Optimal Search and Discovery*

Rafael P. Greminger†

April 9, 2020

Abstract

This paper develops a search problem where a consumer initially is aware of only a few products. To find a good match, the consumer sequentially decides between searching among alternatives he is already aware of and discovering more products. I show that the optimal policy for this search and discovery problem is fully characterized by tractable reservation values. Moreover, I prove that a predetermined index fully specifies the purchase decision of a consumer following the optimal search policy. Finally, a comparison highlights differences to classical random and directed search.

1 Introduction

In many settings, consumers have limited information and first need to search for product information before being able to compare alternatives. The resulting search frictions have received considerable attention in the literature.[1] Under the rational choice paradigm, the analysis of such settings relies on optimal search policies that describe how a consumer with limited information optimally searches among the available alternatives. I add to this literature by developing and solving a sequential search problem that introduces a novel aspect: limited awareness.

To fix ideas, consider a consumer looking to buy a mobile phone. Through advertising or recommendations from friends, the consumer initially is aware of a single available phone and has

*This paper previously was circulated under the title „Optimal Search and Awareness Expansion".

†I am deeply grateful to my advisors, Tobias Klein and Jaap Abbring, for their thoughtful guidance and support. I also thank Bart Bronnenberg and members of the Structural Econometrics Group in Tilburg for excellent comments. Finally, I thank the Netherlands Organisation for Scientific Research (NWO) for financial support through Research Talent grant 406.18.568. Tilburg University, Department of Econometrics & Operations Research, r.p.greminger@tilburguniversity.edu.

[1] For example, Stigler (1961); Diamond (1971); Burdett and Judd (1983); Anderson and Renault (1999); Kuksov (2006); Choi et al. (2018); Moraga-González et al. (2017a,b) study search frictions in equilibrium models and Hortaçsu and Syverson (2004); Hong and Shum (2006); De Los Santos et al. (2012); Bronnenberg et al. (2016); Chen and Yao (2017); Zhang et al. (2018); Jolivet and Turon (2019) study implications of search empirically.

some (but not all) information on what it offers. Given this basic information, the consumer can directly gather more detailed information on this alternative, for example by reading a review online. Besides, there are also phones available that the consumer initially is not aware of. For these alternatives, the consumer knows neither about their existence, nor the features they offer. This precludes the consumer from directly inspecting these phones. Instead, he first needs to discover and become aware of them, for example by getting more recommendations from friends or through a search intermediary. Figure 1 depicts a possible choice sequence for this case.

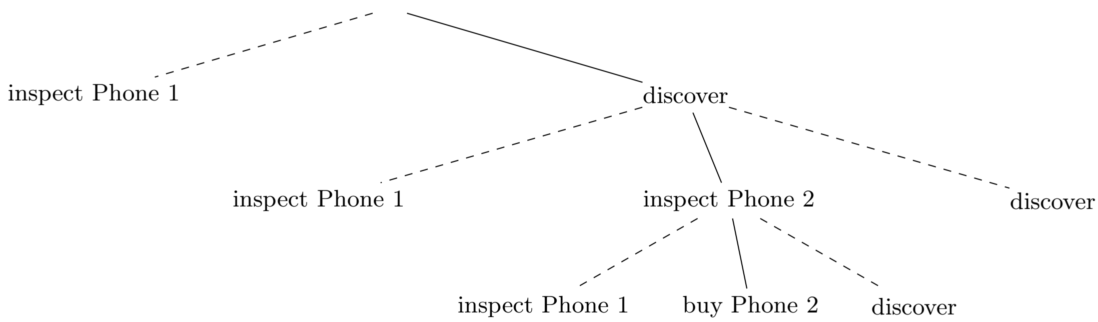

FIGURE 1 – Example of a choice sequence in the search and discovery problem.

The „search and discovery problem" introduced in this paper formalizes the consumer's dynamic decision process in this and similar settings. It nests classical random and directed sequential search as special cases but generalizes both by introducing limited awareness. In random search, a searcher has no prior information, searches randomly across alternatives and decides when to end search (e.g. McCall, 1970; Lippman and McCall, 1976). In directed search, the searcher is aware of all available alternatives and uses partial product information to determine an order in which to inspect products and when to end search (e.g. Weitzman, 1979; Chade and Smith, 2006). In contrast, in the search and discovery problem, the consumer is aware of only a few products. Hence, he not only decides in what order to inspect products and when to end search, but also when to try to discover more alternatives.

The resulting framework allows to study settings that are difficult to accommodate in existing search problems. In particular, neither random nor directed search is well suited to study settings where rational consumers remain oblivious to some, while obtaining only partial information on other products. However, such settings are common in practice. For example, in markets with a large number of alternatives, consumers remain unaware of many alternatives unless they actively set out to discover products beyond those they are already aware of. Similarly, in markets where rapid technological innovations lead to a constant stream of newly available alternatives, few consumers are aware of new releases without exerting effort to remain informed. Moreover, consumers also face a search and discovery problem when buying products online.

Online retailers and search intermediaries present alternatives on a product list that reveals partial information only for some products. Consumers then decide between clicking on products to reveal full information, and browsing further along the list to discover more products.

The contribution of this paper is to show that despite its complexity, optimal search decisions and outcomes in the search and discovery problem remain tractable. First, I prove that the optimal policy is fully characterized by reservation values similar to the well-known reservation prices derived by Weitzman (1979). In each period, a reservation value is assigned to each available action, and it is optimal to always choose the action with the largest value. Each of the reservation values is independent of any other available action and can be calculated without having to consider expectations over a myriad of future periods. Hence, reservation values remain tractable. This allows to determine optimal search behavior under limited awareness without using numerical methods.

Second, I prove that the purchase of a consumer solving the search and discovery problem is equivalent to the same consumer having full information and directly choosing products from a predetermined index. This result generalizes the „eventual purchase theorem" of Choi et al. (2018) to the case of limited awareness.[2] Similar to the eventual purchase theorem, my generalization allows to derive a consumer's expected payoff and market demand without having to consider a multitude of possible choice sequences that otherwise make aggregation difficult.

To provide further details on how the search and discovery problem relates to existing search frameworks, I compare a consumer's stopping decisions, expected payoff and the resulting market demand with classical random and directed sequential search. The comparison highlights several differences and reveals how limited awareness and the availability of partial product information determine search outcomes. Throughout, the comparison focuses on the case where consumers discover products one at a time.

An immediate difference between the three search problems is that having two distinct search actions posits a novel question: Do consumers benefit more from making it easier to discover more alternatives (e.g. through search intermediaries), or from facilitating inspection by more readily providing detailed product information? I show that when the number of available alternatives exceeds a (possibly small) threshold, the expected payoff increases more when facilitating discovery instead of facilitating inspection. This reflects that discovery costs become more important as the number of alternatives grows.[3]

[2] Choi et al. (2018) note that „Our eventual purchase theorem was anticipated by Armstrong and Vickers (2015) and has been independently discovered by Armstrong (2017) and Kleinberg et al. (2017)."

[3] A reduction in inspection costs can always be more beneficial when more than one product is discovered at a time. In this case, the consumer will, in expectation, inspect more products, making inspection costs relatively more important.

The difference in the consumer's initial information between a directed search and a search and discovery problem generates distinct patterns in the resulting market demand. In directed search, more consumers preferring a product based on partial product information will increase market demand. This need not be the case in a search and discovery problem. If consumers remain unaware of a product, its market demand does not increase as it becomes the preferred option for many consumers.

The search and discovery problem generates ranking effects; products that are discovered later are less likely to be bought. This results directly from consumers stopping search before discovering all products. I show that these ranking effects are independent of the number of available alternatives, and decrease as more products are discovered. The latter implies that the demand increase of being discovered first instead of second is larger than when being discovered as the 100th instead of 101st product. This mechanism offers a meaningful interpretation of how advertising that informs consumers of a product's existence is beneficial for a seller:[4] If a seller's marketing efforts make more consumers aware of a product before search or increase the probability of the product being discovered early on, ranking effects directly imply that they will increase the demand.

A directed search problem does not entail a similar mechanism. As consumers are aware of all products from the outset, advertising in the sense of informing consumers about the existence of a product cannot directly be considered in directed search. Nonetheless, a directed search problem can also generate ranking effects by assuming differences in inspection costs. I show that in contrast to the search and discovery problem, this yields that ranking effects increase in the number of available alternatives, as well as in a product's position. Moreover, the analysis reveals that ranking effects generated through differences in inspection costs are sensitive to a cost specification that can be difficult to interpret in practice.

Finally, both a search and discovery and a random search problem imply that consumers may not buy products they initially are unaware of. However, in random search the consumer either has full or no information about a product. He therefore cannot use partial information to decide whether to inspect a product. I show that this reduces the benefits of continuing search and the consumer's expected payoff when total costs of revealing full product information remain the same. This suggests that consumers discover more products and are better off when discovering products on a product list that reveals only partial information instead of inspecting each product in detail before discovering the next.

The remainder of this paper is organized as follows. First, I discuss related search problems.

[4] This relates to the „informative view" of advertising. See e.g. Bagwell (2007) for a summary and comparison to the „persuasive view".

Section 2 introduces the search and discovery problem. The basic setup shares assumptions with recent consumer search models (e.g. Armstrong, 2017; Choi et al., 2018) that facilitate exposition, but are relaxed in several extensions. Section 3 provides the optimal policy and discusses these extensions. In Section 4, I generalize the eventual purchase theorem of Choi et al. (2018) and use this to derive a consumer's expected payoff as well as market demand. Section 5 compares stopping decisions, the resulting ranking effects and the consumer's expected payoff across search problems. Finally, Section 6 concludes. Throughout, proofs are deferred to Appendix A.

1.1 Related Search Problems

The search and discovery problem nests several existing sequential search problems. Most notable is Pandora's problem introduced by Weitzman (1979), which results when the consumer initially is aware of all alternatives, or when there are no costs to discover more alternatives.[5]

To prove the optimality of the reservation value policy, I use a different approach than early contributions to search problems. Instead of directly proving that following the reservation value policy maximizes expected total payoff, I use results from the multi-armed bandit literature to first determine that a Gittins index policy is optimal,[6] and then show that the Gittins index reduces to the simple reservation values. More specifically, I use the results of Keller and Oldale (2003), who showed that a Gittins index policy is optimal in a dynamic decision problem where choosing one action reveals more actions while leaving the state of other available actions unchanged. Introducing a monotonicity condition then allows to generalize their result to the case where with some probability no new actions are revealed and to show that the Gittins index can be calculated based on the myopic comparison of payoffs that defines the reservation values.

Similar monotonicity conditions also apply in some multi-armed bandit problems where they simplify the otherwise difficult calculation of the Gittins index values (see e.g. Section 2.11 in Gittins et al., 2011). The present case differs in that monotonicity is only required for the action of discovering more alternatives, but does not hold when inspecting a product. Furthermore, in a recent working paper, Fershtman and Pavan (2019) independently discovered a similar characterization of the optimal policy when applying a „better-later-than-sooner" property in their multi-armed bandit framework with endogenous arms.

Several other contributions extend Weitzman's (1979) seminal search problem in different directions. Adam (2001) studies the case where the searcher updates beliefs about groups of alternatives during search and finds a similar reservation value policy to be optimal. Olszewski and Weber

[5] Others include classical stopping problems such as those considered in McCall (1970) or Lippman and McCall (1976).

[6] Gittins et al. (2011) provide a textbook treatment of multi-armed bandit problems and the Gittins index policy. As purchasing a product ends search, search problems correspond to stoppable superprocesses as introduced by Glazebrook (1979).

(2015) generalize Pandora's rule to search problems where the final payoff depends on all the alternatives that have been inspected, not only the best one. Finally, Doval (2018) analyzes the optimal policy when a searcher can directly choose alternatives without first inspecting them.

Other studies have estimated structural models based on search problems that are suited for the particular environment. Most closely related is a recent working paper by Choi and Mela (2019), where consumers also decide to reveal more products. The results presented in this paper differ in two important ways. First, I provide a tractable optimal policy based on reservation values, whereas Choi and Mela (2019) use numerical value function iteration to solve for the optimal policy. Second, the search problem I study is more general; it does not limit the decision of the consumer to only inspecting the latest revealed product, accommodates a finite number of alternatives and does not require the consumer to know the number of available alternatives. Note, however, that the optimal policy for the search and discovery problem implies that as long as the consumer has not yet revealed the last alternative, it will never be optimal to go back and inspect a product that was discovered earlier. Hence, this paper provides a justification for studying a reduced search problem where the consumer is precluded from inspecting a product that was revealed earlier.

Besides, Koulayev (2014) estimates a search model where consumers also decide whether to reveal more products. In his model, however, revealing a product also reveals all information on that product. Hence there is no need for inspecting a product as considered in this paper.[7]

Finally, this paper is not the first to consider the concept of limited awareness. Honka et al. (2017) and Morozov (2019) estimate structural search models where consumers also cannot inspect products that they are not aware of. However, in their models, consumers cannot discover additional alternatives. The underlying search problem in these models thus is equivalent to Pandora's problem introduced by Weitzman (1979).

2 The Search and Discovery Problem

A risk-neutral consumer with unit demand faces a market offering a (possibly infinite)[8] number of products gathered in set J. Alternatives are heterogeneous with respect to their characteristics. The consumer has preferences over these characteristics which can be expressed in a utility ranking. To simplify exposition and facilitate a comparison to existing models from the consumer search literature (e.g. Armstrong, 2017; Choi et al., 2018), I assume that the consumer's ex post

[7] Koulayev (2014) solves the dynamic decision problem using numerical backwards induction. For the case where costs are increasing in time (which is the case in his results), the present results suggest that a simple index policy also characterizes the optimal policy for his model.

[8] The problems with infinitely many arms in a multi-armed bandit problem discussed by Banks and Sundaram (1992) do not arise in the present setting.

utility when purchasing alternative j is given by

$$u(x_j, y_j) = x_j + y_j \tag{1}$$

where x_j and y_j are valuations derived from two distinct sets of characteristics. Note, however, that the results presented continue to hold for more general specifications that do not rely on linear additive utility.[9] An outside option of aborting search without a purchase offering u_0 is available.

The consumer has limited information on available alternatives. More specifically, in any period t, the consumer knows both valuations x_j and y_j only for products in a consideration set $C_t \subseteq J$. For products in an awareness set $S_t \subseteq J$, the consumer only knows partial valuations x_j. This captures the notion that if the consumer is aware of a product, he has received some information on the total valuation of the product. Finally, the consumer has no information on any other product $j \in J \backslash (S_t \cup C_t)$.

During search, the consumer gathers information by sequentially deciding which action to take in periods $t = 0, \dots, T$. If the consumer decides to discover more products, n_d alternatives are added to the awareness set. If less than n_d alternatives have not yet been revealed, only the remaining alternatives are revealed. For each of the n_d alternatives, the partial valuation x_j is revealed. To reveal the remaining characteristics of a product j, summarized in y_j, the consumer has to inspect the product. This reveals full information on the product and moves it from the awareness into the consideration set. The latter implies $S_t \cap C_t = \emptyset$.

The order in which products are discovered is tracked by positions $h_j \in \{1, \dots\}$, where a smaller position indicates that a product is discovered earlier. Throughout, it is assumed that products are discovered in increasing order of their index. Whereas this assumption is without loss of generality when focusing on the consumer's search problem, it is worth noting that in equilibrium settings, the order may be determined by sellers' actions, requiring a careful analysis of how these will determine the consumer's beliefs.[10]

Two precedence constraints on the consumer's actions are imposed. First, the consumer can only buy products from the consideration set. Second, the consumer can only inspect products from the awareness set. Whereas the first constraint is used in virtually all search problems,[11] the latter is novel to the proposed search problem. It implies that a product cannot be inspected,

[9] Specifically, suppose that when the consumer becomes aware of alternative j, he reveals a signal on the distribution from which the utility of j will be drawn. Appropriately defining the distribution of signals and the distribution of utilities conditional on these signals then yields an equivalent search problem.

[10] For example, in online settings it is common for sellers to bid on the position at which their product adverts are shown (see e.g. Athey and Ellison, 2011).

[11] Doval (2018) is a notable exception.

unless the consumer is aware of it. In an online setting where a consumer browses through a list of products, this constraint holds naturally: Individual product pages are reached by clicking on the respective link on the list. Hence, unless a product has been revealed on the list, it cannot be clicked on. In other environments, this precedence constraint reflects that, unless a consumer knows whether an alternative exists, he will not be able to direct search efforts and inspect the specific alternative. For example, if a consumer is not aware of a newly released phone model, he will not be able to directly acquire detailed information on it, before becoming aware that it exists.

Given the setting and these constraints, the consumer decides sequentially between the following actions:

i) Purchasing any product from the consideration set C_t and ending the search.

ii) Inspecting any product from the awareness set S_t, thus revealing y_j for that product and adding it to the consideration set.

iii) Discovering n_d additional products, thus revealing their partial valuations x_j and adding them to the awareness set.

These actions are gathered in the *set of available actions*, $A_t = C_t \cup S_t \cup d$, where d indicates discovery. If a consumer chooses an action $a = j \in C_t$, he buys product j, whereas if he chooses an action $a = j \in S_t$, he inspects product j. To clearly differentiate between the different types of actions, this set can also be written as $A_t = \{b0, b3, s4, \ldots, d\}$, where bj indicates purchasing and sj inspecting product j.

Both inspecting a product and discovering more products is costly. Inspection and discovery costs are denoted by $c_s > 0$ and $c_d > 0$ respectively. These costs can be interpreted as the cost of mental effort necessary to evaluate the newly revealed information, or an opportunity cost of the time spent evaluating the new information. In line with this interpretation, I assume that there is free recall: Purchasing any of the products from the consideration set does not incur costs, and c_s is the same for inspecting any of the products in the awareness set.

The consumer has beliefs over the products that he will discover, as well as the valuation he will reveal when inspecting a product j. In particular, x_j and y_j are independent (across j) realizations from random variables X and Y, where the consumer has beliefs over their joint distribution. This implies that the consumer believes that in expectation, products are equivalent. A generalization where the distribution of X depends on index j is discussed in Section 3.

The consumer also has beliefs over the total number of available alternatives. I assume that the consumer believes that with constant probability $q \in [0, 1]$, the next discovery will be the last.[12] As shown in the next section, the optimal policy is independent of the number of remaining discoveries that may be available in the future. Note, however, that this belief specification implicitly assumes that the consumer always knows whether he can reveal n_d more alternatives. An extension presented in Section 3 covers the case where the consumer does not know how many alternatives will be revealed.

All information the consumer has in period t is summarized in the information tuple $\Omega_t = \langle \bar{\Omega}, \omega_t \rangle$. The tuple $\bar{\Omega} = \langle u(x,y), n_d, c_d, c_s, G_X(x), F_{Y|X=x}(y), q \rangle$ represents the consumer's initial knowledge. It contains the utility function, how many products are discovered, and the different costs. It also contains the consumer's beliefs, denoted by the respective cumulative densities $G_X(x)$ and $F_{Y|X=x}(y)$. The latter specifies the cumulative density of Y, conditional on the realization of X, which is observed by the consumer before choosing to inspect a product. As a short-hand notation, I use $G(x)$ and $F(y)$ for these distributions. As a regularity condition, it is assumed that both $G(x)$ and $F(y) \forall x$ have finite mean and variance. Finally, $\bar{\Omega}$ contains the belief q on whether future discoveries will be available.

During search, the consumer reveals valuations x_j and y_j for the various products. This information is tracked in the set ω_t, containing realizations x_j for $j \in S_t \cup C_t$ and y_j for $j \in C_t$. The set of available actions A_t and the information tuple Ω_t capture the state in t. Figure 2 shows their transitions starting from period $t = 0$. The depicted example assumes that there are only two alternatives available and that products are discovered one at a time. If the consumer initially chooses the outside option ($b0$), no new information is revealed, and no further actions remain. If the consumer instead reveals the first alternative, he can inspect it in $t = 1$.

2.1 The Consumer's Dynamic Decision Problem

The setting above describes a dynamic Markov decision process, where the consumer's choice of action determines the immediate rewards, as well as the state transitions. The state in t is given by Ω_t and A_t. As the valuations x_j and y_j can take on any (finite) real values, the state space in general is infinite.[13] Throughout, the state space is defined such that it does not explicitly depend on time.

The consumer's problem consists of finding a feasible sequential policy, which maximizes the expected payoff of the whole decision process. A feasible sequential policy selects actions $\{a_0 \in A_0, a_1 \in A_1, \ldots\}$ for periods $t = 0, 1, \ldots T$ based only on information available in period t.

[12] Alternatively, one can specify these beliefs as a distribution over natural numbers. In this case, however, it becomes necessary to specify how the consumer updates his beliefs during search. In doing so, results hold only if the consumer updates beliefs such that the monotonicity condition (30) presented in Appendix C is satisfied (see also the discussion in Section 3).

[13] An exception is when x_j and y_j are drawn from discrete distributions, which limits the number of possible valuations that can be observed.

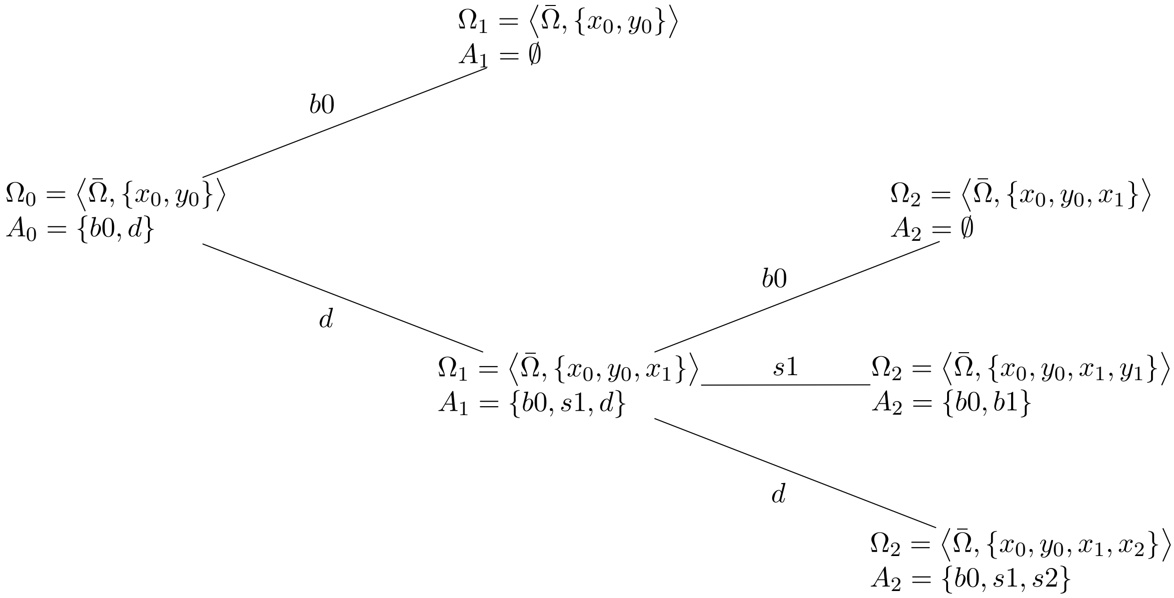

FIGURE 2 – Transition of state variables Ω_t (information tuple) and A_t (set of available actions) for $n_d = 1$ and $|J| = 2$.

Let Π denote the set containing all feasible policies. Formally, the consumer solves the following dynamic programming problem at period $t = 0$

$$\max_{\pi \in \Pi} V_0(\Omega_0, A_0; \pi) \tag{2}$$

where $V_t(\Omega_t, A_t; \pi)$ is the value function in period t, defined as the expected total payoff of following policy π starting in period t, conditional on the state in t. Let

$$[B_a V_t](\Omega_t, A_t; \pi) = R(a) + \mathbb{E}_t [V_{t+1}(\Omega_{t+1}, A_{t+1}; \pi)|a] \tag{3}$$

denote the Bellman operator, where the immediate rewards $R(a)$ either are inspection costs, discovery costs, or the total valuation of a product j if it is bought. Immediate rewards $R(a)$ therefore are known for all available actions. $\mathbb{E}_t [V_{t+1}(I_{t+1}, A_{t+1}; \pi)|a]$ denotes the expected total payoff over the whole future, conditional on policy π and having chosen action a in t.[14] The expectations operator integrates over the respective distributions of X and Y. A purchase in t ends search such that $A_{t+1} = \emptyset$ and $\mathbb{E}_t [V_{t+1}(\Omega_{t+1}, \emptyset; \pi)|a] = 0$ whenever $a \in C_t$. The corresponding Bellman equation is given by

$$V_t(\Omega_t, A_t; \pi) = \max_{a \in A_t} [B_a V_t](\Omega_t, A_t; \pi) \tag{4}$$

[14] In this formulation of the problem, the consumer does not discount future payoffs. This is in line with the consumer search literature, which usually assumes a finite number of alternatives without discounting. However, it is straightforward to show that the results continue to hold if a discount factor $\beta < 1$ is introduced. In this case, the search and discovery values defined in the next section need to be adjusted accordingly.

3 Optimal Policy

The optimal policy for the search and discovery problem is fully characterized by three reservation values. In what follows, I first define these reservation values, before stating the main result. At the end of this section, I discuss possible extensions based on a monotonicity condition, as well as limitations resulting from the condition that available actions need to be independent.

As in Weitzman (1979), suppose there is a *hypothetical* outside option offering utility z. Furthermore, suppose the consumer faces the following comparison of actions: Immediately take the outside option, or inspect a product with known x_j and end search thereafter. In this decision, the consumer will choose to inspect alternative j whenever the following holds:

$$Q_s(x_j, c_s, z) \equiv \mathbb{E}_Y \left[\max\{0, x_j + Y - z\}\right] - c_s \geq 0 \tag{5}$$

$Q_s(x_j, c_s, z)$ defines the expected *myopic* net gain of inspecting product j over immediately taking the outside option. If the realization of Y is such that $x_j + y_j \leq z$, the consumer takes the hypothetical outside option after inspecting j and the gain is zero. When $x_j + y_j > z$, the gain over immediately taking the hypothetical outside option is $x_j + y_j - z$. The expectation operator $\mathbb{E}_Y[\cdot]$ integrates over these realizations.

The *search value* of product j, denoted by z_j^s, then is defined as the value offered by a hypothetical outside option that makes the consumer indifferent in the above decision problem. Formally, z_j^s satisfies

$$Q_s(x_j, c_s, z_j^s) = 0 \tag{6}$$

which has a unique solution (see Lemma 1 in Adam, 2001). The search value can be calculated as

$$z_j^s = x_j + \xi_j \tag{7}$$

where ξ_j solves $\int_{\xi_j}^{\infty} [1 - F(y)]\, \mathrm{d}y - c_s = 0$ (see Appendix B).

The *purchase value* of product j, denoted by z_j^b, is defined as the utility obtained when buying product j:

$$z_j^b = u(x_j, y_j) \tag{8}$$

Based on reservation values given by (6) and (8), Weitzman (1979) showed that it cannot be optimal to inspect a product that does not offer the largest search value, or to stop when the largest remaining search value exceeds the largest purchase value. Hence, for given S_t and C_t, it is optimal to always inspect and buy in decreasing order of search and purchase values. However, this rule does not fully characterize an optimal policy in the search and discovery problem, as the consumer can additionally discover more alternatives.

For this additional action, a third reservation value based on a similar myopic comparison is introduced. Suppose the consumer faces the following comparison of actions: Take a hypothetical outside option offering z immediately, or discover more products and then search among the newly revealed products. The consumer will choose the latter whenever the following holds:

$$Q_d(c_d, c_s, z) \equiv \mathbb{E}_{\boldsymbol{X}} \left[V \left(\left\langle \bar{\Omega}, \omega(\boldsymbol{X}, z) \right\rangle, \{b0, s1, \ldots, sn_d\}; \tilde{\pi} \right) \right] - z - c_d \geq 0 \tag{9}$$

where $\omega(\boldsymbol{X}, z) = \{z, x_1, \ldots, x_{n_d}\}$ denotes the information the consumer has after revealing the n_d more alternatives. Note that with some abuse of notation, product indices were adjusted to the reduced decision problem, such that $j = 0, 1, \ldots, n_d$ indicates the hypothetical outside option and the newly revealed products.

$Q_d(c_d, c_s, z)$ defines the *myopic* net gain of discovering more products and optimally searching among them, over immediately taking the outside option. It is myopic in the sense that it ignores the option to continue searching beyond the products that are discovered. In particular, note that $V \left(\left\langle \bar{\Omega}, \omega(\boldsymbol{X}, z) \right\rangle, \{b0, s1, \ldots, sn_d\}; \tilde{\pi} \right)$ is the value function of having an outside option offering z and optimally inspecting alternatives for which partial valuations in $\boldsymbol{X}$ are known. Possible future discoveries and any products in S_t or C_t are excluded from the set of available actions in this value function. Finally, $\mathbb{E}_{\boldsymbol{X}}[\cdot]$ defines the expectation operator integrating over the joint distribution of the partial valuations in $\boldsymbol{X}$. Formal details on the calculation of the expectations and the value function are provided in Appendix B.

As for the search value, let the *discovery value*, denoted by z^d, be defined as the value of the hypothetical outside option that makes the consumer indifferent in the above decision. Formally, z^d is such that

$$Q_d(c_d, c_s, z^d) = 0 \tag{10}$$

which has a unique solution. In the case where Y is independent of X, the discovery value can be calculated as

$$z^d = \mu_X + \Xi(c_s, c_d) \tag{11}$$

where μ_X denotes the mean of X and $\Xi(c_s, c_d)$ solves (10) for an alternative random variable $\tilde{X} = X - \mu_X$. Further details for the calculation are provided in Appendix B.

Theorem 1 provides the first main result. It states that the optimal policy for the search problem reduces to three simple rules based on a comparison of the search, purchase and discovery values. In particular, the rules imply that in each period t, it is optimal to take the action with the largest reservation value defined in (6), (8), and (10). Hence, despite being fully characterized by myopic comparisons to a hypothetical outside option, these reservation values rank the expected payoffs of actions over all future periods.

Theorem 1. *Let $\tilde{z}^b(t) = \max_{k \in C_t} u(x_k, y_k)$ and $\tilde{z}^s(t) = \max_{k \in S_t} z_k^s$ denote the best products in the consideration and awareness set in period t. An optimal policy for the search and discovery problem is characterized by the following three rules:*

Stopping rule*: Purchase $j \in C_t$ and end search whenever $z_j^b = \tilde{z}^b(t) \geq \max\left\{\tilde{z}^s(t), z^d\right\}$.*

Inspection rule*: Inspect $j \in S_t$ whenever $z_j^s = \tilde{z}^s(t) \geq \max\left\{\tilde{z}^b(t), z^d\right\}$.*

Discovery rule*: Discover more products whenever $z^d \geq \max\left\{\tilde{z}^b(t), \tilde{z}^s(t)\right\}$.*

The proof of Theorem 1 relies on results from the literature on multi-armed bandit problems. In particular, it starts by assuming that the number of alternatives is known prior to search. For this case, Keller and Oldale (2003) showed that a Gittins index policy is optimal. Using a monotonicity condition, it is then shown that the Gittins index is independent of the availability of future discoveries, implying that the policy is optimal independent of the consumer's knowledge of the number of alternatives. Finally, it is shown that the monotonicity condition holds in the proposed search and discovery problem and that the Gittins index is equivalent to the simple reservation values defined above.

Based on Theorem 1, optimal search behavior can be analyzed using only (6), (8) and (10). Weitzman (1979) showed that search values decrease in inspection costs and increase if larger realizations y_j become more likely through a shift in the probability mass of Y. The same applies to the discovery value. It decreases in discovery costs and increases if probability mass of X is shifted towards larger values. The discovery value also depends on inspection costs and the conditional distribution of Y through the value function; it decreases in inspection costs and increases if larger values of Y are more likely.

To see the latter, consider the case where alternatives are discovered one at a time. In this case, the myopic net gain of discovering more products reduces to

$$Q_d(c_d, c_s, z) = \mathbb{E}_X\left[\max\left\{0, Q_s(X, c_s, z)\right\}\right] - c_d \tag{12}$$

For any $c_s' > c_s$, it holds that $Q_s(x, c_s', z) \leq Q_s(x, c_s, z)$ for all finite values of x and z, implying that $Q_d(c_d, c_s', z) \leq Q_d(c_d, c_s, z)$ for all z. As $Q_d(c_d, c_s, z)$ is decreasing in z (see Appendix A.1), it follows that the respective discovery values satisfy $z^{e\prime} \leq z^d$.

The optimal policy being fully characterized by simple rules leads to straightforward analysis of optimal choices for any given awareness and consideration sets. For example, consider a period t where $\max\left\{z^d, \tilde{z}^s(t)\right\} < \tilde{z}^b(t)$ such that the consumer stops searching. When decreasing inspection costs sufficiently in this case, the inequality reverts and the consumer will instead either first discover more products, or inspect the best product from the awareness set.

3.1 Monotonicity and Extensions

For the reservation value policy of Theorem 1 to be optimal, the discovery value needs to fully capture the expected net benefits of discovering more products, including the option value of being able to continue discovering products. The monotonicity condition used in the proof of the theorem ensures that this holds. It states that the expected net benefits of discovering more products do not increase during search. Hence, whenever the consumer is indifferent between taking the hypothetical outside option and discovering more products in t, he will either continue to be indifferent or take the outside option in $t+1$. Whether the consumer can continue to discover products in $t+1$ thus does not affect expected net benefits in t, and the discovery value fully captures the expected net benefits.[15]

For this monotonicity condition to hold, not all assumptions stated in the baseline search problem are required. Below, several such extensions are presented. Formal results and further details are presented in Appendix C.

Ranking in distribution: In various settings, the distribution of partial valuations depends on the position at which a product is discovered. For example, in a market environment where sellers of differentiated products compete in marketing efforts for consumers to become aware of their products early on, sellers offering better valuations may have a stronger incentive to be discovered first.[16] Furthermore, in online settings, consumers also have the option of directly sorting product lists (see e.g. Chen and Yao, 2017), which can lead to the pattern that higher valued products are shown earlier. In both cases, if only the mean of X_j depends on the position, and this mean is decreasing in the position, then monotonicity is satisfied and the optimal policy continues to be characterized by Theorem 1, the only difference being that the mean of X_j in the calculation of the discovery value is updated during search. This leads to a discovery value that decreases during search, making it optimal to recall products in some cases.

Unknown n_d: In some environments a consumer may not know how many alternatives

[15]For the search and purchase values, no monotonicity condition is required. This follows from the fact that in the independent comparison to the hypothetical outside option, both actions do not provide the option to continue searching. After buying a product, search ends, and after having inspected a product, the only option that remains is to either buy the product or choose the hypothetical outside option. Consequently, for inspection and purchase, at most one future period needs to be considered to fully capture the respective net benefits over immediately taking the outside option.

[16]See, for example, the discussion on non-price advertising and the related references cited in Armstrong (2017).

he might discover. For example, a consumer may believe that there are still alternatives he is not aware of and thus try to discover them, only to realize that he already is aware of all the available alternatives. In such cases, a belief over how many alternatives are going to be discovered needs to be specified. The reservation value policy continues to be optimal if these beliefs remain constant during search, or more generally as long as monotonicity is satisfied. The only difference to the baseline is that in $Q_d(c_d, c_s, z)$, expectations are additionally based on beliefs over how many alternatives will be revealed.

Multiple discovery technologies: Finally, there may be multiple technologies through which the consumer can discover more alternatives. In an online setting, for example, each technology may represent a different online shop offering alternatives. Moreover, advertising measures may separate products into different product pools. In such settings, the consumer also decides which technology to use to discover more alternatives. However, by assigning each of the discovery technologies a different discovery value, the optimal policy can be adjusted to accommodate this case.

3.2 Limitations

Though the optimal policy applies to a broad class of search problems, two important limitations exist. The first is that the monotonicity condition discussed above needs to hold for the discovery value to be based on a myopic comparison. The second limitation is that in the dynamic decision process, all available actions need to be independent of each other; performing one action in t should not affect the payoff of any other action that is available in t.

Independence is required to guarantee that the reservation values fully capture the effects of each action. Recall that each reservation value does not depend on the availability of other actions. If independence does not hold, however, the availability of other actions also influences the expected payoff of an action. Choosing actions based only on reservation values that disregard these effects therefore will not be optimal. Appendix D presents alternative search problems that violate this independence assumption.

4 Eventual Purchases, Consumer's Payoff and Demand

In an environment where consumers sequentially inspect products, a consumer's expected payoff and the market demand results from integrating over different possible choice sequences leading to eventual purchases. Conceptually, this poses a major challenge, as the number of possible choice sequences grows extremely fast in the number of available alternatives.[17]

[17] For example, with only one alternative and an outside option, there are four possible choice sequences. With two alternatives, the number of possible choice sequences increases to 20, and with three alternatives, there are

Theorem 2 allows to circumvent this difficulty. It states that the purchase outcome of a consumer solving the search problem is equivalent to a consumer directly buying a product that offers the highest *effective value*. Importantly, a product's effective value does not depend on the various possible choice sequence leading to its purchase.

Theorem 2. *Let*

$$w_j \equiv \begin{cases} x_j + \min\{\xi_j(c_s), y_j\} & \textit{if } x_j + \min\{\xi_j(c_s), y_j\} < z^d \\ z^d + f(h_j) & \textit{else} \end{cases}$$

be the effective value for product j revealed on position h_j, where $f(h_j)$ is a non-negative function and strictly decreasing in h_j. The solution to the search and discovery problem leads to the eventual purchase of the product with the largest effective value.

This result generalizes the „eventual purchase theorem" of Choi et al. (2018) to the case where the consumer has limited awareness. The generalization follows from the following implication of the optimal policy: Whenever both the inspection and the purchase value of a product in the awareness set exceed the discovery value, the consumer will buy the product and end search. This is captured in the effective values by the term $z^d + f(h_j)$, which ranks alternatives based on when during search they are discovered.

The result continues to hold for extensions of the search problem, as long as the discovery values are predetermined. The only difference then is that in the effective value of an alternative j, the discovery value depends on the position at which j is revealed.

4.1 Expected Payoff

Based on Theorem 1, it is now possible to get a simple characterization of a consumer's expected payoff, as shown in Proposition 1. Once a distribution of valuations is specified, it is only necessary to derive the resulting distribution of the effective values without having to explicitly consider different choice sequences.

Proposition 1. *A consumer's expected payoff in the search and discovery problem is equal to $\mathbb{E}_{\hat{\boldsymbol{W}}}[\max\{u_0, \max_{j\in J} \hat{W}_j\}]$, where $\mathbb{E}_{\hat{\boldsymbol{W}}}[\cdot]$ integrates over the distribution of $\hat{\boldsymbol{W}} = \left[\hat{W}_1, \ldots, \hat{W}_{|J|}\right]'$, with $\hat{w}_j = \min\{w_j, z^d\}$ and w_j being the effective value.*

Whereas the search problem already implies that making either inspection or discovery easier leads to an increase in the expected payoff, it is not obvious which of these two changes is more beneficial for a consumer. For the case where $n_d = 1$, Proposition 2 shows that if the number

already more than 100 possible choice sequences.

of alternatives exceeds some threshold, then the consumer benefits more from facilitating the discovery of additional products.[18]

Proposition 2. *If $n_d = 1$, there exists a threshold n^* such that whenever $|J| > n^*$, a consumer benefits more from a decrease in discovery costs than a decrease in inspection costs. This threshold decreases in the value of the outside option.*

The intuition is that when there are only few alternatives available, the consumer is more likely to first discover all alternatives and then start inspecting alternatives. Hence in expectation, he pays the inspection costs relatively often and a reduction in inspection costs can be more beneficial. Similarly, when the value of the outside option is large, the consumer is likely to inspect fewer of the products he discovers, leading to relatively small benefits of a reduction in inspection costs.

For settings where $n_d > 1$, it becomes difficult to obtain similarly general results. In particular, for some distributions and n_d, it is possible that decreasing inspection costs increases the discovery value z^d by more than decreasing the discovery costs by the same amount. In such cases, the consumer will benefit more from making inspection less costly. Nonetheless, the general intuition remains the same in such settings; a reduction in inspection costs is more beneficial, the more likely it is that the consumer inspects relatively many alternatives.

4.2 Market Demand

Using Theorem 1, it is also straightforward to derive market demand functions when heterogeneous consumers optimally solve the search and discovery problem. In particular, let the effective value w_{ij} for each consumer i be a realization of the random variable W_j and gather the random variables in $\boldsymbol{W} = \left[W_0, \dots, W_{|J|}\right]'$. For a unit mass of consumers the market demand for a product j then is given by

$$D_j = \mathbb{E}_h \left[\mathbb{P}_{\boldsymbol{W}} \left(W_j \geq W_k \forall k \in J \backslash j\right)\right] \tag{13}$$

where the expectations operator $\mathbb{E}_h [\cdot]$ integrates over all permutations of the order in which products are discovered by a consumer.

As the effective value decreases in the position at which a product is discovered, (13) reveals that the demand for a product depends on the probability of each position at which it is displayed. Specifically, the demand for a product exhibits ranking effects; products that are more likely to be discovered early are more likely to be bought. As discussed in detail in the next section, this

[18] Note that this threshold can be zero. For example, this is the case when $u_0 = 0$, $c_s = 0.1$ and $c_d = 0.1$, and the valuations are drawn from standard normal distributions.

follows from the structure of the search and discovery problem. As search progresses, it becomes less likely that a consumer has not yet settled for an alternative; hence, fewer consumers discover the products that would be revealed later, leading to a lower demand for such products.

5 Comparison of Search Problems

To highlight how the search and discovery problem differs from existing approaches, I compare it with the two classical sequential search problems; directed search as in Weitzman (1979) and random search as in Lippman and McCall (1976). Both these search problems are nested within the search and discovery problem. Directed search results if the consumer initially has full awareness (i.e. $S_0 = J$), or discovery costs are equal to zero. Random search results if in addition the information received when becoming aware of an alternative is non-informative (e.g. $x_j = 0 \forall j$). More generally, however, directed search differs in that the consumer initially knows all partial valuations and does not discover products, and random search abstracts from the option of not inspecting products with low partial valuations.

For clarity, I focus the comparison on the case where products are discovered one at a time ($n_d = 1$). Furthermore, valuations x_j and y_j are assumed to be realizations of mutually independent random variables X and Y. Assumptions specific to each search problem are described below.

Search and Discovery (SD): The consumer searches as described in Section 2, incurring inspection costs c_s and discovery costs c_d. Without loss of generality, I assume that the consumer discovers products in increasing order of their index, making subscripts for position h and product j interchangeable.

Random Search (RS): When discovering a product j, the consumer reveals both x_j and y_j; hence does not have to pay a cost to inspect the product. Costs to reveal this information are given by c^{rs}. In this case, the consumer optimally stops and buys product j if $x_j + y_j \geq z^{rs}$, where z^{rs} is implicitly defined by $c^{rs} = \mathbb{E}_{X,Y}\left[\max\{0, X + Y - z^{rs}\}\right]$. Products are discovered in the same order as in SD. Furthermore, I assume $u_0 < z^{rs}$ to ensure a non-trivial case.

Directed Search (DS): The consumer initially observes $x_j \forall j$, based on which he chooses to search among alternatives following Weitzman's (1979) reservation value policy. Costs to inspect product j are given by a function $c_j^{ds} = v_{ds}(c_s, h_j)$, where c_s are baseline costs that are adjusted for the position through a function $v_{ds} : \mathbb{R}_+^2 \to \mathbb{R}_+$ which is assumed to be strictly increasing in a product's position h_j. As costs vary across products, reservation values are given by $z_j^s = x_j + \xi_j$, where the assumption on $v_{ds}(c_s, h_j)$ implies that ξ_j decreases in j. I impose this functional form restriction as otherwise the DS problem does not generate similar patterns, which will become clear in the following.

5.1 Stopping decisions

In search settings, consumers' stopping decisions determine which products consumers consider and buy. Stopping decisions therefore shape how firms compete in prices, quality or for being discovered early during search.[19] Moreover, stopping decisions identify search cost parameters estimated in structural search models (e.g. Kim et al., 2017; Ursu, 2018). Hence, comparing stopping decisions across the different search problems provides important insights on how well existing approaches are able to capture the more general setting where consumers are not aware of all alternatives and use partial information to determine whether to inspect products.

In the SD problem, a consumer always stops search at a product k whenever the product is both promising enough to be inspected and offers a large enough valuation to not make it worthwhile to continue discovering more products. Formally, this is given by the condition $x_k + \min\{y_k, \xi\} \geq z^d$. The probability that a consumer will stop searching before discovering product j thus is given by

$$\mathbb{P}_{\boldsymbol{X},\boldsymbol{Y}}(X_k + \min\{Y_k, \xi\} \leq z^d \forall k < j) = 1 - \mathbb{P}(X + \min\{Y, \xi\} \leq z^d)^{j-1} \tag{14}$$

Similarly, in the RS problem, a consumer will always stop search at a product k whenever $x_k + y_k \geq z^{rs}$, hence the probability of stopping search before discovering product j is given by

$$\mathbb{P}_{\boldsymbol{X},\boldsymbol{Y}}(X_k + Y_k \leq z^{rs} \forall k < j) = 1 - \mathbb{P}(X + Y \leq z^{rs})^{j-1} \tag{15}$$

In both search problems, a consumer may stop search before discovering a product j. Consequently, stopping decisions in the SD and the RS problem imply the same feature: Products that a consumer initially has no information on may never be discovered and bought, independent of how the consumer values them.

However, as the consumer has the option of not inspecting products with low partial valuations, the stopping probabilities differ. In particular, in the case where the total cost to reveal all information about a product are the same, stopping probabilities are smaller in the SD problem. This is stated in Proposition 3 and follows from the fact that not having to inspect alternatives with small partial valuations allows to save on inspection costs. This increases the expected benefit of discovering more products, which implies a smaller probability of search stopping, and that on average, more products will be discovered in the SD problem.

[19] The latter is studied by the literature on sponsored search auctions (e.g. Athey and Ellison, 2011). In these models, sellers bid on positions at which their product or advert is shown to a consumer on a website.

Proposition 3. *If costs in the RS problem are given by $c^{rs} = c_s + c_d$, a consumer is less likely to stop before discovering any given product in the SD problem, than he is in the RS problem.*

In contrast, stopping decisions are different in the DS problem. As the consumer initially knows of the existence of all products and can order them based on partial information, there is no stopping decision in terms of discovering products. Instead, the consumer directly compares all partial valuations and the different inspection costs, based on which he decides the order in which to inspect products. Hence, he can directly inspect highly valued products even when they are presented at the last position.

This difference arises from the different assumptions on consumers' initial information and is paramount in the analysis of search frictions. Consider an equilibrium setting where horizontally differentiated alternatives are supplied by firms that compete by setting mean partial valuations (e.g. by setting prices as in Choi et al., 2018). If consumers are aware of all alternatives and search as in the DS problem, all firms will compete directly with each other. In contrast, in an SD problem, the firm that is discovered first initially competes only with the consumers' beliefs of discovering better products later on. This difference is further illustrated in Appendix E, and as it determines how firms compete, will lead to different equilibrium dynamics.[20]

When using structural models to estimate search costs and preference parameters, it is similarly important to consider this difference. For example, a structural search model will use price differences across all products to inform parameter estimates if it abstracts from limited awareness and assumes that consumers initially observe all prices. Consumers not inspecting low-price products at later positions then will be attributed either to a small (or even negative) price sensitivity, or large search cost. However, if instead this is the result of consumers not having discovered these products, then estimates and counterfactuals (e.g. the effects of price changes) will be biased. Whereas inspection costs that increase in position can alleviate this bias, it does not vanish. Instead, the estimation becomes sensitive to the choice of functional form of $v_{sd}(c_s, h_j)$, as this function determines how consumers are assumed to relate differences in partial valuations to differences in position-specific inspection costs.

5.2 Ranking Effects

The above analysis already suggests that the market demand structure differs across the three search problems. To provide further details, I focus on a particular pattern that is generated by all three search problems: Market demand for a product decreases in its position.

[20] To give an example, Anderson and Renault (1999) and Choi et al. (2018) model a similar environment, with the difference that in the former, consumers initially are not aware of any alternatives, whereas in the latter they are aware and observe prices of all alternatives. Whereas in the former, decreasing inspection costs increases lowers the equilibrium price in a symmetric equilibrium, the opposite holds in the latter environment.

Such ranking effects are important as they determine how fiercely sellers compete for their products to be revealed on early positions, for example through informative advertising or position auctions (e.g. Athey and Ellison, 2011). Furthermore, they have received considerable attention in the marketing literature, which has produced ample empirical evidence that suggests their importance in online sales (e.g. Ghose et al., 2014; De los Santos and Koulayev, 2017; Ursu, 2018).

To compare the mechanism producing ranking effects across the search problems, I use the following definition: The ranking effect for a product is the difference in market demand of the product being revealed at position h and at $h+1$, with the corresponding exchange of the product previously revealed at position $h+1$. Formally, the ranking effect is defined as

$$r_k(h) \equiv d_k(h) - d_k(h+1) \tag{16}$$

where $d_k(h)$ denotes the market demand for a product when revealed at position h in search problem $k \in \{SD, RS, DS\}$. For clarity, product specific subscripts are either omitted or exchanged with position subscripts in the following.

To investigate ranking effects, it is first necessary to derive the market demand at a particular position h. For a unit mass of consumers in the SD problem, it is given by

$$\begin{aligned} d_{SD}(h) = \mathbb{P}_W(W < z^d)^{h-1} \Big[& \mathbb{P}_W(W \geq z^d) \\ & + \mathbb{P}_W(W < z^d)^{|J|-(h-1)} \mathbb{P}_{\boldsymbol{W}}(W \geq \max_{k \in J} W_k | W_k < z^d \forall j) \Big] \end{aligned} \tag{17}$$

The expression follows from Theorem 2 which implies that if a consumer discovers a product with $w_j \geq z^d$, he will stop searching and buy a product j. The consumer will only discover and have the option to buy a product on position h if $w_j < z^d$ for all products on better positions. In contrast, when $w_j < z^d$, the consumer will first discover more products, and only recall j if he discovers all products and j is the best among them.

In the latter case, a product's position does not affect market demand; once all products are discovered, products are equivalent in terms of their inspection costs and the order in which they are inspected is only determined based on partial valuations. This implies that the ranking effect in the SD problem is independent of the number of alternatives and simplifies to

$$r_{SD}(h) = \mathbb{P}_W\left(W \geq z^d\right)\left[\mathbb{P}_W(W < z^d)^{h-1} - \mathbb{P}_W(W < z^d)^h\right] \tag{18}$$

This expression reveals that the ranking effect in the SD problem solely results from the difference in the probability of a consumer reaching positions h or $h+1$ respectively. Besides the distribution of valuations and the inspection and discovery costs, Proposition 4 shows that the ranking effect is determined by the position h to which the product is moved. When h is large, fewer consumers will not have already stopped searching before reaching h. Hence, the later a product is revealed, the smaller is the increase in demand when moving one position ahead.

The demand in a random search problem is derived similarly. In RS, a consumer will only be able to buy a product if he has not stopped searching before, which requires that $x + y < z^{rs}$ for all products on better positions. Furthermore, a consumer will also only recall a product if he has inspected all alternatives. Similar to the SD problem, this implies that the ranking effect in the RS problem is given by

$$r_{RS}(h) = \mathbb{P}\left(X + Y \geq z^{rs}\right) \left[\mathbb{P}_{X,Y}\left(X + Y < z^{rs}\right)^{h-1} - \mathbb{P}_{X,Y}\left(X + Y < z^{rs}\right)^{h}\right] \tag{19}$$

Comparing (18) with (19) reveals that ranking effects in both these search problems result from the similar stopping decisions. In both search problems; fewer consumers buy products at later positions due to the increasing the probability of having stopped searching before discovering these products. It follows that in both search problems, ranking effects decrease in the position and are independent of the total number of alternatives.

Though their extent generally differs, Proposition 4 additionally shows that at later positions, ranking effects will be larger in the SD problem. The result is a direct implication of Proposition 3. As a consumer is more likely to reach a product at a later position in the SD problem, ranking effects at later positions will be larger.

Proposition 4. *The ranking effect in both the SD and the RS problem decreases in position h and is independent of the number of alternatives. Furthermore, if $c^{rs} = c_s + c_d$, there exists a threshold h^* such that $r_{SD}(h) \geq r_{RS}(h)$ for all $h > h^*$.*

Given the different stopping decisions, ranking effects in directed search do not result from consumers having stopped searching before reaching products revealed at later positions. Instead, they result from differences in the cost of inspecting products at different positions. To see this, write the ranking effect in the DS problem as

$$r_{DS}(h) = \mathbb{E}_{\tilde{W}_h}\left[\prod_{k \neq h} \mathbb{P}(\tilde{W}_k \leq \tilde{W}_h)\right] - \mathbb{E}_{\tilde{W}_{h+1}}\left[\prod_{k \neq h+1} \mathbb{P}(\tilde{W}_k \leq \tilde{W}_{h+1})\right] \tag{20}$$

This expression reveals that the ranking effect results from two sources in the DS problem. First, by moving a product j one position ahead, the product previously on position h is now more costly to inspect, making it more likely that j is bought for any $\tilde{w}_j$. Second, by making it less costly to inspect j, the distribution of $\tilde{w}_j$ shifts such that larger values $\tilde{w}_j$ become more likely.

In contrast to RS and SD, the ranking effect in the DS problem depends on the number of available alternatives. In RS and SD, ranking effects result from the decreasing probability of a consumer having stopped searching before reaching a particular position, which does not depend on how many alternatives there are in total. In DS, however, a consumer directly compares all alternatives based on partial valuations. Adding more alternatives thus can affect the demand on each position.

Specifically, Proposition (5) shows that ranking effects in the DS problem will be smaller if there are many alternatives. The reason is that as the number of alternatives increases, each product is less likely to be bought and differences in the position-specific market demand decrease. Note, however, that in cases where the probability of consumers buying products on the last positions is very small or exactly zero (e.g. when inspection costs are large), adding more alternatives will not affect ranking effects in the DS problem.

Proposition 5. *The ranking effect in the DS problem is weakly decreasing in the number of alternatives.*

A second difference to the RS and SD problem is that the ranking effect does not necessarily decrease in position. This is possible as there are two counteracting channels through which position affects the ranking effect in. First, as there is lower demand for products at later positions, differences between them will be smaller. Second, if $v_{ds}(c_s, h)$ is such that ξ_h decreases in h at an increasing rate, the difference in the purchase probability at h instead of at $h+1$ increases in the position. When the latter dominates, the ranking effect will first increase in position.

The above comparison highlights that the mechanism producing ranking effects in the DS problem is distinct from the one in the SD and RS problem, leading to a different demand structure. In the former, ranking effects result from differences in inspection costs relative to differences in partial valuations. Hence, a better partial valuation is a substitute for moving positions ahead. In contrast, in an SD or RS problem, a product's large partial valuation does not affect consumers that stop search before discovering it. Hence, offering a larger partial valuation does not substitute for being discovered early in an SD or RS problem.[21]

Moreover, the size of the ranking effects also determines how important it is for products to be

[21] Note, however, that in an equilibrium setting, offering larger partial valuations may indirectly serve as a substitute for being discovered early by raising consumers' expectations and induce them to search longer.

revealed on an early position. As ranking effects are independent of the number of alternatives in SD and RS, so are sellers' incentives to have their products revealed early during search. In contrast, in DS, the demand increase of moving positions ahead becomes smaller when the number of alternatives increases. Hence, sellers can have smaller incentives to be revealed on early positions when there are many, relative to when there are only few alternatives.

Finally, the above comparison between the number of alternatives and ranking effects also suggests the existence of an empirical test to distinguish the search modes in settings where some consumers purchase the last product. If data is available that allows to test whether ranking effects depend on the number of alternatives, then it will be possible to empirically determine whether a DS problem, instead of a RS or SD problem provides a framework that better captures ranking effects in a particular setting. Furthermore, if data is available that allows to test whether a product's partial valuation has an effect on whether it is inspected, it will be possible to distinguish between RS and SD.

5.3 Expected Payoff

If costs are specified such that the total costs of revealing all product information remain the same, then the three search problems differ only in the amount of information the consumer has during search. A comparison of a consumer's expected payoff based on such a specification therefore provides some insight into whether it is always to the consumer's benefit to provide information that helps to direct search towards some alternatives.

For total costs of revealing full information about a product on position h to be the same in the three search problems, inspection costs in the RS and DS problem are specified as $c_{rs} = c_s + c_d$ and $c_j^{ds} = c_s + h_j c_d$ respectively.

The SD problem extends the RS problem by additionally providing the consumer with the option to not inspect products depending on their partial valuations. This allows the consumer to save on inspection costs by not inspecting products with small partial valuations. As stated in Proposition 6, this increases the expected payoff which implies that providing product information across two layers, as done for example by online retailers or search intermediaries, is beneficial for consumers.

Proposition 6. *If $c^{rs} = c_s + c_d$, then a consumer's expected payoff in the SD problem is larger than in the RS problem.*

In contrast to the SD problem, the consumer can use all partial valuations to direct search in the DS problem. Hence, if inspection costs for all products are the same in both problems (i.e. $c_j^{ds} = c_s \forall j$), it is clear that a consumer will have a larger expected payoff in the DS problem.

However, under the assumption that total costs of revealing full information are the same in both search problems, a more detailed analysis is necessary to determine which search problem offers a larger expected payoff.

Denote a consumer's expected payoff in a search problem k as π_k for $k \in \{SD, DS\}$. Proposition 1 implies that

$$\pi_{SD} = \mathbb{E}_{\hat{\boldsymbol{W}}} \left[\max\{u_0, \max_{j \in J} \hat{W}_j\} \right]$$
$$\pi_{DS} = \mathbb{E}_{\tilde{\boldsymbol{W}}} \left[\max\{u_0, \max_{j \in J} \tilde{W}_j\} \right]$$

Furthermore, let $H_k(\cdot)$ denote the cumulative density of the respective maximum value over which the expectation operator integrates in problem k. The difference in expected payoffs of the SD and the DS problem then is given by

$$\pi_{SD} - \pi_{DS} = \int_{z^d}^{\infty} H_{DS}(w) - 1 \mathrm{d}w + \int_{u_0}^{z^d} H_{DS}(w) - H_{SD}(w) \mathrm{d}w \tag{21}$$

The first expression in (21) is negative, capturing the advantage of observing partial valuations for all products and being able to directly inspect a product at a later position. Given $H_{DS}(w) \leq H_{SD}(w)$ on $w \in [u_0, z^d]$, the second expression in (21) is positive, revealing that directly observing all partial valuations x_j does not only yield benefits.

The latter stems from the difference in how inspection and discovery costs are taken into consideration in the two dynamic decision processes. In DS, the total cost of inspecting a product j at a later position is directly weighed against its benefits given the partial valuations. In contrast, in SD, the consumer first weighs the discovery costs against the expected benefits of discovering a product with a larger partial valuation. Once product j is revealed, the accumulated cost paid to discover j (jc_d) is a sunk cost and does not affect the decision whether to inspect j.

Hence, in cases where products on early positions have below-average partial valuations x_j, the optimal policy in SD tends to less often prescribe to inspect these products compared to the direct cost comparison in DS. In some cases, the former can be more beneficial, leading to a larger expected payoff.[22] Directly revealing all partial valuations therefore does not always improve a consumer's benefit, if the consumer continues to incur the same total costs to reveal the full valuation of any given product.[23]

[22] For example, this is the case if $X \sim N\left(0, \frac{1}{3}\right)$, $Y \sim N\left(0, \frac{2}{3}\right)$, $c_s = c_d = 0.05$ and $|J| = 10$.

[23] No threshold result as in Proposition 2 applies in this case. The first expression in (21) decreases whereas the second expression increases in the number of alternatives.

6 Conclusion

This paper introduces a search problem that generalizes existing frameworks to settings where consumers have limited awareness and first need to become aware of alternatives before being able to search among them. The paper's contribution is to provide a tractable solution for optimal search decisions and expected outcomes for this search and discovery problem. Moreover, a comparison with classical random and directed search highlights how limited awareness and the availability of partial product information determine search outcomes

A promising avenue for future research is to build on this paper's results and study limited awareness in an equilibrium setting. This could yield novel insights into how consumers' limited information shapes price competition. Furthermore, the search and discovery problem can serve as a framework to analyze how firms compete for consumers' awareness. For example, informative advertising can make it more likely that consumers are aware of a seller's products from the outset. Ranking effects derived in this paper already suggest that it will be in a seller's best interest to make consumers aware of his product, but further research is needed to determine equilibrium dynamics.

Another avenue for future research entails incorporating the search and discovery problem into a structural model that is estimated with click-stream data. The available actions in the search and discovery problem closely match how consumers scroll through product lists (discovery) and click on products (inspection) on websites of search intermediaries and online retailers. By accounting for the fact that consumers initially do not observe entire list pages, such a model could improve the estimation of consumers' preferences, inspection costs and ranking effects relative to models that abstract from consumers not observing the whole product list.

References

Adam, K. (2001). Learning while searching for the best alternative. *Journal of Economic Theory 101*(1), 252–280.

Anderson, S. P. and R. Renault (1999). Pricing, product diversity, and search costs: A bertrand-chamberlin-diamond model. *The RAND Journal of Economics 30*(4), 719–735.

Armstrong, M. (2017). Ordered consumer search. *Journal of the European Economic Association 15*(5), 989–1024.

Armstrong, M. and J. Vickers (2015). Which demand systems can be generated by discrete choice? *Journal of Economic Theory 158*, 293–307.

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

Bagwell, K. (2007). The economic analysis of advertising. In *Handbook of Industrial Organization*, Volume 3, Chapter 28, pp. 1701–1844.

Banks, J. S. and R. K. Sundaram (1992). Denumerable-armed bandits. *Econometrica 60*(5), 1071–1096.

Banks, J. S. and R. K. Sundaram (1994). Switching costs and the Gittins index. *Econometrica 62*(3), 687–694.

Bronnenberg, B. J., J. B. Kim, and C. F. Mela (2016). Zooming in on choice: How do consumers search for cameras online? *Marketing Science 35*(5), 693–712.

Burdett, K. and K. L. Judd (1983). Equilibrium price dispersion. *Econometrica 51*(4), 955–969.

Chade, H. and L. Smith (2006). Simultaneous search. *Econometrica 74*(5), 1293–1307.

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

Choi, H. and C. F. Mela (2019). Monetizing online marketplaces. *Marketing Science 38*(6), 913–1084.

Choi, M., A. Y. Dai, and K. Kim (2018). Consumer search and price competition. *Econometrica 86*(4), 1257–1281.

De Los Santos, B., A. Hortacsu, and M. R. Wildenbeest (2012). Testing models of consumer search using data on web browsing and purchasing behavior. *The American Economic Review 102*(6), 2955–2980.

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

Diamond, P. A. (1971). A model of price adjustment. *Journal of Economic Theory 3*(2), 156–168.

Doval, L. (2018). Whether or not to open Pandora's box. *Journal of Economic Theory 175*, 127–158.

Fershtman, D. and A. Pavan (2019). Searching for arms. Unpublished manuscript. Available at http://faculty.wcas.northwestern.edu/ apa522/SFA.pdf.

Ghose, A., P. G. Ipeirotis, and B. Li (2014). Examining the impact of ranking on consumer behavior and search engine revenue. *Management Science 60*(7), 1632–1654.

Gittins, J., K. Glazebrook, and R. Weber (2011). *Multi-armed Bandit Allocation Indices* (2nd ed.). John Wiley & Sons, Inc.

Glazebrook, K. D. (1979). Stoppable families of alternative bandit processes. *Journal of Applied Probability 16*(4), 843–854.

Hong, H. and M. Shum (2006). Using price distributions to estimate search costs. *The RAND Journal of Economics 37*(2), 257–275.

Honka, E., A. Hortaçsu, and M. A. Vitorino (2017). Advertising, consumer awareness, and choice: evidence from the u.s. banking industry. *RAND Journal of Economics 48*(3), 611–646.

Hortaçsu, A. and C. Syverson (2004). Product differentiation, search costs, and competition in the mutual fund industry: A case study of S & P 500 index funds. *The Quarterly Journal of Economics 119*(2), 403–456.

Jolivet, G. and H. Turon (2019). Consumer search costs and preferences on the internet. *Review of Economic Studies 86*(3), 1258–1300.

Keller, G. and A. Oldale (2003). Branching bandits: A sequential search process with correlated pay-offs. *Journal of Economic Theory 113*(2), 302–315.

Kim, J. B., P. Albuquerque, and B. J. Bronnenberg (2017). The probit choice model under sequential search with an application to online retailing. *Management Science 63*(11), 3911–3929.

Kleinberg, R., B. Waggoner, and E. G. Weyl (2017). Descending price optimally coordinates search.

Koulayev, S. (2014). Search for differentiated products: identification and estimation. *The RAND Journal of Economics 45*(3), 553–575.

Kuksov, D. (2006). Search, common knowledge, and competition. *Journal of Economic Theory 130*(1), 95–108.

Lippman, S. A. and J. J. McCall (1976). The economics of job search: A survey. *Economic Inquiry 14*(2), 155–189.

McCall, J. J. (1970). Economics of information and job search. *The Quarterly Journal of Economics 84*(1), 113–126.

Moraga-González, J. L., Z. Sándor, and M. R. Wildenbeest (2017a). Nonsequential search equilibrium with search cost heterogeneity. *International Journal of Industrial Organization 50*, 392–414.

Moraga-González, J. L., Z. Sándor, and M. R. Wildenbeest (2017b). Prices and heterogeneous search costs. *The RAND Journal of Economics 48*(1), 125–146.

Morozov, I. (2019). Measuring benefits from new products in markets with information frictions. Unpublished manuscript. Available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3446398.

Olszewski, W. and R. Weber (2015). A more general Pandora rule? *Journal of Economic Theory 160*, 429–437.

Stigler, G. J. (1961). The economics of information. *The Journal of Political Economy 69*(3), 213–225.

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

Weitzman, M. L. (1979). Optimal search for the best alternative. *Econometrica 47*(3), 641–654.

Zhang, X., T. Y. Chan, and Y. Xie (2018). Price search and periodic price discounts. *Management Science 64*(2), 495–510.

Appendix

A Proofs

A.1 Uniqueness of Discovery Value

Proposition 7. (10) *has a unique solution.*

Proof. Differentiating $Q_d(c_d, c_s, z)$ with respect to z yields (see Appendix B)

$$\frac{\partial Q(c_d, c_s, z)}{\partial z} = \begin{cases} +H(z) - 1 & \text{if } z < 0 \\ -2 + H(z) & \text{else} \end{cases} \tag{22}$$

where $H(\cdot)$ denotes the cumulative density of the random variable $\max_{k \in \bar{J}} \tilde{W}_k$. This implies $\frac{\partial Q_d(c_d, c_s, z)}{\partial z} \leq 0$, which combined with continuity, $Q_d(c_d, c_s, \infty) = -c_d$ and $Q_d(c_d, c_s, -\infty) = \infty$ imply that a solution to (10) exists. Finally, uniqueness requires $Q_d(c_d, c_s, z)$ to be strictly decreasing at $z = z^d$. $\frac{\partial Q_d(c_d, c_s, z^d)}{\partial z} = 0$ would require that $H(z^d) = 1$, which contradicts the definition of the discovery value value z^d in (10), as it implies $Q_d(c_d, c_s, z^d) \leq -c_d < 0$. □

A.2 Theorem 1

Proof. Let $\Theta(\Omega_t, A_t, z)$ denote the value function of an alternative decision problem, where in addition to the available actions in A_t, there exists a hypothetical outside option offering value z. Suppose the consumer knows $|J|$. For this case, Theorem 1 of Keller and Oldale (2003) states that a Gittins index policy is optimal, and that the following holds:

$$\Theta(\Omega_t, A_t, z) = b - \int_z^b \Pi_{a \in A_t} \frac{\partial \Theta\left(\Omega_t, \{a\}, w\right)}{\partial w} \mathrm{d}w \tag{23}$$

where b is some finite upper bound of the immediate rewards.[24] The Gittins index of action d (discovering products) is defined by $g_t^d = \mathbb{E}_{\boldsymbol{X}}\left[\Theta(\Omega_{t+1}, A_{t+1} \backslash A_t, g_t^d)\right]$. Now consider a period t in which more discoveries will be available in $t+1$ (known due to knowledge of $|J|$). In this case

[24] Note that immediate rewards $R(a) \geq -\max\{c_s, c_d\}$, and that finite mean and variance of the distributions of X and Y imply that for all realizations x, y, there exists some b such that $R(a) = x_j + y_j \leq b$.

we have

$$g_t^d = \mathbb{E}_{\boldsymbol{X}}\left[\Theta(\Omega_{t+1}, \{d, s1, \ldots, sn_d\}, g_t^d)\right] - c_d \tag{24}$$
$$= \mathbb{E}_{\boldsymbol{X}}\left[b - \int_{g_t^d}^{b} \frac{\partial \Theta\left(\Omega_{t+1}, \{d\}, w\right)}{\partial w} \prod_{k=1}^{n_d} \frac{\partial \Theta\left(\Omega_{t+1}, \{s_k\}, w\right)}{\partial w} \mathrm{d}w\right] - c_d$$

where $s_k \in S_{t+1} \backslash S_t \forall k$. $\Theta(\Omega_t, \{s_k\}, z)$ is the value of a search problem with an outside option offering z and the option of inspecting product k (with known partial valuation x_k). $\Theta\left(\Omega_{t+1}, \{e\}, w\right)$ is the value of a search problem with an outside option offering z, and the option to discover more products. Finally, $\mathbb{E}_{\boldsymbol{X}}[\cdot]$ is the expectation operator integrating over the joint distribution of the n_d random variables in $\boldsymbol{X} = [X, \ldots, X]$.

Optimality of the Gittins index policy then implies that when $z \geq g_{t+1}^d$, the consumer will choose the outside option in $t+1$. Hence $\Theta\left(\Omega_t, \{e\}, w\right) = w \forall w \geq g_{t+1}^d$ which yields $\frac{\partial \Theta(\Omega_t, \{e\}, w)}{\partial w} = 1 \forall w \geq g_{t+1}^d$. This implies that for $g_t^d \geq g_{t+1}^d$, g_t^d does not depend on whether more products can be discovered in the future, and the optimal policy is independent of the beliefs over the number of available alternatives. As a result, as long as the Gittins index is weakly decreasing during search, i.e. $g_t^d \geq g_{t+1}^d \forall t$, a Gittins index policy remains optimal in the case where the consumer does not know $|J|$ and has beliefs over the number of available alternatives.

It remains to show that $g_t^d \geq g_{t+1}^d \forall t$ holds in the proposed search problem. When $|J| = \infty$, it is clear that $g_t^d = g_{t+1}^d$ as in both periods infinitely many products remain to be discovered, and q (the consumer's belief on whether more products can be discovered in the next period) is constant. For $|J| < \infty$, backwards induction yields that this condition holds: Suppose that in period $t+1$, no further products can be discovered. In this case, the Gittins index is given by

$$g_{t+1}^d = \mathbb{E}_{\boldsymbol{X}}\left[b - \int_{g_{t+1}^d}^{b} \prod_{k=1}^{n_d} \frac{\partial \Theta\left(\Omega_{t+1}, \{s_k\}, w\right)}{\partial w} \mathrm{d}w\right] - c_d \tag{25}$$

As $0 \leq \frac{\partial \Theta(\Omega_t, \{e\}, w)}{\partial w} \leq 1$ and $\frac{\partial \Theta(\Omega_t, \{s_k\}, w)}{\partial w} \geq 0$, it holds that

$$\mathbb{E}_{\boldsymbol{X}}\left[b - \int_{g_t^d}^{b} \prod_{k=1}^{n_d} \frac{\partial \Theta\left(\Omega_{t+1}, \{s_k\}, w\right)}{\partial w} \mathrm{d}w\right] \leq \mathbb{E}_{\boldsymbol{X}}\left[b - \int_{g_t^d}^{b} \frac{\partial \Theta\left(\Omega_t, \{d\}, w\right)}{\partial w} \prod_{k=1}^{n_d} \frac{\partial \Theta\left(\Omega_t, \{s_k\}, w\right)}{\partial w} \mathrm{d}w\right] \tag{26}$$

which implies $g_t \geq g_{t+1}$.

Finally, $\Theta(\Omega_{t+1}, \{d, s1, \ldots, sn_d\}, g_t^d) = V\left(\left\langle \bar{\Omega}, \omega(x, z)\right\rangle, \{b0, s, \ldots, sn_d\}; \tilde{\pi}\right)$ of the discovery value implies $z^d = g_t^d$. Similarly, the definition of the inspection and purchase values (in 6 and 10) are equivalent to the definition of Gittins index values for these actions and it follows that the reservation value policy is the Gittins index policy. □

A.3 Theorem 2

Proof. As a product (incl. the outside option) always is bought, it suffices to show that product j is not bought whenever there exists another product k with $w_k > w_j$. First, consider the case where k is revealed before j ($h_k < h_j$). In this case, $w_k > w_j$ if and only if either (i) $\tilde{w}_k \equiv \min\left\{z_k^s, z_k^b\right\} \geq z^d$ or (ii) $z^d > \tilde{w}_k > \tilde{w}_j$. In the former, the consumer will not discover products beyond k and never buy product j. This follows from $z_k^s \geq z^d$ guaranteeing that k is inspected, and $z_k^b \geq z^d$ implying that search ends with buying k. In the latter, $w_j = \tilde{w}_j < w_k = \tilde{w}_k$, and the consumer continues discovering such that both products are in the awareness set. Hence the eventual purchase theorem of Choi et al. (2018) applies and product j will not be bought. Finally, consider the case where k is discovered after or at the same time as j ($h_k \geq h_j$). In this case, $w_k > w_j$ if and only if $z^d > \tilde{w}_k > \tilde{w}_j$, which is the same as (ii) above. □

A.4 Proposition 1

Proof. Let $\tilde{w}_j \equiv x_j + \min\{y_j, \xi_j\}$, $\bar{w} \equiv \max_{j\in J} \hat{w}_j$, $\bar{w}_r \equiv \max_{k\in J_{1:r-1}} \hat{w}_k$ and $\tilde{\bar{w}}_{r,j} \equiv \max_{k\in J_r\setminus j} \tilde{w}_k$ where $J_{a:b}$ denotes the set of products discovered on position $r \in \{a, \ldots, b\}$, and J_r is shorthand for $J_{r:r}$. Furthermore, let $1(\cdot)$ denote the indicator function and $\sum_r$ sum up over possible product positions r.

For $z^d \geq u_0$, the payoff of a consumer given values x_j and y_j then is given by

$$\sum_r 1(\bar{w}_r < z^d) \left[\sum_{j\in J_r} 1(\tilde{w}_j \geq \max\left\{z^d, \tilde{\bar{w}}_{r,j}\right\})(x_j + y_j) - 1(x_j + \xi_j \geq \max\left\{z^d, \tilde{\bar{w}}_{r,j}\right\})c_s \right] - \sum_r 1(\bar{w}_r < z^d)c_d + 1(\bar{w} < z^d)\pi \quad (27)$$

which follows from Theorem 2: (i) If $\bar{w}_r < z^d$, the consumer continues beyond position $r-1$, hence pays discovery costs. (ii) If $\tilde{w}_j \geq \max\left\{z^d, \tilde{\bar{w}}_{r,j}\right\}$, then the consumer buys j and does not continue beyond position r. (iii) If $x_j + \xi_j \geq \max\left\{z^d, \tilde{\bar{w}}_{r,j}\right\}$, then the consumer inspects j and incurs costs c_s. (iv) If $\bar{w} < z^d$, the consumer reveals all products, and has payoff π, which denotes the payoff of a directed search problem over products $\left\{j | x_j + \xi_j < z^d\right\}$ with outside option $\bar{u}_0 = \max\{u_0, \max_{k\in\left\{j|x_j+\xi_j\geq z^d, x_j+y_j\leq\xi_j\right\}} x_k + y_k\}$.

Let $\mathbb{E}[\cdot]$ integrate over the distribution of $X_j, Y_j \forall j \in J$. Using the definition of the reservation values to rewrite the costs as $c_s = \mathbb{E}[1(Y_j \geq \xi_j)(Y_j - \xi_j)] \,\forall j$ and $c_d = \mathbb{E}\left[1(\tilde{\bar{W}} \geq z^d)(\tilde{\bar{W}} - z^d)\right]$

(see Appendix B), this yields an expected payoff of:

$$
\begin{aligned}
&\sum_r \mathbb{E}\left[1(\bar{W}_r < z^d)\left(\sum_{j\in J_r} 1(\tilde{W}_j \geq \max\{z^d, \tilde{\tilde{W}}_{r,j}\})(X_j + Y_j)\right.\right.\\
&\quad \left.\left. - 1(X_j + \xi_j \geq \max\{z^d, \tilde{\tilde{W}}_{r,j}\})1(Y_j \geq \xi_j)(Y_j - \xi_j)\right)\right]\\
&\quad - \sum_r \mathbb{E}\left[1(\bar{W}_r < z^d)1(\tilde{\tilde{W}}_r \geq z^d)(\tilde{\tilde{W}}_r - z^d)\right] + \mathbb{E}\left[1(\bar{w} < z^d)\pi\right]\\
&= \sum_r \mathbb{E}\left[1(\bar{W}_r < z^d)\left(\sum_{j\in J_r} 1(\tilde{W}_j \geq \max\{z^d, \tilde{\tilde{W}}_{r,j}\})(X_j + \xi_j)\right)\right]\\
&\quad - \sum_r \mathbb{E}\left[1(\tilde{\tilde{W}}_r \geq z^d)(\tilde{\tilde{W}}_r - z^d)\right] + \mathbb{E}\left[1(\bar{W} < z^d)\pi\right]\\
&= \sum_r \mathbb{E}\left[1(\bar{W}_r < z^d)1(\tilde{\tilde{W}}_r \geq z^d)\tilde{\tilde{W}}_r\right]\\
&\quad - \sum_r \mathbb{E}\left[1(\tilde{\tilde{W}}_r \geq z^d)(\tilde{\tilde{W}}_r - z^d)\right] + \mathbb{E}\left[1(\bar{W} < z^d)\pi\right]\\
&= \sum_r \mathbb{E}\left[1(\bar{W}_r < z^d)1(\tilde{\tilde{W}}_r \geq z^d)z^d\right]\\
&\quad + \mathbb{E}\left[1(\bar{W} < z^d)\max\{u_0, \max_{j\in J} \tilde{W}_j\}\right]\\
&= \mathbb{E}\left[\max\{u_0, \max_{j\in J} \hat{W}_j\}\right]
\end{aligned}
$$

where the second-to-last step follows from Corollary 1 in Choi et al. (2018). Finally, if $z^d < u_0$, the payoff is equal to u_0. As in this case $\hat{W}_j \leq z^d$ by definition, the above expression continues to apply. □

A.5 Proposition 2

Proof. Consider a situation where $u_0 < z^d$, and we decrease initial costs c_s and c_d to either $c_s' = c_s - \Delta$ or $c_d' = c_d - \Delta$, while keeping the other cost constant. Let $H_1(\cdot)$ and $H_2(\cdot)$ denote the cumulative density of $\bar{W} \equiv \max\left\{u_0, \max_{j\in J} \hat{W}_j\right\}$ of the former and the latter case respectively. Similarly, let z_1^e and z_2^e denote the associated discovery values. Given $n_d = 1$, we have $\frac{\partial Q_d(c_d,c_s,z)}{\partial c_d} < \frac{\partial Q_d(c_d,c_s,z)}{\partial c_s}$; hence $\left|\frac{\partial z^d}{\partial c_d}\right| > \left|\frac{\partial z^d}{\partial c_s}\right|$ and $z_2^e > z_1^e$. The difference in a consumer's expected payoff following the two changes thus can be written as

$$
\int_{z_1^e}^{z_2^e} 1 - H_2(w)\mathrm{d}w - \int_{u_0}^{z_1^e} H_2(w) - H_1(w)\mathrm{d}w \tag{28}
$$

Whereas the first part is strictly positive, the second part is negative. This follows as $\bar{W} = \max_{j\in J} X_j + \min\{Y_j, \xi\}$ and $\frac{\partial \xi}{\partial c_s} < 0$ imply $H_1(w) \leq H_2(w)$ for $w \in [u_0, z_1^e]$. Since valuations are independent across products, we have $H_k(w) = \mathbb{P}_{X,Y}\left(X + \min\{Y, \xi_k\} \leq w\right)^{|J|}$; hence, as $|J|$ increases, $H_2(w) - H_1(w)$ and $H_2(w)$ decrease for $w \in [u_0, z_2^e]$.[25] Consequently, for all $\Delta > 0$

[25] Note that if $\mathbb{P}_{X,Y}\left(X + \min\{Y, \xi_k\} \leq w\right)$ is large, then $H_1(w) - H_2(w)$ will first increase in $|J|$, before starting to decrease.

there exists some threshold n^* for $|J|$ such that

$$\int_{z_1^e}^{z_2^e} 1 - H_2(w)\mathrm{d}w > \int_{u_0}^{z_1^e} H_2(w) - H_1(w)\mathrm{d}w \tag{29}$$

For the case where $u_0 \geq z^d$, having $z_2^e > z_1^e$ immediately yields the result. Finally, note that if $z_2^e < u_0$, then neither change affects the expected payoff which remains at u_0. □

A.6 Proposition 3

Proof. At $c_s = 0$, we have $z^d = z^{rs}$. $\left|\frac{\partial z^{rs}}{\partial c_s}\right| \geq \left|\frac{\partial z^d}{\partial c_s}\right|$ then implies $z^d \geq z^{rs}$. Using this in (14) and (15) immediately yields the result. □

A.7 Proposition 4

Proof. The first two statements immediately follow from (18) and (19). To see the latter, rewrite (18) as $\mathbb{P}_W\left(W < z^d\right)^{h-1}\mathbb{P}_W\left(W \geq z^d\right)^2$, and (19) in a similar way. $c^{rs} = c_s + c_d$ then implies $z^d \geq z^{rs}$. Hence, $\mathbb{P}_W\left(W < z^d\right) = \mathbb{P}_{X,Y}(X + \min\{Y, \xi\} < z^d) \geq \mathbb{P}_{X,Y}\left(X + Y < z^{rs}\right)$ which directly implies the existence of the threshold. □

A.8 Proposition 5

Proof. Write the first expression in (20) (demand at position h) as $\mathbb{E}_{\tilde{W}_h}\left[\mathbb{P}\left(\tilde{W}_{h+1} \leq \tilde{W}_h\right)\prod_{k\notin\{h,h+1\}}\mathbb{P}\left(\tilde{W}_k \leq \tilde{W}\right.\right.$ When $|J|$ decreases, this expression decreases through the product term, which is weighted by the first term $\mathbb{P}\left(\tilde{W}_{h+1} \leq \tilde{W}_h\right)$. As $\mathbb{P}\left(\tilde{W}_{h+1} \leq t\right) \geq \mathbb{P}\left(\tilde{W}_h \leq t\right) \forall t$, the first expression in (20) decreases by more than the second one when the number of alternatives increases. □

A.9 Proposition 6

Proof. The RS problem is equivalent to a policy in the SD problem that commits on inspecting every product that is discovered, conditional on which the consumer chooses to stop optimally. However, as the optimal policy in the SD problem is not this policy, it must yield a (weakly) larger payoff. □

B Further details on Search and Discovery Values

The search value of a product j is defined by equation (6) and sets the myopic net gain of the inspection over immediately taking a hypothetical outside option offering utility z to zero. This myopic net gain can be calculated as follows:

$$\begin{aligned}
Q_s(x_j, c_s, z) &= \mathbb{E}_Y\left[\max\{0, x_j + Y - z\}\right] - c_s \\
&= \int_{z-x_j}^{\infty} (x_j + y - z)\mathrm{d}F(y) - c_s \\
&= \left[(x_j + y - z)F(y)\right]_{y=z-x_j}^{y_j=\infty} - \int_{z-x_j}^{\infty} F(y)\mathrm{d}y - c_s \\
&= \int_{z-x_j}^{\infty} \left[1 - F(y)\right]\mathrm{d}y - c_s
\end{aligned}$$

Substituting $\xi_j = z - x_j$ then yields (7).

The discovery value is defined by equation (10) and sets the expected myopic net gain of discovering more products over immediately taking a hypothetical outside option offering utility z to zero. Corollary 1 in Choi et al. (2018) then implies that:

$$
\begin{aligned}
Q_d(c_d, c_s, z) &= \mathbb{E}_{\boldsymbol{X},\boldsymbol{Y}}\left[\max\left\{z, \max_{k\in\{1,\ldots,n_d\}} \tilde{W}_k\right\}\right] - z - c_d \\
&= \int_{\max\{0,z\}}^{\infty} 1 - H(w)\mathrm{d}w - \int_{\min\{z,0\}}^{0} H(w)\mathrm{d}w - z - c_d
\end{aligned}
$$

where $H(\cdot)$ denotes the cumulative density of the random variable $\max\left\{z, \max_{k\in\tilde{J}} \tilde{W}_j\right\}$. The above also implies that in the case where Y is independent of X, a change in variables in the ingegration yields that the discovery value is linear in the mean of X, denoted by μ_X:

$$
z^d = \mu_X + \Xi(c_s, c_d)
$$

where $\Xi(c_s, c_d)$ solves (10) for an alternative random variable $\tilde{X} = X - \mu_X$.

C Monotonicity and Extensions

Monotonicity of the Gittins index values ($g_t^d \geq g_{t+1}^d \forall t$) is satisfied whenever the following holds:

$$
\begin{aligned}
0 \leq &\mathbb{E}_{\boldsymbol{X},Y,n_d,J,t}\left[\Theta(\Omega_{t+1}, \tilde{A}_{t+1}, g_t^d)\right] \\
&- \mathbb{E}_{\boldsymbol{X},Y,n_d,J,t+1}\left[\Theta(\Omega_{t+2}, \tilde{A}_{t+2}, g_{t+1}^d)\right]
\end{aligned} \tag{30}
$$

where g_t^d is the Gittins index of discovering products (defined by (24)), and $\tilde{A}_{t+1} \equiv \{e, s1, \ldots, sn_d\}$ is the set of actions available in $t+1$ containing the newly revealed products and (if available) the possible future discoveries. The expectation operator $\mathbb{E}_{\boldsymbol{X},Y,n_d,J,t}[\cdot]$ integrates over the following random realizations, where the respective joint distribution may be time-dependent: (i) Partial valuations drawn from $\boldsymbol{X} = [X_1, \ldots, X_{n_d}]$; (ii) conditional distributions $F_{Y|X=x}(y)$; (iii) the number of revealed alternatives (n_d); (iv) whether more products can be discovered in future periods determined by the beliefs over $|J|$.

It goes beyond the scope of this paper to determine all possible specifications of the joint distribution which satisfy this condition. However, Proposition 8 provides several specifications that can be of interest and for which (30) holds (see also Section 3).

Proposition 8. (30) *holds for the below deviations from the baseline model:*

i) Y is independent of X. The revealed partial valuations in $\boldsymbol{X}$ are i.i.d. with time-dependent cumulative density $G_t(x)$ such that $G_t(x) \geq G_{t+1}(x) \forall x \geq z^d - \xi$.

ii) The consumer does not know how many alternatives he will discover. Instead, he has beliefs such that with each discovery, at most the same number of alternatives are revealed as in previous periods ($n_{e,t+1} \leq n_{e,t}$).

Proof. Each part is proven using slightly different arguments.

i) Let $\tilde{x} \equiv \max_{k\in\{1,...,n_d\}} x_k$. If $\tilde{z}^s = \tilde{x} + \xi \leq z^d$, $\Theta(\Omega_{t+1}, \tilde{A}_{t+1}, z^d) = 1$, whereas for $\tilde{x} > z^d - \xi$, $\frac{\partial\Theta(\Omega_{t+1},\{e,s1,...,sn_d\},z^d)}{\partial\tilde{x}} \geq 0$. Independence implies that the cumulative density of the maximum $\tilde{x}$ is $\tilde{G}_t(x) = G_t(x)^{n_d}$. Consequently, whenever the distribution of X shifts such that $G_t(x) \geq G_{t+1}(x) \forall x \geq z^d - \xi$, larger values of $\Theta(\Omega_{t+1}, \tilde{A}_{t+1}, g_t^d)$ become less likely in $t+1$, and hence (30) holds.

ii) Since $\frac{\partial\Theta(\Omega_{t+1},\{s_k\},w)}{\partial w} \leq 1$, we have $\frac{\partial\Theta(\Omega_{t+1},\tilde{A}_{t+1},g_t^d)}{\partial n_d} \geq 0$. Hence (30) holds given $n_{e,t+1} \leq n_{e,t}$.

□

Based on this monotonicity condition, Proposition 9 generalizes Theorem 1. It implies that whenever (30) holds, the discovery value can be calculated based on the expected myopic net gain of discovering products over immediately taking the hypothetical outside option. Hence, whenever (30) holds, the optimal policy continues to be fully characterized by reservation values that can be obtained without having to consider many future periods.

Proposition 9. *Whenever* (30) *is satisfied, Theorem 1 continues to hold (with appropriate adjustment of the discovery value's time-dependence).*

Proof. Follows directly from the proof of Theorem 1. □

D Violations of independence assumption

Costly recall: Consider a variation to the search problem, where purchasing a product in the consideration set is costly unless it is bought immediately after it is inspected. If in period t product j is inspected, then inspecting another product or discovering more products in $t+1$ will change the payoff of purchasing product j by adding the purchase cost. In the context of a multi-armed bandit problem, this case arises if there are nonzero costs of switching between arms. Banks and Sundaram (1994), for example, provide a more general discussion on switching costs and the nonexistence of optimal index-based strategies. The same reasoning also applies in a search problem where inspecting a product is more costly if the consumer first discovers more products. The exception is if there are infinitely many alternatives. In this case, the optimal policy never prescribes to recall an alternative.

Learning: Independence is also violated for some types of learning. Consider a variation of the search problem, where the consumer updates his beliefs on the distribution of Y. In this case, by inspecting a product k and revealing y_{ik}, the consumer will update his belief about the distribution of Y, thus affecting the expected payoffs of both discovering more and inspecting other products. Independence therefore is violated and the reservation value policy is no longer optimal.[26] Note, however, that as long as learning is such that only payoffs of actions that will be available in the future are affected, independence continues to hold.

[26] Adam (2001) studies a similar case where independence continues to hold across groups of products. However, his results do not extend to the case with limited awareness, as the beliefs of Y also determine the expected benefits of discovering more products.

Purchase without inspection: A final setting where independence does not hold is when a consumer can buy a product without first inspecting it. In this case, the consumer has two actions available for each product he is aware of. He can either inspect a product, or directly purchase it. Clearly, when the consumer first inspects the product, the information revealed changes the payoff of buying the product. Independence therefore is violated and the reservation value policy is not guaranteed to be optimal. Doval (2018) studies this search problem for the case where a consumer is aware of all available alternatives, and characterizes the optimal policy under additional conditions

E Sellers' decisions

To illustrate the difference in sellers' decision making across the SD and DS problem, we can compare the market demand generated by the SD problem with the one from the DS problem when there are infinitely many alternatives. Given a unit mass of consumers, market demand for a product discovered at position h is given by

$$d_{SD}(h) = \mathbb{P}_{\boldsymbol{W}}\left(W_k < z^d \forall k < h\right) \mathbb{P}_{W_h}\left(W_h \geq z^d\right) \tag{31}$$

where W_h is the random effective value of a product on position h. The expression immediately follows from the stopping decision which implies that if a consumer discovers a product with $w_j \geq z^d$, he will stop searching and buy a product j. Hence, the consumer will only discover and have the option to buy a product on position h if $w_h < z^d$ for all products on earlier positions.

For the DS problem, Choi et al. (2018) showed that the market demand is given by

$$d_{DS}(h) = \mathbb{P}_{\boldsymbol{W}}\left(\tilde{W}_h \geq \max_{k \in J} \tilde{W}_k\right) \tag{32}$$

where $\tilde{W}_k = X_k + \min\{Y_k, \xi_k\}$.

Now suppose that the seller of a product on position h sets the mean of X_h, for example by choosing a price. In the SD problem, this is equivalent to choosing $\mathbb{P}_{W_h}\left(W_h \geq z^d\right)$; the probability that the consumer inspects and then stops search by buying the seller's product. Importantly, this does not directly depend on partial valuations of both products at earlier, and products at later positions. This results from the stopping decisions, and given the infinite number of products a consumer will never recall a product discovered earlier.

In contrast, in the SD problem, choosing the mean of X_h influences demand through the joint distribution of all products. As consumers are aware of all products, they compare all partial valuations. Hence, each seller's choice of partial valuations affects all other sellers demand, and sellers do not make independent decisions.