Preventing Search with Defaults*

Andreas Ortmann[†], Dmitry Ryvkin[‡], Tom Wilkening[§], Jingjing Zhang[¶]

January 21, 2026

Abstract

We explore how firms might use defaults in environments where buyers' search is costly. A *search-preventing default* (SPD) selects a product within a buyer's acceptance set that is profit-maximizing for the firm. SPDs reduce buyers' search costs but distort product choices—benefiting low-cost buyers and harming high-cost buyers. In a lab experiment, we find that SPDs affect buyers' payoffs as predicted, although sub-optimal search leads to welfare losses relative to the theoretical benchmark. In a follow-up experiment, we show that outcomes for high-cost buyers can be improved by an empirically informed default-setting algorithm calibrated on historical data.

Keywords: defaults, costly search, experiment, algorithmic defaults
JEL Codes: C91, D21, D83

*We are grateful to Zachary Breig, Ran Spiegler, Jakub Steiner, the participants of the 2022 UNSW-UQ Economic Theory Festival, the 2022 North-American ESA Meeting, the 2023 ANZWEE, the 2023 and 2024 ARC Workshops on Behavioral Market Design, the 2024 FUR Conference, the 2024 Australian Conference of Economists, and seminar participants at various universities for their insightful comments. We thank Lunzheng Li for programming the experiment in oTree. We gratefully acknowledge the financial support of the Australian Research Council through ARC Discovery Award DP19013475. The experimental protocol was reviewed and approved by the University of Technology Sydney Human Research Ethics Committee.

[†]Department of Economics, UNSW Sydney Business School, a.ortmann@unsw.edu.au.

[‡]Corresponding author: School of Economics, Finance and Marketing, RMIT University, d.ryvkin@gmail.com.

[§]Department of Economics, University of Melbourne, tom.wilkening@unimelb.edu.au.

[¶]School of Economics, Finance and Marketing, RMIT University, jingjing.zhang@rmit.edu.au.

1 Introduction

Defaults have emerged as an important policy tool for nudging individuals into socially beneficial choices (Sanders, Snijders and Hallsworth, 2018; DellaVigna and Linos, 2022). In the hands of benevolent "choice architects," defaults (i) do not restrict the choices of consumers who are informed and active in their decision making, but (ii) help improve the choices made by individuals who find decision making difficult and who make bad (or no) decisions as a result. As such, defaults are seen as light-touch "Libertarian Paternalistic" policies that can improve the outcomes of confused or uninformed individuals while respecting the freedom of others (Thaler and Sunstein, 2003; Sunstein, 2014; Jachimowicz et al., 2019).

Despite defaults being celebrated as an important policy tool, not all choice architects are benevolent. Stores, brokers, salespeople, websites, and shopping platforms regularly curate products as part of their service. As curation often involves the location or sequencing of observed products, there is substantial scope for selecting defaults that nudge consumers into products that are preferred by the firm.

We study defaults theoretically and experimentally in the context of a profit-maximizing firm that can use defaults to strategically influence the initial information acquired by its consumers.[1] Dating back to Simon (1955), it has been posited that many individuals make choices using an item-by-item sequential search process that stops when individuals identify a choice with a sufficiently high level of utility. Such "satisficing" behavior appears to be a common approach to complex decision making (Conlisk, 2003; Bhatia et al., 2021), and has been documented experimentally (e.g., Caplin, Dean and Martin, 2011). We are interested in understanding how defaults might be used to influence the choices made by individuals who exhibit satisficing behavior and the welfare consequences of these defaults. We are also interested in understanding how individuals respond to defaults that might be extractive in nature.

We begin by developing a theoretical model that shows how firms can use defaults to increase their profits when search is costly. Our search model is in the tradition of McCall (1970) and Weitzman (1979) where buyers must search sequentially over a finite set of products to learn their underlying values. When buyers know about the distribution of values and evaluating each new product is costly, optimal search requires a threshold stopping rule so that search continues until a threshold utility is achieved or all objects are

[1]Throughout the paper, we concentrate on the case where the firm can only set the default product but not the entire search sequence of an individual. As discussed below, setting the default is a relatively subtle nudge that is easy to implement and enforce. More aggressive strategies, such as trying to order multiple products, may lead buyers to actively modify their search behavior in response and are likely to have fewer real-world applications.

evaluated.

We show that firms can increase their profits by strategically offering defaults that fall within buyers' acceptance set when values are drawn from a known distribution. These firm-optimal *search-preventing defaults* prevent additional search and push buyers to products that are more profitable for the firm to sell. Buyers purchase products that are worse on average than what they would have bought had they searched on their own in a random order, but economize on search costs. On net, firms will always be better off setting a search-preventing default because they are able to extract the expected rents associated with search. Buyers may be worse off under the search-preventing default if (i) there is a strong misalignment in the ordering of products between buyers and the firm and (ii) if buyers' costs of search are sufficiently high.

Our model provides a rational mechanism by which profit-maximizing firms can use defaults to increase profits. The range of economic decision-making problems where such defaults might be used is wide and includes any setting in which information acquisition is costly and decision making involves a threshold stopping rule. The channel is distinct from the rationales for defaults offered in the literature, which commonly have behavioral underpinnings such as status-quo bias or quasi-hyperbolic discounting (e.g., Choi et al., 2003, 2004; Carroll et al., 2009; Dinner et al., 2011).

Although our theoretical results are compelling, the experimental literature on search (discussed below) has documented that individuals adopt a wide variety of strategies when faced with sequential search and often deviate from optimal search (Bhatia et al., 2021). Therefore, it is important to understand whether there are behavioral responses to defaults that are not captured by our theoretical model. We provide experimental evidence that defaults do indeed influence the behavior of individuals faced with sequential search problems and that the main comparative statics of the model hold. However, we find substantial variation in how individuals respond to search-preventing defaults, and we explore how defaults might be modified to address this heterogeneity.

We report the results of two experiments. The first experiment studies the impact of search-preventing defaults, directly testing the implications of our model. The second experiment explores empirically-informed defaults based on profit maximization, taking into account participants' search strategies in training data.

Participants in our first experiment faced a total of 80 sequential search problems. In each problem, an individual faces a sequence of 10 options with values drawn from a known distribution. The first option—the default—is always revealed, and the rest are initially hidden. Participants can either accept the default option or reveal one of the hidden options at a cost. They can keep searching or stop at any point and accept the best revealed offer so far. We assume extreme misalignment, with the firm's profit determined as a constant less the value of the offer the buyer eventually accepts. In half of the rounds (assigned in a random order), the entire search sequence, including the default, is random, while in the other half we modify the search order by initially offering participants the search-preventing option, i.e., the lowest-valued option in the theoretical acceptance set. In addition to this within-subject variation in default settings, we conduct between-subject treatments varying participants' search costs. We, therefore, used a baseline 2×2 design that allows us to study the comparative statics of our model.

In this baseline 2×2 design, we find, in agreement with our theory, stark negative effects of optimal defaults on buyers' welfare when search costs are high; and small, but statistically significant, positive effects when search costs are low. Our buyers mostly search excessively and earn below the theoretical predictions. Because of these deviations from optimal search behavior, the firm does not benefit from optimal defaults when dealing with low-cost buyers. The firm does, however, benefit from offering optimal defaults to high-cost buyers, albeit to a lesser extent than theoretically predicted.

Our data from the first experiment suggest that buyer search behavior is relatively consistent with theory in the low-cost treatment but that buyers adopt a wide variety of cutoff rules in the high-cost treatment. Some buyers have a cutoff rule near the theoretical benchmark and the firm is able to extract rents from these individuals by introducing the theoretically optimal search-preventing default. However, many buyers have a search threshold that is far above the theoretical prediction, and the search-preventing default results in excessive search for these individuals since they never accept the defaults offered. Buyers in the high-cost treatment on average search more with the search-preventing default than with the random default, and this negatively impacts their earnings (and the firm's profit).

Motivated by the observed sub-optimal search behavior and heterogeneous responses to defaults, we designed a second experiment that explored whether the firm could increase its profits by using the empirical search data to better select its default. For this second experiment, we focus on high-cost buyers and use data from additional random and search-preventing default treatments to estimate the probability that a given default option would be accepted. The estimation is based on a simple logistic choice model for the likelihood that the first and subsequent offers would be accepted by each individual in the training set. Next, we calculate the expected profit of the firm for each potential default, adjusting for the fact that individuals who reject the initial offer are a selected sample who are more likely to continue searching until a high-valued offer is found. Selecting the default that maximizes the empirically estimated expected-profit function yields a default level above the theoretically optimal search-preventing default (SPD) in our high-cost setting. Relative to the SPD treatment, this group-estimated default results in significantly less buyer search and significantly higher firm profits and buyer welfare.

We also explored whether it was possible to offer individual-specific defaults based on each participant's search behavior in the training data. Here, our results are more mixed: while overall profits are higher than in rounds where defaults are random, firms' profits under individually estimated defaults are not significantly different from the search-preventing default treatment. The data suggest that participants learn not to search over time. As such, the efficiency of the search-preventing default improves, while the empirical algorithms set defaults that are slightly too high.

To summarize, the main contributions of our paper are in (i) proposing a simple mechanism through which strategic defaults allow firms to extract surplus from consumers by preventing search; (ii) showing that such defaults can harm some consumers while enhancing welfare for others; (iii) testing these predictions in an experiment; and (iv) examining whether empirically informed defaults, constructed from observed consumer search, can perform better given that consumers search suboptimally. The modeling contributions in (i) and (ii), though simple, are novel in the literature, and the resulting comparative statics with respect to search costs are nontrivial. Testing these predictions in a stylized environment is, therefore, a natural step. Recognizing that consumers often search suboptimally, we then demonstrate that empirically informed (data-driven) defaults can outperform theoretical benchmarks on both profits and buyer welfare.

Importantly, our focus is not on how consumers search, but on the impact of strategically chosen defaults. We find it neither surprising nor novel that our participants search suboptimally. The contribution here is to show that strategic defaults matter, and largely as the theory predicts. Given the generality of the mechanism, this is a robust finding that extends beyond the current setting.

Our framework also allows us to explore the interaction between defaults and information about the objectives of the default setter. In our main experiments, we do not tell participants how defaults are set, which is a reasonable approximation of reality given the light-touch nature of our manipulation. However, in general, there is considerable evidence in the literature (e.g., from guessing games, Bosch-Domènech et al., 2002; Bhatia et al., 2021) that participants do take into account who they interact with. Specific to defaults, Altmann, Falk and Grunewald (2025) explore the extent to which defaults can influence behavior in an environment where defaults provide signals about the state of the world and where the default setter's and decision makers' preferences may be partially misaligned. They show that in this setting, (i) decision makers take their strategic environment into account and (ii) default effects are less pronounced in cases where the interests of the default setter and decision maker are more misaligned. Thus, there is an a priori reason that buyers' awareness of the seller's default-setting algorithm may influence their search behavior.[2]

To evaluate whether awareness of optimal defaults influences behavior, we ran additional treatments—a 2×2 design that paralleled the first experiment—in which we informed participants explicitly in each round whether the default was random or set in a profit-maximizing manner.[3] Our theoretical model predicts that the decision maker accepts the default even if the preferences of parties are known to be fully misaligned. Thus, in contrast to Altmann, Falk and Grunewald (2025), we predict that default-accepting behavior in our setting is not affected by information regarding the objective of the default setter. Consistent with this prediction, we do not find evidence that awareness of the firm's default-setting decision rule affects buyers' behavior. As such, the firm continues to benefit from search-preventing defaults when search costs are high, even when its objective is revealed.

The remainder of the paper is structured as follows. We review the related literature in Section 2 and develop our theoretical model in Section 3. The baseline design and results of our first experiment with search-preventing defaults are discussed in Section 4. The second experiment using empirically-driven defaults is presented in Section 5. The results on the effects of awareness are relegated to Appendix F. Section 6 concludes.

2 Related literature

Our paper is related to the broader literature on defaults and the issues of misaligned incentives and information acquisition. Gigerenzer (2015) provides several examples of potentially misaligned incentives in the health-care environment where practitioners order tests or recommend procedures due to concerns over their own liability (e.g., Studdert et al., 2005) or the revenue the procedures generate (e.g., Gøtzsche and Jørgensen, 2013; Woloshin and Schwartz, 2012; Gigerenzer, Mata and Frank, 2009). In the economics literature, Altmann, Falk and Grunewald (2025) explore how defaults might be used to shape decisions in environments where incentives are potentially misaligned and where defaults convey information about the underlying state of the world. Dobrescu et al. (2016) explore empirically how defaults in pension plans and the type of investment allocation influence wealth accumulation

[2] Brown and Krishna (2004) and Campbell (2007) also find that decision makers alter their response to the default based on their beliefs about the default setter.

[3] As discussed in the design section, we use a computerized seller in all of our experiments so that we can offer theoretically optimal defaults. This design choice is motivated by the observation that many real-world default and pricing decisions are made at the firm level, or by an algorithm, and are hard to attribute to a human agent that the buyer interacts with, if any. Therefore, behavioral motives such as reciprocity and other-regarding preferences that are likely to arise in a lab setting with one-to-one buyer-seller interactions are unlikely to be major factors in evaluating how buyers respond to defaults outside the lab.

in a superannuation setting and highlight suboptimal default settings. Ortmann et al. (2023) explore how defaults might be used to nudge individuals into insurance contracts when the seller has an informational advantage about customers' true risk.

The issue of information acquisition and nudge policies has also been discussed in the economics literature in settings where the choice architect is benevolent. Carlin, Gervais and Manso (2013) explore how defaults offered by a benevolent planner can reduce the information collected by individuals in the context of investment plans. The paper emphasizes how defaults may lead to information externalities when individuals can observe the investment plans selected by others. Caplin and Martin (2017) provide experimental evidence that the introduction of defaults causes decision makers to "drop out" of the decision process and accept the default without deliberation. We add to the literature by explicitly exploring the exploitative use of defaults by profit-maximizing firms and showing how these defaults can reduce consumer welfare even in single-person decision problems where there are no information spillovers.

Our paper is also related to the literature on consumer search (for a review see, e.g., Lippman and McCall, 1976; Ratchford, 2009), and optimal stopping problems more generally (e.g., Rapoport, Seale and Spiliopoulos, 2022). Our model is a variation of early models of sequential search (McCall, 1970; Weitzman, 1979), in which the decision maker searches as long as the expected marginal benefit of one additional search exceeds the marginal search cost, and the optimal stopping rule is based on a reservation value. Parallel strands of literature consider non-sequential, or fixed sample size, search where the decision maker first decides how many options to sample and then chooses the best one among those (Stigler, 1961), and variable sample size search, which is a generalization of the other two (Morgan and Manning, 1985). In combination with price competition, these models have been used to explain price dispersion in markets for homogeneous products (e.g., Stahl, 1989; Burdett and Judd, 1983). A number of empirical studies using natural market search data, including online search, have tried to identify which of the search algorithms better explains behavior, with mixed results (e.g., Hong and Shum, 2006; De los Santos, Hortaçsu and Wildenbeest, 2012).

The early experimental economics literature on search focused mainly on testing sequential search models (Schotter and Braunstein, 1981; Grether, Schwartz and Wilde, 1988) as well as comparing potential search algorithms (Harrison and Morgan, 1990; Davis and Holt, 1996; Sonnemans, 1998; Cason and Friedman, 2003). Generally, the observed behavior in these studies is consistent with predicted comparative statics but deviates from the predictions of optimal sequential search. Subjects tend to search less than predicted (Karle et al., 2025), and the (imputed) reservation value declines with the number of searches, which is consistent with sunk-cost fallacy or elements of a fixed-sample size approach to search. Prices also tend to be lower than predicted, i.e., sellers do not fully exploit their informational advantage.[4] Gabaix et al. (2006) explored search in more complex settings where option values are drawn from different distributions and are obtained by aggregating multiple attributes (in the latter case, the fully optimal search algorithm is too complex to be derived). They find that a boundedly-rational "direct cognition" search heuristic explains subjects' behavior well, and better than the fully optimal solution when it is available.

The introduction of a default implies that options in our environment are not sampled in a uniformly random way. As such, there are links between our paper and the literature on ordered search (Arbatskaya, 2007; Armstrong, 2017). The ordered-search literature explores buyer search behavior and firm competition in environments where buyers use characteristics such as location, advertisement recall, or observed prices to guide the order in which they search over options (Armstrong, 2017).[5] A main interest in this literature is the question of how firms set prices when consumers have heterogeneous valuations (Armstrong, 2017) or search costs (Arbatskaya, 2007; Armstrong, Vickers and Zhou, 2009; Xu, Chen and Whinston, 2010).

Distinct from the ordered-search literature, our setting can be interpreted as the decision problem of a multi-product intermediary who is fully informed about the search cost of the consumer and who can alter the default but not the price-quality characteristics of each offering. This is common in environments where an informed intermediary, such as a broker, is offering products to clients and where the broker may receive product-specific bonuses.[6] Our interest is in understanding how consumer welfare is affected by (search-preventing) defaults and empirically exploring search behavior under those conditions. Similar to the

[4]Similar results have also been obtained by Ryvkin, Serra and Tremewan (2017) and Ryvkin and Serra (2019, 2020) in their experiments on competition among corrupt public officials demanding bribes from citizens who could search for lower bribes across multiple offices, and by Cason and Datta (2006) and Cason and Mago (2010) in their experiments on costly search where sellers could advertise prices (at a cost), and where excessive advertisement was found.

[5]A subset of this literature studies environments where firms strategically order search. For example, Armstrong and Zhou (2011) study cases where competing firms can pay for prominence via commission or advertising while Varian (2007), Athey and Ellison (2011) and Chen and He (2011) study position auctions for websites.

[6]This broker interpretation of our seller is related to the work of Janssen and Williams (2023) who consider consumer search in a market with an influencer recommending products to consumers. The influencer's incentives are not fully aligned with those of consumers, and firms pay the influencer to recommend their product. The influencer samples some number of firms and recommends the one that gives her the highest value, effectively offering a default. Consumers then examine the recommended firm and buy its product if it is good enough, and search otherwise. In the environment considered in Janssen and Williams (2023), it is shown that consumers' welfare is increased by the presence of the influencer even when incentives are misaligned, because it is optimal for the influencer to recommend high-quality products. Our paper studies an environment where defaults can improve or hurt consumer welfare depending on search costs.

existing literature (Wilson, 2010; Armstrong and Zhou, 2011; Armstrong, 2017), our results imply that firms have an incentive to damage the product environment by driving up search costs since this increases the set of potential products that fall within the buyers' acceptance set.

The role of defaults in settings with restricted search is also studied in the behavioral industrial organization literature on the impact of bounded rationality on competition (Spiegler, 2011, 2016). Spiegler (2015) explores how the default policy chosen by regulators influences obfuscation in an environment first studied by Varian (1980) where firms compete on price but can endogenously influence the comparability of choices by either (i) investing in obfuscation (Carlin, 2009; Chioveanu and Zhou, 2013) or (ii) selecting a pricing format that limits comparability (Piccione and Spiegler, 2012). He shows that firms' incentives to obfuscate can be strongly influenced by the default policies set by regulators.

Finally, we note a relation of our model to the model of manipulated attention by Gossner, Steiner and Stewart (2021). In their model, a decision maker receives a sequence of signals about multiple items by focusing attention on, and updating beliefs about, one item at a time. The authors show that for a wide class of attention allocation strategies, i.e., rules by which the decision maker switches attention from one item to the next, manipulating attention by focusing it initially on a particular item unambiguously raises the likelihood of that item being selected.[7] Default setting in our model can then be viewed as an attention manipulation where the seller-preferred (and incentive-compatible) item is placed at the start of the decision maker's consideration sequence.

3 The model

In this section, we propose a simple model of how firms might select defaults when facing buyers whose decision making involves costly search. Our main interest is in cases where the firm is an intermediary, such as a store, website, or broker, that curates products offered by upstream suppliers. As curation often involves the location or sequencing of observed products, the firm in our model has the power to set the default object observed by individuals, but cannot directly affect relative values through prices.

In the main text, we consider the case where the buyer is unaware that the default option is being manipulated. This is a reasonable assumption to make given the relatively subtle intervention of influencing the first object observed through search. In Section 3.1, we

[7] The marketing literature has documented the importance of location in the selection of goods offered by profit-maximizing firms, which is likely related to attention. For example, Chandon et al. (2009) show that changing the location of products influences their demand while Atalay, Bodur and Rasolofoarison (2012) find that horizontal product placement increases demand but does not change beliefs about product quality.

show that it is optimal for the firm to use a search-preventing default in this environment. We further explore the welfare implications of these defaults in Section 3.2 and argue that buyers may be helped or hurt by search-preventing defaults depending on their search costs.

Finally, in Section 3.3 we briefly describe how the model extends to settings where the buyers are aware of how the default is chosen and cases where buyers are risk averse or use alternative search strategies. More formal treatments of these extensions are included in Appendices B-E.

3.1 Search-preventing defaults with unaware buyers

A risk-neutral buyer faces a finite set of $n \geq 2$ options indexed by $i = 1, \ldots, n$, where each option i yields a benefit b_i to her and a profit of Π_i to the firm. The buyer does not observe the benefits initially, but believes they are i.i.d. draws from a known distribution $F(\cdot)$ with support $[\underline{b}, \overline{b}] \subseteq \mathbb{R}_+$. For simplicity, we assume that F is absolutely continuous, with a bounded density $f(\cdot)$.

The buyer's goal is to select one option. She starts a search process by costlessly[8] considering option 1 and finding out about b_1; at this point the buyer may stop and select option 1 as her final choice, or continue searching by sequentially considering other options in a random order. Every additional search costs $c > 0$. We assume perfect recall: The buyer remembers the benefits of all previously encountered options and can costlessly go back to select any of them. The buyer must select one option eventually.

The setup follows a standard sequential search model (see, e.g., McCall, 1970; Weitzman, 1979). The optimal search strategy for the buyer is to search until the best available option has a value greater than or equal to a reservation value, $z^*(c)$, determined by the condition

$$\int_z^{\overline{b}} (b - z) \mathrm{d}F(b) = c. \tag{1}$$

The left-hand side of (1), which can be written as $\mathbb{E}(\max\{b - z, 0\})$, is the expected benefit of one additional search given the best option encountered so far is $z \in [\underline{b}, \overline{b}]$, and the right-hand side is the cost of one additional search.

It is easy to see that the left-hand side of (1) is a strictly decreasing function of z that is $\overline{c} \equiv \mathbb{E}(b) - \underline{b}$ when $z = \underline{b}$ and 0 when $z = \overline{b}$. Thus, a nontrivial reservation value will exist, and the buyer will search with positive probability, if $c \in (0, \overline{c})$, which is an assumption we maintain throughout. It also follows that $z^*(c)$ is a strictly decreasing function of c in this range.

[8] It could be argued that considering option 1 should also be costly; however, in our setting the buyer has to select one option eventually, and hence treating option 1 as costless is just a normalization.

We consider two settings. In the first one, referred to as *random default* (RD), the ordering of options is fully random. Let $A = \{i : b_i \geq z^*(c)\}$ denote the buyer's (possibly empty) set of *acceptable* options with benefits above the reservation value. The buyer's optimal strategy is to continue search until he finds an element in A and then to stop. If $A = \varnothing$, the buyer will search through all n objects and then select the object with the highest value in the offer set.

Next, consider the case of the firm-optimal *search-preventing default* (SPD) where the firm can observe the realization of the benefits $(b_1, \ldots, b_n)$ and profits $(\Pi_1, \ldots, \Pi_n)$ of each offer and can, unbeknowst to the buyer, designate any of the options as option 1—the starting point of search. Since the buyer is unaware that the default is manipulated, she will accept any offer in A without further search and reject anything outside of A. It follows immediately that the firm's profit-maximizing choice is to select the option in A that maximizes the firm's value.

Proposition 1 *If $A \neq \varnothing$ and the buyer is unaware, the search-preventing default $d \in \arg\max_{i \in A} \Pi_i$ is profit-maximizing.*

3.2 Comparative statics

We now explore the consequences of setting the search-preventing default. There are two competing effects for the buyer's welfare: The default helps the buyer save on search costs but precludes finding better than seller-optimal acceptable options. The latter effect becomes larger the more the seller's profit and the buyer's benefit are misaligned. To simplify matters, we consider an extreme case of fully misaligned incentives.[9]

Assumption 1 *All options $i = 1, \ldots, n$ have a fixed total surplus S such that $b_i + \Pi_i = S$.*

Under the fixed surplus assumption, the seller would like the buyer to choose the option with the lowest possible benefit; therefore, the optimal search-preventing default is simply $d \in \arg\min_{i \in A} b_i$. We can write the difference in the buyer's payoffs between the search-preventing default and random default as

$$\pi_{\text{SPD}} - \pi_{\text{RD}} = \sum_{k=1}^{n} \binom{n}{k} F(z^*(c))^{n-k}[1 - F(z^*(c))]^k \left[cs_k - \int_{z^*(c)}^{\bar{b}} (\tilde{F}_{(1:k)}(b) - \tilde{F}(b)) \mathrm{d}b \right]. \quad (2)$$

[9]The mechanism we identify here exists in all cases where incentives are at least partially misaligned, although, of course, the effect is strongest under full misalignment. Strong misalignment emerges, for example, in an insurance setting where the insurer cannot fully price discriminate, such as employer-provided group health insurance.

Here, summation goes over the number of acceptable options $k = |A|$, and s_k is the expected number of searches under RD conditional on k (see Appendix A for details). Equation (2) shows explicitly the two competing effects of the search-preventing default described previously. Both effects are zero in the absence of search costs.

The first effect is the savings on search costs. This effect is always small because when c is low, there is not much to save. On the other hand, when c is high, $F(z^*(c))$ is low, and hence the terms with larger k contribute more to the sum in (2); however, s_k—the expected number of searches—declines with k.

The second effect is the loss in expected benefits due to the optimal choice of default. It arises due to first-order stochastic dominance between $\tilde{b}$—a random benefit above the reservation value, and $\tilde{b}_{(1:k)}$—the lowest order statistic among k draws of such benefits. This effect is increasing in c. Indeed, as mentioned above, $F(z^*(c))$ falls with c and hence the terms with higher values of k contribute more to the sum in (2). But larger values of k produce stochastically larger differences in benefits $\tilde{b} - \tilde{b}_{(1:k)}$.

In general, while the difference $\pi_{\text{SPD}} - \pi_{\text{RD}}$ may depend on c in a complicated way determined by the shape of F, we show that it is always positive for low enough c and negative for high enough c. Formally, the result is as follows.[10]

Proposition 2 *When Assumption 1 holds, there exist search costs $c_L, c_H \in (0, \bar{c})$ such that $\pi_{\text{SPD}} < \pi_{\text{RD}}$ for all $c \in (c_H, \bar{c}]$ and $\pi_{\text{SPD}} > \pi_{\text{RD}}$ for all $c \in (0, c_L)$.*

Obviously, the seller is always weakly better off in offering a default since the seller can select a profit-maximizing option in A. Moreover, if $|A| > 1$ with positive probability, the seller is strictly better off in expectation. Further, when total surplus is fixed and preferences are fully misaligned, the seller's additional profit from optimal default setting is strictly increasing in c. This can be seen by noting that the extra surplus extracted by the seller is exactly the second term in (2), and this term is strictly increasing in c as explained above.

Corollary 1 *When Assumption 1 holds, the seller's additional profit from optimal default setting is strictly increasing in c.*

3.3 Model robustness

Our main analysis assumed that buyers were unaware of the way defaults were set, risk neutral, and searched sequentially. We analyze the consequences of relaxing these assumptions

[10] We note that there is no universal characterization of what happens for intermediate values of c. In particular, $\pi_{\text{SPD}} - \pi_{\text{RD}}$ can cross zero multiple times. An example is available upon request.

in Appendices B-E. In Appendix B, we discuss additional issues that arise when buyers are aware how the default is being set by the firm. We first argue that if the firm continues to use a search-preventing default, the cutoff selected by the firm is not affected by buyers' awareness for a wide class of value distributions (a formal result on buyers' belief updating is presented in Appendix C). Therefore, the only substantive change that arises with awareness has to do with buyers' search behavior after observing a default outside the acceptable set. When the number of options is large, these events are rare, and more complicated algorithms that seek to exploit awareness can generate only minimal (asymptotically vanishing) additional profits for the firm.

We also consider the case of risk averse buyers in Appendix D. While a general characterization is too complex, our exploratory calculations for CARA utility show that under the parameters of our experiment, the predictions of Proposition 2 are qualitatively unchanged. One caveat is that the magnitude of the positive effect of the SPD for low-cost buyers declines with risk aversion. This is expected because more risk averse buyers search less, and hence their gain from savings on search costs under SPD is lower.

Finally, we acknowledge that sequential search is not the only possible search algorithm, and it may be of interest to what extent the results of this section are robust to variation in search strategies. In Appendix E, we consider an alternative model—fixed sample size search, where buyers observe the default option, b_1, and commit to a fixed number of options $s(b_1)$ they would like to reveal, at cost $cs(b_1)$. We show the existence of a generically unique optimal strategy, $s^*(b_1)$, and find that under the optimal search-preventing default identified in this section (i.e., based on the reservation value $z^*(c)$), a result similar to Proposition 2 continues to hold. The main reason is that $s^*(b_1) = 0$ for any $b_1 > z^*(c)$, i.e., such defaults shut down fixed sample size search as well.

4 Experiment 1: Search-preventing defaults

Our theory framework suggests that defaults can be used by sellers to prevent additional search and guide consumers to products that are optimal for the firm to sell. Consumers buy products that are worse on average than what they would have received had they searched on their own in a random order, but they economize on search costs. These competing forces imply that consumers can be better or worse off under a search-preventing default depending on search costs.

Although these theoretical results are compelling, they rely on the assumption that buyers optimally respond to the default and are deterred from searching. Given that the empirical literature, using both controlled experiments and observational data, has documented a wide variety of strategies individuals adopt when faced with sequential search, it is important to understand whether there are behavioral responses to defaults that are not captured by our theoretical model and whether the main predictions of our theory are likely to hold. We explore these issues with a lab experiment.

As a framework for testing, we use individual choice problems where a human buyer faces an automated, computerized seller. This environment allows us to offer the theoretically predicted search-preventing default. Our design does not explore whether *human* sellers are able to extract rents by correctly choosing defaults, nor does it explore behavioral forces that may arise due to interactions between buyers and sellers. We note that in many real-life settings, defaults (and prices) are set at a high level in an organization or by algorithms. In these settings, the strategic decisions of the firm are made at a distance and reciprocity and other-regarding preferences—forces common in laboratory settings with interactions between human buyers and sellers—may not be significant considerations. As such, we believe that our simple setting captures key aspects of the decision problem faced by buyers outside the lab.

4.1 Treatments

We implemented a 2×2 design in which we generated within-subject variation in the way that defaults were assigned and between-subject variation in the cost of search. Across sessions, individuals were assigned to either a **High-Cost (HC)** environment or a **Low-Cost (LC)** environment. Within a session, participants played a total of 80 rounds, 40 of which were randomly assigned to a **Random-Default (RD)** round and the remaining 40 rounds were assigned to a **Search-Preventing-Default (SPD)** round.

For each subject, we used a pre-drawn 80-round sequence of decision problems. To minimize noise, we generated 30 different sequences for 30 subjects and re-used them across experimental sessions. Each decision problem was generated in two steps. First, we generated 10 "cards" by randomly drawing integer values uniformly on $\{0, 1, \ldots, 100\}$. Next, we assigned one of the cards to be the initial offer and arranged the remaining 9 cards in a random sequence. The way the initial offer card was selected differed by treatment. In RD rounds, each card had an equal chance of being assigned as the initial offer. In SPD rounds, we assigned the card with the lowest value in the interval $[z^*(c), 100]$ to be the initial offer.[11]

Participants were endowed with 200 tokens in each round and were shown the value

[11] Following the model with unawareness, we randomly selected a default if the acceptable offer set was empty. The probability of this occurring for a given participant in any given round is $F(z^*(c))^n = 3.39 \times 10^{-7}$ in the HC environment and 0.0223 in the LC environment. The latter implies on average $0.0223 \times 9 \approx 0.2$ searches in SPD rounds, see Table 1.

Table 1: Parameters, theoretical predictions, and number of subjects in Experiment 1

Treatments	Search Cost, c	Cutoff, $z^*(c)$	# of Searches	Buyer's Payoff	Seller's Payoff	Subjects
LC-RD	5	68.4	2.09	73.2	16.3	30
HC-RD	30	22.5	0.29	52.5	38.7	30
LC-SPD	5	68.4	0.20	75.5	23.5	30
HC-SPD	30	22.5	0.00	31.6	68.4	30

Values in the experiment are drawn from a uniform distribution on $[0, 100]$ in each treatment cell, and the total surplus is set to $S = 100$ tokens.

of the card containing the initial offer. They then had the option of "accepting the best offer" and ending the round or paying to reveal the next card in the sequence. Revealing an additional offer incurred a fixed cost of 5 tokens in the LC environment and 30 tokens in the HC environment. The participant could continue to reveal cards until either all cards were revealed or they chose to accept the best offer.

Throughout a round, the computer automatically highlighted the card containing the best offer revealed so far and the value of this card was prominently displayed. Subjects also were shown the total cost of previous searches and the payoff that the subject would receive if they accepted the best offer at that point.

To implement unawareness, participants were told that the first card would be chosen as the initial offer. However, we did not explicitly tell participants how this card was chosen nor alert participants as to whether they were in a RD round or a SPD round. The 2×2 treatment cells in our design are denoted as **HC-RD**, **HC-SPD**, **LC-RD**, and **LC-SPD**.

Parameters of the experiment and theoretical predictions are summarized in Table 1. We set total surplus to $S = 100$ and follow Assumption 1 that buyers' and sellers' incentives are fully misaligned, i.e., if a buyer chooses an option with benefit b, the seller's profit is $100 - b$. As noted above, the distribution of benefits is uniform on $\{0, 1, \ldots, 100\}$ and search costs are set to $c = 30$ and $c = 5$ in the HC and LC environments, respectively.

We pre-registered Experiment 1 on the Wharton Credibility Lab AsPredicted platform under #96405. Prior to this registration, we conducted two pilot sessions, one for HC with 11 participants and one for LC with 13 participants, and performed a power analysis for the difference in buyers' payoffs between the RD and SPD rounds in the two treatments. We obtained Cohen's $d = 2.04$ in HC and $d = 0.72$ in LC, indicating "huge" and "medium" effect sizes, respectively, following the classification of Sawilowsky (2009). The corresponding number of subjects based on a power calculation for paired data with $\alpha = 0.05$ and $\beta = 0.1$ produced $N = 5$ and $N = 23$ for HC and LC, respectively. Therefore, conservatively, we recruited a sample size of 30 subjects per treatment cell in the baseline treatments.

4.2 Experimental protocols

All of the sessions were run online with participants from the pool of pre-registered student subjects at the Behavioural Lab of the University of Technology Sydney in May of 2022. The experiment was implemented using oTree (Chen, Schonger and Wickens, 2016). A total of 60 subjects participated in 7 sessions. All participants were recruited from the subject pool using ORSEE (Greiner, 2015).

Participants were provided with instructions (see Appendix H) which described the search game and included a quiz that required the individual to identify the best revealed offer and the cost of search from a hypothetical round where the participant searched three additional times. After successfully completing the quiz, participants then went through the 80 rounds of the search game.

After the main experiment, subjects participated in three additional lottery tasks where participants' risk aversion, loss aversion, and ambiguity aversion were elicited using list methods. In each of these tasks, participants were presented with a list of 21 choices between a lottery and a sure amount of money. Each list was constructed so that if a participant's preference respected monotonicity, there was a unique point at which they were willing to switch from the lottery to the sure amount. In the risk task, individuals chose between a lottery $(0, \$2.00; 0.5, 0.5)$ and a sure amounts of money that increased from zero to $2.00 in 10 cent increments. In the loss task, the lotteries were $(-\$x, \$2.00; 0.5, 0.5)$, where x changed from 0 to 2.00 in 10 cent increments, and the sure amount of money was always 0. Finally, in the ambiguity task the lottery was $(0, \$2.00; p, 1-p)$, where, unbeknownst to subjects, p was generated randomly from the uniform distribution on $[0, 1]$, and the sure amounts were the same as in the risk task. The three tasks were presented to participants in a random order, without feedback, and one of them was randomly selected for actual payment.

We randomly selected 4 rounds from the search game and one of the three lottery tasks for payment. Participants were informed about this earnings mode. Participants were paid AU$1 for every 100 tokens they earned in those 4 rounds as well as an additional $5 show-up payment. On average, participants spent about 35 minutes completing the experiment and earned $15.86.

4.3 Hypotheses

The following hypotheses are based on the model predictions summarized in Table 1. We formulate hypotheses in terms of comparative statics, rather than point predictions, because participants are likely to deviate from payoff-maximizing search behavior, and our interest is in how robust the model predictions are to such deviations.

Our experimental design generates between-subject variation in costs (LC or HC) and within-subject variation in defaults (RD or SPD). We first describe our hypotheses related to the cost dimension since strong deviations in this dimension would make it more difficult to understand how behavior is influenced by the default.

Hypothesis 1 *Search is lower in the HC environment than in the LC environment in RD rounds but not in SPD rounds. Buyers' payoffs are decreasing in costs regardless of how the default is set. Likewise, sellers' profits are increasing in costs regardless of how the default is set.*

The first part of this hypothesis follows the comparative statics of search behavior. For participants using the reservation value stopping rule, the lower bound of the acceptance set falls as costs increase. As such, we would predict more search when defaults are random and costs are low relative to the case where defaults are random and costs are high. In search-preventing-default rounds, the default was set to prevent search in both the high-cost and low-cost treatments. Thus there is no predicted difference in search except for the rare cases where all options happen to be below the reservation value.[12]

An increase in search costs will influence the default offered in SPD rounds and reduce search in RD rounds. As both of these responses will lead to lower-valued options being selected, buyers' payoffs are predicted to be decreasing in search costs, and sellers' payoffs are predicted to be increasing in search costs.[13]

We next turn to our comparative static predictions with relation to the way that defaults are set. We have three hypotheses related to the comparison of the Random-Default treatments and Search-Preventing-Default treatments.

Hypothesis 2 *Search in SPD rounds is lower than in RD rounds in both the LC and HC environments.*

Hypothesis 3 *Buyers' payoffs are lower in SPD rounds than in RD rounds in the HC environment. In the LC environment, the opposite holds.*

Hypothesis 4 *Firm's profits are higher in SPD rounds than in RD rounds in both the LC environment and the HC environment.*

[12] This occurs in 2% of cases in the LC environment and never occurs in the HC environment. These proportions are similar to the theoretical proportions discussed in footnote 11.

[13] The expected search cost paid in the low-cost random-default rounds is 10.46 while the expected search cost paid in the high-cost random-default rounds is 8.73. As such, the difference in search costs is small and the comparative statics in buyers' payoffs is driven primarily by lower-valued options being selected in the high-cost environment.

Hypothesis 2 follows directly from the choice of the search-preventing default, which is predicted to always prevent further search except in a small number of cases under LC. Hypothesis 3 is based on our choice of parameters for the high-cost and low-cost environments. We selected $c = 5$ for the low-cost environment since this cost was in the range where buyers benefit from search-preventing defaults, and $c = 30$ for the high-cost environment since this is in the range of costs where the buyers are predicted to be worse off. The former effect is always relatively small, as discussed in Section 3.[14]

Hypothesis 4 is due to the assumption that buyer and seller preferences are perfectly misaligned and the fact that the search-preventing default is selected to be the lowest option that is in the buyer's acceptance set. Since this is always lower than the average acceptable option in expectation, the seller is predicted to be better off under the search-preventing default regardless of the cost of search.

4.4 Results

We start our analysis with an overview of the data, a comparison of the treatment outcomes to their point predictions, and the comparative statics related to costs. Table 2 reports summary statistics for each of the four treatment cells alongside the theoretical point predictions, which are reproduced from Table 1 for convenience.

One general observation across all treatments is that participants search too much as compared to the average optimal number of searches. The difference is statistically significant in all treatment conditions and, importantly, participants search substantially in rounds with search-preventing defaults. This is in contrast to the theoretical prediction that search should only occur in the search-preventing default rounds when all options fall below the reservation value (which, as noted in footnote 11, is very unlikely).[15] The deviations from optimal search behavior have consequences for participants' payoffs and firms' profits, with both the buyers' payoffs and the firms' profits generally lower than predicted under optimal search.

Despite there being too much search, the comparative statics with respect to costs are all strongly consistent with theory.

Result 1 *Consistent with Hypothesis 1, subjects search less in the HC environment than in the LC environment in RD rounds, but in SPD rounds there is no difference. Buyers'*

[14] As discussed at the end of Section 4.1, we based our power analysis on pilot sessions comparing buyers' payoffs between RD and SPD rounds in the HC and LC environments. Despite the predicted difference being small in the LC environment, the effect size in the pilot is large enough for our experiment to be adequately powered.

[15] These deviations may reflect that (i) "errors" in search behavior are truncated and show up in one direction only; and (ii) participants have a preference to do something rather than nothing.

Table 2: Summary statistics and theoretical predictions in Experiment 1

Treatments	# of searches		Buyer's payoff		Seller's profit		Subjects
	data	theory	data	theory	data	theory	
LC-RD	2.51	2.09	68.84	73.24	18.60	16.30	30
	(0.22)		(0.81)		(1.02)		
HC-RD	1.08	0.29	37.08	52.54	30.54	38.73	30
	(0.12)		(2.64)		(1.57)		
LC-SPD	1.76	0.20	72.61	75.48	18.61	23.52	30
	(0.28)		(0.84)		(0.82)		
HC-SPD	1.51	0.00	16.74	31.63	38.01	68.37	30
	(0.14)		(2.42)		(2.35)		

Robust standard errors clustered by subject in parentheses.

profits are lower in the HC environment than in the LC environment, while sellers' profits are higher in the HC environment than in the LC environment.

Support for the first part of Result 1 is based on nonparametric Wilcoxon rank-sum tests shown in Panel A of Table 3. To perform these tests, we calculated the average number of searches in RD rounds and in SPD rounds for each subject and treated the two treatments as separate samples[16] As seen in the panel, the search is significantly lower in the HC treatment than in the LC treatment for RD rounds, but there is no significant decrease in search in SPD rounds.

We use the same nonparametric Wilcoxon rank-sum test to assess buyers' payoffs and sellers' profits. As seen in Panel B, buyers' profits are significantly lower in all four pair-wise comparisons of the LC environment and the HC environment. Likewise, as seen in Panel C, the sellers' profits are significantly higher in all four pair-wise treatment comparisons.

Having established that the comparative statics with respect to costs are consistent with theory, we now turn our attention to search under random and search-preventing defaults. We find mixed support for our second hypothesis.

Result 2 *Consistent with Hypothesis 2, subjects search more in RD rounds than SPD rounds in the LC environment. However, in contrast to Hypothesis 2, the opposite is true in the HC environment.*

[16] We, however, do not compare RD rounds and SPD rounds to each other at this point; rather, we compare RD rounds to RD rounds and SPD rounds to SPD rounds across the two cost environments, i.e., between subjects.

Table 3: Statistical Tests for Hypothesis 1

Treatments Compared	z-score	p-value	N_1	N_2	Statistical Test
Panel A: Comparison of Search by Search Cost					
LC-RD vs HC-RD	4.89	< 0.01	30	30	Rank-Sum
LC-SPD vs HC-SPD	0.07	0.95	30	30	Rank-Sum
Panel B: Comparison of Buyers' Payoffs by Search Cost					
Treatments Compared	z-score	p-value	N_1	N_2	Statistical Test
LC-RD vs HC-RD	6.65	< 0.01	30	30	Rank-Sum
LC-SPD vs HC-SPD	6.65	< 0.01	30	30	Rank-Sum
Panel C: Comparison of Sellers' Profits by Search Cost					
Treatments Compared	z-score	p-value	N_1	N_2	Statistical Test
LC-RD vs HC-RD	-5.10	< 0.01	30	30	Rank-Sum
LC-SPD vs HC-SPD	-6.17	< 0.01	30	30	Rank-Sum

Statistical tests for Hypothesis 1. Panel A is a comparison of search; Panel B is a comparison of buyers' payoffs; Panel C is a comparison of sellers' profits. All tests are a nonparametric Wilcoxon rank-sum test with the two treatments treated as independent samples.

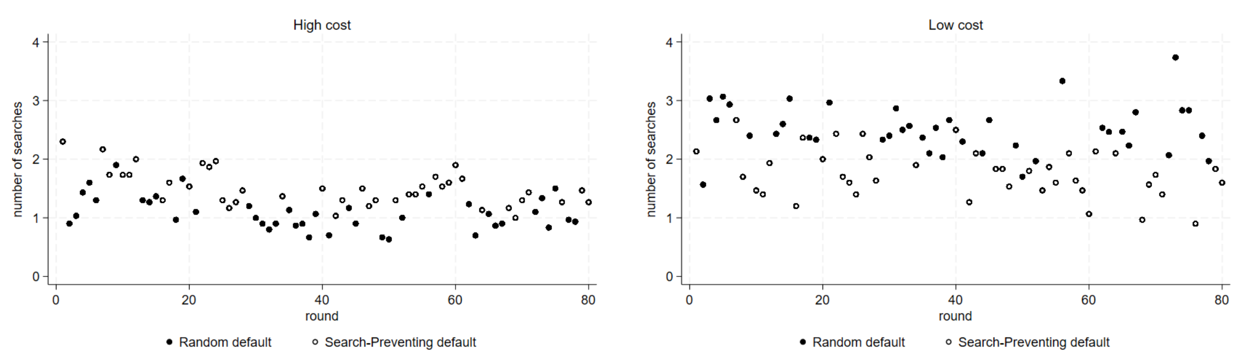

Figure 1: The average number of searches per round, by treatment.

Figure 1 shows the average number of searches per round in each treatment. As seen from the figure, when the search cost is high, subjects consistently search more in rounds with search-preventing default than in rounds with random default; however, the opposite holds when the search cost is low.[17] As seen in Panel A of Table 4, the differences are significant based on the nonparametric Wilcoxon matched-pairs signed-rank test where we calculated the average number of searches in RD rounds and in SPD rounds for each subject.

Since our theoretically predicted cutoffs are payoff-maximizing, the excessive search observed in the SPD rounds leads to a decrease in payoffs relative to the theoretical predictions, especially in the HC treatment. Interestingly, it does not invalidate the predicted reversal of payoff comparisons between SPD and RD rounds under high and low costs.

Result 3 *Consistent with Hypothesis 3, buyers' payoffs are lower in SPD rounds than in RD rounds in the HC environment. In the LC environment, the opposite holds.*

Support for Result 3 is based on the same nonparametric Wilcoxon matched-pairs signed-rank test that we used above with buyers' payoffs as the variable of interest. As seen in Panel B of Table 4, buyers' payoffs are significantly higher in SPD rounds than RD rounds in the LC treatment and significantly lower in SPD rounds than RD rounds in the HC treatment.

Finally, we look at sellers' profits. In both the HC environment and LC environment, buyers tend to search excessively in SPD rounds, which tends to reduce overall seller profits. In the HC environment, excessive search also occurs in RD rounds. On net, however, the seller is better off in SPD rounds relative to RD rounds since a subset of buyers accept the search-preventing default. In the LC environment, because defaults in SPD rounds have high values to begin with, the excessive search observed in SPD rounds results in similar quality products being purchased in RD and SPD rounds. Consequently, there is no significant difference in seller profits between RD rounds and SPD rounds in the LC environment. These findings are summarized in the following result, and the corresponding tests are shown in Panel C of Table 4.

Result 4 *Consistent with Hypothesis 4, seller's payoffs are higher in SPD rounds than in RD rounds in the HC environment. However, there is no significant difference in sellers' profits between SPD and RD rounds in the LC environment.*

[17] We note that there is a small downward trend in search intensity in the high-cost treatments (in pooled OLS regressions of the number of searches on round number, the slope estimates are -0.0052, $p = 0.011$, in HC), but not in the low-cost treatments. This behavior is transient: None of the treatments has a significant time trend when the first 10 rounds of the data are excluded. Following our pre-analysis plans, we use data from all rounds in all of the analysis. Excluding the first 10 rounds does not materially affect any of our results.

Table 4: Statistical Tests for Hypotheses 2-4

Panel A: Comparison of Search by Default				
Treatments Compared	z-score	p-value	N	Statistical Test
LC-RD vs LC-SPD	3.64	< 0.01	30	Sign-rank test
HC-RD vs HC-SPD	-3.91	< 0.01	30	Sign-rank test
Panel B: Comparison of Buyers' Payoffs by Default				
Treatments Compared	z-score	p-value	N	Statistical Test
LC-RD vs LC-SPD	-3.59	< 0.01	30	Sign-rank test
HC-RD vs HC-SPD	4.68	< 0.01	30	Sign-rank test
Panel C: Comparison of Sellers' Profits by Default				
Treatments Compared	z-score	p-value	N	Statistical Test
LC-RD vs LC-SPD	-0.46	0.64	30	Sign-rank test
HC-RD vs HC-SPD	-4.52	< 0.01	30	Sign-rank test

Statistical tests for Hypotheses 2-4. Panel A is a comparison of search by default; Panel B is a comparison of buyers' payoffs by default; Panel C is a comparison of sellers' profits by default. All tests are a nonparametric Wilcoxon matched pairs signed rank test with RD rounds paired to SPD rounds by individual buyer.

Taken together, our data suggest that buyers' behavior is partially consistent with theory and that the comparative statics of the model hold. However, we document more search by buyers than predicted by theory, particularly in the HC treatment with the search-preventing default. This is likely due to many subjects viewing the search-preventing defaults under HC (with $z^* = 22.5$) as too low to accept, even though from the payoff maximization point of view it is optimal for them to do so. In the next section, we analyze subjects' search behavior in more detail.

4.5 Empirical cutoff analysis

To better understand the deviations in search behavior in the HC environment of Experiment 1, we explored the cutoff rules employed by individual subjects in an ex-post analysis. For each subject, we plotted the distribution of values that were accepted in each of the 80 decision rounds. These values are an upper bound of what an individual's reservation value is since an individual is likely to be offered a value that is above their cutoff in each round. However, by finding a low quantile of this distribution, we can approximate a cutoff rule for each individual.

The left-hand side of Figure 2 shows the distribution of cutoff rules when we use the 10th percentile of accepted values to accommodate a small amount of noise. As seen, buyers have heterogeneous cutoffs in the high-cost environment, and many use a cutoff rule that is between 30 and 50. By contrast, most buyers in the low-cost environment have a cutoff near the theoretical value of 68.4.

We also estimated the cutoff used by each buyer via a local logistic regression analysis. For each individual, we broke each decision problem where a decision maker searched $s-1$ times into a series of s binary decision problems where the decision maker rejected the best available offer in the first $s-1$ decisions and accepted the best available offer for the s-th decision. Next, using a local logistic regression, we estimated the point where the decision maker was indifferent between accepting and rejecting the offer and used this point as the estimated cutoff.[18] The right-hand side of Figure 2 shows the distribution of these estimated cutoff rules. Similar to the simple analysis, search cutoffs in the low-cost treatments are centered close to the theoretical cutoff and narrowly dispersed. By contrast, the cutoff for the high-cost treatments is widely dispersed, with a mean close to 50.[19]

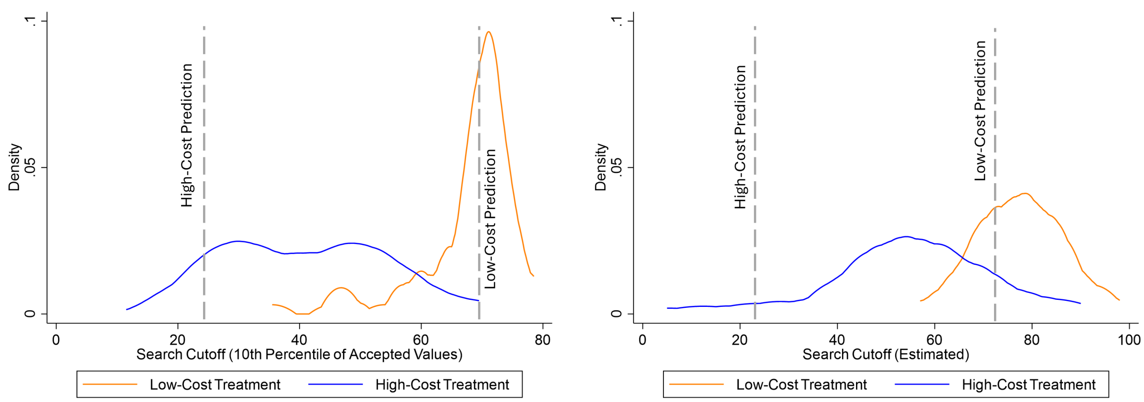

Figure 2: Epanechnikov kernel plots of the search cutoffs employed by buyers in HC treatments of Experiment 1. Optimal bandwidth is used. The left panel uses the the 10th percentile of accepted offers as the estimated cutoff. The right panel shows the cutoffs estimated from a local logistic regression.

The cutoff analysis shows that there is substantial heterogeneity in search strategies

[18]In Appendix G, we also use a duration analysis to explore whether other aspects of the decision problem impact search. We show in this appendix that the cutoff rule does not appear to change with the search round, consistent with sequential search strategy.

[19]We also explored alternative specifications where we use the lowest value in the distribution of accepted values as the cutoff, a specification where we estimate a lower bound cutoff by using the highest value rejected by each buyer, and a specification where we use the cutoff rules reported in the questionnaire at the end of the experiment. The results from these alternative analyses are similar to the results shown here with (i) substantial dispersion in cutoff rules in the high-cost environment and estimated or reported reservation values between 30 and 50 and (ii) less dispersed cutoff rules in the low-cost environment with modal estimated or reported reservation values around 70.

employed by buyers in the high-cost treatments. While some buyers have a cutoff near the theoretical benchmark, most buyers have a higher cutoff and, therefore, reject the search-preventing default. In the HC environment, this leads to more search and lowers both the buyers' payoffs and the sellers' profits relative to the theoretical predictions. Note that while higher cutoffs imply more search, the converse is not necessarily true. Indeed, while there is excessive search in the SPD rounds of the LC environment, the imputed cutoffs are mostly aligned with theory because defaults in those rounds are high and there is little room for improvement in accepted values.

The data also suggest that there may be scope to improve both the buyers' payoffs and the sellers' profits by setting a higher threshold for defaults. To test this hypothesis, we conducted Experiment 2 where we sought to construct a better default selection rule based on (i) the behavior of an out-of-sample set of buyers and (ii) each buyer's own past behavior in a series of "training" rounds.

5 Experiment 2: Empirically-informed defaults

5.1 Experimental design

Participants in this experiment made decisions in an environment identical to the HC environment of Experiment 1, with the only differences being that (i) sessions were conducted in the Experimental Economics Laboratory at the University of Melbourne in November 2024; (ii) 20 new sequences of value draws were generated and re-used across treatments; and (iii) participants played a different number of decision rounds and faced different default-setting rules, as explained in detail below.[20]

In addition to random defaults and search-preventing defaults, across the treatments of Experiment 2 we used two types of empirically-informed defaults: **Group-estimated defaults (GED)** and **individual-estimated defaults (IED)**. The former was based on out-of-sample training data from separate treatments, whereas the latter relied on within-subject training data from a set of initial decision rounds. We conducted a total of six treatments, summarized in Table 5.

To generate training data for GED, we conducted two treatments consisting of 80 rounds. In one treatment (**RD**), a random default was used in each round while in the other (**SPD**) the theoretically optimal search-preventing default was used.

[20] The experiment was pre-registered on AsPredicted under #195964. On average, participants spent about 35 minutes completing the experiment and earned $29.52.

Table 5: Treatments and number of subjects in Experiment 2

Treatments	Rounds	Defaults	Subjects
RD	80	random	42
SPD	80	search-preventing	43
GED	80	group-estimated	45
mixed-SPD	160	RD/SPD/GED-RD/SPD	46
mixed-GED	160	RD/SPD/GED-RD/GED	46
mixed-IED	160	RD/SPD/GED-RD/IED	46

The RD/SPD/GED mix in rounds 1-80 of the mixed treatments included 40 rounds of RD, 20 rounds of SPD and 20 rounds of GED. The RD/$\cdot$ mix in rounds 81-160 of the mixed treatments included 40 rounds of RD and 40 rounds of $\cdot$.

Group-estimated default (GED) In order to identify a better default setting rule, we used a two-step algorithm to predict the expected profit that would arise empirically from each potential default. Using data from treatments RD and SPD, we estimated the probability for each individual buyer to accept a given best offer. To this end, we estimated two logistic search models to allow for some variation in how buyers treated the initial offer and subsequent offers. The first model used data only from the first acceptance/rejection decision in each round. For the second model, we used data from rounds where buyers searched $s > 0$ times and broke each decision problem into s binary acceptance/rejection decisions for the best offer in each case. Thus, the second logistic choice model predicted the probability that the buyer would accept a given best offer conditional on searching at least once.

Next, for each potential default d, we used the first model to estimate $\mathrm{P}_i(A|d)$—the predicted probability that buyer i accepts the default. We then used the second model to estimate $V_i(Search|d)$—the expected value of the offer eventually accepted by the buyer initially facing default d, conditional on the buyer searching at least once. This value was estimated via simulations, by generating 20,000 draw sequences and finding the average value of the accepted offer over these draws. We used these estimates to compute the expected profit of the firm from buyer i if the default offer was d as

$$\mathrm{E}[\Pi_i(d)] = \mathrm{P}_i(A|d)[100 - d] + [1 - \mathrm{P}_i(A|d)][100 - V_i(Search|d)]. \tag{3}$$

Finally, we calculated the expected profit from each default by averaging over all buyers in the hold-out sample.

Figure 3 shows the resulting profit function along with the profit function from the

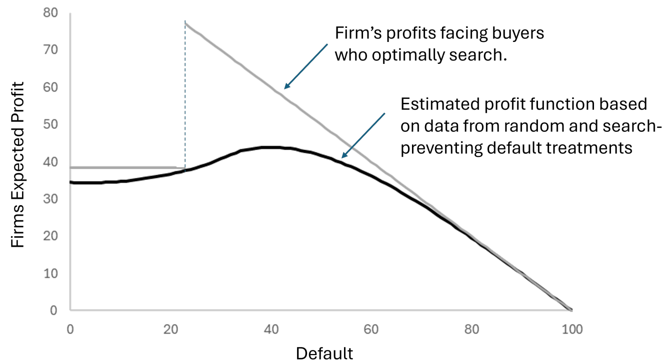

Figure 3: Profit function with fully rational buyers and the profit function estimated from behavior in the search-preventing default and random treatments.

theoretical model with a homogeneous set of fully rational buyers. As can be seen, while the theoretical benchmark generates a cutoff rule at $d = 23$, the empirically estimated profit function is smooth with a maximum around $d = 40$.

In treatment GED, subjects faced in each round the default with the highest value from this empirical profit function.

Individual-estimated default (IED) To explore individually-estimated defaults, we conducted treatment **mixed-IED** consisting of 160 decision rounds. In rounds 1-80, serving to generate training data, subjects faced a mixture of 40 RD rounds, 20 SPD rounds, and 20 GED rounds, with the latter based on the estimated profit function in Figure 3.[21] This training data was then used at the individual level for each subject to estimate their profit function (3). In rounds 81-160, subjects faced a mixture of 40 RD rounds, and 40 **IED** rounds in which the default was chosen as the offer with the highest value from the individual's estimated profit function.[22] For comparison, we also ran two additional control treatments, **mixed-SPD** and **mixed-GED**, where the first 80 rounds were identical to those in mixed-IED and rounds 81-160 contained a 40/40 mixture of RD with SPD and GED, respectively.

5.2 Results: Group-estimated defaults

We compare profits and buyers' payoffs in the GED treatment with the same measures in the control (and training data generating) treatments RD and SPD, where subjects faced

[21] The first block thus allows for a within-subject comparison between group-estimated defaults, SPD defaults, and random defaults, while also providing a wide variety of defaults to estimate individual cutoffs.

[22] We chose not to continuously update the default selection function with the data arriving each round to make the entire block of 80 choices comparable to the mixed-GED treatment where 80 rounds of data was used in the estimation.

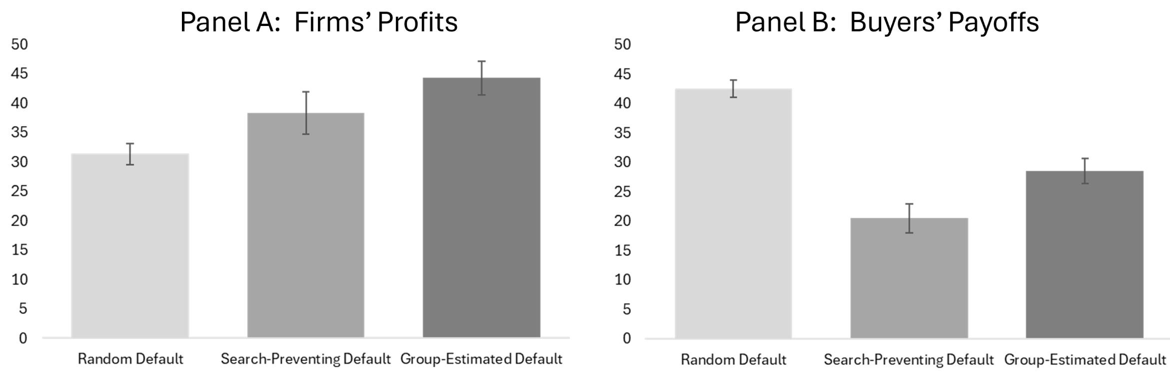

Figure 4: Average Sellers' Profits and Buyers' Payoffs in follow-up Experiment 2 where participants played 80 rounds under either the random default, the search-preventing default, or the group-estimated default. Error bars represent 95% confidence intervals with errors clustered at the individual level.

either a random default or the theoretically optimal search-preventing default of 23.

Result 5 *Firms' profits are highest in the group-estimated default treatment and lowest in the random default treatment. Buyers' payoffs are higher in the group-estimated default treatment than the search-preventing default treatment but lower than in the random default treatment.*

The support for Result 5 is provided in Figure 4, which shows the firms' average profit (left) and buyers' average payoff (right) in the three treatments. As can be seen in the left panel, the group-estimated default treatment has an average firm profit of 44.29, which is significantly higher than the firms' profits in the search-preventing default treatment ($p = 0.01$) and the random default treatment ($p < 0.01$), in a comparison of means with errors clustered at the individual level. As can be seen in the right panel, the buyer's average payoff is also significantly higher in the group-estimated default treatment compared to the SPD treatment but lower than under a random default treatment. The increase in buyers' payoffs in the group-estimated default treatment relative to the SPD treatment is due to the fact that buyers in the former search 0.91 times on average, compared to 1.37 times in the latter. The difference in search is also significant ($p < 0.01$).

In our high-cost environment, at least, variation in search strategies causes the firm to raise its default relative to the theoretical benchmark in order to entice a larger subset of buyers to accept the default rather than search. On net, the adjusted default rule helps both the firm and the buyer since it reduces search costs but leads the buyer to purchase goods that are worse than they are predicted to receive if the first offer were random.

5.3 Results: Individual-estimated defaults

The previous section suggests that it is possible to increase profits relative to the theoretically optimal default-setting strategy by using out-of-sample data to better predict search behavior. In this section we explore whether it is possible for a firm to use individually-tailored cutoff rules if significant data is available from past decisions.

The first block allows for a within-subject test of the relationship between group-estimated defaults, SPD defaults, and random defaults while also providing a wide variety of defaults to estimate individual cutoffs. Consistent with the results in Section 5.2, we find:

Result 6 *In Rounds 1-80 of treatments mixed-SPD, mixed-GED and mixed-IED:*

1. *Sellers' profits and buyers' payoffs are significantly higher on average in rounds with a group-estimated default relative to rounds with a search-preventing default.*

2. *Sellers' profits are significantly lower in rounds with a random default compared to rounds using the SPD default or the group-estimated default.*

3. *Buyers' profits are significantly lower in rounds with a random default compared to rounds using the SPD default or the group-estimated default.*

Support for Result 6 is based on the same comparison of means test reported in the previous section. Similar to the results there, average sellers' profits in the group-estimated default rounds is 41.6. This is significantly higher than sellers' profits in search-preventing default rounds (profit = 40.0; p-value < 0.01) and in rounds where the default is random (profit = 32.0; p-value < 0.01).

In contrast, average buyers' payoffs in rounds with a random default is 39.7, which is significantly higher than buyers' payoffs in group-estimated default rounds (payoff = 25.9; p-value < 0.01). Buyers' payoffs in group-estimated default rounds is, in turn, significantly larger than buyers' payoffs in search-preventing default rounds (payoff = 20.8; p-value < 0.01).

Having replicated our previous results in a similar environment, we now turn to behavior in rounds 81-160 where we introduce individual-estimated defaults in some of the sessions. We find:

Result 7 *In Rounds 81-160 of the experiment:*

1. *There are no significant differences in the average firms' profits in rounds played under the search-preventing default, the group-estimated default, and the individually-estimated default.*

2. *Buyers' payoffs are significantly higher in rounds played under the individually-estimated default and the group-estimated default when compared to the search-preventing default. However, there are no significant difference in buyers' payoffs when comparing the group-estimated default to the individually-estimated default.*

Support for Result 7 is provided in Figure 5, which shows the firms' average profit (left) and buyers' average payoff (right) in the four treatments. As can be seen on the left-hand side of the figure, the average profit of the sellers in rounds using the SPD default is 48.4, which is not significantly different from the average profit in rounds using the individually-estimated default (profit = 44.0; p-value = 0.15) and rounds using the group-estimated default (profit = 46.4; p-value = 0.48).

As can be seen in the right panel, buyers' payoffs in the individually-estimated default rounds is 34.4, which is significantly higher than buyers' payoffs under the search-preventing default (payoff = 27.9; p-value < 0.01), but not significantly different from buyers' payoffs under the group-estimated default (payoff = 33.1; p-value = 0.52). Buyers' payoffs in the group-estimated default are also significantly higher than payoffs under the search-preventing default (p-value < 0.01).[23]

In an ex-post analysis, we also explored why the group-estimated default increases the average profits of the sellers relative to the search-preventing default in rounds 1-80 but not in rounds 81-160. The data suggest that the difference is due to buyers' search decisions becoming more similar to optimal behavior over time. Figure 6 shows the average probability of rejection for each potential initial offer with the data split between rounds 1-80 and 81-160. As seen, initial offers between 30 and 50 are more likely to be accepted by the buyer in the second block. As such, the group-estimated default is set too high in these rounds, and the performance of the theoretically selected default improves.

Taken together, our data suggest that empirically-estimated cutoff rules can improve on the theoretically optimal rules by accounting for suboptimal search behavior, but that over time the gains to an empirical cutoff diminish as participants learn. Individually optimal default performs similarly to the group-estimated default, and we find limited additional gains in trying to tailor the algorithm to the individual in our environment. We note, however, that we are studying an environment in which search costs are homogeneous. It is an open question as to how tailored algorithms perform in environments where search costs are heterogeneous and unknown to the default designer.

[23] As seen in the graph, the random treatment generates significantly lower profit for the firm (profit = 33.5; p-value < 0.01 for both comparisons) and significantly higher payoffs for the buyer (payoff = 46.8, p-value < 0.01 for all three pairwise comparisons).

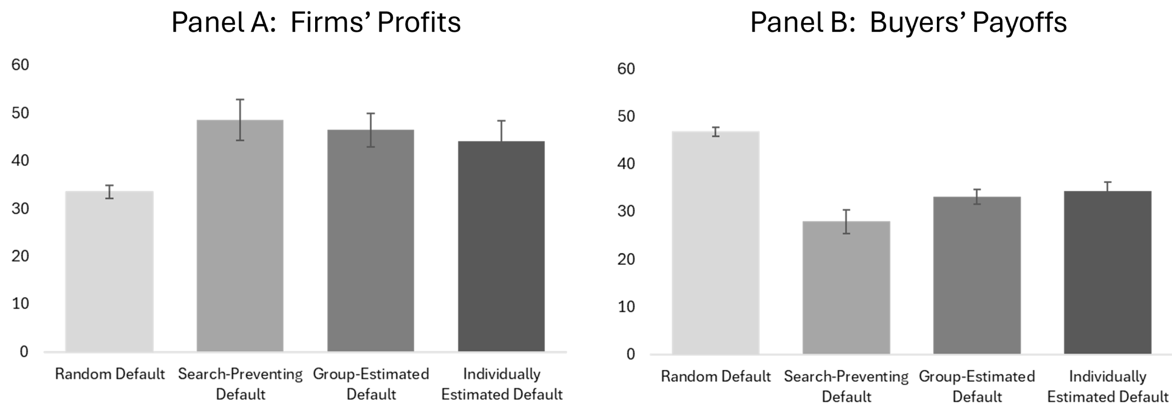

Figure 5: Average Sellers' Profits and Buyers' Payoffs in rounds 81-160 of follow-up Experiment 3 where participants played 40 rounds under a random default and 40 rounds under a search-preventing default, group-estimated default, or individually-estimated default. Error bars represent 95% confidence intervals with errors clustered at the individual level.

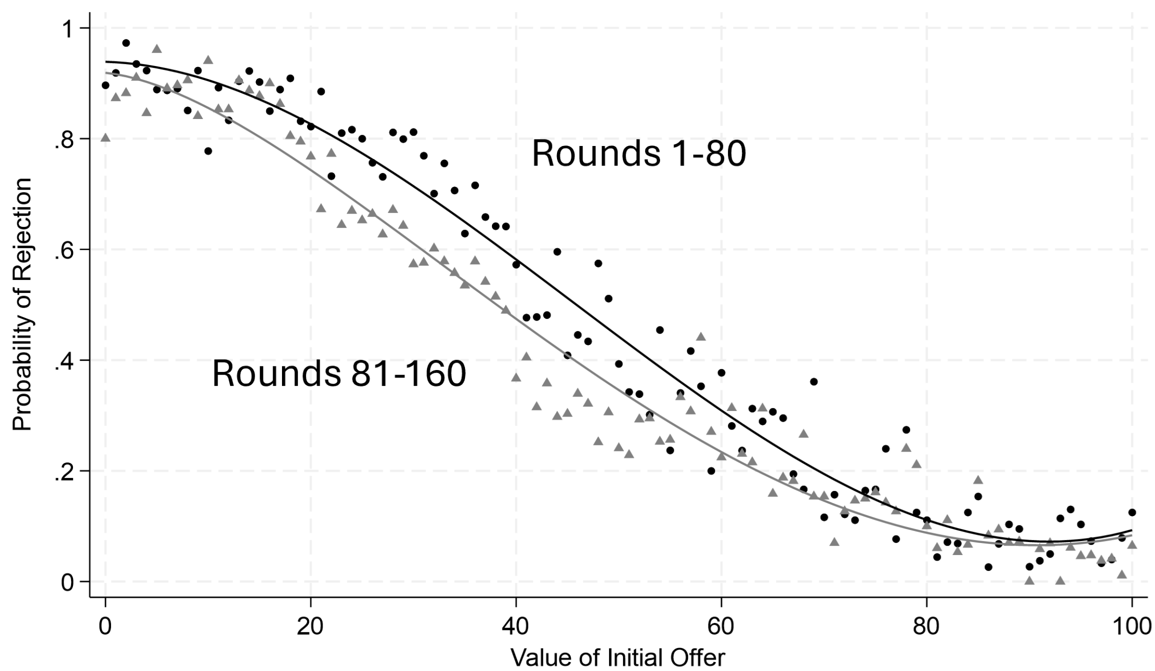

Figure 6: Probability of Rejection in Rounds 1-80 (black) and Rounds 81-160 (grey). The probability of rejection in later rounds is lower than in the original block leading to greater efficiency in the search-preventing default and lower efficiency in the group-estimated default.

6 Conclusions

We proposed a simple model of how a profit-maximizing seller can benefit from strategically setting defaults in environments where buyers choose one of several available options whose values are initially unknown but can be discovered through costly search. Theoretically, search-preventing defaults prevent search and sellers always benefit from setting them, whereas the expected impact on buyers depends on search costs. Buyers with low enough search costs benefit from strategic default-setting because they economize on search, and it is optimal for the firm to offer them high-quality defaults. However, high-cost buyers suffer because their defaults are inferior, yet good enough for them to refrain from searching.

We find that the main comparative statics of our model are realized in an experimental environment where we vary the way in which the default is set between rounds and the cost of search between individuals. Search-preventing defaults are beneficial to buyers when search costs are low but harmful to buyers when search costs are high. However, in the high-cost treatments, buyers tend to search more than predicted, and this leads to even worse outcomes for the buyers than our theoretical model would predict.

We then explored empirically-informed alternative mechanisms that accommodate for variation in search behavior. We found that a group-estimated default tends to improve the profits of firms and also improves buyer welfare when search costs are high. This is due to the empirically driven search algorithm selecting higher defaults that reduce the likelihood of search. Individually tailored defaults do not improve on the group-estimated defaults in our setting, where participants have homogeneous search costs.

Our theoretical model is stylized, but the issue of sequential information acquisition in the context of a better-informed principal is quite general. Examples of such environments abound in industries where sellers of goods or services have asymmetric expertise. Such sellers include financial advisors, car mechanics, IT specialists, medical consultants, lawyers, real-estate agents, and contractors.

The negative consequences of search-preventing defaults are more severe the larger the search costs, and the more the incentives of buyers and seller are misaligned. To the extent search costs are correlated with characteristics such as age, cognitive ability, income, or education, it appears that the most vulnerable segments of the population may be disproportionately affected.

Our results suggest that search-preventing defaults may be an effective strategy for surplus extraction by monopoly sellers even when buyers' and sellers' incentives are known to be misaligned. In future research we will consider richer environments involving multiple firms, to study the effects of competition, or the availability of feedback and information systems such as consumer ratings or product reviews, as possible mechanisms mitigating the negative impact of strategic default settings. Other possible extensions are considering boundedly rational buyers, buyers who learn about the distribution of values, or sellers who face uncertainty about buyers' preferences or search costs.

References

Altmann, Steffen, Armin Falk, and Andreas Grunewald. 2025. "Communicating through defaults." *Review of Economics and Statistics*, 107(1): 256–268.

Arbatskaya, Maria. 2007. "Ordered Search." *RAND Journal of Economics*, 38(1): 119–126.

Armstrong, Mark. 2017. "Ordered Consumer Search." *Journal of the European Economic Association*, 15(5): 989–1024.

Armstrong, Mark, and Jidong Zhou. 2011. "Paying for Prominence." *The Economic Journal*, 121(556): F368–F395.

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

Atalay, A. Selin, H. Onur Bodur, and Dina Rasolofoarison. 2012. "Shining in the Center: Central Gaze Cascade effect on Product Choice." *Journal of Consumer Research*, 39(4): 848–866.

Athey, Susan, and Glenn Ellison. 2011. "Position Auctions with Consumer Search." *Quarterly Journal of Economics*, 126(3): 1213–1270.

Bhatia, Sudeep, Lisheng He, Wenjia J. Zhao, and Pantelis P. Analytis. 2021. "Cognitive Models of Optimal Sequential Search with Recall." *Cognition*, 210: 104595.

Bosch-Domènech, Antoni, José G. Montalvo, Rosemarie Nagel, and Albert Satorra. 2002. "One, Two, (Three), Infinity, ...: Newspaper and Lab Beauty-Contest Experiments." *American Economic Review*, 92(5): 1687–1701.

Brown, Christina L., and Aradhna Krishna. 2004. "The Skeptical Shopper: A Metacognitive Account for the Effects of Default Options on Choice." *Journal of Consumer Research*, 31(3): 529–539.

Burdett, Kenneth, and Kenneth L. Judd. 1983. "Equilibrium Price Dispersion." *Econometrica*, 51(4): 955–969.

Campbell, Margaret C. 2007. ""Says Who?!" How the Source of Price Information and Affect Influence Perceived Price (Un) Fairness." *Journal of Marketing Research*, 44(2): 261–271.

Caplin, Andrew, and Daniel Martin. 2017. "Defaults and Attention: The Drop Out Effect." *Revue Économique*, 68(5): 747–755.

Caplin, Andrew, Mark Dean, and Daniel Martin. 2011. "Search and Satisficing." *American Economic Review*, 101(7): 2899–2922.

Carlin, Bruce I. 2009. "Strategic Price Complexity in Retail Financial Markets." *Journal of Financial Economics*, 91(3): 278–287.

Carlin, Bruce Ian, Simon Gervais, and Gustavo Manso. 2013. "Libertarian Paternalism, Information Production, and Financial Decision Making." *The Review of Financial Studies*, 26(9): 2204–2228.

Carroll, Gabriel, James Choi, David Laibson, Brigitte Madrian, and Andrew Metrick. 2009. "Optimal Defaults and Active Decisions." *Quarterly Journal of Economics*, 124: 1639–1674.

Cason, Timothy N., and Daniel Friedman. 2003. "Buyer Search and Price Dispersion: A Laboratory Study." *Journal of Economic Theory*, 112(2): 232–260.

Cason, Timothy N., and Shakun Datta. 2006. "An Experimental Study of Price Dispersion in an Optimal Search Model with Advertising." *International Journal of Industrial Organization*, 24(3): 639–665.

Cason, Timothy N., and Shakun D. Mago. 2010. "Costly Buyer Search in Laboratory Markets with Seller Advertising." *The Journal of Industrial Economics*, 58(2): 424–449.

Chandon, Pierre, J. Wesley Hutchinson, Eric T. Bradlow, and Scott H. Young. 2009. "Does In-Store Marketing Work? Effects of the Number and Position of Shelf Facing on Brand Attention and Evaluation at the Point of Purchase." *The Journal of Industrial Economics*, 73(6): 1–17.

Chen, Daniel L., Martin Schonger, and Chris Wickens. 2016. "oTree—An Open-source Platform for Laboratory, Online, and Field Experiments." *Journal of Behavioral and Experimental Finance*, 9(C): 88–97.

Chen, Yongmin, and Chuan He. 2011. "Paid Placement: Advertising and Search on the Internet." *The Economic Journal*, 121(556): F309–F328.

Chioveanu, Ioana, and Jidong Zhou. 2013. "Price Competition with Consumer Confusion." *Management Science*, 59(11): 2450–2469.

Choi, James, David Laibson, Brigitte Madrian, and Andrew Metrick. 2003. "Optimal Defaults." *American Economic Review*, 93: 180–185.

Choi, James, David Laibson, Brigitte Madrian, and Andrew Metrick. 2004. "For Better of for Worse Default Effects and 401(K) Savings Behavior." In *Perspectives on the Economics of Aging.* , ed. D. A. Wise, 81–125. University of Chicago Press, Chicago.

Conlisk, John. 2003. "A Note on Convergence of Adaptive Satisficing to Optimal Stopping." *Journal of Political Economy*, 111(6): 1353–1360.

Davis, Douglas D., and Charles A. Holt. 1996. "Consumer Search Costs and Market Performance." *Economic Inquiry*, 34(1): 133–151.

DellaVigna, Stefano, and Elizabeth Linos. 2022. "RCTs to Scale: Comprehensive Evidence from Two Nudge Units." *Econometrica*, 90(1): 81–116.

De los Santos, Babur, Ali Hortaçsu, and Matthijs R. Wildenbeest. 2012. "Testing Models of Consumer Search Using Data on Web Browsing and Purchasing Behavior." *American Economic Review*, 102(6): 2955–2980.

Dinner, Isaac, Eric J. Johnson, Daniel G. Goldstein, and Kaiya Liu. 2011. "Partitioning Default Effects: Why People Choose Not to Choose." *Journal of Experimental Psychology: Applied*, 17(4): 332–341.

Dobrescu, Loretti, Xiaodong Fan, Hazel Bateman, Ben Newell, Andreas Ortmann, and Susan Thorp. 2016. "Retirements Savings: A Tale of Decisions and Defaults." *Economic Journal*, 128(610): 1047–1094.

Gabaix, Xavier, David Laibson, Guillermo Moloche, and Stephen Weinberg. 2006. "Costly Information Acquisition: Experimental Analysis of a Boundedly Rational Model." *American Economic Review*, 96(4): 1043–1068.

Gigerenzer, Gerd. 2015. "On the Supposed Evidence for Libertarian Paternalism." *Review of Philosophy and Psychology*, 6: 361–383.

Gigerenzer, Gerd, Jutta Mata, and Ronald Frank. 2009. "Public Knowledge of Benefits of Breast Cancer Screening in Europe." *Journal of the National Cancer Institute*, 101(17): 1216–1220.

Gossner, Olivier, Jakub Steiner, and Colin Stewart. 2021. "Attention Please!" *Econometrica*, 89(4): 1717–1751.

Gøtzsche, Peter C, and Karsten Juhl Jørgensen. 2013. "Screening for Breast Cancer with Mammography." *Cochrane Database of Systematic Reviews*, 2013(6): CD001877.

Greiner, Ben. 2015. "Subjects with ORSEE." *Journal of the Economic Science Association*, 1(1): 114–125.

Grether, David M., Alan Schwartz, and Louis L. Wilde. 1988. "Uncertainty and Shopping Behaviour: An Experimental Analysis." *Review of Economic Studies*, 55(2): 323–342.

Harrison, Glenn W., and Peter Morgan. 1990. "Search Intensity in Experiments." *Economic Journal*, 100(401): 478–486.

Hong, Han, and Matthew Shum. 2006. "Using Price Distributions to Estimate Search Costs." *RAND Journal of Economics*, 37(2): 257–275.

Jachimowicz, Jon M., Shannon Duncan, Elke U. Weber, and Eric J. Johnson. 2019. "When and Why Defaults Influence Decisions: A Meta-analysis of Default Effects." *Behavioural Public Policy*, 3(2): 159–186.

Janssen, Maarten C.W., and Cole Williams. 2023. "Influencing Search." *Working paper*. https://colerandallwilliams.com/wp-content/uploads/2023/02/influencingsearch.pdf.

Jenkins, Stephen P. 2005. "Survival Analysis." *Unpublished manuscript, Institute for Social and Economic Research, University of Essex, Colchester, UK*. https://citeseerx.ist.psu.edu/document?repid=rep1&type=pdf&doi=9bb46b98492c0d8e33ffbddab4a0f99d84f3f0c0.

Karle, Heiko, Florian Kerzenmacher, Heiner Schumacher, and Frank Verboven. 2025. "Search Costs and Context Effects." *American Economic Journal: Microeconomics*, 17(2): 127–161.

Lippman, Steven A., and John J. McCall. 1976. "The Economics of Job Search: A Survey." *Economic Inquiry*, 14(2): 155–189.

McCall, John Joseph. 1970. "Economics of Information and Job Search." *Quarterly Journal of Economics*, 113–126.

Morgan, Peter, and Richard Manning. 1985. "Optimal Search." *Econometrica*, 53(4): 923–944.

Ortmann, Andreas, Dmitry Ryvkin, Tom Wilkening, and Jingjing Zhang. 2023. "Defaults and cognitive effort." *Journal of Economic Behavior & Organization*, 212: 1–19.

Piccione, Michele, and Ran Spiegler. 2012. "Price competition Under Limited Comparability." *Quarterly Journal of Economics*, 127(1): 97–135.

Rapoport, Amnon, Darryl A. Seale, and Leonidas Spiliopoulos. 2022. "Progressive Stopping Heuristics that Excel in Individual and Competitive Sequential Search." *Theory and Decision*. https://link.springer.com/article/10.1007/s11238-02-09881-0.

Ratchford, Brian T. 2009. "Consumer Search and Pricing." In *Handbook of Pricing Research in Marketing*. 91–107. Edward Elgar Publishing.

Ryvkin, Dmitry, and Danila Serra. 2019. "Is More Competition Always Better? An Experimental Study of Extortionary Corruption." *Economic Inquiry*, 57(1): 50–72.

Ryvkin, Dmitry, and Danila Serra. 2020. "Corruption and Competition Among Bureaucrats: An Experimental Study." *Journal of Economic Behavior & Organization*, 175: 439–451.

Ryvkin, Dmitry, Danila Serra, and James Tremewan. 2017. "I Paid a Bribe: An Experiment on Information Sharing and Extortionary Corruption." *European Economic Review*, 94: 1–22.

Sanders, Michael, Veerle Snijders, and Michael Hallsworth. 2018. "Behavioural Science and Policy: Where Are We Now and Where Are We Going?" *Behavioural Public Policy*, 2(2): 144–167.

Sawilowsky, Shlomo S. 2009. "New Effect Size Rules of Thumb." *Journal of Modern Applied Statistical Methods*, 8(2): 467–474.

Schotter, Andrew, and Yale M. Braunstein. 1981. "Economic Search: An Experimental Study." *Economic Inquiry*, 19(1): 1–25.

Simon, Herbert A. 1955. "A Behavioral Model of Rational Choice." *Quarterly Journal of Economics*, 70(5, pt.2): 99–118.

Sonnemans, Joep. 1998. "Strategies of Search." *Journal of Economic Behavior & Organization*, 35(3): 309–332.

Spiegler, Ran. 2011. *Bounded Rationality and Industrial Organization.* Oxford University Press.

Spiegler, Ran. 2015. "On the Equilibrium Effects of Nudging." *The Journal of Legal Studies*, 44(2): 389–416.

Spiegler, Ran. 2016. "Choice Complexity and Market Competition." *Annual Review of Economics*, 8: 1–25.

Stahl, Dale O. 1989. "Oligopolistic Pricing with Sequential Consumer Search." *American Economic Review*, 700–712.

Stigler, George J. 1961. "The Economics of Information." *Journal of Political Economy*, 69(3): 213–225.

Studdert, David M., Michelle M. Mellow, William M. Sage, Catherine M. DesRoches, Jordon Peugh, Kinga Zapert, and Troyen A. Brennan. 2005. "Defensive Medicine Among High-risk Specialist Physicians in a Volatile Malpractice Environment." *Journal of the American Medical Association*, 293: 2609–2617.

Sunstein, Cass R. 2014. *Why Nudge? The Politics of Libertarian Paternalism.* Yale University Press.

Thaler, Richard H., and Cass R. Sunstein. 2003. "Libertarian Paternalism." *American Economic Review*, 93(2): 175–179.

Varian, Hal. 2007. "Position Auctions." *International Journal of Industrial Organization*, 25(6): 1163–1187.

Varian, Hal R. 1980. "A Model of Sales." *American Economic Review*, 70(4): 651–659.

Weitzman, Martin L. 1979. "Optimal Search for the Best Alternative." *Econometrica*, 47(3): 641–654.

Wilson, Chris M. 2010. "Ordered Search and Equilibrium Obfuscation." *The Economic Journal*, 28(5): 496–506.

Woloshin, Steven, and Lisa M. Schwartz. 2012. "How a Charity Oversells Mammography." *British Medical Journal*, 345(e5132).

Xu, Lizhen, Jianqing Chen, and Andrew Whinston. 2010. "Oligopolistic Pricing with Online Search." *Journal of Managment Information Systems*, 27(3): 111–142.

A Proofs

We start by calculating s_k—the expected number of searches in the random default condition when the acceptable set contains k options.

Lemma 1 *The expected number of searches for a given* $k = |A| \geq 1$ *is*

$$s_k = \frac{n-k}{k+1}. \tag{4}$$

Proof of Lemma 1 The expected number of searches for a given $k = |A| \geq 1$ is

$$\begin{aligned} s_k = &\frac{k}{n} \cdot 0 + \frac{n-k}{n} \cdot \frac{k}{n-1} \cdot 1 + \frac{n-k}{n} \cdot \frac{n-k-1}{n-1} \cdot \frac{k}{n-2} \cdot 2 \\ &+ \ldots + \frac{n-k}{n} \cdot \frac{n-k-1}{n-1} \cdot \ldots \cdot \frac{1}{n-k} \cdot \frac{k}{k} \cdot (n-k) = \sum_{l=0}^{n-k} \frac{(n-k)!(n-l)!}{(n-k-l)!n!} \frac{k}{n-l} l. \end{aligned}$$

The last expression is equivalent to:

$$\frac{k!(n-k)!}{n!} \sum_{l=0}^{n-k} \frac{(n-l-1)!l}{(k-1)!(n-k-l)!}.$$

We will now show that this expression simplifies to (4). Consider the following:

$$\sum_{l=0}^{n-k} \frac{(n-1-l)!l}{(k-1)!(n-k-l)!} = \sum_{l=0}^{n-k} \binom{n-1-l}{k-1} l = \sum_{l=0}^{n-k} \binom{k-1+l}{k-1} (n-k-l).$$

It is, therefore, sufficient to prove that

$$\sum_{l=0}^{n-k} \binom{k-1+l}{k-1} (n-k-l) = \frac{(n-k)n!}{(k+1)!(n-k)!}.$$

We prove by induction over $n = k, k+1, \ldots$. For $n = k$, the left-hand side is 0 and so is the right-hand side. Suppose the identity is valid for some $n \geq k$, and consider the same for $n+1$:

$$\sum_{l=0}^{n+1-k} \binom{k-1+l}{k-1}(n+1-k-l) = \sum_{l=0}^{n-k} \binom{k-1+l}{k-1}(n+1-k-l)$$
$$= \sum_{l=0}^{n-k} \binom{k-1+l}{k-1}(n-k-l) + \sum_{l=0}^{n-k} \binom{k-1+l}{k-1} = \frac{(n-k)n!}{(k+1)!(n-k)!} + \binom{n}{k}$$
$$= \frac{n!}{k!(n-k)!}\left(\frac{n-k}{k+1}+1\right) = \frac{(n+1)!}{(k+1)!(n-k)!} = \frac{(n+1-k)(n+1)!}{(k+1)!(n+1-k)!}.$$

The resulting expression has the desired form. In the second line, we used the induction hypothesis and the well-known identity $\sum_{l=0}^{n-k} \binom{k+l}{k} = \binom{n+1}{k+1}$. ■

Deriving (2) The buyer's expected payoff under the search-preventing default is

$$\pi_{\text{SPD}} = \int_{\underline{b}}^{z^*} b\mathrm{d}F(b)^n - F(z^*)^n c(n-1) + \sum_{k=1}^{n} \binom{n}{k} F(z^*)^{n-k}[1-F(z^*)]^k \int_{z^*}^{\bar{b}} b\mathrm{d}\tilde{F}_{(1:k)}(b), \quad (5)$$

where (i) function $\tilde{F}(b) = \frac{F(b)-F(z^*)}{1-F(z^*)}$ is the updated distribution of benefits above z^* and (ii) function $\tilde{F}_{(1:k)}(b) = 1-(1-\tilde{F}(b))^k$ is the distribution of the lowest order statistic among k i.i.d. draws from $\tilde{F}$. The first two terms of (5) represent the expected payoff if $A = \varnothing$; the buyer will search $n-1$ times in this case and select the best option. The third term is the expected payoff when $A \neq \varnothing$. The sum is over $k = |A|$. For each possible k, the default is set at (and the buyer selects, without searching) the option with the lowest benefit above z^*.

The buyer's payoff in the random default setting can similarly be written as

$$\pi_{\text{RD}} = \int_{\underline{b}}^{z^*} b\mathrm{d}F(b)^n - F(z^*)^n c(n-1) + \sum_{k=1}^{n} \binom{n}{k} F(z^*)^{n-k}[1-F(z^*)]^k \left[\int_{z^*}^{\bar{b}} b\mathrm{d}\tilde{F}(b) - cs_k\right], \quad (6)$$

where, as before, $\tilde{F}(b)$ is the updated distribution of beliefs about the benefits above z^* and s_k is the expected number of searches for a given k. The first two terms are the same as in (5) and correspond to $A = \varnothing$. The third term is again written as a sum over $k = |A|$, but this time the buyer's expected payoff for a given k is the expectation of b above the reservation value less the expected search cost. Equation (2) obtains by combining (5) and (6) and integrating by parts.

Proof of Proposition 2 Throughout this proof, we drop the asterisk in z^* for brevity. We begin by simplifying (2). Using Lemma 1, the coefficient on c in (2) becomes

$$\sum_{k=1}^{n}\binom{n}{k}F(z)^{n-k}[1-F(z)]^{k}\frac{n-k}{k+1}=\sum_{k=1}^{n}\binom{n}{k}F(z)^{n-k}[1-F(z)]^{k}\left(\frac{n+1}{k+1}-1\right)$$

$$=\sum_{k=1}^{n}\binom{n+1}{k+1}F(z)^{n-k}[1-F(z)]^{k}-[1-F(z)^{n}]$$

$$=\sum_{k=2}^{n+1}\binom{n+1}{k}F(z)^{n+1-k}[1-F(z)]^{k-1}-[1-F(z)^{n}]$$

$$=\frac{1-F(z)^{n+1}-(n+1)F(z)^{n}(1-F(z))}{1-F(z)}-1+F(z)^{n} \tag{7}$$

$$=1+F(z)+F(z)^{2}+\ldots+F(z)^{n}-(n+1)F(z)^{n}-1+F(z)^{n} \tag{8}$$

$$=F(z)[1+F(z)+\ldots+F(z)^{n-1}-nF(z)^{n-1}]$$

$$=F(z)\left[\frac{1-F(z)^{n}}{1-F(z)}-nF(z)^{n-1}\right]. \tag{9}$$

To obtain (7), we multiplied and divided the binomial sum by $1-F(z)$ and used the identity

$$\sum_{k=0}^{n}\binom{n}{k}F(z)^{n-k}[1-F(z)]^{k}=1, \tag{10}$$

with n replaced by $n+1$. To obtain (8), we used the identity

$$\frac{1-F(z)^{n+1}}{1-F(z)}=\sum_{k=0}^{n}F(z)^{k}.$$

Finally, to obtain (9) we used a similar identity with $n+1$ replaced by n.

The next term in (2) can be written as

$$\sum_{k=1}^{n}\binom{n}{k}F(z)^{n-k}[1-F(z)]^{k}\int_{z}^{\bar{b}}\tilde{F}_{(1:k)}(b)\mathrm{d}b$$

$$=\sum_{k=1}^{n}\binom{n}{k}F(z)^{n-k}[1-F(z)]^{k}\int_{z}^{\bar{b}}[1-(1-\tilde{F}(b))^{k}]\mathrm{d}b$$

$$=\int_{z}^{\bar{b}}[1-F(z)^{n}]\mathrm{d}b-\int_{z}^{\bar{b}}\sum_{k=1}^{n}\binom{n}{k}F(z)^{n-k}[1-F(b)]^{k}\mathrm{d}b \tag{11}$$

$$=\int_{z}^{\bar{b}}[1-F(z)^{n}]\mathrm{d}b-\int_{z}^{\bar{b}}[(1-F(b)+F(z))^{n}-F(z)^{n}]\mathrm{d}b. \tag{12}$$

To obtain (11), we used (10) to write the binomial sum over $k=1,\ldots,n$ as $1-F(z)^{n}$. We also used that

$$1-\tilde{F}(b)=\frac{1-F(b)}{1-F(z)},\quad b\in[z,\bar{b}] \tag{13}$$

holds for the updated distribution of benefits. Finally, to obtain (12), we used the binomial formula

$$\sum_{k=0}^{n}\binom{n}{k}[1-F(b)]^{k}F(z)^{n-k}=[1-F(b)+F(z)]^{n}. \tag{14}$$

The last term in (2) can be simplified as

$$\sum_{k=1}^{n}\binom{n}{k}F(z)^{n-k}[1-F(z)]^{k}\int_{z}^{\bar{b}}\tilde{F}(b)\mathrm{d}b=[1-F(z)^{n}]\int_{z}^{\bar{b}}\tilde{F}(b)\mathrm{d}b.$$

Next, we integrate (1) by parts to rewrite it in the form

$$c=\bar{b}-z-\int_{z}^{\bar{b}}F(b)\mathrm{d}b=\int_{z}^{\bar{b}}[1-F(b)]\mathrm{d}b. \tag{15}$$

Plugging these into (2), we obtain

$$\begin{aligned}
&\pi_{\mathrm{SPD}}-\pi_{\mathrm{RD}}\equiv\Delta(z)=\frac{F(z)[1-F(z)^{n}]}{1-F(z)}\int_{z}^{\bar{b}}[1-F(b)]\mathrm{d}b-nF(z)^{n}\int_{z}^{\bar{b}}[1-F(b)]\mathrm{d}b\\
&-\int_{z}^{\bar{b}}[1-F(z)^{n}]\mathrm{d}b+\int_{z}^{\bar{b}}[(1-F(b)+F(z))^{n}-F(z)^{n}]\mathrm{d}b+[1-F(z)^{n}]\int_{z}^{\bar{b}}\tilde{F}(b)\mathrm{d}b\\
&=-[1-F(z)][1-F(z)^{n}]\int_{z}^{\bar{b}}[1-\tilde{F}(b)]\mathrm{d}b-nF(z)^{n}\int_{z}^{\bar{b}}[1-F(b)]\mathrm{d}b\\
&+\int_{z}^{\bar{b}}[(1-F(b)+F(z))^{n}-F(z)^{n}]\mathrm{d}b\\
&=\int_{z}^{\bar{b}}[(1-F(b)+F(z))^{n}-F(z)^{n}]\mathrm{d}b-[1+(n-1)F(z)^{n}]\int_{z}^{\bar{b}}[1-F(b)]\mathrm{d}b.
\end{aligned}$$

To show the existence of c_H, note that at $c=\bar{c}$ we have $z=\underline{b}$, $F(z)=0$, and hence

$$\Delta(\underline{b})=\int_{\underline{b}}^{\bar{b}}[(1-F(b))^{n}-(1-F(b))]\mathrm{d}b<0.$$

The existence of c_H then follows by continuity.

Finally, we show the existence of c_L. Recall that $z(c)$ is decreasing; it is, therefore, sufficient to show that $\Delta(z)$ is decreasing in z in a neighborhood of $z=\bar{b}$. Differentiating

$\Delta(z)$, obtain

$$\Delta'(z) = -[1 - F(z)^n - (1 + (n-1)F(z)^n)(1 - F(z))] \tag{16}$$
$$+ nf(z)\int_z^{\bar{b}} [(1 - F(b) + F(z))^{n-1} - F(z)^{n-1} - (n-1)F(z)^{n-1}(1 - F(b))]\mathrm{d}b.$$

The first term in (16) can be further transformed as

$$-[1 - F(z)]\left[\sum_{k=0}^{n-1} F(z)^k - 1 - (n-1)F(z)^n\right] = -[1 - F(z)]\sum_{k=1}^{n-1}[F(z)^k - F(z)^n]$$
$$= -[1 - F(z)]\sum_{k=1}^{n-1} F(z)^k[1 - F(z)^{n-k}] = -[1 - F(z)]^2 \sum_{k=1}^{n-1} F(z)^k \sum_{l=0}^{n-k-1} F(z)^l. \tag{17}$$

The second term in (16) can be transformed as

$$nf(z)\int_z^{\bar{b}}[1 - F(b)]\left[\sum_{k=0}^{n-2}(1 - F(b) + F(z))^{n-2-k}F(z)^k - (n-1)F(z)^{n-1}\right]\mathrm{d}b$$
$$= nf(z)\int_z^{\bar{b}}[1 - F(b)]\sum_{k=0}^{n-2} F(z)^k\left[(1 - F(b) + F(z))^{n-2-k} - F(z)^{n-1-k}\right]\mathrm{d}b \tag{18}$$
$$= nf(z)\int_z^{\bar{b}}[1 - F(b)]\sum_{k=0}^{n-2} F(z)^k\left[\sum_{l=1}^{n-2-k}\binom{n-2-k}{l}(1 - F(b))^l F(z)^{n-2-k-l}\right.$$
$$\left. + F(z)^{n-2-k} - F(z)^{n-1-k}\right]\mathrm{d}b \tag{19}$$
$$= nf(z)[1 - F(z)]^2\int_z^{\bar{b}}[1 - \tilde{F}(b)]\sum_{k=0}^{n-2} F(z)^k$$
$$\times\left[\sum_{l=1}^{n-2-k}\binom{n-2-k}{l}(1 - \tilde{F}(b))(1 - F(b))^{l-1}F(z)^{n-2-k-l} + F(z)^{n-2-k}\right]\mathrm{d}b. \tag{20}$$

To obtain (18), we represented $(n-1)F(z)^{n-1}$ as $\sum_{k=0}^{n-2} F(z)^{n-1}$ to bring $F(z)^{n-1}$ inside the common sum and then extracted the multiplier $F(z)^k$. For (19), we used the binomial representation similar to (14). Finally, for (20), we used (13).

Combining (17) and (20), we can write (16) as

$$\Delta'(z) = -[1 - F(z)]^2\left[\sum_{k=1}^{n-1} F(z)^k \sum_{l=0}^{n-k-1} F(z)^l - S(z)\right],$$

where $S(z)$ is the term multiplying $[1 - F(z)]^2$ in (20). Importantly, $S(z)$ converges to zero for $z \to \bar{b}$, whereas the double sum in brackets converges to a positive number; therefore, there exists a $z_0 < \bar{b}$ such that $\Delta'(z) < 0$ for all $z \in (z_0, \bar{b})$, which implies the existence of c_L. ■

B Additional considerations with buyer awareness

In our analysis in Section 3, we maintained the assumption that the buyer was unaware that the firm strategically sets the default. In this section, we explore how awareness changes buyer behavior if the seller commits to a search-preventing default and discuss the equilibrium that results from the game between the firm and the buyer when the firm seeks to exploit awareness.

Suppose first that the firm commits to using a search-preventing default and that a fully rational buyer is aware of the strategy employed by the firm. Would the buyer change her strategy in response to the observed default value?

To answer this question, it is useful to divide the analysis into two cases. First, suppose that the default offer b_1 is above $z^*(c)$, i.e., it is in the buyer's *ex ante* acceptance set. In this case, the buyer would realize that there is a gap in the distribution of values between $z^*(c)$ and b_1 and would update her beliefs about the value distribution accordingly, which can potentially affect her optimal search strategy. However, we show in Appendix C that such updating does not lead to additional search for value distributions that have an increasing hazard rate. As such, awareness does not affect behavior after observing a default in the acceptance set for a wide class of distributions.

Next, suppose that the default offered is below $z^*(c)$. Under a search-preventing default, a cutoff below $z^*(c)$ would arise only in the case where the acceptance set is empty. Therefore, the buyer should optimally switch her strategy and accept the default. As such, awareness when combined with a search-preventing default is welfare-improving for the buyer since it cuts down on search costs in cases where an unaware buyer would have to exhaust the entire offer set and eventually accept an offer below the cutoff. We note, however, that the probability that the acceptance set is empty equals $F(z^*(c))^n$, which is typically small even for a moderate n. Thus, the observable implications of awareness when a search-preventing default algorithm is used are likely to be minimal.

The above results make clear that if the firm is able to *commit* to a search-preventing default, awareness by buyers is not predicted to have an effect on buyers' behavior outside of the edge cases where the acceptance set is empty. We view this as a reassuring robustness check since there is likely to be heterogeneous levels of awareness in real-world settings.

However, the existence of the edge cases also implies that the search-preventing default is not the optimal mechanism when buyers are aware, and this is common knowledge. In particular, a profit-maximizing firm that knows that the buyer will accept offers outside the acceptance set will have an incentive to "lie" by making offers below the acceptance threshold even when the acceptance set is non-empty. This naturally leads to one or more mixed strategy equilibria in the extensive-form game between the profit-maximizing firm and the buyer. Fully specifying this equilibrium is difficult because the seller's strategy depends on the full vector of value draws. However, we note that the profit of the firm under this mixed equilibrium is only marginally higher than the profit it obtains by committing to the search-preventing default described above when n is not small. This is due to the fact that, in equilibrium, the buyer must be made indifferent between searching and not searching for each option outside the acceptance set. Thus, the probability that the seller "lies" by selecting a default outside of the acceptance set when the acceptance set is non-empty cannot be much higher than the probability that the acceptance set is empty.[24] Since the probability that A is empty is equal to $F(z^*(c))^n$, the probability of lies is typically small even for a moderate n.[25]

Finally, in our model with unawareness, the firm cannot do better by manipulating the order of search beyond setting the search-preventing default.[26] This is not the case when awareness exists: the profit-maximizing sequence would just be ordering products in the descending order of Π_i (i.e., from worst to best for the buyer if Assumption 1 holds). This type of manipulation is likely to be unrealistic in practice. Buyers faced with such aggressive strategies are likely to change their search behavior and firms who curate their offerings in such an aggressive way are unlikely to survive competition. By contrast, the search-preventing default we consider is a relatively subtle nudge and is often welfare-enhancing for

[24]Indeed, suppose the probability that the acceptance set is empty is ε, and the probability that the firm lies, i.e., offers a default below the cutoff when the acceptance set is non-empty, is λ. Then, having observed an offer below the cutoff, the buyer's belief that the acceptance set is non-empty will be updated to $\frac{\lambda(1-\varepsilon)}{\lambda(1-\varepsilon)+\varepsilon}$, which converges to 1 as $\varepsilon \to 0$ unless λ is at most $O(\varepsilon)$. Thus, for ε small, either the seller almost never lies or the buyer almost never accepts an offer below the cutoff. In either case, the welfare consequences are minimal.

[25]For illustration, consider an environment where each option is drawn independently from three equally likely values: 30, 80, and 100. If the cost of search is 20, the buyer's acceptance set is $A \subseteq \{80, 100\}$. Thus, if there are $n = 10$ options, the probability that no option is acceptable is equal to $(\frac{1}{3})^{10} \approx 0.0000167$. In the mixed strategy equilibrium of the game between the firm and the buyer, the probability that the buyer receives an offer of 30 when $|A| > 0$ is equal to 0.0000992. This is roughly 6 times larger than the probability that A is empty but still exceptionally small.

[26]Indeed, under any other ordering of the options the buyer would search until she encounters an option in A, and d is the best for the seller among those. For example, the seller could place inferior options ahead of d, which would make the buyer worse off, but the seller's profit would remain the same. From a behavioral viewpoint, considering incentives outside our model, choosing d as the default is optimal for a seller who cares, in a lexicographic manner, about the buyer's welfare alongside the seller's own. This is especially important in settings with repeated interactions where the seller may be interested in cultivating the buyer's trust.

buyers who do not lose too much value and economize on search costs.

C Belief updating under awareness

As mentioned in Section 3, throughout the paper we assume that buyers are boundedly rational in that they do not update their beliefs about the underlying distribution of values having observed the default option. In this section, we explore the consequences of such updating in more detail. While in general the answer depends on the shape of the distribution of values, for a large class of distributions, including the one used in our experiment, the consequences are minimal and confined to the cases where none of the n options are above the cutoff $z^*(c)$ (see footnote 11).

We note first that if a buyer observes a default with value below $z^*(c)$, she can conclude that all n options have benefits below $z^*(c)$ and refrain from further search. This would not be an equilibrium outcome if the seller were strategic; however, it is the buyer's optimal response in our experiment where sellers are automated.

Consider now a setting where a buyer observes a default option such that $b_1 > z^*(c)$. She can then conclude that there are no options with values in the interval $[z^*, b_1]$, and hence her updated distribution of values becomes

$$\hat{F}(b; c, b_1) = \frac{1}{1 - F(b_1) + F(z^*)} \times \begin{cases} F(b), & b \in [\underline{b}, z^*] \\ F(z^*), & b \in [z^*, b_1] \\ F(b) - F(b_1) + F(z^*), & b \in [b_1, \bar{b}] \end{cases} \tag{21}$$

This updated distribution produces a modified optimal cutoff $\hat{z}(c, b_1)$ that is, in general, different from $z^*(c)$, and depends on the observed realization of optimal default. Our main interest is in whether it is possible to have $\hat{z}(c, b_1) > b_1$, i.e., whether the belief updating can lead to additional search. The answer is negative in a wide range of cases, as stated in the following proposition.

Proposition 3 *Suppose the distribution of benefits F has an increasing hazard rate. Then $\hat{z}(c, b_1) \leq b_1$.*

Proof Assume, by contradiction, that $\hat{z} > b_1$. Using representation (15) for the optimal cutoff, we obtain that $\hat{z}$ must satisfy

$$\int_{\hat{z}}^{\bar{b}} [1 - \hat{F}(b; c, b_1)] \mathrm{d}b = c.$$

Equation (21) then gives

$$\int_{\hat{z}}^{\bar{b}} \left[1 - \frac{F(b) - F(b_1) + F(z^*)}{1 - F(b_1) + F(z^*)}\right] \mathrm{d}b = c,$$

which, using (15) to represent c, can be written as

$$\int_{\hat{z}}^{\bar{b}} [1 - F(b)]\mathrm{d}b = [1 - F(b_1) + F(z^*)] \int_{z^*}^{\bar{b}} [1 - F(b)]\mathrm{d}b. \tag{22}$$

The integral in the left-hand side of (22) is decreasing in $\hat{z}$ and equals zero for $\hat{z} = \bar{b}$. A necessary and sufficient condition for (22) to have a solution $\hat{z} > b_1$ is, therefore,

$$\int_{b_1}^{\bar{b}} [1 - F(b)]\mathrm{d}b > [1 - F(b_1) + F(z^*)] \int_{z^*}^{\bar{b}} [1 - F(b)]\mathrm{d}b. \tag{23}$$

We will now show that the opposite inequality holds. Splitting the integral in the right-hand side,

$$\int_{z^*}^{\bar{b}} [1 - F(b)]\mathrm{d}b = \int_{z^*}^{b_1} [1 - F(b)]\mathrm{d}b + \int_{b_1}^{\bar{b}} [1 - F(b)]\mathrm{d}b,$$

we further rewrite (23) as

$$[F(b_1) - F(z^*)] \int_{b_1}^{\bar{b}} [1 - F(b)]\mathrm{d}b > [1 - F(b_1) + F(z^*)] \int_{z^*}^{b_1} [1 - F(b)]\mathrm{d}b. \tag{24}$$

We can show, however, that

$$[F(b_1) - F(z^*)] \int_{b_1}^{\bar{b}} [1 - F(b)]\mathrm{d}b \leq [1 - F(b_1)] \int_{z^*}^{b_1} [1 - F(b)]\mathrm{d}b, \tag{25}$$

which contradicts (24). Rewrite (25) as

$$\int_{z^*}^{b_1} f(x)\mathrm{d}x \int_{b_1}^{\bar{b}} [1 - F(y)]\mathrm{d}y \leq \int_{b_1}^{\bar{b}} f(y)\mathrm{d}y \int_{z^*}^{b_1} [1 - F(x)]\mathrm{d}x,$$

or, equivalently,

$$\int_{z^*}^{b_1} \left[\int_{b_1}^{\bar{b}} [f(x)(1 - F(y)) - f(y)(1 - F(x))] \, \mathrm{d}y \right] \mathrm{d}x \leq 0. \tag{26}$$

Note that $x \leq y$ in the domain of integration. The increasing hazard rate property then implies

$$\frac{f(x)}{1-F(x)} \leq \frac{f(y)}{1-F(y)},$$

which implies the integrand in (26) is negative, and hence the inequality holds. ■

The class of increasing hazard rate distributions (also known in the reliability theory as *increasing failure rate*, or IFR) includes many standard distributions such as the uniform, normal, exponential, logistic, or Gumbel. Any log-concave distribution is IFR (but not the other way around). The IFR property is associated with sub-exponential (or *light*) tails. In contrast, distributions violating the IFR property are *heavy-tailed* distributions such as the Pareto (power laws) family.

Intuitively, heavy-tailed distributions are relatively more likely to produce large realizations away from the mean. In our context, this implies the updated distribution may have a larger weight shifted to the upper tail $[b_1, \bar{b}]$, and hence a buyer may expect sufficient benefits to justify additional search.

We conclude that in our experiment, where the distribution of benefits is uniform, belief updating would not affect our predictions about subjects' search behavior except in rare cases when the buyer's acceptance set is empty.

D Risk averse buyers

In this section, we consider risk averse buyers and discuss how the comparative statics predictions of Proposition 2 are affected by risk aversion. We provide numerical illustrations for CARA utility.

Suppose a buyer has a concave utility of money $u(\cdot)$ and initial endowment $w > 0$. Having searched s times, the buyer will search again if the best option revealed so far is below the threshold z that solves the equation

$$\int_z^{\bar{b}} [u(w-cs+b) - u(w-cs+z)]\mathrm{d}F(b) = u(w-cs) - u(w-cs-c). \tag{27}$$

Let z_s^* denote the solution, assuming it exists, which may now depend on the number of searches. For CARA utility, $u(x) = \frac{1}{r}(1-e^{-rx})$, (27) simplifies to

$$\frac{1}{r}\int_z^{\bar{b}} (e^{-rz} - e^{-rb})\mathrm{d}F(b) = \frac{1}{r}(e^{rc}-1),$$

i.e., z_s^* is independent of w and s. Following the derivation of (2) in Appendix A, we have the expected utilities under SPD and RD in the form

$$U_{\text{SPD}} = \int_{\underline{b}}^{z^*} u(w+b-(n-1)c)\mathrm{d}F(b)^n + \sum_{k=1}^{n} \binom{n}{k} F(z^*)^{n-k}[1-F(z^*)]^k \int_{z^*}^{\bar{b}} u(w+b)\mathrm{d}\tilde{F}_{(1:k)}(b)$$

and

$$U_{\text{RD}} = \int_{\underline{b}}^{z^*} u(w+b-(n-1)c)\mathrm{d}F(b)^n + \sum_{k=1}^{n} \binom{n}{k} F(z^*)^{n-k}[1-F(z^*)]^k \int_{z^*}^{\bar{b}} \mathbf{E}[u(w+b-cS_k)]\mathrm{d}\tilde{F}(b),$$

where the expectation in the second equation is over the number of searches S_k for a given $k = |A|$.[27] From the proof of Lemma 1, we have

$$\mathbf{E}[u(w+b-cS_k)] = \sum_{s=0}^{n-k} \frac{(n-k)!(n-s)!}{(n-k-s)!n!} \frac{k}{n-s} u(w+b-cs).$$

Therefore, for the CARA utility,

$$\begin{aligned}
&U_{\text{SPD}} - U_{\text{RD}} \\
&= \sum_{k=1}^{n} \binom{n}{k} F(z^*)^{n-k}[1-F(z^*)]^k \left[\int_{z^*}^{\bar{b}} \frac{1-e^{-r(w+b)}}{r} \mathrm{d}\tilde{F}_{(1:k)}(b) - \int_{z^*}^{\bar{b}} \frac{1-\mathbb{E}[e^{-r(w+b-cS_k)}]}{r} \mathrm{d}\tilde{F}(b) \right] \\
&= \frac{e^{-rw}}{r} \sum_{k=1}^{n} \binom{n}{k} F(z^*)^{n-k}[1-F(z^*)]^k \left[\mathbb{E}[e^{rcS_k}] \int_{z^*}^{\bar{b}} e^{-rb}\mathrm{d}\tilde{F}(b) - \int_{z^*}^{\bar{b}} e^{-rb}\mathrm{d}\tilde{F}_{(1:k)}(b) \right].
\end{aligned}$$

This expression has a structure similar to (2) showing the two competing effects: savings on search costs versus benefits from finding options with values exceeding the SPD. Figure 7 shows the rescaled utility difference for $r = 0.001$, 0.3 and 0.8. As expected, there is a crossing similar to the one predicted by Proposition 2; however, the positive effect of SPD for low-cost buyers declines with risk aversion.

E Defaults with fixed sample size search

In this section, we analyze the effect of the optimal default characterized in Proposition 1 when the buyer's search is non-sequential. As we show, the effect of default is qualitatively very similar to the case of sequential search, and hence our predictions are to a certain extent robust to assumptions about search.

[27] Note that for CARA utility z_s^* is independent of s and hence the distribution of the number of searches is the same as under risk neutrality.

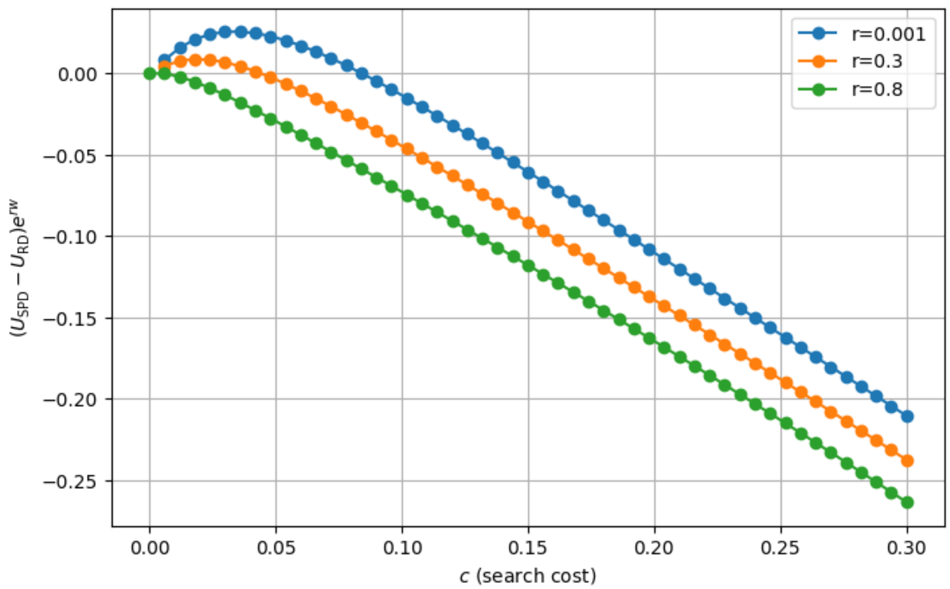

Figure 7: The (rescaled) utility difference $(U_{\text{SPD}} - U_{\text{RD}})e^{rw}$ as a function of search cost c for risk averse buyers with CARA utility. Parameters: $n = 10$, F is uniform on $[0, 1]$.

Specifically, we consider a fixed sample size search algorithm where, having observed a default option with benefit b_1, the buyer chooses a sample size $s \in \{0, \dots, n-1\}$, randomly draws s other options, at cost cs, and selects the option with the largest benefit among the $s+1$ options (including the default).

Conditional on b_1, the buyer's expected payoff from sampling s other options is

$$\pi_s(b_1) = \mathbb{E}(\max\{b_1, b_{(s:s)}\}|b_1) - cs = b_1 F(b_1)^s + \int_{b_1}^{\bar{b}} b \mathrm{d}F(b)^s - cs. \tag{28}$$

For $s \geq 1$, let $\Delta\pi_s(b_1) = \pi_s(b_1) - \pi_{s-1}(b_1)$ denote the payoff's first difference with respect to s. Integrating (28) by parts, we can write it as

$$\Delta\pi_s(b_1) = \int_{b_1}^{\bar{b}} F(b)^{s-1}[1 - F(b)]\mathrm{d}b - c. \tag{29}$$

As seen from (29), $\Delta\pi_s(b_1)$ is decreasing in s, implying that $\pi_s(b_1)$ is discrete concave in s. Additionally, $\Delta\pi_s(b_1)$ is decreasing in b_1 and in c, and hence, from the monotone comparative statics, so does $\arg\max_s \pi_s(b_1)$. The payoff-maximizing choice of s is, therefore, characterized as follows.

Proposition 4 *(i) There exists a generically unique optimal sample size*

$$s^*(b_1) = \begin{cases} 0, & \text{if } \Delta\pi_1(b_1) < 0 \\ \max\left\{s \in \{1, \ldots, n-1\} : \Delta\pi_s(b_1) \geq 0\right\}, & \text{otherwise} \end{cases}$$

(ii) $s^*(b_1)$ *is decreasing in* b_1 *and in* c.
(iii) $s^*(b_1) = 0$ *for* $b_1 > z^*$.

To verify part (iii), it is easy to show via integration by parts that $\Delta\pi_1(z^*) = 0$, where z^* is the solution of (1).

In the baseline setting where the default is random, the buyer's *ex ante* expected payoff is simply $\pi_{\text{RD}} = \int \pi_{s^*(b)}(b)\mathrm{d}F(b)$. When the default is chosen optimally as in Proposition 1, almost surely either (i) $b_{(n:n)} < z^*$, i.e., all the realized benefits are below the reservation value, and the default is chosen randomly; or (ii) $b_{(n:n)} > z^*$, in which case a default $b_1 > z^*$ is chosen, and hence, from part (iii) of Proposition 4, the optimal number of searches is zero as in the sequential search case. The buyer's expected payoff is then

$$\begin{aligned} \pi_{\text{SPD}} = F(z^*)^n \int_{\underline{b}}^{z^*} \left[b\underset{\sim}{F}(b)^{s^*(b)} + \int_b^z b' \mathrm{d}\underset{\sim}{F}(b')^{s^*(b)} - cs^*(b) \right] \mathrm{d}\underset{\sim}{F}(b) \\ + \sum_{k=1}^{n} \binom{n}{k} F(z^*)^{n-k}[1 - F(z^*)]^k \int_{z^*}^{\bar{b}} b\mathrm{d}\tilde{F}_{(1:k)}(b). \end{aligned} \tag{30}$$

Here, $\underset{\sim}{F}(b) = \frac{F(b)}{F(z^*)}$ is the conditional distribution of benefits below the reservation value. The first term represents the expected payoff from fixed sample size search when all options are below z^* (the probability of this event is $F(z^*)^n$), and the default is chosen randomly from $\underset{\sim}{F}$. However, the buyer is unaware of that and still selects the sample size according to the rule in part (i) of Proposition 4. The second term, similar to (5), sums up over all possible numbers of options, k, with benefits above z^*, for each of which $\tilde{b}_{(1:k)}$—the lowest order statistic from the updated distribution $\tilde{F}(b) = \frac{F(b)}{1-F(z^*)}$—is set as the optimal default.

As seen from (30), the same two effects are at play here as in the case of sequential search. When the cost of search is low, the optimal default is high, and the buyer benefits from it due to savings on search costs. That benefit, however, is small because as soon as the search cost rises enough for the probability of having more than one option above z^* to become nontrivial, those savings are outweighed by losses from the unavailability of the best options. As the cost of search approaches $\bar{c}$, the buyer does not search and the optimal default is set at the worst possible option, implying $\pi_{\text{SPD}} < \pi_{\text{RD}}$. By continuity, a c_H similar to the one identified in Proposition 2 then exists as well.

F Additional treatments to study the effect of awareness

As discussed in Section B, a buyer who is aware that the search-preventing default is being selected by the firm has no incentive to change their behavior for a large class of decision problems, which includes the uniform distribution of values that was used in Experiment 1. This differs from the effect of defaults that are designed to convey information or act as an endorsement for one action or choice over another.

To study the effect of awareness, we ran additional treatments duplicating the 2×2 design of Experiment 1, in which we **announced (A)** how the default was set in each round.[28] Comparing the baseline treatments of Experiment 1 to the corresponding Announcement treatments (encoded as HC-RD-A, HC-SPD-A, LC-RD-A and LC-SPD-A), we predict and find the following.

Hypothesis 5 *Announcements alerting the buyers to the default setting strategy of the seller do not have an effect.*

Result 8 *Consistent with Hypothesis 5, there are no significant changes in search, buyers' payoffs, or firms' profits when buyers are informed about the search-preventing default choice rule.*

Support for Result 8 is based on the comparison of Table 2, which provides summary statistics for the unawareness treatments of Experiment 1, and Table 6, which provides summary statistics for the treatments with awareness. Table 7 reports Wilcoxon nonparametric rank-sum tests comparing each Announcement treatment to the corresponding treatment of Experiment 1.[29] As with the analysis of costs, we average the outcome variable at the individual level and compare the resulting treatment-level distributions.

As seen in Panel A of Table 7, there are no significant differences in search between the treatments with and without Announcement. Likewise, as seen in Panels B and C, there are no significant differences in buyers' payoffs or sellers' profits.

[28]These treatments were pre-registered on AsPredicted under #96863. In the instructions, we stated that "in the other 50% of rounds, the initial offer is carefully chosen from among the 10 random offers by a profit-maximizing seller. The seller's profit is negatively related to the value of the offer that you eventually accept. The seller has selected the lowest offer that it believes you will accept without searching." We then informed participants in each round whether the default in that round was random or selected by a profit-maximizing seller.

[29]An alternative test would be to pair the data by value sequence and perform a Wilcoxon nonparametric sign-rank test across the pairs. As with the tests reported here, we cannot reject the null in any of the 12 tests under this alternative specification.

Table 6: Summary statistics and theoretical predictions in treatments with awareness

Treatments	# of searches		Buyer's payoff		Seller's profit		Subjects
	data	theory	data	theory	data	theory	
LC-RD-A	2.08	2.09	67.20	73.24	22.43	16.30	30
	(0.16)		(1.29)		(1.80)		
HC-RD-A	1.16	0.29	34.31	52.54	30.91	38.73	30
	(0.20)		(5.04)		(1.57)		
LC-SPD-A	1.49	0.20	73.14	75.48	19.39	23.52	30
	(0.22)		(0.64)		(0.67)		
HC-SPD-A	1.58	0.00	13.31	31.63	39.41	68.37	30
	(0.23)		(5.07)		(2.70)		

Robust standard errors clustered by subject in parentheses.

G Duration analysis of search

In this section we present a detailed analysis of subjects' search behavior using a duration model. Recall that sequential search with perfect recall entails choosing, at each stage, whether to search further at a cost c or accept the best available option discovered so far, which constitutes the state variable of search.

Following a standard approach to duration analysis with individual heterogeneity (see, e.g., Jenkins, 2005), we model each subject's hazard rate, i.e., the probability of acceptance in the current state conditional on not having accepted so far, using a proportional hazard model with *frailty*. The hazard rate of acceptance is modeled as $\theta(t, X_t|v) = \theta_0(t)\exp(\beta' X_t)v$, where t is a discrete time index corresponding to the rounds of search, X_t is a vector of (possibly time-varying) covariates, $\theta_0(t)$ is a baseline hazard that is the same for all individuals, and $v > 0$ is an unobserved individual effect (frailty) with unit mean and a finite variance. Assuming $\log(v)$ is normally distributed, we use a random effects cloglog regression, with robust standard errors clustered by subject, to estimate the vector of parameters β. The results are shown in Table 8.

As explanatory variables, we include *Search round* (the duration of search), *Best option* (the state), and the treatment controls. We estimate the model separately for the treatments in Experiment 1 (under unawareness) and in the additional treatments with awareness (with the announcement of SPD rounds). There is evidence of significant heterogeneity (frailty) in the data, with parameter ρ (a measure of heterogeneity variance) estimated at 0.288 and 0.323 under unawareness and Announcement, respectively ($p < 0.01$ in both cases).

As seen from Table 8, the results are very similar between the two conditions. Impor-

Table 7: Statistical Tests for Hypothesis 5

Panel A: Comparison of Search in Experiments 1 and 3					
Treatments Compared	z-score	p-value	N_1	N_2	Statistical Test
LC-RD vs LC-RD-A	1.27	0.20	30	30	Rank-Sum test
LC-SPD vs LC-SPD-A	0.34	0.73	30	30	Rank-Sum test
HC-RD vs HC-RD-A	0.49	0.63	30	30	Rank-Sum test
HC-SPD vs HC-SPD-A	0.59	0.55	30	30	Rank-Sum test
Panel B: Comparison of Buyers' Payoffs in Experiments 1 and 3					
Treatments Compared	z-score	p-value	N_1	N_2	Statistical Test
LC-RD vs LC-RD-A	0.86	0.39	30	30	Rank-Sum test
LC-SPD vs LC-SPD-A	−0.10	0.92	30	30	Rank-Sum test
HC-RD vs HC-RD-A	−0.27	0.79	30	30	Rank-Sum test
HC-SPD vs HC-SPD-A	−0.51	0.61	30	30	Rank-Sum test
Panel C: Comparison of Sellers' Profits in Experiments 1 and 3					
Treatments Compared	z-score	p-value	N_1	N_2	Statistical Test
LC-RD vs LC-RD-A	−1.37	0.17	30	30	Rank-Sum test
LC-SPD vs LC-SPD-A	−0.47	0.64	30	30	Rank-Sum test
HC-RD vs HC-RD-A	0.04	0.97	30	30	Rank-Sum test
HC-SPD vs HC-SPD-A	−0.25	0.80	30	30	Rank-Sum test

Statistical tests for Hypothesis 5. Panel A is a comparison of search; Panel B is a comparison of buyers' payoffs; Panel C is a comparison of sellers' profits. All tests are a nonparametric Wilcoxon rank-sum test with the two treatments treated as independent samples.

tantly, there is no evidence of duration affecting the hazard rate conditional on the state, which supports the "memoryless" search hypothesis. As expected, acceptance is more likely the higher the value of the best option encountered so far, and the higher the search cost. The search-preventing default condition has no significant effect on search, for a given search cost.

H Experimental instructions (Low Cost, treatments with awareness)

Welcome and thank you for participating in todays experiment. Please turn off your phone now and put it away. Please do not talk during the experiment. **If you have a question, please type it in the Zoom chatbox and send it only to the experimenter who will answer it**.

Your earnings in this experiment will depend on your own choices and on chance moves. Understanding the instructions is likely to increase your earnings. Earnings are private. You

Table 8: Duration Analysis of Search

	Unawareness	Announcement
Search round	-0.036	-0.049
	(0.034)	(0.034)
Best option	0.050***	0.050***
	(0.004)	(0.005)
High cost	0.912***	0.891***
	(0.216)	(0.239)
search-preventing default	0.025	-0.100
	(0.079)	(0.096)
(High cost)×(search-preventing default)	0.142	0.235*
	(0.125)	(0.133)
Intercept	-4.248***	-3.977***
	(0.394)	(0.426)
Observations	13,026	12,365
Subjects	60	60

Random effects complementary log-log (cloglog) regressions, normally distributed frailty, robust standard errors clustered by subject in parentheses. Each round of search is an observation. The dependent variable is an indicator equal one if the best option is accepted and zero otherwise. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

will receive a $5.00 participation fee. You will be asked to submit your PayID at the end of the experiment to receive the online transfer of your earnings.

The Choice Tasks and Payoffs: There are **80 decision rounds** in this experiment.

In each round, you will collect points depending on your choice and on chance. At the end of the experiment, <u>4 of the 80 rounds</u> will be chosen randomly and the points you earn in those rounds will be paid out to you with an exchange rate of $1 for every 100 points.

<u>Endowment</u>

At the beginning of each round, you will receive an endowment of 200 points.

<u>Offers</u>

For each round, the values for **10** offer cards were randomly generated by the computer.

All values are integers between and including 0 and 100. Each of these values is equally likely, and independent of others. These values differ from round to round

One of the offer cards is always chosen as the initial offer. The 9 remaining cards are reshuffled and placed in a random order every round. The initial offer card is always shown on the top.

If you like the initial offer, you can accept it immediately.

If you do not like the initial offer, you can reveal further offers.

How the initial offer is determined[30]

In 50% of rounds, selected randomly, the initial offer is random just like the remaining 9 (hidden) offers.

In the other 50% of rounds, the initial offer is carefully chosen from among the 10 random offers by a profit-maximizing seller. The seller's profit is negatively related to the value of the offer that you eventually accept. **The seller has selected the lowest offer that it believes you will accept without searching**.

Cost of revealing other offers

Revealing additional offers incurs a fixed cost of 5 per offer.

So your COST in a round equals (number of offers you reveal) × 5 cents

Accepting an offer and the round termination

You can always accept the best offer revealed so far. Accepting an offer terminates the round.

Your Payoffs

Your payoff in a round equals:

200 + BEST OFFER at the moment you accept − YOUR COST in that round

In this experiment you will get information about the best offer so far, about the cost of the offers so far and about what your payoff would be if you accepted at that moment. Here is an example:

[30] This section is absent in Experiment 1.

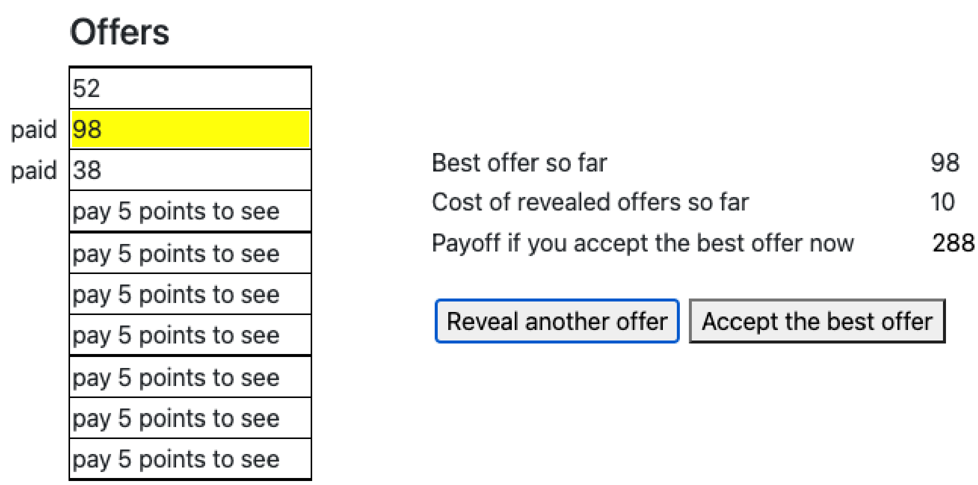

If you have any questions at this time, please message one of the experimenters using the chat function in Zoom. If there are no questions, please go ahead and complete a quiz.

Note that the instructions will remain on your screen for future reference.

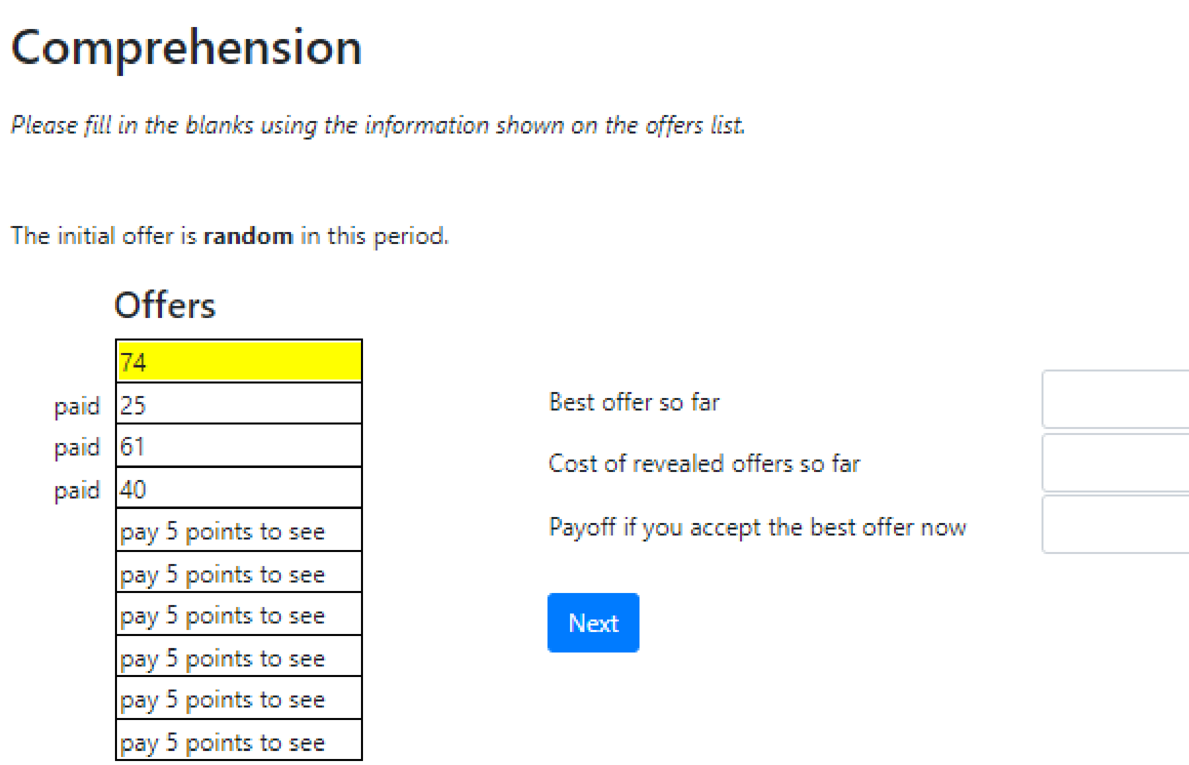

Actual decision screens (a SPD round with announcement)

(a) The beginning of a round

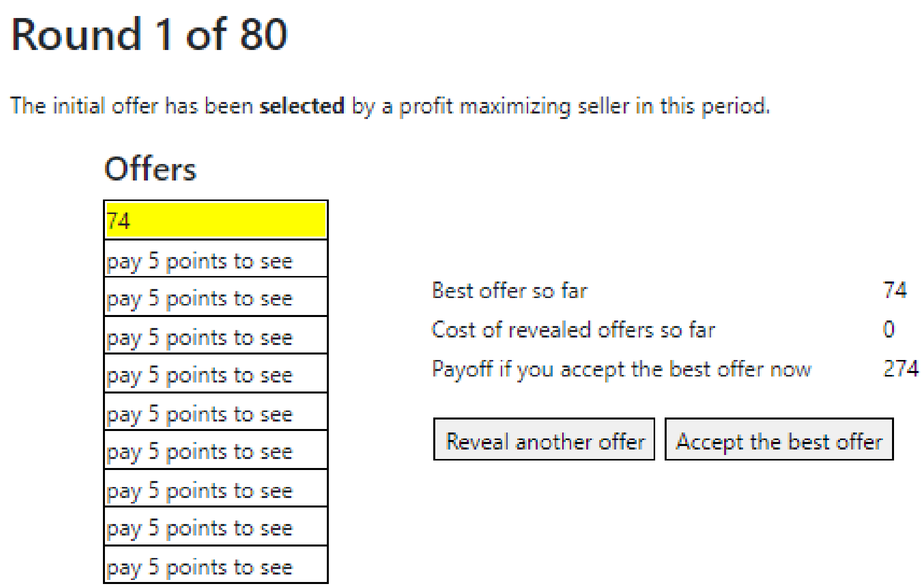

(b) After 3 searches