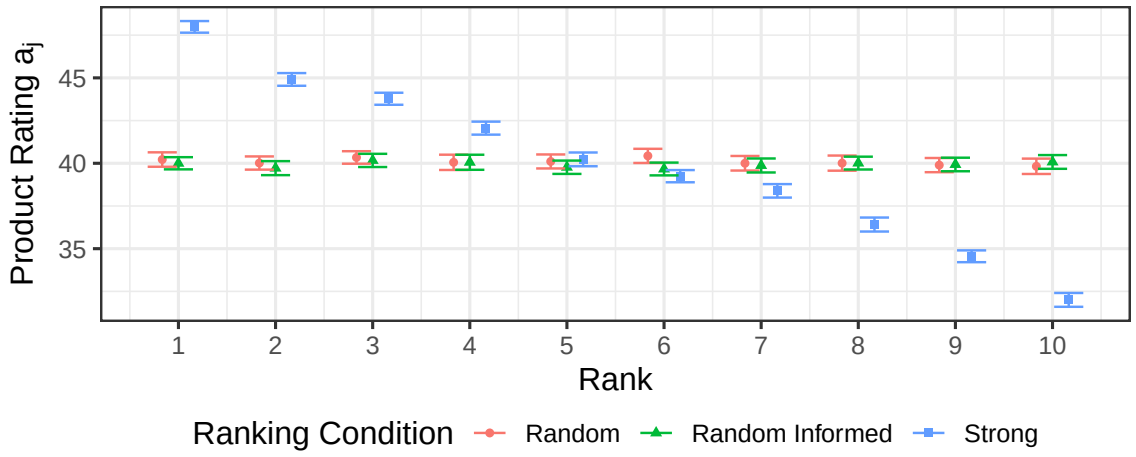
Appendix

A Study 1B: Correlated Attributes

A.1 Design

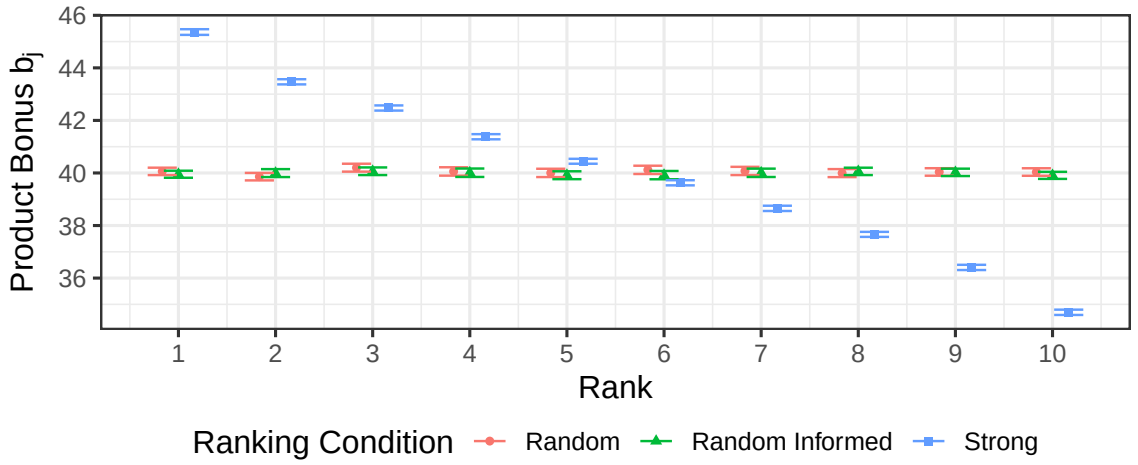
We constructed an alternative version of our first search task paradigm—Study 1B—which incorporates an additional feature of many search markets: informative pre-search signals. We redesign the attribute space to achieve this. We draw bonus b from a Normal distribution with mean 40 and standard deviation of 4. We eliminate bonus a as a component of the incentives and replace it with an informative rating (also called a here). The rating is a noisy but unbiased prediction of bonus b, which we create by adding $N(0, 10)$ draws. We show the average values by rank and condition of the rating a_j and the bonus payoff b_j in Figures A.1 and A.2.

Figure A.1: Rating by Rank by Condition (Study 1B)



Notes: Error bars represent 95% confidence intervals, which are based on standard errors clustered at the participant level.

Figure A.2: Bonus by Rank by Condition (Study 1B)



Notes: Error bars represent 95% confidence intervals, which are based on standard errors clustered at the participant level.

A.2 Results

Participants in Study 1B are more likely to start searching at the highest rating (max a) option first. Because options with high ratings (a_j) also have high payoffs (b_j), this difference is also reflected in shorter search sequences relative to Study 1.

Table A.1: Summary Statistics across Conditions (Study 1B)

	Condition Mean (SD)		
	Random	Random Informed	Strong
Num Searches	1.32 (0.71)	1.4 (0.93)	1.38 (0.87)
Bonus Earned	40.77 (3.75)	40.62 (3.8)	40.92 (3.98)
Click Max(a) First	0.69 (0.46)	0.71 (0.46)	0.69 (0.46)
N(Tasks)	2690	2860	2810
N(Participants)	269	286	281

Notes: Table reports mean and standard deviation by condition for each variable. Num Searches is the number of products searched in a task. Bonus Earned is the net bonus earned in a task (bonus of selected product net of nominal search costs). Click Max(a) First is a dummy variable which is equal to one if the participant clicked on the product with the highest rating first in a task. The prior statistics summarize across participants and tasks within a condition.

Despite searching less overall, participants in Study 1B display similar position effects to participants in Study 1 across conditions (Table A.2). We conclude from these results that the primitive rank-specific search costs are likely robust to alternative information structures.

Table A.2: Study 1B: Effect of Rank on Search by Condition

Dependent Variable:	1(Product Searched)					
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
$r \times$ Condition = Random	-0.0037***	-0.0032***	-0.0037***	-0.0032***	-0.0037***	-0.0032***
	(0.0008)	(0.0007)	(0.0008)	(0.0007)	(0.0008)	(0.0007)
$r \times$ Condition = Random Informed	-0.0030***	-0.0030***	-0.0030***	-0.0030***	-0.0030***	-0.0030***
	(0.0009)	(0.0008)	(0.0009)	(0.0008)	(0.0009)	(0.0008)
$r \times$ Condition = Strong	-0.0282***	-0.0075***	-0.0282***	-0.0075***	-0.0282***	-0.0074***
	(0.0009)	(0.0009)	(0.0009)	(0.0009)	(0.0009)	(0.0009)
a_j (demeaned)		0.0129***		0.0129***		0.0130***
		(0.0002)		(0.0002)		(0.0002)
Fixed-effects						
Condition (3)	Yes	Yes	Yes	Yes		
Search Task (10)			Yes	Yes	Yes	Yes
Participant (836)					Yes	Yes
Fit statistics						
Observations	83,600	83,600	83,600	83,600	83,600	83,600
R^2	0.01933	0.18157	0.01939	0.18165	0.05091	0.21431
Within R^2	0.01925	0.18150	0.01925	0.18153	0.01988	0.18862

Clustered (Participant) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

B Bayesian Estimation Details

B.1 Estimator Details

We outline the implementation of our estimator for Random-Informed and for all other conditions. In all cases, the following inequalities must hold for reservation utilities:

1. Search order: $z_{i1} > z_{i2}, > ... > z_{iJ}$ (where $j = 1$ is the first searched, and J is the last searched)

2. Continuation: $max(a_{ik} + b_{ik}) < z_{ij}$ for all $k < j$, all $j \in \{2, ..J\}$ (max bonus of searched is less than the reservation value of next searched)

3. Stopping: $max_{k \in \{1,..J-1}(a_{ik} + b_{ik}) > z_{iJ'}$ (max bonus of searched is greater than the reservation values of all unsearched)

B.1.1 Random-Informed

To sample from the distribution of β and σ_ϵ, we augment the data by drawing the search propensities δ's, which allows us to have values for z which satisfy the above conditions. We do not impose any sign restrictions on β, so that it is possible that the position effects go in either direction. Omitting subscript t, let w_{ij} denote the vector of attributes that shift search propensities for product j, such that

$$\delta_{ij} = \beta_i^0 + \beta_i^r \cdot (r_{ij} - 1) + \epsilon_{ij} \tag{B.1}$$
$$= \beta_i' w_{ij} + \epsilon_{ij}. \tag{B.2}$$

In this specification, w_{ij} is a vector of 1 and $(r_{ij} - 1)$, where r_{ij} is the rank of product j for user i. In the alternative specification noted in Equation 11, w_j includes 1, an indicator for whether rank of j for user i equals 1, and $(r_{ij} - 2) \cdot \mathbb{1}\{r_{ij} \neq 1\}$. Also, we denote pre-search component of the reservation utility by $\theta_j = a_j + E[b_j]$. Recall that in Random-Informed, $E[b_j] = 40$ for all j.

Since we have imposed a normal distribution on this error term, we use conjugate priors (Normal for the mean parameters, Inverse Gamma for the variance). We estimate this model using a Gibbs sampler, with each step outlined below.

1. **Augment δ and update z.** Given β and σ_ϵ, the distribution of δ is fully defined. However, we have to impose censoring on the draws such that the search sequences are optimal at that set of draws. For ease of exposition, we drop subscripts i and t, as this holds for all tasks. We also omit the superscript m for objects outside the step of the Gibbs sampler (e.g., when drawing δ, parameters β and σ_ϵ are already at their "most recent" value).

   (a) **Case: Search 1 product** For the searched product 1, we require that $z_1 > max(z_k|k > 1) \iff \delta_1^m > max(\theta_k + \delta_k^{m-1}|k > 1) - \theta_1$. Thus, draw δ_1^m from a Truncated Normal $TN(\beta' w_1, \sigma_\epsilon^2|lb = max(z_k^{m-1}|k > 1) - \theta_1, ub = \infty)$.
   For unsearched products 2 through 10, the stopping rule provides the upper bound on the reservation utilities: $z_j < a_1 + b_1 \iff \delta_j < a_1 + b_1 - a_j - 40$, thus

$\delta_j^m \sim TN(\beta' w_j, \sigma_\epsilon^2 | lb = -\infty, ub = a_1 + b_1 - a_j - 40)$.
Before proceeding, set $z_j^m = \theta_j + \delta_j^m$

(b) **Case: Searches 2 through 9 products** Let K be the number of searches. For the first searched product 1, we require that $z_1 > max(z_k | k > 1) \iff \delta_1^m > max(\theta_k + \delta_k^{m-1} | k > 1) - \theta_1$. Thus, draw from a Truncated Normal $\delta_1^m \sim TN(\beta' W_1, \sigma_\epsilon^2 | lb = max(z_k^{m-1} | k > 1) - \theta_1, ub = \infty)$. Set $z_1^m = \theta_1 + \delta_1^m$

Next, for searched products j in 2 through $K-1$, $z_j^m < z_{j-1}^{m-1} \iff \delta_j^m < z_{j-1}^m - \theta_j$. In addition, there is a lower bound condition to satisfy: that the subsequent product searched has a lower z. This is $\delta_j^m > z_{j+1}^{m-1} - \theta_j$, so we have the following $\delta_j^m \sim TN(\beta' w_j, \sigma_\epsilon^2 | lb = z_{j+1}^{m-1} - \theta_j, ub = z_{j-1}^{m-1} - \theta_j)$. Set $z_j^m = \theta_j + \delta_j^m$

For the last-searched product K we have the same style upper bound but a different lower bound (continuation restriction + search order restriction). $z_K^m < z_{K-1}^{m-1} \iff \delta_K^m < z_{K-1}^{m-1} - \theta_K$, $z_K > max(z_k | k > K) \iff \delta_K^m > max(\theta_k + \delta_k^{m-1} | k > K) - \theta_K$ and $z_K^m > max_{j<K}(a_j + b_j) \iff \delta_K^m > max_{j<K}(a_j + b_j) - \theta_K$. This yields $\delta_K^m \sim TN(\beta' w_K, \sigma_\epsilon^2 | lb = max(max_{j<K}(BA_j + BB_j) - \theta_K, max(\theta_k + \delta_k^{m-1} | k > K) - \theta_K), ub = z_{K-1}^{m-1} - \theta_K)$. Set $z_K^m = \theta_K + \delta_K^m$.

Finally, for unsearched products l, the stopping rule and search order rule provide the upper bound $z_l^m < max_{j \le K}(a_j + b_j) \iff \delta_l^m < max_{j \le K}(a_j + b_j) - \theta_l$ and $z_l^m < z_K^{m-1} \iff \delta_l^m < z_K^{m-1} - \theta_l$. Draw $\delta_l^m \sim TN(\beta' w_l, \sigma_\epsilon^2 | lb = -\infty, ub = max(max_{j \le K}(a_j + b_j) - \theta_l, z_K^{m-1} - \theta_l)$ and update $z_l^m = \theta_l + \delta_l^m$

(c) **Case: Searches 10 products** This is the same as above case, but just uses the search orders for products 1 through 9. For the final searched product, the upper bound is based on product 9. The lower bound is based on only the continuation condition. After each product, update z with the new δ.

2. **Draw** β. Given the search propensities δ^m and the prior draws of σ_ϵ^{m-1},Δ^{m-1} and V_β^{m-1}, we update the regression coefficients β_i's using a Bayes regression step. The posterior distribution is

$$\beta_i^m | \sigma_\epsilon^2, w_i, V_\beta, \Delta \sim N((w_i' w_i / \sigma_\epsilon^2 + V_\beta^{-1})^{-1} (w_i' \delta_i / \sigma_\epsilon^2 + V_\beta^{-1} \Delta' X_i), (w_i' w_i / \sigma_\epsilon^2 + V_\beta^{-1})^{-1}) \tag{B.3}$$

where w_i is a matrix of values w_{itj} with rows for each j, t corresponding to user i.

3. **Draw** σ_ϵ.

$$\sigma_\epsilon^2 \sim IG(a + \frac{N}{2}, (\frac{1}{b} + \sum_{i,j,t} (\frac{(\delta_{itj} - w_{itj} \beta_i)^2}{2}))^{-1}) \tag{B.4}$$

Here N is the total number of observations, and prior value of $\sigma_\epsilon^2 \sim IG(a, b)$

4. **Draw Δ, V_β.**

$$vec(\Delta) \sim N(vec(\tilde{\Delta}), V_\beta \otimes (X'X + A)^{-1}) \tag{B.5}$$

$$V_\beta \sim IW(\nu_0 + Nusers, V_0 + (\beta - X\tilde{\Delta})'(\beta - X\tilde{\Delta}) + (\tilde{\Delta} - \bar{\Delta})'A(\tilde{\Delta} - \bar{\Delta})) \tag{B.6}$$

Here, β is a matrix that stacks β_i's across users, and $\tilde{\Delta} = (X'X + A)^{-1}(X'\beta + A\bar{\Delta})$. Prior value of $V_\beta \sim IW(\nu_0, V_0)$ and prior value of $vec(\Delta) \sim N(vec(\bar{\Delta}), V_\beta \otimes A^{-1})$.

B.1.2 Other Conditions

The previous subsection described estimation for Random-Informed only. We jointly estimate across all conditions in Study 1 by re-parameterizing the reservation utilities to be condition specific. Recall that $z_{itj} = \theta_{itj} + \delta_{itj}$. We still let $\theta_{itj} = a_{itj} + E[b] = a_{itj} + 40$ for all conditions. We allow for differences in beliefs across ranks to be captured in the parameters β. The mean baseline and rank-specific search propensity coefficients are specified as below.

$$\begin{aligned}\beta_i^0 = \Delta_0^0 + \Delta_S^0 \mathbb{1}\{Cond_i = S\} + \Delta_R^0 \mathbb{1}\{Cond_i = R\} + \Delta_{comp}^0 Comp_i + \\ \Delta_{shop}^1 ShopFreq_i\end{aligned} \tag{B.7}$$

$$\begin{aligned}\beta_i^r = \Delta_0^r + \Delta_S^r \mathbb{1}\{Cond_i = S\} + \Delta_R^r \mathbb{1}\{Cond_i = R\} + \Delta_{comp}^r Comp_i + \\ \Delta_{shop}^r ShopFreq_i\end{aligned} \tag{B.8}$$

B.1.3 Priors

Let p denote the length of vector β_i. Thus $p = 2$ in the linear specification, and $p = 3$ in the alternative specification noted in Equation 11. The values of the prior parameters we used for estimation are given in Table B.1.

Table B.1: Priors

Parameter	Prior Variable	Value
Search Propensity(Δ)	$\bar{\Delta}$	$\mathbf{0}_{5\times p}$
	A	I_5
Heterogeneity(V_β)	ν_0	5
	V_0	$0.1 \times I_p$
Search Propensity Shocks(σ_ϵ^2)	a	10
	b	1

Notes: I_n denotes the identity matrix of dimension $n \times n$. p refers to the length of vector β_i. $\mathbf{0}_{5\times p}$ denotes a matrix of zeroes of dimension $5 \times p$.

B.2 Simulations

B.2.1 Study 1

To demonstrate that the proposed estimator can recover the correct parameters, we simulate search behavior based on the stimuli used in Random-Informed in Study 1. We take the actual data used ($N = 317$ participants, $T = 10$ tasks of $J = 10$ products each), remove outcomes, and simulate outcomes under a set of known parameters. We choose these parameters to approximately match the distribution of the number of searched products we observe in the data. We focus on the Random-Informed condition to show that the model works well to recover both baseline costs across positions and to recover the noise in behavior σ_ϵ.

We conduct three simulations. First, we simulate data and recover estimates under a DGP with exclusively unobserved heterogeneity. Second, we add observed shifters of mean search propensity and rank effect which match our empirical application: dummy variables for $Comp_i$ and $ShopFreq_i$. Third, we allow for a distinct position effect in position 1 (as opposed to exclusively a linear effect of rank).

In the first simulation, $z_{itj} = a_{itj} + E[b] + \beta_i^0 + \beta_i^r r_{itj} + \eta_{itj}$, where $\eta_{itj} \sim N(0, 15^2)$, $\beta_i^0 \sim N(-11, 20^2)$, and $\beta_i^1 \sim TN(-0.5, 1.5, \leq 0)$. All draws are i.i.d. across i, j, t or across i, depending on the unit over variation. This yields 55% of observations with only one search (versus 54% in our data).

In the second simulation, we allow for observable differences in search propensities (e.g., as in Equations 24 and 25 above). We simulate with parameters reported in Table B.3. Baseline propensity coefficients (equivalent to position 1) are drawn from a normal distribution, rank coefficients are drawn from a truncated normal distribution (with no values greater than zero). In the third simulation, we keep the existing parameters from the second simulation, and we draw the position 1 fixed effects as normal draws with parameters reported in Table B.4.

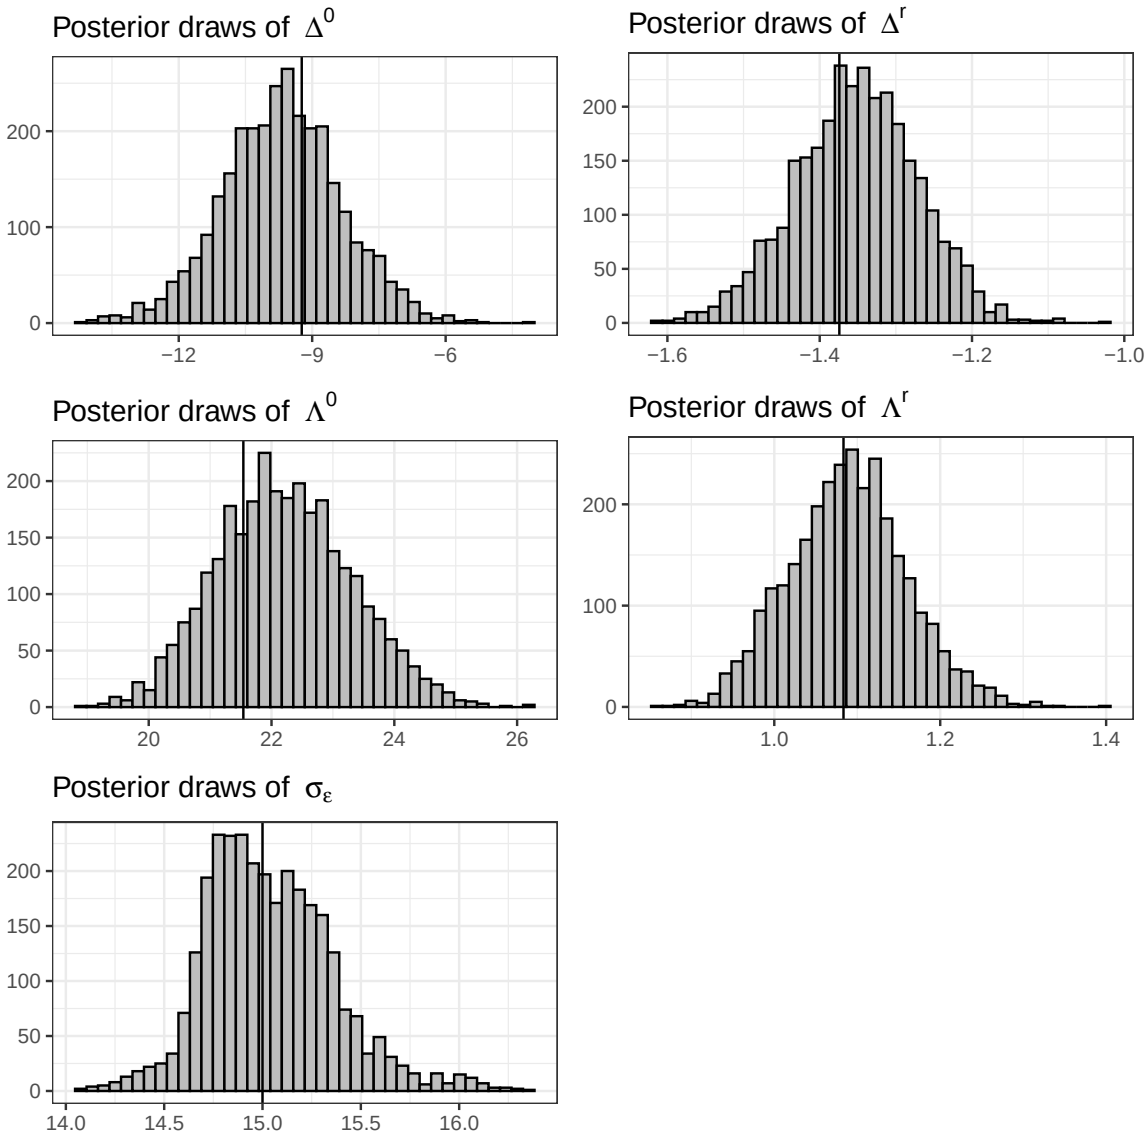
We estimate the model using 5,000 draws from the Gibbs sampler. We apply diffuse priors, and start at random draws for each value. In practice, the starting draws do not impact the chain. We discard the first 2,000 draws, and we thin the subsequent chain by using every fifth draw. Both simulations have large degrees of heterogeneity across consumers.

We report the plotted posterior distributions in Figures B.1, B.2 and B.3. We also conduct inference on the thinned chain using every 10th draw. We are able to recover all parameters, even with large degrees of across-person heterogeneity.

Table B.2: Simulation 1 Parameter Estimates

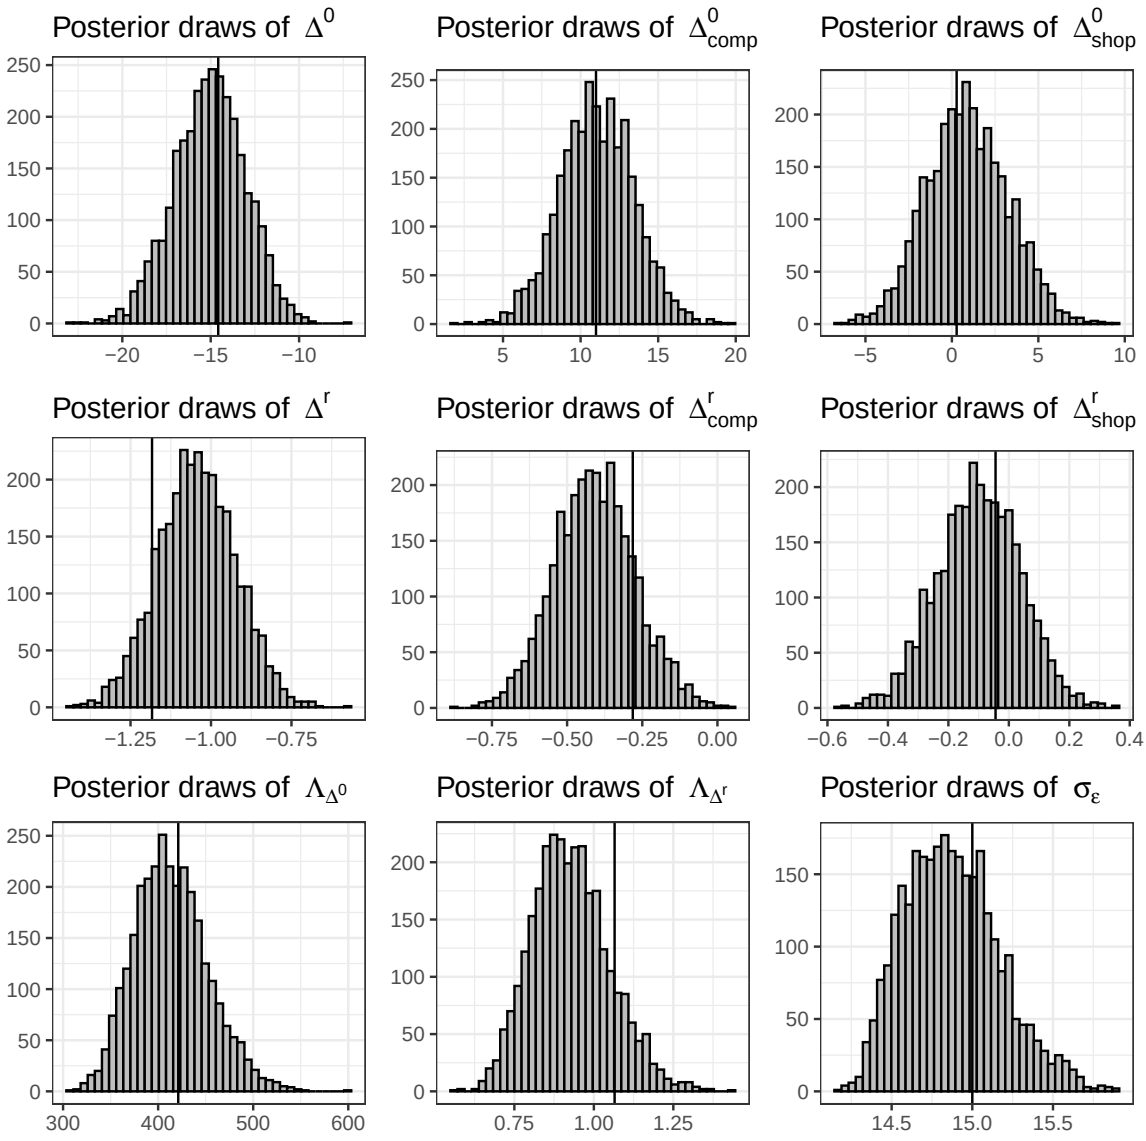
Parameter	Coef	Truth	Posterior Mean	Posterior SD
Baseline search propensity	Δ^0	-9.235	-9.502	1.332
Rank Coefficient	Δ^r	-1.374	-1.341	0.077
Unobs heterogeneity in (β^0)	λ_{Δ^0}	21.541	21.91	1
Unobs heterogeneity in (β^r)	λ_{Δ^r}	1.083	1.073	0.068
Standard deviation of ϵ	σ_ϵ	15	14.893	0.288

Figure B.1: Simulation Posterior Draws — Unobservable Heterogeneity Only



Notes: Each plot provides a histogram of posterior draws for mean search propensity (Δ^0), mean rank effect (Δ^1), the variance in search propensities (Λ^0), the variance in position effects (Λ^1), and the noise term on search decisions (σ^ϵ). The vertical black line corresponds to true value of the parameter.

Figure B.2: Simulation Posterior Draws — Observable and Unobservable Heterogeneity



Notes: Each plot provides a histogram of posterior draws for mean search propensity (Δ^0), mean search propensity FEs (Δ^0_{comp}, Δ^0_{shop}), mean rank effect (Δ^r), mean rank effect fixed slope differences (Δ^r_{comp}, Δ^r_{shop}), the variance in search propensities (Λ_{Δ^0}), the variance in position effects (Λ_{Δ^r}), and the noise term on search decisions (σ^ϵ). The vertical black line corresponds to true value of the parameter.

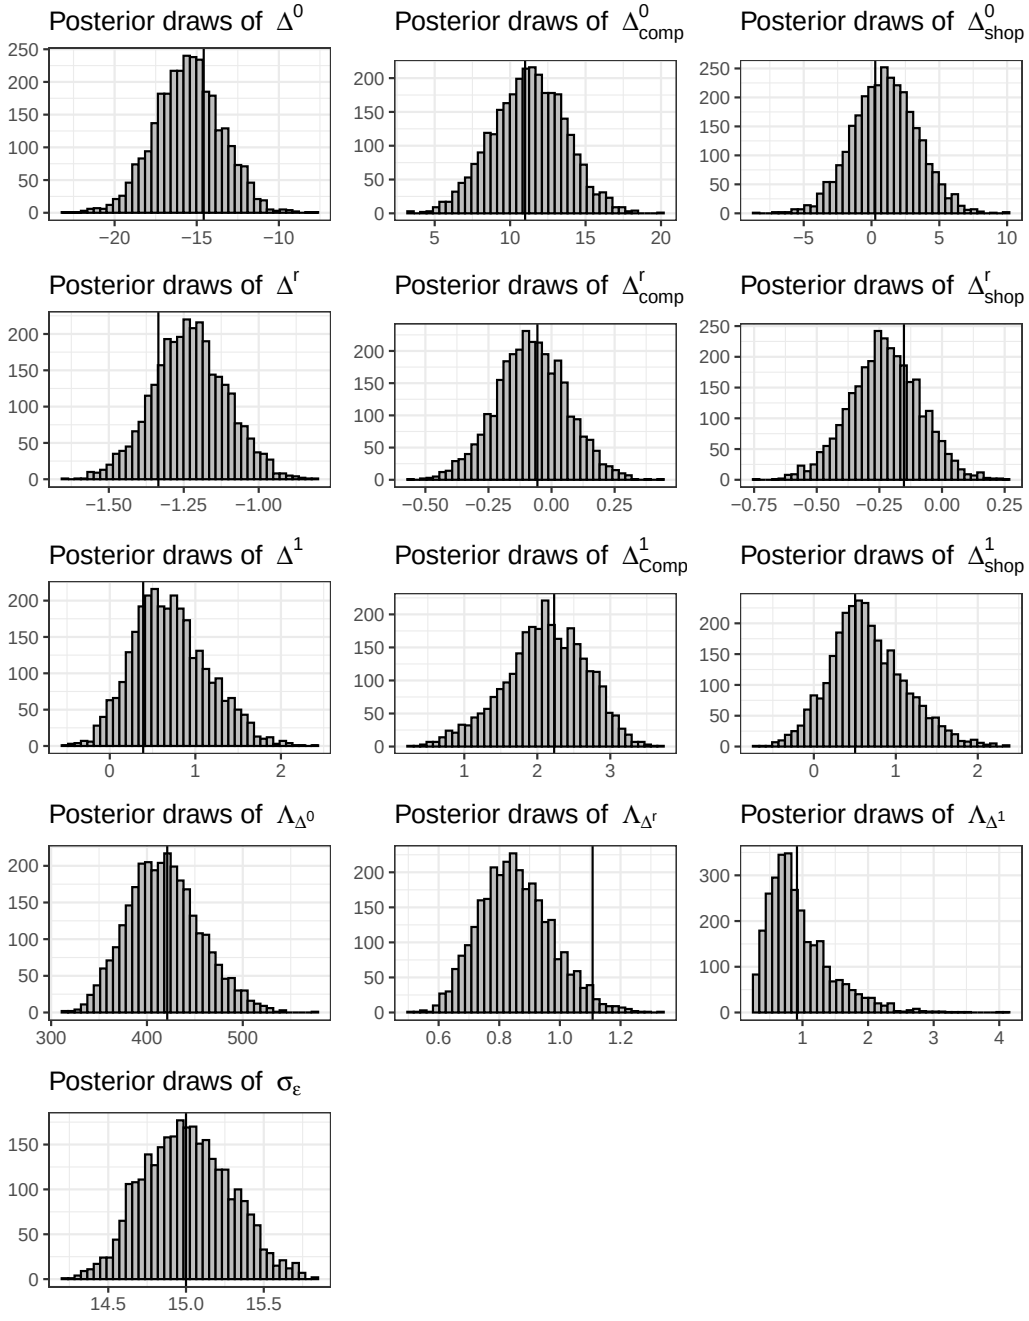
Table B.3: Simulation 2 Parameter Estimates

Parameter	Coef	Truth	Posterior Mean	Posterior SD
Baseline search propensity	Δ^0	-14.575	-14.634	1.948
Baseline search propensity Comp FE	Δ^0_{comp}	10.993	10.865	2.287
Baseline search propensity Shop FE	Δ^0_{shop}	0.276	0.689	2.36
Rank Coefficient	Δ^r	-1.183	-1.022	0.114
Rank Coefficient Comp FE	Δ^r_{comp}	-0.282	-0.429	0.143
Rank Coefficient Shop FE	Δ^r_{shop}	-0.044	-0.1	0.119
Unobs heterogeneity in (β^0)	λ_{Δ^0}	20.522	20.216	0.963
Unobs heterogeneity in (β^r)	λ_{Δ^r}	1.032	0.947	0.063
Standard deviation of ϵ	σ_ϵ	15	14.784	0.292

Table B.4: Simulation 3 Parameter Estimates

Parameter	Coef	Truth	Posterior Mean	Posterior SD
Baseline search propensity	Δ^0	-14.575	-15.759	2.253
Baseline search propensity Comp FE	Δ^0_{comp}	10.993	11.569	2.47
Baseline search propensity Shop FE	Δ^0_{shop}	0.276	0.961	2.375
Rank Coefficient	Δ^r	-1.335	-1.24	0.125
Rank Coefficient Comp FE	Δ^r_{comp}	-0.056	-0.084	0.141
Rank Coefficient Shop FE	Δ^r_{shop}	-0.152	-0.216	0.139
Position 1 Effect	Δ^1	0.39	0.692	0.377
Position 1 Comp FE	Δ^1_{comp}	2.231	2.047	0.55
Position 1 Shop FE	Δ^1_{shop}	0.504	0.788	0.487
Unobs heterogeneity in (β^0)	λ_{Δ^0}	20.522	20.543	0.963
Unobs heterogeneity in (β^r)	λ_{Δ^r}	1.053	0.937	0.066
Unobs heterogeneity in (β^1)	λ_{Δ^1}	0.957	0.889	0.239
Standard deviation of ϵ	σ_ϵ	15	15.075	0.357

Figure B.3: Simulation Posterior Draws — Separate Postion 1 Effect



Notes: Each plot provides a histogram of posterior draws for mean search propensity (Δ^0), mean search propensity FEs (Δ^0_{comp}, Δ^0_{shop}), mean rank effect (Δ^r), mean rank effect fixed slope differences (Δ^r_{comp}, Δ^r_{shop}), position 1 rank effects (Δ^1, Δ^1_{comp}, Δ^1_{shop}), the variance in search propensities (Λ_{Δ^0}), the variance in position effects (Λ_{Δ^r}), and the noise term on search decisions (σ^ϵ). The vertical black line corresponds to true value of the parameter.

B.2.2 Study 2

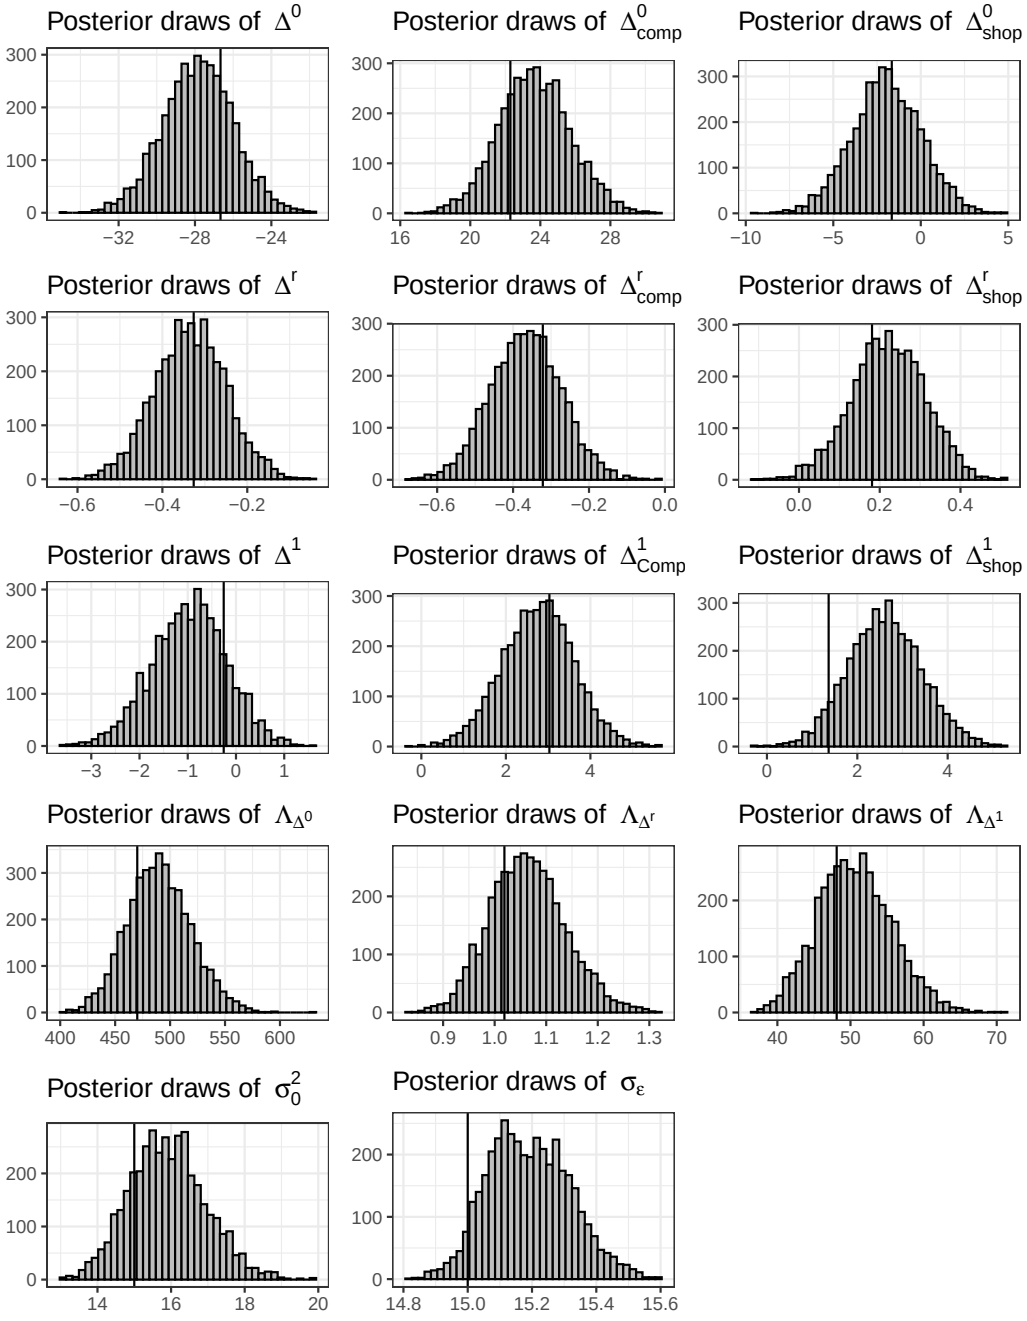
We simulate search behavior as per the learning model described in Study 2. We take the actual data used in the Strong condition ($N = 613$ participants, $T = 15$ tasks of $J = 20$ products each), remove outcomes, and simulate search under a set of known parameters. We choose the search propensity parameters to approximately match the estimates from Study 1 and select the prior variance parameter to reasonably replicate the rate of learning we observe in the data. We focus on the Strong condition to show that the model can recover the prior variance parameter (which governs the rate of learning), which is the additional parameter to be estimated in this model.

We estimate the model using 5,000 draws from the Gibbs sampler. We discard the first 1,000 draws and also conduct inference on the thinned chain using every 10th draw. We are able to recover all parameters, even with large degrees of across-person heterogeneity. The posterior distribution is plotted in Figure B.4 and the estimated mean (SD) of the parameters are reported in Table B.5.

Table B.5: Simulation – Study 2 Learning Model: Parameter Estimates

Parameter	Coef	Truth	Posterior Mean	Posterior SD
Baseline search propensity	Δ^0	-26.662	-27.383	1.857
Baseline search propensity Comp FE	Δ^0_{comp}	22.288	23.199	2.023
Baseline search propensity Shop FE	Δ^0_{shop}	-1.658	-1.959	1.859
Rank Coefficient	Δ^r	-0.326	-0.331	0.081
Rank Coefficient Comp FE	Δ^r_{comp}	-0.322	-0.37	0.092
Rank Coefficient Shop FE	Δ^r_{shop}	0.181	0.219	0.082
Position 1 Effect	Δ^1	-0.255	-0.879	0.776
Position 1 Comp FE	Δ^1_{comp}	3.03	2.716	0.842
Position 1 Shop FE	Δ^1_{shop}	1.365	2.591	0.774
Unobs heterogeneity in (β^0)	λ_{Δ^0}	21.683	21.987	0.802
Unobs heterogeneity in (β^r)	λ_{Δ^r}	1.009	1.029	0.035
Unobs heterogeneity in (β^1)	λ_{Δ^1}	6.936	7.097	0.341
Prior Variance	σ_0^2	15	15.963	1.062
Standard deviation of ϵ	σ_ϵ	15	15.181	0.149

Figure B.4: Simulation Posterior Draws — Study 2 Learning Model



Notes: Each plot provides a histogram of posterior draws for mean search propensity (Δ^0), mean search propensity FEs (Δ^0_{comp}, Δ^0_{shop}), mean rank effect (Δ^r), mean rank effect FEs (Δ^r_{comp}, Δ^r_{shop}), position 1 rank effects (Δ^1, Δ^1_{comp}, Δ^1_{shop}), the variance in search propensities (Λ_{Δ^0}), the variance in position effects (Λ_{Δ^r}), the prior variance (σ_0^2) and the noise term on search decisions (σ^ϵ). The vertical black line corresponds to true value of the parameter.

C Study 1: Heterogeneity by Participant Attributes

Table C.1: Search Intensity and Position Effects by Participant Attributes (Study 1)

Dependent Variables:	N Products Searched	1(Product Searched)
Model:	(1)	(2)
Variables		
Comp Checks Correct	7.749***	0.1083***
	(0.8681)	(0.0128)
Comp Checks Correct × r		-0.0058***
		(0.0013)
ShopFreq	0.7155	0.0150
	(0.9183)	(0.0135)
ShopFreq × r		-0.0013
		(0.0014)
Condition = Random	1.230	0.0121
	(1.057)	(0.0112)
Condition = Strong	-0.1996	-0.0019
	(1.068)	(0.0097)
Bonus A (demeaned)		0.0235***
		(0.0008)
r		-0.0055***
		(0.0013)
Constant	17.13***	0.2017***
	(1.040)	(0.0127)
Fit statistics		
Observations	961	96,100
R^2	0.07849	0.06599
Adjusted R^2	0.07464	0.06591

*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: Table reports OLS estimates. The data in column 1 is at participant level, and at the participant-product-task level in column 2. In columns 1 and 2, the dependent variables are the number of products that the participant searches across all ten tasks and an indicator for whether a product is searched, respectively. Comp Checks Correct is a dummy variable which equals one if the participant passed comprehension checks on the first attempt. ShopFreq is a dummy variable which equals one if the user shops online at least once a week (self-reported). The variable b_A (demeaned) is the demeaned value of the product's bonus A. Standard errors are clustered at the participant level in Column 2.

Table C.1 reports how the total number of products searched (column 1) and whether the participant searches a particular product (column 2) vary by whether the user passed all comprehension checks on the first attempt and whether the user self-reports that they shop online at least once a week. There is significant heterogeneity in search behavior by the former attribute. Participants who passed all comprehension checks on the first attempt search eight more products (or 44% more) over all ten tasks than those who did not.[47] They also exhibit stronger position effects, as demonstrated by the statistically significant

[47] Fourty-four percent of participants answered all comprehension checks correctly on the first try.

negative effect of the interaction between rank and whether all the comprehension checks were answered correctly in Column 2. One might expect that frequent online shoppers might demonstrate different position effects compared to an infrequent shopper due to greater familiarity with an online retail setting. Here, the interaction between whether the participant is a high-frequency shopper and the product's rank is negative, indicating that these participants have stronger position effects, but it is not statistically significant in this specification.

Table C.2: Position Effects by Participant Attributes in Random and Strong (Study 1)

Dependent Variable:	1(Product Searched)	
Model:	(1)	(2)
Variables		
Comp Checks Correct	0.1092^{***}	0.1180^{***}
	(0.0247)	(0.0206)
Comp Checks Correct $\times$ r	-0.0024	-0.0099^{***}
	(0.0022)	(0.0021)
ShopFreq	-0.0141	0.0047
	(0.0279)	(0.0223)
ShopFreq $\times$ r	-6.53×10^{-5}	0.0003
	(0.0025)	(0.0023)
Bonus A (demeaned)	0.0214^{***}	0.0236^{***}
	(0.0014)	(0.0013)
r	-0.0060^{**}	-0.0085^{***}
	(0.0024)	(0.0021)
Constant	0.2234^{***}	0.2234^{***}
	(0.0262)	(0.0204)
Fit statistics		
Observations	33,000	31,400
R^2	0.05902	0.06934
Adjusted R^2	0.05885	0.06917
Condition	Random	Strong

Clustered (Participant) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: Table reports OLS estimates. The data in both columns are at the participant-product-task level, and the dependent variable is an indicator for whether a product is searched. Comp Checks Correct is a dummy variable which equals one if the participant passed comprehension checks on the first attempt. ShopFreq is a dummy variable which equals one if the user shops online at least once a week (self-reported). Columns 1 and 2 report estimates only for participants in the Random condition and the Strong condition, respectively. The variable b_A (demeaned) is the demeaned value of the product's bonus A.

Table C.2 reports whether position effects differ by whether the participant answers all comprehension checks correctly in the first attempt and is a high frequency online shopper for the Random (column 1) and Strong (column 2) conditions. Participants who answer all comprehension checks correctly on the first try search more in both conditions, but only exhibit strong position effects in Strong. This suggests that participants in Strong learned more quickly about the search environment. The coefficients for whether the participant

is a high frequency online shopper are not statistically significant and are close to zero in magnitude.

These results demonstrate that participants' search behavior is correlated with their observable attributes, especially whether they pass the comprehension checks on the first try. Allowing the model to account for heterogeneity among these dimensions can improve the model's fit by better explaining across-participant differences in search costs and position effects.

D Study 1 Learning: Task-Specific Estimates

Recall that participants were not informed about the ranking scheme in the Strong and Random conditions and therefore had to learn the underlying algorithm through search experience. Consistent with learning through experience, we find that position effects in the first task do not differ significantly between these conditions (Figure I.3 in this appendix). Therefore, the presence of position effects in Strong when averaged over all search tasks suggests that participants are learning that positions are informative. Learning has not yet been captured in our model estimates, as thus far, they have been averaged across all tasks. In this section, we investigate the possibility and speed of learning by estimating a model which has task-specific marginal search propensities and position effects. We assume that δ_{itj} is given by

$$\begin{aligned}\delta_{itj} = \beta_i^1 + \beta_i^2 \mathbb{1}\{Task_i \geq 6\} + \beta_i^3 \mathbb{1}\{r_{itj} = 1\} + \beta_i^4 \mathbb{1}\{r_{itj} = 1\} \cdot \mathbb{1}\{Task_i \geq 6\} + \\ + \beta_i^5 (r_{itj} - 2) + \beta_i^6 (r_{itj} - 2) \cdot \mathbb{1}\{Task_i \geq 6\} + \epsilon_{itj},\end{aligned} \tag{D.1}$$

where the terms β_i^1 represents i's search propensity for tasks 1–5, and β_i^2 represents the *additional* search propensity of tasks 6–10, relative to the first five tasks. Similarly, the terms β_i^3 and β_i^5 represent the effect of *rank* on i's search propensity for tasks 1–5, and β_i^4 and β_i^6 are the additional effect of rank of tasks 6 – 10 relative to the first five. The set of participant attributes (e.g., whether i is in the Strong condition, whether i got the comprehension checks correct on the first try, etc.) remains the same as in Equation 9.

The results of this estimation are presented in Table D.1. We omit reporting the heterogeneity of the βs and σ_ϵ for brevity. As predicted, the rank effect due to *beliefs* in the Strong condition increases in magnitude over tasks. This is most evident for the rank effect of the first item (Δ_S^1): the interaction with the Strong condition increases by 66% between the first and second set of tasks (1.5 versus $1.5 + 1$). We also find that participants do not need all ten tasks to learn that positions are informative. The rank effect of the first item appears within the first five tasks in the Strong condition (Δ_S^1 for tasks 1–5).

We also see that Δ_0^r decreases over tasks, indicating that in all conditions, position effects get stronger. This could possibly be due to fatigue—participants in latter tasks might have preferred the "easier-to-search" top-ranked products more, even in Random-Informed.

Table D.1: Estimation Results – Task Specific Search Model (Study 1)

		Estimate (SD)	
Variable	Coefficient	Tasks 1-5	Tasks 6-10
Rank Coefficient β^r			
Intercept	Δ_0^r	-0.1 (0.094)	-0.17 (0.078)**
Strong	Δ_S^r	-0.026 (0.093)	-0.082 (0.077)
Random	Δ_R^r	-0.05 (0.097)	-0.11 (0.064)
Comp	Δ_{comp}^r	-0.37 (0.089)***	0.044 (0.064)
ShopFreq	Δ_{shop}^r	-0.044 (0.084)	-0.082 (0.056)
Pos 1 Coefficient β^1			
Intercept	Δ_0^1	-0.2 (0.5)	-0.24 (0.34)
Strong	Δ_S^1	1.4 (0.53)**	1.2 (0.46)***
Random	Δ_R^1	0.01 (0.59)	0.047 (0.41)
Comp	Δ_{comp}^1	2.1 (0.5)***	0.71 (0.37)**
ShopFreq	Δ_{shop}^1	0.68 (0.45)	0.27 (0.38)
Baseline Search Propensity β^0			
Intercept	Δ_0^0	-24 (1.9)***	4.1 (0.86)***
Strong	Δ_S^0	1.2 (2.1)	1.2 (0.92)
Random	Δ_R^0	-0.081 (2.1)	0.78 (0.89)
Comp	Δ_{comp}^0	20 (1.7)***	-3.3 (0.8)***
ShopFreq	Δ_{shop}^0	-1.8 (1.7)	-0.64 (0.75)

Notes: Table reports posterior means and standard deviation from the search propensity specified in Equation D.1. Posterior means and standard deviations are based on the thinned chain which drops the first 2,000 draws of the chain and keeps every tenth draw thereafter. We omit variance of heterogeneity results for brevity. Signif. Codes: $***: 0.01, **: 0.05, *: 0.1$

E Study 2B: Zero Marginal Cost

In our studies in the main paper, we impose an explicit marginal search cost of 1 point. As we discussed, we impose this cost because participants who are paid to take a survey might weigh their outside option differently from those conducting a real search. However, a concern with imposing this marginal cost is that it creates another difference between our experimental setting and the field: in field settings, consumers do not pay an explicit search cost to click on products, but rather make the trade-off between searching again and choosing either the outside option or the best-searched alternative. In this section, we present results from a pilot study where we implemented a version of Study 2 that has zero explicit marginal search cost (i.e., it does not cost any points to search a product).

This pilot study includes 89 participants, randomized across 4 experimental conditions (Random-Informed, Strong, Strong-to-Random5, and Strong-to-Pref5). Like in Study 2, each participant performs each search task fifteen times, with twenty products in each search task. Table E.1 reports the summary statistics across conditions. The first row shows that participants in this pilot search more than those in the main Study 2, where participants search between 1.78–1.99 products on average. (Table 5, row 1). As a result, they spend longer per search task. In addition, pilot participants are less likely to search the product with the highest bonus A first; in Random-Informed, where this is the optimal strategy, participants search the product with the highest bonus A 20% of the time (row 3), as compared to 31% of the time in Study 2 (Table 5, row 3). This suggests that without the explicit marginal search cost, participants deviate more from optimal sequential search. Furthermore, in the pilot, participants searched all twenty products in 12% of the search tasks, compared to only 0.1% of tasks in the main Study 2. A high proportion of tasks where all products are searched presents a challenge for estimating search costs, as these cases only allow us to determine the upper bound of the search costs.

Table E.1: Summary Statistics across Conditions (Study 2, Zero Marginal Cost)

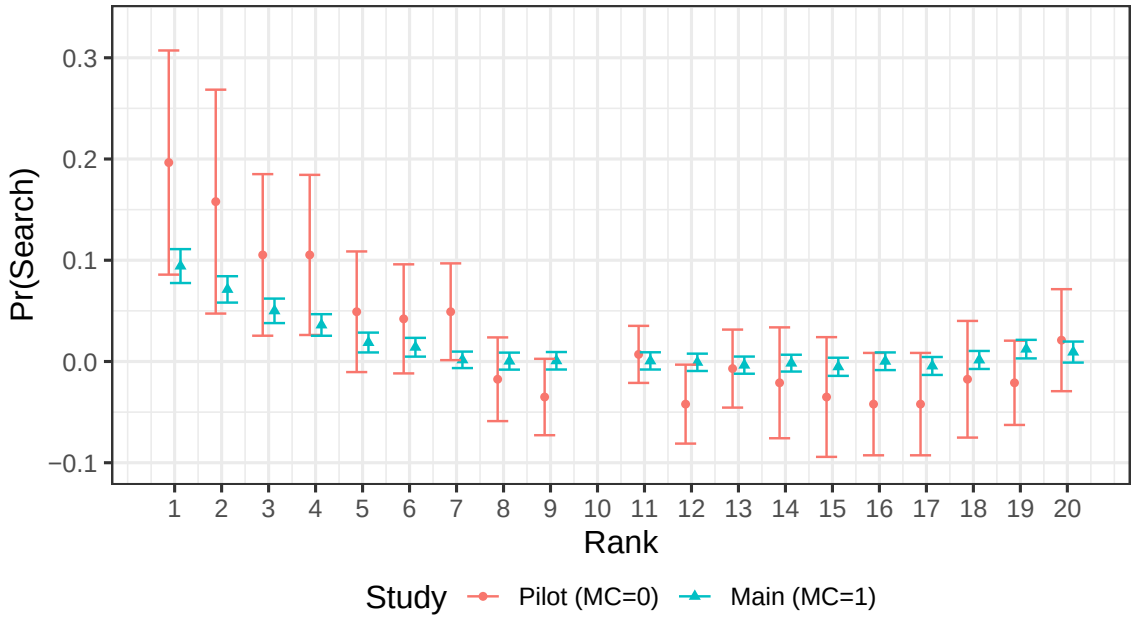
	Condition Mean (SD)			
	RI	Str	Str to P (5)	Str to R (5)
Num Searches	7.4 (7.92)	4.55 (5.15)	6.34 (6.84)	5.07 (5.58)
Bonus Earned	67.2 (10.52)	73.15 (10.43)	70.47 (12.03)	69.8 (10.43)
Click Max(a) First	0.2 (0.4)	0.23 (0.42)	0.18 (0.39)	0.22 (0.41)
Time Spent per Task	24.25 (24.57)	15.92 (16.8)	21.79 (21.6)	19.06 (18.3)
N(Tasks)	285	315	405	330
N(Participants)	19	21	27	22

Notes: "Str" stands for Strong, "P" for Pref, "R" for Random, and "RI" for Random-Informed. Table reports mean and standard deviation by condition for each variable. Num Searches is the number of products searched in a task. Bonus Earned is the total bonus of the selected product in a task. Click Max(a) First is a dummy variable which is equal to one if the participant searched the product with the highest ex-ante bonus value first in a task. Time Spent per Task is the seconds spent on a task from page load until final click. The prior statistics summarize across participants and tasks within a condition.

Figure E.1 plots the average probability of search by rank for the pilot and for main study 2. For ease of comparison, we include only the Random-Informed condition. Because

the average probability of search for each item is higher in the pilot, we plot the probability of search relative to the tenth-ranked product. The pilot's sample size is relatively small, so we do not have the statistical power to detect differences. However, it appears that position effects are stronger in magnitude when the explicit marginal cost is zero. That is, participants are more likely to search the first product relative to the tenth in the pilot than in the main study. A possible explanation is that in the pilot study, participants rely less on Bonus A to guide their search (consistent with the summary statistics above) because they do not feel constrained by how many products to search.

Figure E.1: Search Probability by Rank for Random-Informed



Notes: This figure plots the average probability of searching each product by its rank, relative to the tenth product, for participants in the Random-Informed condition. Error bars represent 95% confidence intervals, which are based on standard errors clustered at the participant level.

Our results suggest that having zero explicit marginal search cost leads to search behavior that is less consistent with the search patterns observed in real-world contexts. In particular, it leads to greater search intensity compared to that reported in prior research (e.g., Morozov, 2023; Ursu, 2018), and it leads to search behavior that is less consistent with optimal behavior predicted by the Weitzman sequential search model.

Notably, the introduction of even a minimal marginal search cost—such as one point or 0.25 cents—has a surprisingly large impact on search intensity. We believe this effect is unique to the transition from zero to one point. Introducing a marginal cost of one point makes the trade-off between stopping and continuing search more salient to participants, thereby resulting in search behavior more consistent with Weitzman's model. We anticipate smaller differences in search behavior when increasing the marginal search cost from one to two points.

F Study 2C: Higher Stakes

We implement a robustness check where we raised the incentives in the search task. This study is similar to Study 2, with a few exceptions: (1) we increased the payoffs from search by using a conversion rate of 200 points per dollar (instead of 500), (2) participants conduct ten search tasks (instead of fifteen) and only include the Strong and Random-Informed conditions (for cost considerations), (3) decreased the point value of the marginal search cost to 0.5 (instead of 1), and (4) payoffs were summed over all tasks (instead of eight out of fifteen). Otherwise, everything else, such as the distributions of the bonuses (pointwise), are the same as in Study 2.

We compare the search behaviors with that of Study 2 by combining the data from the "higher stakes" study with the data from each participants' first ten tasks in Study 2. We then estimate the following regression:

$$y_{it} = \alpha_t + \beta_1 \mathbb{1}\{HigherStakes_i\} + \beta_2(\mathbb{1}\{HigherStakes_i\} \times Condition_i) + \epsilon_{it}. \quad \text{(F.1)}$$

We consider two dependent variables for participant i in task t: the number of products searched, and an indicator for whether i searches the product with the highest bonus A first. Recall that searching the product with the highest bonus A first is the optimal strategy in Random-Informed. Note that we do not include participant fixed effects in this specification because they are multicollinear with the $\mathbb{1}\{HigherStakes_{it}\}$. We report the estimates in Table F.1. Here, we find that participants search 0.27 more products per search task in the higher stakes study (column 1). This is expected because the returns to search are greater. However, we do not see differences in treatment effects on search depth, as indicated by the null interaction effect between condition and $\mathbb{1}\{HigherStakes_{it}\}$. This shows that having greater incentives does not have a differential effect on our treatment effects, indicating that our treatment effects should generalize to settings with higher incentives.[48]

The second column shows the effect on whether the participant searches the product with the highest bonus A first, as an indicator of optimal search (for Random Informed). Again, we do not find differences between the higher stakes study and Study 2. This suggests search behavior is as "optimal" in our main studies as when stakes are greater. That is, raising the incentives for attention does not seem to change the degree to which participants pay attention to bonus A, only their willingness to search further (since it translates to higher payoffs in dollar terms). However, we cannot rule out that even higher stakes would not generate behavior that more closely aligns with the Weitzman search model.

We also measure whether higher stakes differentially impacts position effects by estimating the following:

$$\mathbb{1}\{Searched_{itj}\} = \alpha_t + \beta(r_{itj} \times Condition_i \times \mathbb{1}\{HigherStakes_i\}) + \epsilon_{itj}. \quad \text{(F.2)}$$

This regression has a three-way interaction between the product j's rank r_{itj}, the participant's condition, and whether the participant is in the higher stakes study. We report the

[48] Note that this is not explicitly an experimental manipulation, since whether an observation was generated with higher stakes is not randomized within a single study. Instead, we combine non-overlapping samples from the same participant pool and compare the same experimental manipulation (RI vs Strong) across samples.

Table F.1: Differences in Search Behavior Between Study 2 and Study 2C (Higher Stakes)

Dependent Variables:	Search Depth	Searched Highest b_A First
Model:	(1)	(2)
Variables		
HigherStakes	0.2660***	-0.0187
	(0.0897)	(0.0180)
Condition = Strong	-0.1157**	-0.0228
	(0.0572)	(0.0171)
Condition = Strong × HigherStakes	-0.0312	-0.0141
	(0.1234)	(0.0246)
Fixed-effects		
Search Task (10)	Yes	Yes
Fit statistics		
Observations	21,910	21,910
R^2	0.00596	0.00243
Within R^2	0.00567	0.00190

Clustered (Participant) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: This table presents the OLS estimates of Equation F.1.

estimates in Table F.2. We find that the higher stakes study does not show significantly different position effects than Study 2; the coefficients for the interaction between whether the participant is in the higher stakes study and rank, and with "Condition = Strong" are not statistically significant and close to zero. This demonstrates that our conclusions from our main study are robust to higher stakes.

Table F.2: Differences in Search Behavior Between Study 2 and Study 2C (Higher Stakes)

Dependent Variable:	1(Product Searched)
Model:	(1)
Variables	
r	-0.0029***
	(0.0003)
r × Condition = Strong	-0.0013***
	(0.0004)
r × HigherStakes	-0.0004
	(0.0004)
r × Condition = Strong × HigherStakes	-0.0002
	(0.0006)
Condition = Strong	0.0083
	(0.0058)
HigherStakes	0.0174**
	(0.0073)
Condition = Strong × HigherStakes	5.5×10^{-5}
	(0.0101)
Fixed-effects	
Search Task	Yes
Fit statistics	
Observations	438,200
R^2	0.00621
Within R^2	0.00618

Clustered (Participant) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: This table presents the OLS estimates of Equation F.2.

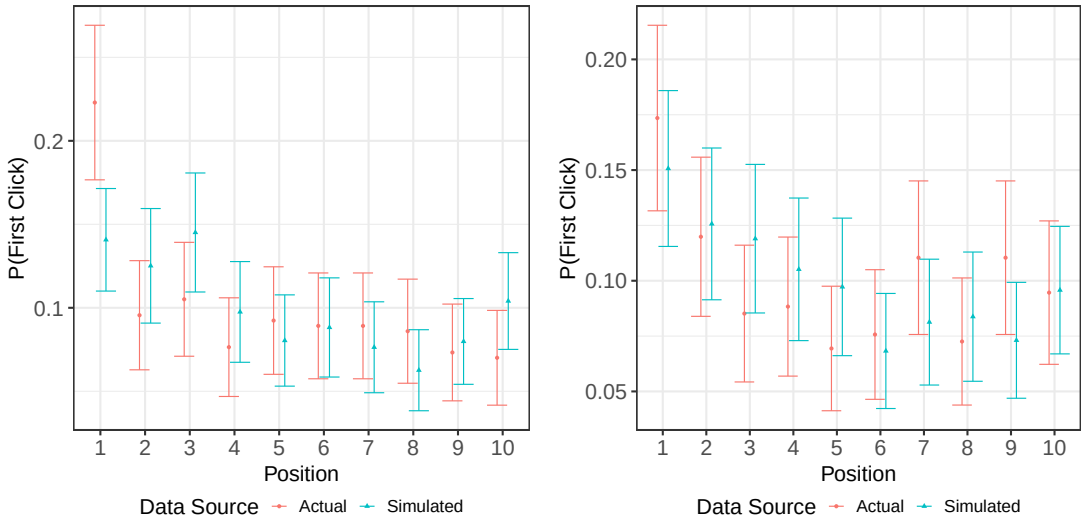
G Model Fit

G.1 Study 1

Simulating behavior from Random-Informed shows that we cannot match the moments in the Strong condition using the variation generated by random rankings. This is by design, since there are no ex ante observable differences in how products appear to participants across rankings (bonus A and bonus B are independent).

To demonstrate this, we re-estimate the 'Separate Positions 1' model using only Random-Informed tasks 1–9. We then simulate the implied search patterns in task 10 for Random-Informed and for Strong. Figure G.1 shows these simulated click probabilities against the data. Simulated search from Random-Informed fails to capture the higher probability of searching the first item (Panel (a)). The model fits the data moments well for Random-Informed (Panel (b)).

Figure G.1: Predicting First Search in Task 10 (Study 1)



(a) Strong

(b) Random-Informed

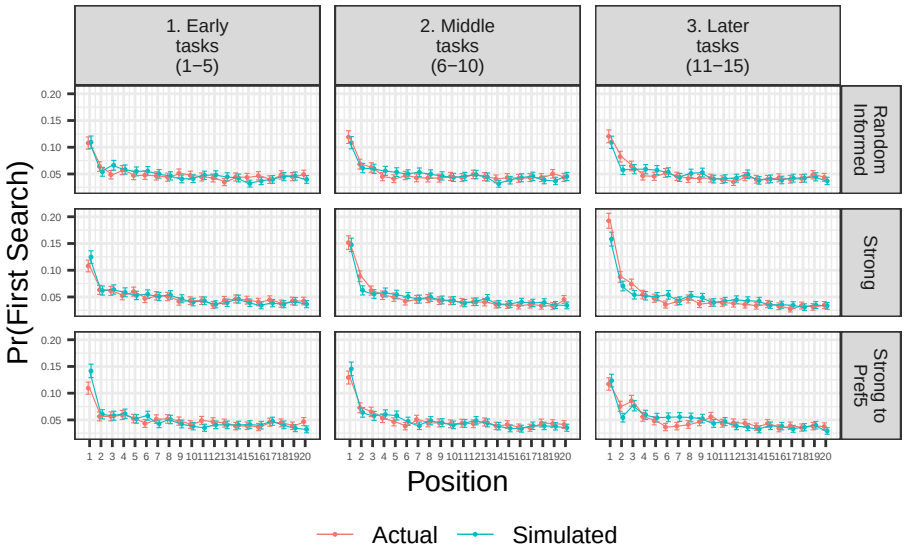
Notes: Panels (a) and (b) show the predicted versus actual first click probabilities based on the baseline model estimated on the first nine tasks in Random-Informed for the tenth task in Strong and Random-Informed, respectively.

G.2 Study 2

How well do the model predictions fit the observed search patterns? We simulate search behavior based on the model described in Section 4.3 using the estimated parameters and then compare the first search probability by rank with the corresponding probabilities observed in the actual data. Figure G.2 plots this comparison, for conditions Random-Informed, Strong and Strong-to-Pref5. The figure shows that the model recovers underlying pa-

rameters reasonably well. The simulated and actual search patterns are very similar for Random-Informed. In Strong, the simulated search predicts an increase in probability of first searching the top-ranked product—matching what we see in the actual data.[49] Finally, in Strong-to-Pref5, the model is also able to predict the 'peak' beginning to form at rank 3 in tasks $11-15$, consistent with the observed search patterns and the underlying algorithm. The plots for the other conditions are in Figure G.3.

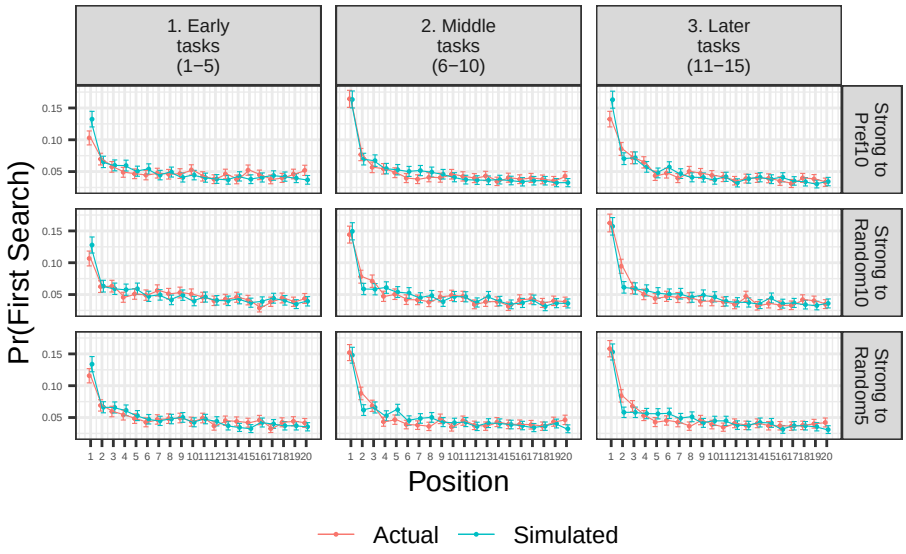
Figure G.2: Actual versus Model-Predicted Search Patterns (Study 2)



Notes: The above figures plot the actual observed first click probability by rank versus the same probability as predicted by the model described in the preceding section, based on estimated parameters. The left, middle and right panels show the search probabilities for tasks 1–5, 6–10 and 11–15 respectively, to illustrate the evolution in search probabilities as participants learn about the underlying algorithm and update their beliefs.

[49]The probability of searching rank 1 in tasks 11–15 in *Strong* is slightly understated by the model (middle right panel in Figure G.2). This suggests that there may be some form of learning that occurs which is not captured by our model. However, given the trade-off between accuracy and model tractability, we feel that our model reasonably captures the essential elements of the consumer search process.

Figure G.3: Actual versus Model-Predicted Search Patterns (Study 2)



Notes: The above figures plot the actual observed first click probability by rank vs the same probability as predicted by the model described in the preceding section, based on estimated parameters. The left, middle and right panels show the search probabilities for tasks $1-5$, $6-10$ and $11-15$ respectively, to illustrate the evolution in search probabilities as participants learn about the underlying algorithm and update their beliefs.

H Counterfactuals

In this section, we describe in detail how we obtain our surplus predictions.

H.1 Comparing Surplus with and without Belief Adjustments

We simulate behavior for a simulated set of 594 consumers with the same attributes (e.g., comprehension check accuracy and shopping frequency) as those in our Study 2 sample. We then predict consumer surplus for thirty search tasks for each consumer using the estimates from models that account or do not account for position-specific beliefs. In the first five of the thirty tasks, the ranking algorithm is *Strong*, and in the remainder of the tasks, the algorithm is *Pref*. We generate the bonuses for each of these tasks by sampling from the pool of tasks under the same algorithm that generated the bonuses in the experiment.

We first describe how we generate the surplus predictions with beliefs (and thus, learning) and without beliefs.

1. Estimate the model described in Section 3.3 using the data generated from tasks 6 and 7 in the Strong-to-Random5 condition. We assume away the possibility that beliefs have adjusted over the course of the first five tasks under *Strong* rankings: the expected belief about bonus B is 40 for all ranks. We refer to these estimates as the "incorrect" estimates (i.e., conflating beliefs with costs).

2. Estimate the model described in Section 3.3 using the data generated from tasks 6 and 7 in the Random-Informed condition. Note that because all the data comes from Random-Informed, $\Delta_S = \Delta_R = 0$. In other words, this is the sequential search model in which all position effects are attributed to search costs, and the expected belief about bonus B is 40 for all ranks. We refer to these estimates as the "correct" estimates." We re-estimate this model using the same set of tasks, as opposed to using the estimates from Random-Informed in Section 4.3, to provide a fair comparison to the model in Step 1.

3. For each product in each task, draw ϵ_{itj} using the estimate of σ_ϵ from the correct estimates. We do this to keep the ϵ_{itj} draws constant when comparing the two model predictions.

4. Using the *correct* estimates, draw $\beta_i^0, \beta_i^1, \beta_i^r$ for each user.

5. Simulate search behavior for thirty tasks for each user under the draws from the correct estimates in Step 4. We assume that at the start of the first task, the prior mean of each position is 40, as in the experiment, and σ_0, the variance of their prior belief is 11, which is the estimate from Study 2 (Table 9). Position-specific beliefs are updated after each task based on sampled (searched) products.

6. Using the *incorrect* estimates, draw $\beta_i^0, \beta_i^1, \beta_i^r$ for each user.

7. Simulate search behavior for all thirty tasks for each user under the draws from the incorrect estimates in Step 6. We simulate under the typical assumption that there are no position-specific beliefs ($E[b_r] = 40 \forall r$).

8. Calculate the consumer's surplus for each task for the search behavior generated under the correct and incorrect estimates. This is defined as the bonus A + bonus B of the selected product minus the incurred search costs. With the incorrect estimates, we compute costs by inverting the full search propensity, including any residual effects of beliefs: $c_{itj} = \sigma_b \zeta^{-1}(\frac{\delta_{itj} - \epsilon_{itj}}{\sigma_b})$

9. Take the average of consumer over consumers for each task.

10. Repeat Steps 3–9 200 times.

11. Take the average surplus per task across all 200 iterations. The confidence intervals are the 2.5 and 97.5 percentiles.

H.2 Auxiliary Model Estimates

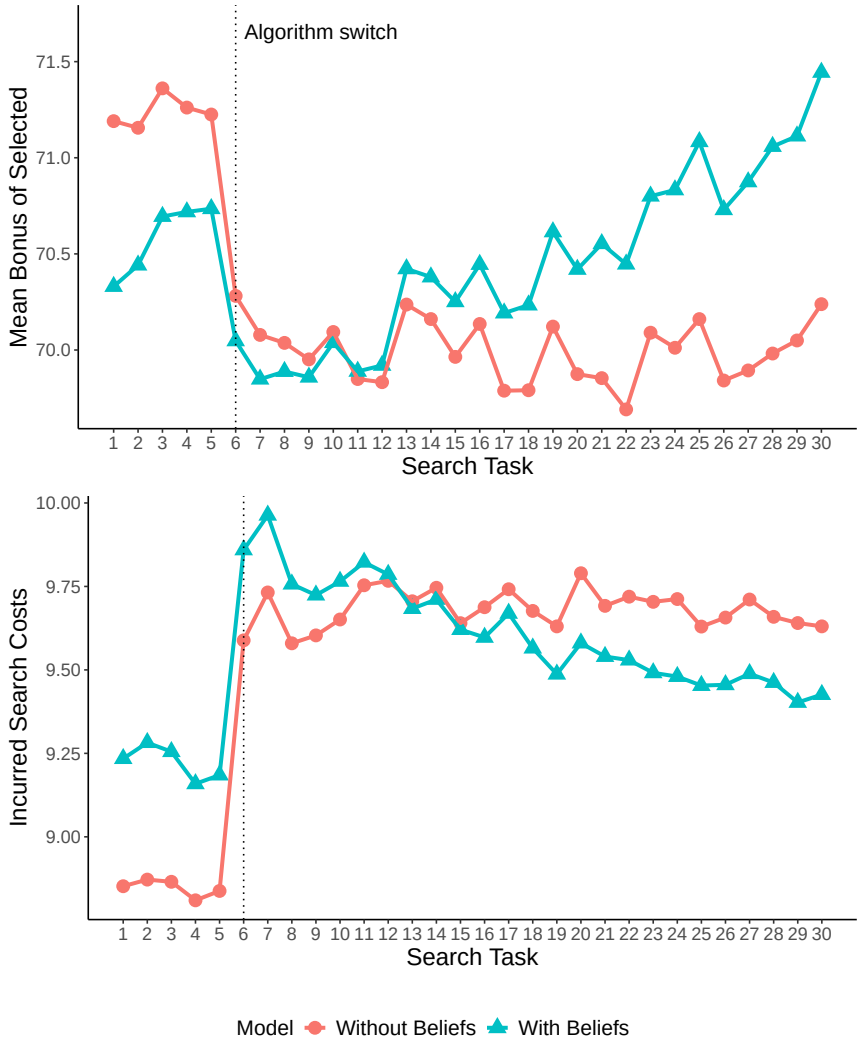
Table H.1: Estimation Results—Auxiliary Models Estimated on Tasks 6-7 (Study 2)

		Estimate (SD)		
Variable	Coefficient	Random Informed	S2R5	S2R5 + Rat. Exp.
Rank Coefficient β^r				
Intercept	Δ_0^r	-0.11 (0.059)*	-0.21 (0.063)***	1.5 (0.073)***
Comp	Δ_{comp}^r	-0.059 (0.062)	-0.05 (0.064)	-0.051 (0.081)
ShopFreq	Δ_{shop}^r	-0.023 (0.047)	-0.052 (0.061)	-0.052 (0.077)
Pos 1 Coefficient β^1				
Intercept	Δ_0^2	2.3 (1)***	3.8 (0.94)***	-0.088 (0.83)
Comp	Δ_{comp}^2	0.34 (0.99)	0.43 (0.94)	-0.61 (0.61)
ShopFreq	Δ_{shop}^2	-0.12 (0.81)	-0.48 (0.95)	0.17 (0.48)
Baseline Search Propensity β^0				
Intercept	Δ_0^0	-23 (1.5)***	-23 (1.5)***	-39 (1.6)***
Comp	Δ_{comp}^0	8.9 (1.6)***	10 (1.7)***	10 (1.8)***
ShopFreq	Δ_{shop}^0	-2.7 (1.4)**	-0.85 (1.5)	-0.73 (1.6)
Heterogeneity				
Rank Coefficient	$\sqrt{V}_{\beta^r}$	0.41 (0.033)	0.5 (0.037)	0.63 (0.052)
Pos1 Coefficient	$\sqrt{V}_{\beta^1}$	4.1 (6.9)	5.3 (9.2)	0.8 (0.77)
Mean Search Propensity	$\sqrt{V}_{\beta^0}$	15 (18)	16 (21)	17 (23)
Reservation Utility	σ_ϵ	11 (8.2)	12 (8.2)	13 (11)

Notes: Table reports posterior means and standard deviations for three models, all estimated only on tasks 6 and 7. Random-Informed refers to estimates using only Random-Informed participants. "S2R5" refers to estimates which use only Strong-to-Random5 participants. These two columns assume that there is no belief component to search. "S2R5 + Strong Beliefs" uses the same Strong-to-Random5 participants and assumes that participants have correct (in expectation) pre-switch beliefs based on the *Strong* algorithm from tasks 1–5. Posterior means and standard deviations are based on the thinned chain which drops the first 1,000 draws of the chain and keeps every tenth draw thereafter. We omit variance of heterogeneity results for brevity. Signif. Codes: $***: 0.01, **: 0.05, *: 0.1$

H.3 Counterfactual Search Behavior

Figure H.1: Bonus and Costs with and without Accounting for Beliefs



Notes: Figure plots the simulated bonus earned (top panel) and total search costs incurred (bottom panel) per consumer for each search task. We calculate the average surplus (bonus chosen - incurred search costs) per consumer per task, given a draw of consumer parameters. We repeat this 200 times. The error bars represent the 95% confidence intervals, which are taken over these iterations. The algorithm switches from *Strong* to *Pref* after task 5.

I Additional Figures and Tables

I.1 Reduced-Form Position Effects

I.1.1 Study 1

The main paper reports the position effects as the probability that a participant searches an item. Position effects are starker when we examine the first product searched. The top-ranked product is significantly more likely to be searched in all conditions (Figure I.1). In Table I.2, we document that these position effects on the first click are strongest in the Strong condition. For participants in all conditions, the position effects on the first click *get larger* over the ten tasks. This is not driven by large changes in the average search depth over tasks (Appendix Figure I.2).

Table I.1: Study 1: Effect of Rank on Search by Condition

Dependent Variable:	1(Product Searched)					
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
r × Condition = Random	-0.0072***	-0.0071***	-0.0072***	-0.0071***	-0.0072***	-0.0071***
	(0.0011)	(0.0011)	(0.0011)	(0.0011)	(0.0011)	(0.0011)
r × Condition = Random Informed	-0.0075***	-0.0073***	-0.0075***	-0.0073***	-0.0075***	-0.0073***
	(0.0011)	(0.0011)	(0.0011)	(0.0011)	(0.0011)	(0.0011)
r × Condition = Strong	-0.0127***	-0.0130***	-0.0127***	-0.0130***	-0.0127***	-0.0130***
	(0.0011)	(0.0011)	(0.0011)	(0.0011)	(0.0011)	(0.0011)
a_j (demeaned)		0.0236***		0.0236***		0.0237***
		(0.0008)		(0.0008)		(0.0008)
Fixed-effects						
Condition (3)	Yes	Yes	Yes	Yes		
Search Task (10)			Yes	Yes	Yes	Yes
Participant (961)					Yes	Yes
Fit statistics						
Observations	96,100	96,100	96,100	96,100	96,100	96,100
R^2	0.00446	0.05726	0.00452	0.05734	0.11900	0.17196
Within R^2	0.00434	0.05715	0.00434	0.05717	0.00491	0.06472

Clustered (Participant) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

The first search is not the only decision impacted by rank. In Table I.3, we condition on products having been searched (which includes more top-ranked items, as in Figure 2). Top-ranked products are more likely to be searched earlier. If we condition on the length of search sequence, for the Random and Strong conditions, the rank effect on search order becomes weaker, highlighting that some of the rank effect in these conditions is coming from the duration of search. However, in all three conditions, lower ranked products are searched later, even if they have a high bonus A.

Table I.2: Effect of Rank on First Search by Condition (Study 1)

Dependent Variable:	1(Product Searched First)			
Model:	(1)	(2)	(3)	(4)
Variables				
$r \times$ Condition = Random	-0.0055***	-0.0054***	-0.0025**	-0.0024**
	(0.0009)	(0.0008)	(0.0010)	(0.0010)
$r \times$ Condition = Random Informed	-0.0050***	-0.0049***	-0.0027***	-0.0025**
	(0.0009)	(0.0009)	(0.0010)	(0.0010)
$r \times$ Condition = Strong	-0.0068***	-0.0070***	-0.0038***	-0.0041***
	(0.0010)	(0.0010)	(0.0011)	(0.0011)
a_j (demeaned)		0.0158***		0.0158***
		(0.0005)		(0.0005)
$r \times t-1 \times$ Condition = Random			-0.0007***	-0.0007***
			(0.0001)	(0.0001)
$r \times t-1 \times$ Condition = Random Informed			-0.0005***	-0.0005***
			(0.0001)	(0.0001)
$r \times t-1 \times$ Condition = Strong			-0.0007***	-0.0006***
			(0.0001)	(0.0001)
Fixed-effects				
Condition (3)	Yes	Yes	Yes	Yes
Search Task (10)			Yes	Yes
Fit statistics				
Observations	96,100	96,100	96,100	96,100
R^2	0.00307	0.04795	0.00337	0.04825
Within R^2	0.00307	0.04795	0.00337	0.04825

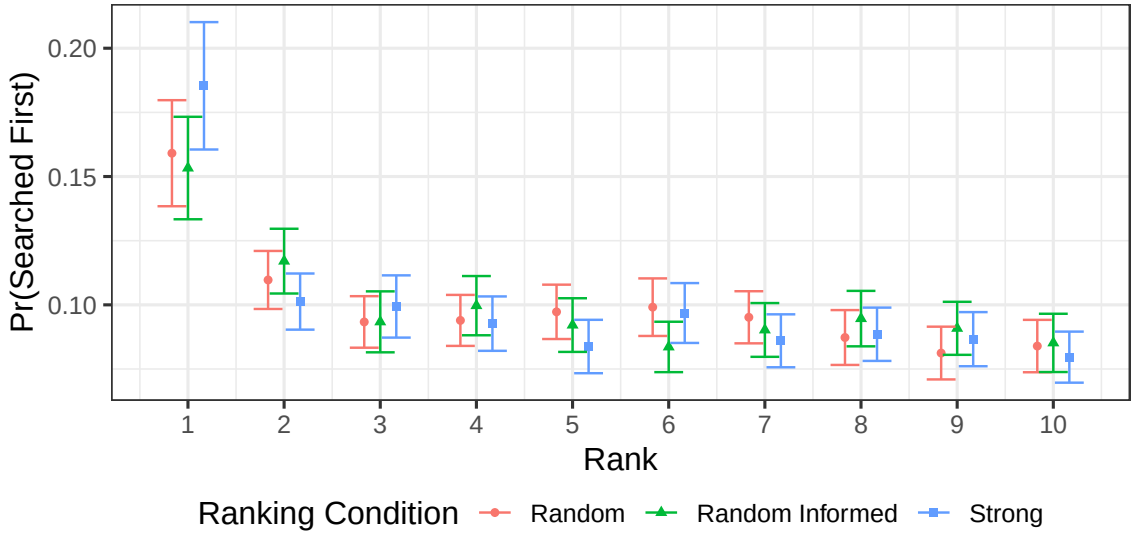
Clustered (Participant) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Table I.3: Effect of Rank on Search Order by Condition

Dependent Variable:	Click Order Given Search							
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Variables								
$r \times$ Condition = Random	0.1405***	0.1427***	0.1403***	0.1425***	0.1578***	0.1582***	0.0940***	0.0949***
	(0.0276)	(0.0272)	(0.0276)	(0.0273)	(0.0279)	(0.0277)	(0.0143)	(0.0142)
$r \times$ Condition = Random Informed	0.0785***	0.0808***	0.0784***	0.0808***	0.0871***	0.0872***	0.0798***	0.0815***
	(0.0209)	(0.0207)	(0.0209)	(0.0207)	(0.0213)	(0.0211)	(0.0121)	(0.0121)
$r \times$ Condition = Strong	0.0974***	0.1061***	0.0973***	0.1059***	0.1117***	0.1150***	0.0764***	0.0782***
	(0.0208)	(0.0209)	(0.0208)	(0.0209)	(0.0209)	(0.0209)	(0.0121)	(0.0122)
a_j (demeaned)		-0.0990***		-0.0991***		-0.0626***		-0.0239***
		(0.0062)		(0.0062)		(0.0053)		(0.0032)
Fixed-effects								
Condition (3)	Yes	Yes	Yes	Yes				
Search Task (10)			Yes	Yes	Yes	Yes	Yes	Yes
Participant (961)					Yes	Yes		
Search Depth (10)							Yes	Yes
Fit statistics								
Observations	20,734	20,734	20,734	20,734	20,734	20,734	20,734	20,734
R^2	0.02957	0.07580	0.02990	0.07617	0.38687	0.40119	0.54808	0.55063
Within R^2	0.02667	0.07304	0.02656	0.07299	0.04913	0.07133	0.03456	0.04000

Clustered (Participant) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Figure I.1: Pr(Search First) by Rank by Condition (Study 1)



Notes: Error bars represent 95% confidence intervals, which are based on standard errors clustered at the participant level.

Figure I.2: Depth of Search across Tasks (Study 1)

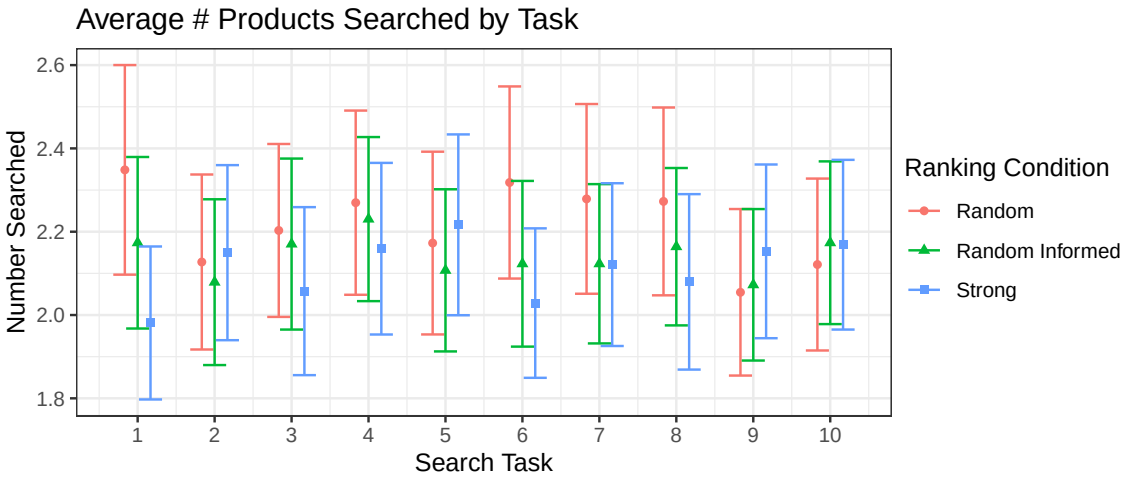
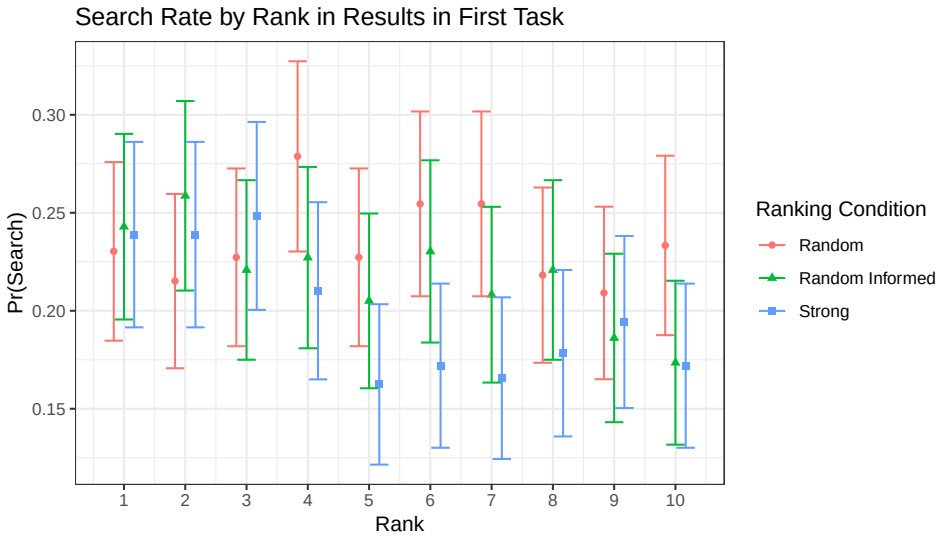


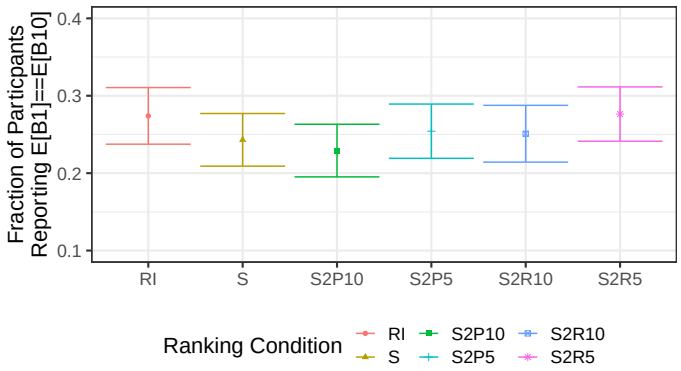
Figure I.3: Search Probability by Rank and Condition for First Task (Study 1)



Notes: Figure shows the probability each item is searched by rank and condition across participants for only the first task. The 95% confidence intervals are based on standard errors clustered at the participant level.

I.1.2 Study 2

Figure I.4: Fraction of Flat Beliefs Reported at End of Tasks (Study 2)



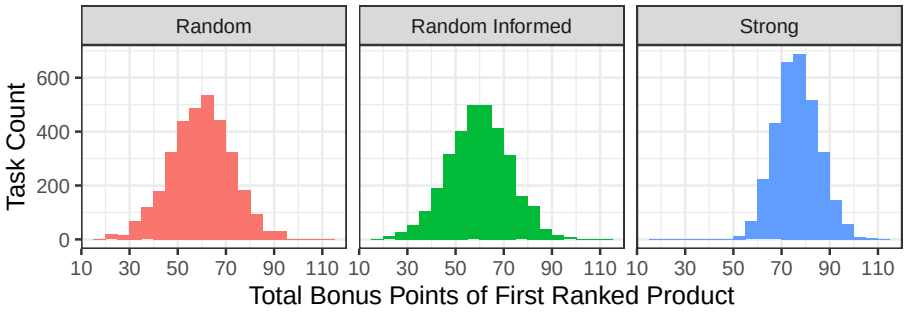
Notes: Error bars represent 95% confidence intervals, which are based on i.i.d. standard errors. "RI" refers to Random-Informed, "S" refers to Strong, "S2P10" and "S2P5" refer to Strong-to-Pref10 and Strong-to-Pref5, respectively, and "S2R10" and "S2R5" refer to Strong-toRandom10 and Strong-to-Random5, respectively.

I.2 Randomization and Manipulation Checks

I.2.1 Study 1

Figure I.5 presents a manipulation check: the first ranked product on average has a higher total bonus in the Strong condition compared to Random and Random-Informed. Figure I.6 shows that the average bonus A does not vary significantly by rank or condition, and the product with the highest bonus A can be found at any rank with similar probability.

Figure I.5: Total Bonus of First-Ranked Products (Study 1)



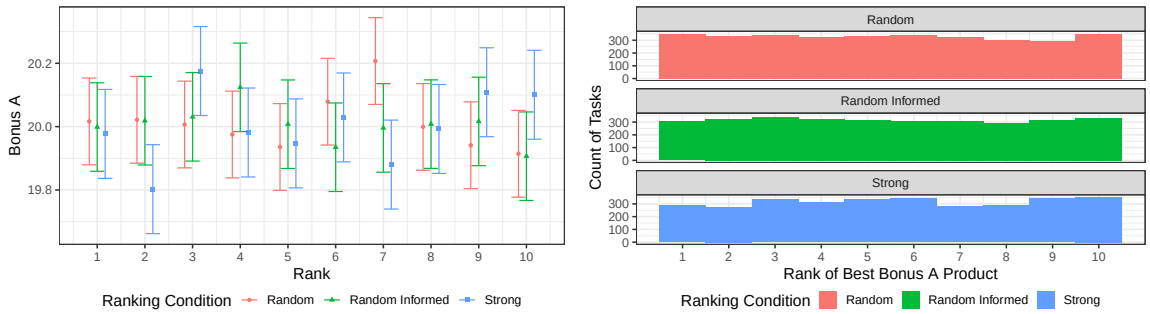
Notes: This figure plots the histograms of the total bonus (bonus A + bonus B) of the first-ranked product for all search tasks and for all participants, for each condition.

Table I.4: Randomization Check—Stimuli Selection (Study 1)

Dependent Variable:	Stimuli Row Number
Model:	(1)
Variables	
Constant	497.3***
	(15.73)
Condition = Random Informed	-21.91
	(22.47)
Condition = Strong	-13.41
	(22.52)
Fit statistics	
Observations	961
R^2	0.00101
Adjusted R^2	-0.00107

IID standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*
Base condition is Random

Figure I.6: Randomization Check of Bonus A (Study 1)



(a) Average Bonus A by Rank

(b) Rank of Best Bonus A

I.2.2 Study 2

Table I.5: Randomization Check—Stimuli Selection (Study 2)

Dependent Variable:	Stimuli Row Number
Model:	(1)
Variables	
Constant	243.6***
	(6.027)
Condition = Strong	10.92
	(8.383)
Condition = Strong2P10	10.02
	(8.466)
Condition = Strong2P5	-0.8277
	(8.448)
Condition = Strong2R10	7.361
	(8.661)
Condition = Strong2R5	1.182
	(8.341)
Fit statistics	
Observations	3,533
R^2	0.00113
Adjusted R^2	-0.00029

IID standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*
Base condition is Random Informed

J Study Instructions

This section provides the instructions given to participants in the Random Informed conditions of Studies 1 and 2. The instructions for the other conditions are the same, except they do not include the last sentence "The products are displayed in a random order."

J.1 Study 1 Instructions

On the following pages, you will complete 1 practice round and 10 bonus-eligible rounds of a product search task.

In each round, you will be presented with a product listing page with 10 product options. You'll need to click on and eventually pick one product. Each product has two bonus values: Bonus A and Bonus B. Your performance bonus will depend on Bonus A and Bonus B of the product that you ultimately pick and how many products you click on.

You can see each product's Bonus A before you click. However, Bonus B is hidden until you click on the product.

Clicking on each new product to reveal Bonus B costs 1 point (clicking to open a box you've already opened previously is free and doesn't cost any points).

Bonus B has an average value of 40 points, and a standard deviation of 12. Put otherwise, 95% of products have their Bonus B between 16 and 64.

The value of Bonus A is not informative of the value of Bonus B. In other words, Bonus A and B are not correlated. Products with high Bonus A can have low Bonus B, and vice versa.

Your earnings towards your bonus for each round will be Bonus A plus Bonus B from the product you selected minus the points you lost from clicking. After all 10 rounds, your total points will be added up over all rounds. Your points will convert directly to your performance bonus. 400 points equals $1.00.

The products are displayed in a random order.

J.2 Study 2 Instructions

On the following pages, you will complete 1 practice round and 15 bonus-eligible rounds of a product search task.

In each round, you will be presented with a product listing page with 20 product options. You'll need to click on one or more products and eventually pick one product. Each product has two bonus values: Bonus A and Bonus B. Your performance bonus will depend on Bonus A and Bonus B of the product that you ultimately pick.

You can see each product's Bonus A before you click. However, Bonus B is hidden until you click on the product.

Clicking on each new product to reveal Bonus B costs 1 point (clicking to open a box you've already opened previously doesn't cost any points).

The value of Bonus A is not informative of the value of Bonus B. In other words, Bonus A and B are not correlated. Products with high Bonus A can have low Bonus B, and vice versa.

Bonus B has an average value of 40 points, and a standard deviation of 12. In most rounds, the highest Bonus B is at least 61.

Because Bonus B has a higher average and standard deviation than Bonus A, Bonus B contributes to a larger share of your payment than Bonus A.

Your earnings towards your bonus for each round will be Bonus A plus Bonus B from the product you selected minus the points you lost from clicking. We will randomly select 8 rounds to generate your additional bonus. Your total points will be added up over these 8 selected rounds and will convert directly to your performance bonus. 500 points equals $1.00.

The products are displayed in a random order.