Tilburg University

Optimal Search and Awareness Expansion

Greminger, Rafael

Publication date:
2019

Document Version
Early version, also known as pre-print

[Link to publication in Tilburg University Research Portal](#)

Citation for published version (APA):
Greminger, R. (2019). *Optimal Search and Awareness Expansion*. (CentER Discussion Paper; Vol. 2019-034). CentER, Center for Economic Research.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

No. 2019-034

OPTIMAL SEARCH AND AWARENESS EXPANSION

By

Rafael P. Greminger

18 November 2019

ISSN 0924-7815
ISSN 2213-9532

Optimal Search and Awareness Expansion

Rafael P. Greminger*

November 18, 2019

Abstract

This paper introduces a search problem where a consumer has to first become aware of an alternative, before being able to search it. Initially, the consumer is aware of only a few alternatives. During search, the consumer sequentially decides between searching alternatives he is already aware of and expanding awareness to discover more products. I show that the optimal policy for this search problem is fully characterized by simple reservation values. Moreover, I prove that the purchase outcome of a consumer optimally solving the search problem is equivalent to the consumer simply choosing the product offering the largest value on a predetermined index.

1 INTRODUCTION

The existence and importance of search frictions have been well established.¹ Under the rational choice paradigm, search frictions continue to be studied relying on optimal search policies that apply in either of two settings: (i) searchers have no prior information and search randomly across alternatives (e.g. McCall, 1970; Lippman and McCall, 1976), and (ii) they are aware of all available alternatives and order their searches based on partial information (e.g. Weitzman, 1979; Chade and Smith, 2006). In this paper, I add to this literature by developing and solving a search problem where a consumer has limited awareness, and sequentially decides whether to become aware of more alternatives or search any of the alternatives he is already aware of.

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 few available phones and has some (but not all) information on what these alternatives offer. Given this information, the

*Tilburg University, Department of Econometrics & Operations Research, r.p.greminger@tilburguniversity.edu. 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 generous financial support through Research Talent grant 406.18.568.

¹See Chade et al. (2017) for a recent review of the theoretical search literature, and Jolivet and Turon (2019) for recent references on empirical search studies.

consumer can directly gather more information on these alternatives (e.g. by reading reviews online). Besides, there are also other phones available that the consumer initially is not aware of. For these alternatives the consumer has no information and does not know whether they exist. This precludes the consumer from directly inspecting these phones. Instead, he first needs to become aware of them (e.g. by getting more recommendations from friends or online sites).

In this case, limited awareness introduces a novel search problem, where the consumer not only decides in what order to search products and when to stop searching, but also when to become aware of more alternatives. Figure 1 provides an example of a choice sequence in such a search and expansion problem. Initially, the consumer is only aware of Product 1 and decides between either searching it or expanding awareness. After becoming aware of Product 2, the consumer decides whether to expand further or search either of the two products. Finally, once Product 2 has been searched, the consumer decides whether to buy Product 2 and end the search, continue expanding, or search Product 1.

FIGURE 1 – Search and awareness expansion. The consumer first expands awareness, then searches and finally buys Product 2.

Environments where consumers face a search problem of this type often emerge in practice. For example, in markets with a large number of alternatives, consumers usually will not have any information on some alternatives. Similarly, in markets where rapid technological innovations lead to a constant stream of newly available alternatives, few consumers are aware of new releases without spending effort to remain informed. Consumers also face the same type of search problem when buying products online. Online retailers and search intermediaries present visitors with 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.

Despite the complexity of the dynamic decision problem, the optimal policy is fully characterized by simple reservation values. More specifically, the optimal search policy is a generalization of Pandora’s rule 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 explicitly considering a myriad of possible future choice sequences. Based on these reservation values, it is then straightforward to analyze optimal search behavior under limited awareness.

The second contribution of this paper is to prove that the purchase outcome of a consumer solving the search and awareness expansion problem is equivalent to the same consumer having full information and directly choosing the product on a predetermined index. This allows deriving market demand without having to consider the multitude of possible choice sequences that otherwise make aggregation difficult. Besides, this result generalizes the eventual purchase theorem of Choi et al. (2018) to the case of limited awareness.²

Based on these technical results, I discuss two implications of limited awareness. First, as search progresses, it becomes less and less likely that a consumer has not yet stopped searching. Consequently, aggregate demand under limited awareness exhibits ranking effects; it decreases in the position at which products are revealed.³ As I show using a simple example, these ranking effects differ both from those generated by standard (sequential) random and directed search.

Second, a consumer's search path is independent of alternatives the consumer is not aware of, implying that the consumer fails to always first search the alternative that looks most promising from partial information. The rationale is similar to how consumers fail to buy the preferred alternative in discrete choice models when they have limited information (e.g. Goeree, 2008); hence can have similar implications when using observed search paths to infer preference and search cost parameters. To see this, consider again the example depicted in Figure 1, and suppose a third product is available. At the time the consumer decides to search Product 2, no information on the third product is available. Consequently, even when Product 3 yields a larger expected valuation than Product 2, the consumer continues to first search Product 2.

To keep exposition clear and facilitate a comparison with recent consumer search models (e.g. Armstrong, 2017; Choi et al., 2018), I present the search and awareness expansion problem using assumptions that then are relaxed in several extensions. Cases where a reservation value policy continues to be optimal include environments where (i) alternatives are ranked in decreasing order of expected utility, (ii) awareness expansion reveals an unknown number of alternatives, and (iii) multiple awareness expansion technologies are available.

The remainder of this paper is organized as follows. First, I briefly discuss related search problems. The search and awareness expansion problem is introduced in Section 2. Section 3 provides the optimal policy and discusses extensions. In Section 4, I generalize the eventual purchase theorem and show how it can be used to derive aggregate demand. Section 5 compares search outcomes from alternative search problems using a simple example, and Section 6

²Choi 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).”

³There is ample empirical evidence of ranking effects from online stores and search intermediaries, see e.g. Ursu (2018) and references therein.

concludes.

RELATED SEARCH PROBLEMS

A range of sequential search problems is nested within the present framework. 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 becoming aware of more alternatives.⁴ Besides, the extensions presented also accommodate search problems where the consumer learns while sampling from an unknown distribution (Kohn and Shavell, 1974; Rothschild, 1974; Rosenfield and Shapiro, 1981; Bikhchandani and Sharma, 1996).

To prove the optimality of the reservation value policy, I take 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,⁵ 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 problems where choosing one action reveals more actions while leaving the state of other available actions unchanged. Introducing monotonicity then allows me 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). However, the present case differs in that monotonicity is only required for the action of expanding awareness, but does not hold when searching a product.⁶

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. His search problem is also nested in an extension presented in this paper. Olszewski and Weber (2015) generalize Pandora’s rule to search problems where the final payoff depends on all the alternatives that have been searched, not only the best one. Finally, Doval (2018) analyzes the optimal policy when a searcher can directly choose alternatives without first searching them.

Other studies have estimated structural models that were based on search problems that are suited for the particular environment. Most closely related is a recent working paper by Choi

⁴Others include classical stopping problems such as those considered in McCall (1970) or Lippman and McCall (1976).

⁵Gittins 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).

⁶The utility revealed when searching a product with some probability is larger than the reservation value of searching that product.

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 searching 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 my results imply that as long as the consumer has not yet revealed the last alternative, it will never be optimal to go back and search a product that was revealed earlier. Hence my results provide a justification for studying a reduced search problem where the consumer is precluded from searching 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 search 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 search products that they are not aware of. However, in their models, consumers cannot become aware of more alternatives. The underlying search problem thus is equivalent to Pandora's problem introduced by Weitzman (1979).

2 THE SEARCH AND AWARENESS EXPANSION PROBLEM

A risk-neutral consumer with unit demand faces a market offering a (possibly infinite)⁸ number of alternatives 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 utility when purchasing alternative j is given by

$$u(x_j, y_j) = x_j + y_j$$

where x_j and y_j are partial 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 $j = 0$ of aborting search without a purchase

⁷Note also that 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.

⁸Note that 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.

⁹In particular, 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. By appropriately defining the distribution of signals and the distribution of utilities, conditional on these signals, the search problem then is equivalent to when the consumer

is available.

Initially, the consumer has limited information on the available alternatives. More specifically, in period t the consumer knows both 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 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 $j \in J \setminus (S_t \cup C_t)$.

During search, the consumer gathers information by sequentially deciding which action to take next in periods $t = 0, \dots, T$. If the consumer decides to expand awareness, he discovers n_e more alternatives which are added to the awareness set. If less than n_e alternatives have not yet been revealed, only the remaining alternatives are revealed. For each of the n_e alternatives, partial product information summarized in the realization x_j is revealed. Without loss of generality, I assume that products are revealed in increasing order of their index.¹⁰ To reveal the remaining characteristics of a product j , summarized in partial valuation y_j , the consumer has to search the product. This reveals full information on the product and moves it from the awareness into the consideration set. The latter implies $S_t \cap C_t = \emptyset$.

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 search products from the awareness set. Whereas the first constraint is used in virtually all search problems,¹¹ the later is novel to the proposed search problem. It implies that a product cannot be searched, 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 search a 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) Searching any product from the awareness set S_t , thus revealing y_j for that product and adding it to the consideration set.
- iii) Expanding awareness, thus revealing x_j for n_e additional products and adding them to the awareness set.

has linear additive utility.

¹⁰I.e. if $n_e = 2$, the first expansion reveals x_1 and x_2 , the second x_3 and x_4 etc.

¹¹Doval (2018) is a notable exception.

These actions are gathered in the *set of available actions*, $A_t = C_t \cup S_t \cup e$, where e indicates expanding the awareness set S_t . 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 searches product j .¹² To clearly differentiate between the different types of actions, this set can also be written as $A_t = \{b0, b3, s4, \dots, e\}$, where bj indicates purchasing and sj searching product j .

Both searching a product and expanding awareness is costly. Search and expansion costs are denoted by $c_s > 0$ and $c_e > 0$ 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 searching any of the products in the awareness set.

The consumer has beliefs over the products that he will discover when expanding awareness, as well as the partial valuations he will reveal when searching 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.¹³ 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 current expansion will be the last.¹⁴ As shown in the next section, the optimal policy is independent of the number of expansions that may be available in the future. Notice, however, that this belief specification implicitly assumes that the consumer always knows whether he can reveal n_e more alternatives. An extension presented in Section 3 covers the case where the consumer does not know how many alternatives are revealed when expanding awareness.

All information the consumer has in period t is summarized in the *information tuple* $\Omega_t = \langle \bar{\Omega}, \omega_t \rangle$. The tuple $\bar{\Omega} = \langle u(x, v), n_e, c_e, c_s, G_X(x), F_{Y|X=x}(y), q \rangle$ represents the consumer's initial knowledge. It contains the utility function, the number of products revealed when expanding the list, 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 search a product and realize y . 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 expansions will be available.

¹²Note that $C_t \cap S_t = \emptyset$.

¹³Note that the proposed policy is optimal conditional on beliefs. When beliefs are not correct, i.e. the true distribution from which the valuations are drawn is not the same as the beliefs, then an alternative policy where beliefs are updated during search could yield higher expected utility.

¹⁴Alternatively, 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 (12) in the Appendix is satisfied (see also the discussion in Section (3)).

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 the 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 are the two state variables. Figure 2 shows how they transition depending on the actions taken starting from period $t = 0$. The depicted example assumes that there are only two alternatives available and that an expansion only reveals a single alternative. If the consumer initially chooses the outside option ($b0$), no new information is revealed, and no further actions remain. If the consumer instead expands awareness, he reveals the first alternative, which then can be searched in $t = 1$. Furthermore, as $|J| = 2$, he also can expand the awareness set once more.

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

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.¹⁵ 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 $\{a_0 \in A_0, a_1 \in A_1, \dots\}$ for periods $t = 0, 1, \dots, T$ based only on information available in period t . Let Π denote the set containing all feasible policies. Formally, the consumer solves the following

¹⁵An exception is when x_j and y_j are drawn from discrete distributions, which limits the number of possible valuations that can be observed.

dynamic programming problem at period $t = 0$

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

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

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

denote the Bellman operator, where the immediate rewards $R(a)$ either are search costs, expansion costs, or the total valuation of a product j if it is bought. Hence immediate rewards $R(a)$ 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 .¹⁶ 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

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

3 OPTIMAL SEARCH AND AWARENESS EXPANSION POLICY

The optimal policy for the search and awareness problem is fully characterized by three reservation values. In what follows, I first define these reservation values, before then 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 and end the search, or search a product with known x_j , and end the search thereafter. In this decision, the consumer will choose to search whenever the following holds:

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

$Q_s(x_j, c_s, z)$ defines the expected *myopic* net gain of searching 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 outside

¹⁶In 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 expansion values defined in the next section need to be adjusted accordingly.

option after search and the gain is zero. When $x_j + y_j > z$, the gain over immediately taking the outside option is $x_j + y_j - z$. The expectation operator $\mathbb{E}_Y[\cdot]$ integrates over these outcomes.

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

$$Q_s(x_j, c_s, z_j^s) = 0 \quad (1)$$

which has a unique solution (see Lemma 1 in Adam, 2001). Conveniently, this can further be simplified to

$$z_j^s = x_j + \xi_j \quad (2)$$

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

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

$$z_j^b = u(x_j, y_j) \quad (3)$$

Based on reservation values given by (1) and (3), Weitzman (1979) showed that it cannot be optimal to search 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 search and buy in decreasing order of search and purchase values. However, this rule does not fully characterize an optimal policy for the search and awareness expansion problem, as the consumer can additionally expand awareness.

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: (i) Take a hypothetical outside option offering z immediately, or (ii) expand awareness, and then search through and buy one of the newly discovered products. The consumer will choose the latter whenever the following holds:

$$Q_e(c_e, c_s, z) \equiv \mathbb{E}_{\mathbf{X}} [V(\langle \bar{\Omega}, \omega(\mathbf{X}, z) \rangle, \{b0, s1, \dots, sn_e\}; \tilde{\pi})] - z - c_e \geq 0$$

where $\omega(\mathbf{X}, z) = \{z, x_1, \dots, x_{n_e}\}$ denotes the information the consumer has after expanding awareness, derived from the n_e revealed partial valuations gathered in $\mathbf{X} \equiv [X, \dots, X]'$. Note that in this notation, product indices were adjusted to the reduced decision problem, such that $j = 0, 1, \dots, n_e$ indicates the hypothetical outside option and the products revealed from the expansion.

$Q_e(c_e, c_s, z)$ defines the *myopic* net gain of expanding the search and optimally searching the revealed products over immediately taking the outside option. It is myopic in the sense that it ignores any possible future expansions, or searches beyond the products that are revealed. In particular, note that $V(\langle \bar{\Omega}, \omega(\mathbf{X}, z) \rangle, \{b0, s1, \dots, sn_e\}; \tilde{\pi})$ is the value function of having an

outside option offering z and optimally searching alternatives for which partial valuations in $\mathbf{X}$ are known. Future expansions and any products in S_t or C_t are excluded from the set of available actions in this value function. Finally, $\mathbb{E}_{\mathbf{X}}[\cdot]$ defines the expectation operator integrating over the joint distribution of the partial valuations in $\mathbf{X}$. Formal details on the calculation of the expectations and the value function are provided in Appendix B.

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

$$Q_e(c_e, c_s, z^e) = 0 \quad (4)$$

which has a unique solution (see Appendix A).

In the case when X is independent of Y , the expansion value defined by (4) is linear in the mean of the distribution of X .¹⁷ In particular, denote the mean of X by μ_X , and suppose Ξ solves (4) for an alternative random variable $\bar{X} = X - \mu_X$. As shown in Appendix B, the expansion value then can conveniently be written as

$$z^e = \mu_X + \Xi(c_s, c_e)$$

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 expansion values. In particular, the rules imply that in each period t , it is optimal to take the action with the largest reservation value defined by (1), (3), and (4). What is remarkable about this result is that these reservation values rank the expected payoffs of the actions over all future periods, despite being fully characterized by myopic comparisons to a hypothetical outside option.

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 expansion problem is characterized by the following three rules:*

STOPPING RULE: *Purchase $j \in C_t$ and end the search whenever $z_j^b = \tilde{z}^b(t) \geq \max\{\tilde{z}^s(t), z^e\}$.*

SEARCH RULE: *Search $j \in S_t$ whenever $z_j^s = \tilde{z}^s(t) \geq \max\{\tilde{z}^b(t), z^e\}$.*

EXPANSION RULE: *Expand awareness whenever $z^e \geq \max\{\tilde{z}^b(t), \tilde{z}^s(t)\}$.*

Proof. See Appendix A. □

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.

¹⁷A similar expression can be derived when only the mean of Y depends on X , but its variance remains constant.

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 expansions, 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 expansion 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 (1), (3) and (4). Weitzman (1979) showed that search values decrease in search costs, and increase if larger realizations y_j become more likely through a shift in the probability mass of Y .¹⁸ The same applies to the expansion value. It decreases in expansion costs, and increases if probability mass of X is shifted towards larger values. The expansion value also depends on search costs and the conditional distribution of Y through the value function. Hence the expansion value decreases in search costs, and increases if larger values of Y are more likely.

To see the latter, consider the case where each expansion only reveals a single product ($n_e = 1$). In this case, the myopic net gain of expansion reduces to (see Appendix B):

$$Q_e(c_e, c_s, z) = \mathbb{E}_X [\max \{0, Q_s(X, c_s, z)\}] - c_e$$

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 . Hence $Q_e(c_e, c'_s, z) \leq Q_e(c_e, c_s, z)$ for all z . As $Q_e(c_e, c_s, z)$ is decreasing in z (see Appendix A), it follows that the respective expansion values satisfy $z^{e'} \leq z^e$.

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

MONOTONICITY AND EXTENSIONS

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

¹⁸The latter implies that the distribution $F(y)$ is shifted such that $\tilde{F}(y) \leq F(y) \forall y$.

net benefits.¹⁹

For this monotonicity condition to hold, not all assumptions stated in the baseline search problem are required. A range of alternative specifications that adjust these assumptions while continuing to satisfy monotonicity can be conceived of. Below, I discuss those cases that I believe to be most relevant for applied work. 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. In this case, $x_1, \dots, x_J$ are realizations from (non-identical) random variables $X_1, \dots, X_J$. For example, in a market environment where sellers compete in marketing efforts for consumers to become aware of their products early on, sellers with on average better valuations may have a stronger incentive to be discovered first.²⁰ Furthermore, in online settings, consumers also have the option of directly sorting product lists (see e.g. Chen and Yao, 2017), which leads to the pattern that higher valued products are shown early on.

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 the reservation values, the only difference being that μ_X in the expansion value is updated during search.

Unknown n_e : A consumer may also not know how many alternatives he discovers when expanding awareness. For example, the consumer may believe that there are still alternatives he is not aware of and thus try to expand awareness, only to discover that he already is aware of all the available alternatives. In these cases, a belief over how many alternatives are revealed per expansion 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 in the baseline is that in $Q_e(c_e, c_s, z)$, expectations now additionally are taken based on the beliefs over how many alternatives will be revealed.

Multiple awareness technologies: Finally, there may be multiple technologies through which the consumer can become aware of more alternatives. In an online setting, for example, each technology may represent a different online shop offering alternatives. In such settings, the consumer also decides which technology to use to become aware of more alternatives. However, by assigning each of the expansion actions an expansion value, the optimal policy can be adjusted to accommodate this case.

For example, suppose that two such technologies are available in a period t . The technologies

¹⁹For 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 searched a product, the only option that remains is to either buy the product or choose the hypothetical outside option. Hence for search and purchase actions, at most one future period needs to be considered to fully capture the respective net benefits over immediately taking the outside option.

²⁰See for example the discussion on non-price advertising and the related references cited in Armstrong (2017).

can differ both in the expected benefits and costs they offer per expansion. The set of available actions then is $A_t = C_t \cup S_t \cup e_1 \cup e_2$, where e_l indicates expanding awareness through technology l . Let z_l^e denote the expansion value defined by (4) for technology l , accounting for the respective beliefs and costs underlying each technology. If both technologies satisfy the monotonicity condition, then as in Theorem 1, the optimal policy for this search problem selects in each period t the action with the largest reservation value from A_t .

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 expansion 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. More specifically, performing one action in t does not affect the payoff of any other action that is available in t .²¹

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. Hence choosing actions based only on reservation values that disregard these effects will not be optimal. Below, I present alternative search problems that violate independence to show this explicitly.

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 searched. If in period t product j is searched, then searching another product or expanding in $t + 1$ will change the payoff of purchasing product j by adding the purchase cost.²² Similarly, in a search problem where searching a product is more costly if the consumer first expanded further, expanding will change the action of searching a product in the overview set.²³ Appendix D discusses this in more detail using a numerical example.

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 searching 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 expanding awareness and searching other products. Hence independence is violated and the reservation value policy is no longer optimal.²⁴ Note, however, that as long as learning is such that only payoffs of actions that will

²¹Formally, this guarantees that the independence assumption in the branching bandits framework of Keller and Oldale (2003) is satisfied.

²²In the context of a multi-armed bandit problem, this case arises if there are nonzero costs of switching between arms. See also Banks and Sundaram (1994) for a more general discussion of switching costs and the nonexistence of optimal index-based strategies.

²³The exception is if there are infinitely many alternatives. In this case, the optimal policy never prescribes to recall an alternative.

²⁴Adam (2001) studies a similar case where independence continues to hold across groups of products. However,

be available in the future are affected, independence continues to hold.

Purchase without search: A final setting where independence does not hold is when a consumer can buy a product without first inspecting it. In this case, when a consumer is aware of a particular product he has two actions available. Either, he can search the product, or directly purchase it and end the search. Clearly, when the consumer first searches the product, the information revealed changes the payoff of buying the product. Hence independence 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.

4 EVENTUAL PURCHASES AND MARKET DEMAND

In an environment where consumers sequentially search products, market demand results from aggregating different possible choice sequences leading to individual purchases. Conceptually, this poses a major challenge, as the number of possible choice sequences grows faster than exponentially in the number of available alternatives.²⁵

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. Hence the market demand function can be obtained without having to explicitly consider the myriad of possible choice sequences.

Theorem 2. *Let*

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

be the effective value for product j revealed after h_j expansions, where $f(h_j)$ can be any non-negative function that is strictly decreasing in h_j . The solution to the search and awareness expansion problem leads to the eventual purchase of the product with the largest effective value.

Proof. As a product 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 discovered before j ($h_k < h_j$). In this case, $w_k > w_j$ if and only if either (i) $\tilde{w}_k \equiv \min \{ z_k^s, z_k^b \} \geq z^e$ or (ii) $z^e > \tilde{w}_k > \tilde{w}_j$. In the former, the consumer will not expand beyond product k and hence never buy product j . This follows from $z_k^s \geq z^e$ guaranteeing that k is searched, and $z_k^b \geq z^e$ implying

his results do not extend to the case with limited awareness, as the beliefs of Y also determine the expected benefits of expanding awareness.

²⁵For example, with only one alternative and an outside option, there are already four possible choice sequences. With two alternatives, the number of possible choice sequences increases to 20.

that k then is bought before expanding further. In the latter, $w_j = \tilde{w}_j < w_k = \tilde{w}_k$, and the consumer expands 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 $z^e > \tilde{w}_k > \tilde{w}_j$, which is the same as (ii) above. $\square$

This result generalizes the *eventual purchase theorem* of Armstrong (2017) and 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 search and the purchase value of a product in the awareness set exceed the expansion value, the consumer will buy the product and end the search. This is captured in the effective values through the term $z^e + f(h_j)$, which ranks alternatives based on when during search they are discovered.

The first term in the effective value then is equivalent to what Choi et al. (2018) define as effective value. As Choi et al. (2018) show, it fully characterizes the purchase decision conditional on the awareness set. Intuitively, this term reflects that unless both the search and the purchase value of an alternative are large relative to those of other alternatives, it will not be searched and/or bought. For example, if the search value of an alternative j is smaller than both the search and purchase value of another alternative k , then the optimal policy requires the consumer to first search and then buy k , without first searching j .

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

Based on Theorem 1, it is straightforward to derive market demand functions when heterogeneous consumers optimally solve the search and expansion problem. Once a distribution of consumers' valuations and/or search costs is specified, it is only necessary to derive the resulting distribution of the effective values, which then yields the market demand. 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 $\mathbf{W} = [W_0, \dots, W_{|J|}]'$. For a unit mass of consumers the market demand for a product j then is given by

$$D_j = \mathbb{P}_{\mathbf{W}} (W_j \geq W_k \forall k \in J \setminus j) \quad (5)$$

An important property of the demand resulting from consumers having limited awareness is that it exhibits ranking effects; products that are revealed early on are more likely to be bought. To gain some intuition, I now discuss the market demand for a particular environment, and then show how ranking effects emerge.

Proposition 1 provides an explicit market demand function for a specific environment.²⁶ The

²⁶The intuition in more general cases remains the same. Having $n_e > 1$ would add complexity in notation without

first term in (6), $\mathbb{P}_{\tilde{\mathbf{W}}}(\tilde{W}_k < z^e \forall k \in J_h)$, captures that a share of consumers will stop searching before having become aware of j . The second term, $\mathbb{P}(\tilde{W}_j > z^e)$, reflects those consumers that reveal j and do not continue expanding. This follows directly from the optimal policy, which implies that consumers will stop searching as soon as both the search and purchase values exceed the expansion value. Finally, the last term reflects those consumers which continue expanding until they have reached the end of the list, and then return to buy product j .

Proposition 1. *Let products $j = 1, 2, \dots$ be horizontally differentiated such that the valuations x_j and y_j are independent realizations of X and Y . Furthermore, let all consumers reveal only one product per expansion ($n_e = 1$) in the same order and be homogeneous with respect to their search and expansion costs. The market demand for product j revealed after h expansions then is given by*

$$D_{j|h} = \mathbb{P}_{\tilde{\mathbf{W}}}(\tilde{W}_k < z^e \forall k \in J_h) \left[\mathbb{P}_{\tilde{\mathbf{W}}_j}(\tilde{W}_j \geq z^e) \right. \\ \left. + \mathbb{P}_{\tilde{\mathbf{W}}}(\tilde{W}_k < z^e \forall k \in J \setminus J_h) \mathbb{P}_{\tilde{\mathbf{W}}}(\tilde{W}_j \geq \tilde{W}_k \forall k \in J \setminus J_h | \tilde{W}_k \leq z^e \forall k \in J) \right] \quad (6)$$

where $\tilde{W}_j = X + \min\{\xi(c_s), Y\}$, and $J_h \equiv \{k | k \in \{1, \dots, J\}, h_k \leq h\}$ is the set of products revealed up to and including expansion h .

Proof. Four cases need to be considered. (i) If $\tilde{w}_k > z^e$ for some $k \in J_h$, then $w_j < w_k$ and Theorem 2 implies that j is not bought. (ii) If (i) does not apply, then $\tilde{w}_j \geq z^e$ implies that $w_j \geq w_k \forall k \in J \setminus J_h$ and j is bought. (iii) If (i) and (ii) do not apply and if $\tilde{w}_k \geq z^e$ for some $k \in J \setminus J_h$, then $w_j < w_k$ and j is not bought. Finally, if (i)–(iii) do not apply, then $w_k = \tilde{w}_k \forall k$ and j is bought if and only $\tilde{w}_j \geq \tilde{w}_k \forall k \in J \setminus j$. Appropriately conditioning the joint probability in (5) across these four cases then yields the expression. $\square$

To see the ranking effects in more detail, suppose that there are infinitely many alternatives available. In this case, the probability of returning to product j will be zero, and the ranking effect is equal to

$$D_{j|h} - D_{j|h-1} = \left[\mathbb{P}(\tilde{W}_k \leq z^e \forall k \in J_h) - \mathbb{P}(\tilde{W}_k \leq z^e \forall k \in J_{h-1}) \right] \mathbb{P}(\tilde{W}_j > z^e) < 0$$

This expression reveals that the market demand for product j when moved to a later position decreases due to an increase in the probability of search having stopped before having reached j . In particular, some consumers will not buy the products at later positions, independent of how they value them. This is the case whenever $w_k > z^e$ for a product k revealed after h_k expansions, such that the consumer does not continue expanding. As a consequence, he does not search or buy products that would be revealed at later expansions, independent of their valuation.

providing additional insights.

Intuitively, limited awareness leads to ranking effects through the probability of a consumer not having stopped searching before becoming aware of a particular alternative. As search progresses, it becomes less likely that the consumer has not yet bought an alternative and thus ended the search. This leads to a lower demand for products the consumer discovers later on. Though similar ranking effects can be induced by either random search where the order by which products are searched is exogenously given, or by introducing product specific search costs into directed search (e.g. Ursu, 2018). As the next section highlights, however, both approaches lead to distinct search paths and market demand.

When there are only a finite number of products available, the above derivation does not hold. However, Proposition 3 in Appendix A shows that negative ranking effects persist when $|J| < \infty$.

Finally, note that in the above environment, all consumers discover products in the same order. In most settings, however, not all consumers will discover products in this way. For such cases, (6) defines the market demand conditional on a particular product order. The unconditional market demand then is obtained by taking expectations over the possible orders in which products are revealed. Hence ranking effects persist in the more general case, except for the case where all orders are equally likely.

5 COMPARISON OF SEARCH OUTCOMES

A range of sequential search problems are nested within the search and awareness expansion problem. Directed search as in Weitzman (1979) results if the consumer has full awareness (i.e. $S_0 = J$), or expansion costs are equal to zero. The latter makes the consumer first expand awareness until no further alternatives are left, after which he searches based on the selection and stopping rule only. Classical random search (e.g. Lippman and McCall, 1976) results if in addition the information received when becoming aware of an alternative is non-informative (e.g. $x_j = 0 \forall j$). In this case, the consumer will randomly search any of the available alternatives.

In more general settings, both random and directed search lead to different search paths and market demand. To highlight these differences, I now present a simple numerical example.

The environment is as follows: A unit mass of consumers searches in a market with 2 horizontally differentiated alternatives and no outside option. Consumer i values alternative j at $x_{ij} + y_{ij}$, where x_{ij} and y_{ij} are i.i.d. realizations of random variables X and Y respectively. X and Y follow discrete uniform distributions with respective support $x_{ij} \in \{0, 0.5, 1\}$ and $y_{ij} \in \{-0.5, 0, 0.5\}$. Except for their valuations, consumers are homogeneous and have search costs $c_s = 0.2$ and expansion costs $c_e = 0.2$. Each expansion reveals a single product ($n_e = 1$), hence search and expansion values are given by $z_j^s = x_j - 0.05$ and $z^e = 0.3$ respectively.

Initially, consumers are aware of only the first alternative, and then decide whether to search

alternative 1, or become aware of alternative 2. Several search paths are possible. For example, all consumers with $x_{i1} = 0$ will first expand awareness before searching either product. Consumers with $x_{i1} = 0.5$ will first search alternative 1, before becoming aware of alternative 2. Hence all consumers with $x_{i1} = 0.5$ and $y_{i1} \in \{0, 0.5\}$ will end the search before becoming aware of the second alternative.

These search paths differ from what the solutions to standard random and directed search problems imply for consumers with the same valuations.

Random search: Under classical random search, consumers ignore the information contained by x_{i1} , and first search either alternative 1 or alternative 2 with equal probability. More comparable to the search problem with limited awareness is a stopping problem where the order by which products are revealed is predetermined and outside the consumers' control. For example, all consumers may first search alternative 1, before then deciding whether to continue searching alternative 2. To make it comparable, I assume that searching the next product incurs costs $c_s + c_e$, which captures that the same amount of information is revealed as when both expanding and searching the second product. In this case, the consumer stops searching if $x_1 + y_1 \geq z_2^{s'} \approx 0.23$.

The solution to this stopping problem with a predetermined search order still differs from the optimal search path under limited awareness. When solving the search and awareness expansion problem, consumers with bad partial valuations ($x_{i1} = 0$) will not immediately search alternative 1, and instead first become aware of the second alternative. In contrast, in the stopping problem, whether a consumer first searches alternative 1 is independent of x_{i1} .

Directed search: Under directed search, consumers directly compare x_{i1} and (after searching) y_{i1} with $x_{i2} - 0.05$ to determine whether to first search and then buy alternative 1.²⁷ Hence consumers with $x_{i1} = 0.5$ and $x_{i2} = 1$ will always first search alternative 2. This contrasts the optimal outcomes of the search problem with limited awareness, where consumers with $x_{i1} = 0.5$ and $y_{i1} \in \{0, 0.5\}$ will never search alternative 2 because they stop searching before becoming aware of it. Hence even when $x_{i2} = 1$, these consumers will never search it.

The distinct search paths lead to differences in aggregate demand. Based on the search problem with limited awareness, the demand for alternative 1 is approximately $D_1 \approx 0.65$. This highlights the ranking effect resulting from consumers initially being aware only of alternative 1.²⁸ In classical random or directed search, this ranking effect will not occur as consumers are equally likely to buy either product.

When all consumers first search alternative 1 due to exogenous factors, there is also a ranking effect in the demand. In particular, in this case the demand will be $D_1 \approx 0.72$, indicating a more pronounced ranking effect. The reason is that those consumers with $x_{i1} = 0$ and

²⁷The reservation value in directed search is equivalent to the search value, hence the consumer searches based on reservation values $x_{ij} - 0.05$.

²⁸Note that $D_2 = 1 - D_1 < D_1$.

$y_{i1} = 0.5$ will now buy alternative 1 before searching alternative 2.²⁹ In contrast, under limited awareness, consumers with $x_{i1} = 0$ will first become aware of alternative 2, and end up buying it if $x_{i2} \in \{0.5, 1\}$. Hence search and expansion leads to less pronounced ranking effects, as products with low partial valuations will not be searched initially.

The differences in outcomes across search problems highlighted in the above example indicate two areas where accounting for limited awareness may be important. First, limited awareness can lead to biased inference when using a revealed preference approach to infer preference parameters and search costs from observed search paths.³⁰ As the example shows, consumers not searching some alternatives can be unrelated to the partial valuations (x_{ij}) of these products. A search model that disregards limited awareness then may spuriously attribute this lack of search either to consumers having low partial valuations or large search costs.

Given the prevalence of ranking effects in click-stream data, empirical studies have proposed an approach to alleviate this problem (e.g. Chen and Yao, 2017; Ursu, 2018). In this approach, search costs are modeled such that they depend on the position at which the alternatives are revealed to the consumer. This, however, still implies that consumers compare partial valuations across all products to inform their search decisions. For the consumer to buy the first product before searching the second one then would require that $(x_{i1} - x_{i2}) + y_{i1} \geq \xi(c_{s2})$, where c_{s2} is the cost of searching product 2. If valuations are observed by the researcher, observing a consumer with $x_{i1} + y_{i1} = 0.5$ and $x_{i2} = 1$ requires that $c_{s2} \geq 0.5$. In other words, rationalizing this observed path based on the alternative formulation of the search problem requires very large search costs.

In a more realistic setting, valuations are not observed by the researcher and a full estimation approach will consider additional search decisions. Nonetheless, the above example suggests that even when allowing search costs to depend on position, it is not guaranteed that such an approach fully accounts for consumers having limited awareness. It thus remains unclear how well this approach performs in settings where consumers have limited awareness.

Second, aggregate demand and ranking effects generated by consumers solving the search problem under limited awareness are distinct from those generated by solutions to other sequential search problems. This suggests that limited awareness can lead to distinct equilibrium outcomes in competitive settings where sellers set prices that are part of the valuations. For example, an interesting result of Choi et al. (2018) is that as search costs increase, sellers lower their prices. This happens as when consumers observe all prices prior to search, sellers have an incentive to lower prices to be searched first. As search costs increase, the benefit of being searched first increases, thus leading to lower prices. Limited awareness, however, mutes this channel as consumers would observe only prices of the products that they are aware of. Hence

²⁹Note

³⁰Such an approach is used for example by De Los Santos et al. (2012); Honka (2014); Chen and Yao (2017); Ursu (2018).

under limited awareness, sellers' demand functions depend on whether they are in the awareness set of consumers, leading to different price setting incentives and equilibrium outcomes.

MYOPIC POLICY

Besides comparing search outcomes across search problems, they can also be compared across different feasible policies. Most interesting is a comparison of the optimal to a myopic policy. The latter does not take into account choices available in the future. It thus always prescribes to perform the action that offers the largest expected myopic net benefit over taking the currently best option in the consideration set. If no action offers a positive expected myopic net benefit, the myopic policy prescribes the consumer to end the search.³¹

For $n_e = 1$, the myopic (expected) net benefits of expanding awareness are given by $Q_e(c_e, c_s, z)$, where z denotes the utility offered by the currently best option. This does not define myopia in a narrow sense; it takes into account future searches of the newly revealed product. I apply this definition as a policy based on a narrower definition of myopia will never prescribe to expand, given that an expansion generates positive net benefits only if a newly revealed product is searched and then bought. Without any future actions, expanding awareness thus only generates costs c_e , but no benefits.

For $n_e > 1$ it is less clear how myopic net benefits should be defined; they could take into account all of the revealed products, only the best one, or ignore that $n_e > 1$ all together. However, in all these cases, the myopic net benefits will understate the full value of an expansion, as the option value of being able to search the other revealed products is not taken into account. Hence the myopic policy tends not to prescribe an expansion when it would be optimal, leading to a premature stop of search.

Focusing on the case where $n_e = 1$, the following reveals an additional channel how myopia affects search decisions. Based on the above definition, a myopic policy prescribes the following expansion rule: Expand whenever $Q_e(c_e, c_s, z(t)) \geq \max \{\max_{j \in S_t} Q_s(x_j, c_s, z(t)), z(t)\}$, where $Q_s(x_j, c_s, z(t))$ is the expected myopic net benefit of searching j and $z(t) \equiv \max_{j \in C_t} u_j$ is the value offered by the best alternative in the consideration set in t .

This expansion rule differs from the optimal one in that both $Q_e(c_e, c_s, z(t))$ and $Q_s(x_j, c_s, z(t))$ depend on $z(t)$. Hence depending on $z(t)$, the expansion rule may either prescribe to expand, or search j . This contrasts the optimal expansion rule where the current best alternative enters only through the purchase value, but not the search or expansion value. Hence it is optimal to expand instead of searching j independent of the current best alternative.

This is highlighted in Figure 3, which depicts the myopic net benefits of searching an alternative and expanding depending on the current best valuation. In the given example, the optimal policy prescribes to expand, independent of $z(t)$.³² The myopic policy will prescribe to

³¹This policy is similar to what Gabaix et al. (2006) refer to as directed cognition algorithm.

³²More specifically, in the depicted case we have $z^e \approx 0.55$ and $z_j^s \approx 0.49$, hence $z^e > z_j^s$.

expand only when $z(t) > 0$. If for example $z(t) = -0.1$, the myopic policy will contradict the optimal one and prescribe to search j .

FIGURE 3 – Myopic net benefits of search and expansion for standard normal distributions, $c_s = c_e = 0.2$, $x_j = 0$ and $n_e = 1$.

The reason is that the myopic net benefits of expanding ignore the option to go back and search j . Hence if an expansion reveals a low partial valuation x_k , the consumer is left with $z(t)$ and cannot search j . Consequently, when $z(t)$ is small relative to x_j , the myopic net benefits of expanding are smaller and the myopic policy selects to search j , instead of expanding. However, when $z(t)$ is large relative to x_j , the option of being able to search j after expanding has no value as it will be better to take $z(t)$ instead of going back and searching j .

Finally, note that the definition of the search and expansion values implies that the myopic is equivalent to the optimal stopping rule (in the case of $n_e = 1$). As $Q_e(c_e, c_s, z) > 0$ holds if and only if $z^e > z$, it is only optimal to stop searching if the myopic net benefits of expanding are negative.³³

This is a common property of optimal stopping rules in sequential search problems with free recall (e.g. Weitzman, 1979). The myopic net benefits provide a lower bound to the total net benefits of an action. Hence whenever it is optimal to stop searching, a myopic policy will also prescribe to stop. Besides, whenever the myopic net benefits are negative, the monotonicity condition guarantees that they remain negative in all remaining periods. Thus the reverse holds as well, making the two stopping rules equivalent.

LEARNING

Finally, it is also worth pointing out that the search and awareness expansion problem nests search problems where a consumer samples from an unknown distribution and updates beliefs during search (e.g. Rothschild, 1974; Rosenfield and Shapiro, 1981; Bikhchandani and Sharma,

³³The same holds for $Q_s(x_j, c_s, z)$.

1996). Such learning problems result when the search and expansion problem is specified as follows: Partial valuations x_j are non-informative (e.g. $x_j = 0 \forall j$), $n_e = 1$, and the consumer updates beliefs about the distribution of Y during search. In this case, each expansion reveals a single product, which will be searched immediately.³⁴ $n_e = 1$ then implies $A_t = C_t \cup e$ in all periods where the consumer can discover an additional product. Learning in this case does not violate independence of actions, as it only occurs for future expansions, but does not affect the valuations of the alternatives in the consideration set. Hence as long as beliefs are updated such that monotonicity is satisfied, an optimal stopping rule will be based on the expansion value defined in (4).

Adam (2001) extended these learning problems to the case where a consumer not only decides when to stop searching (and learning), but also selects from which of multiple unknown distributions to draw from. This generalized search problem is not directly nested within the search and awareness expansion problem. However, it can be represented within the extension that adds multiple awareness technologies discussed in Section (3). In particular, let each available awareness technology represent an unknown distribution the consumer can draw from. This search problem then is equivalent to the one studied by Adam (2001), if each technology is as defined as above for the case where the consumer is drawing from a single unknown distribution. The optimal policy then again is a reservation value policy as described in Section (3).

6 CONCLUSION

This paper provides two results that apply to search problems where consumers have limited awareness and first need to become aware of alternatives before being able to search them. First, I show that the optimal policy is fully characterized by simple reservation values that can be calculated based on a myopic comparison. Second, by proving that the search problem can be transformed into a discrete choice problem, I show that market demand resulting from consumers optimally searching can be derived without having to explicitly consider a myriad of possible choice sequences.

Whereas the search problem introduced in this paper nests several classical search problems, a numerical example reveals that limited awareness leads to search paths and market demand that are distinct from those generated by existing approaches. In particular, I show that the order in which products are searched can be independent of products the consumer initially is not aware of, and that this leads to ranking effects in aggregate demand. The result suggests that consumers' limited awareness can affect equilibrium outcomes in competitive settings, as well as influence inference based on observed search paths. It is left for future research to study

³⁴As no new information is received when revealing a new product, expanding but not searching the newly revealed product cannot be part of an optimal policy.

the effects of limited awareness more generally.

APPENDIX

A PROOFS

UNIQUENESS OF EXPANSION VALUE

Proposition 2. $Q_e(c_e, c_s, z)$ is continuous, weakly decreasing and differentiable. Furthermore, (4) has a unique solution.

Proof. Let $\tilde{G}(\tilde{\mathbf{x}})$ denote the joint distribution of $\tilde{\mathbf{X}} = [X_{(1)}, \dots, X_{(n_e)}]$, where $X_{(1)}, \dots, X_{(n_e)}$ are order statistics of X (which has cumulative density $G(x)$), such that the corresponding (random) search values satisfy $Z_{(1)}^s \leq Z_{(2)}^s \leq \dots \leq Z_{(n_e)}^s$. Differentiating (10) with respect to z yields

$$\frac{\partial Q_e(c_e, c_s, z)}{\partial z} = - \sum_{q=0}^{n_e-1} \int_{\chi_q(z)} \left[1 - \prod_{k=0}^q F(z - x_{(n_e-k)}) \right] d\tilde{G}(\tilde{\mathbf{x}}) \leq 0$$

where $\chi_q(z) = \left\{ \tilde{\mathbf{X}} | Z_{(n_e-q-1)}^s \leq z \leq Z_{(n_e-q)}^s \right\}$. $Q_e(c_e, c_s, \infty) = -c_e$ and $Q_e(c_e, c_s, -\infty) = \infty$ then imply that a solution to (4) exists. Finally, uniqueness requires $Q_e(c_e, c_s, z)$ to be strictly decreasing at $z = z^e$. $\frac{\partial Q_e(c_e, c_s, z^e)}{\partial z} = 0$ requires that

$$\int_{\chi_q(z^e)} \left[1 - \prod_{k=1}^q F(z^e - X_{(q-k)}) \right] d\tilde{G}(\tilde{\mathbf{x}}) = 0$$

holds for all q . This contradicts the definition of the expansion value z^e in equation (4), as it implies $Q_e(c_e, c_s, z^e) \leq -c_e < 0$ (see Appendix B). Hence $\frac{\partial Q_e(c_e, c_s, z^e)}{\partial z} < 0$. $\square$

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 (for a detailed discussion of Gittins index policies see e.g. Gittins et al., 2011) is optimal, and that the following holds:

$$\Theta(\Omega_t, A_t, z) = b - \int_z^b \Pi_{a \in A_t} \frac{\partial \Theta(\Omega_t, \{a\}, w)}{\partial w} dw \quad (7)$$

where b is some finite upper bound of the immediate rewards.³⁵ The Gittins index of action e (expanding awareness) is defined by $g_t^e = \mathbb{E}_{\mathbf{X}} [\Theta(\Omega_{t+1}, A_{t+1} \setminus A_t, g_t^e)]$. Now consider a period t in which a future expansion will be available in $t+1$ (known due to knowledge of $|J|$). In this case we have

$$\begin{aligned} g_t^e &= \mathbb{E}_{\mathbf{X}} [\Theta(\Omega_{t+1}, \{e, s_1, \dots, s_{n_e}\}, g_t^e)] - c_e \\ &= \mathbb{E}_{\mathbf{X}} \left[b - \int_{g_t^e}^b \frac{\partial \Theta(\Omega_{t+1}, \{e\}, w)}{\partial w} \prod_{k=1}^{n_e} \frac{\partial \Theta(\Omega_{t+1}, \{s_k\}, w)}{\partial w} dw \right] - c_e \end{aligned} \quad (8)$$

where $s_k \in S_{t+1} \setminus S_t \forall k$. $\Theta(\Omega_t, \{s_k\}, z)$ is the value of a search problem with an outside option offering z and the option of searching product k (with known partial valuation x_k). $\Theta(\Omega_{t+1}, \{e\}, w)$ is the value of a search problem with an outside option offering z and the only option to expand the search. Finally, $\mathbb{E}_{\mathbf{X}} [\cdot]$ is the expectation operator integrating over the joint distribution of the n_e random variables in $\mathbf{X} = [X, \dots, X]$.

³⁵Note that immediate rewards $R(a) \geq -\max\{c_s, c_e\}$, 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 $R(a) = x_j + y_j \leq b$.

Optimality of the Gittins index policy then implies that when $z \geq g_{t+1}^e$, the consumer will choose the outside option in $t + 1$. Hence $\Theta(\Omega_t, \{e\}, w) = w \forall w \geq g_{t+1}^e$ which yields $\frac{\partial \Theta(\Omega_t, \{e\}, w)}{\partial w} = 1 \forall w \geq g_{t+1}^e$. This implies that for $g_t^e \geq g_{t+1}^e$, g_t^e does not depend on whether future expansions are available, 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. $g_t^e \geq g_{t+1}^e \forall t$, 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 $g_t^e \geq g_{t+1}^e \forall t$ holds in the proposed search problem. When $|J| = \infty$, it is clear that $g_t^e = g_{t+1}^e$ as in both periods infinitely many expansions remain and q (the consumer's belief on whether a future expansion will be available) is constant. For $|J| < \infty$, backwards induction yields that this condition holds: Suppose that in period $t + 1$, no future expansions are available. In this case, the Gittins index is given by

$$g_{t+1}^e = \mathbb{E}_{\mathbf{X}} \left[b - \int_{g_{t+1}^e}^b \prod_{k=1}^{n_e} \frac{\partial \Theta(\Omega_{t+1}, \{s_k\}, w)}{\partial w} dw \right] - c_e$$

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

$$\mathbb{E}_{\mathbf{X}} \left[b - \int_{g_t^e}^b \prod_{k=1}^{n_e} \frac{\partial \Theta(\Omega_{t+1}, \{s_k\}, w)}{\partial w} dw \right] \leq \mathbb{E}_{\mathbf{X}} \left[b - \int_{g_t^e}^b \frac{\partial \Theta(\Omega_t, \{e\}, w)}{\partial w} \prod_{k=1}^{n_e} \frac{\partial \Theta(\Omega_t, \{s_k\}, w)}{\partial w} dw \right]$$

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

Finally, as $\Theta(\Omega_{t+1}, \{e, s_1, \dots, s_{n_e}\}, g_t^e) = V(\langle \bar{\Omega}, \omega(x, z) \rangle, \{b_0, s, \dots, s_{n_e}\}; \bar{\pi})$ of the expansion value implies $z^e = g_t^e$. Similarly, the definition of the search and purchase values (in 1 and 4) 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. $\square$

NEGATIVE RANKING EFFECT

Proposition 3. $D_{j|h} - D_{j|h-1} < 0$ holds for the market demand given by Proposition (1).

Proof. Independence implies

$$\mathbb{P}_{\tilde{\mathbf{W}}}(\tilde{W}_k \leq z^e \forall k \in J_h) \mathbb{P}_{\tilde{\mathbf{W}}}(\tilde{W}_k \leq z^e \forall k \in J \setminus J_h) = \mathbb{P}_{\tilde{\mathbf{W}}}(\tilde{W}_k \leq z^e \forall k \in J)$$

Hence

$$\begin{aligned} D_{j|h} - D_{j|h-1} &= \left[\mathbb{P}(\tilde{W}_k \leq z^e \forall k \in J_h) - \mathbb{P}(\tilde{W}_k \leq z^e \forall k \in J_{h-1}) \right] \left[\mathbb{P}(\tilde{W}_j > z^e) \right. \\ &\quad \left. + \mathbb{P}_{\tilde{\mathbf{W}}}(\tilde{W}_k \leq z^e \forall k \in J) \mathbb{P}(\tilde{W}_j \geq \tilde{W}_k \forall k \in J \setminus j | \tilde{W}_k \leq z^e \forall k \in J) \right] \\ &< 0 \end{aligned}$$

$\square$

B SEARCH AND EXPANSION VALUES

The search value of a product j is defined by equation (1) and sets the myopic net gain of the search over immediately taking a hypothetical outside option offering utility z to zero. This myopic net gain can be calculated as follows:

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

Substituting $\xi_j = z - x_j$ then yields (2).

The expansion value is defined by equation (4) and sets the expected myopic net gain of expanding the search over immediately taking a hypothetical outside option offering utility z to zero. Using (7), this myopic net gain can be rewritten as:

$$\begin{aligned}
Q_e(c_e, c_s, z) &= \mathbb{E}_{\mathbf{X}} [V(\langle \bar{\Omega}, \omega(\mathbf{X}, z) \rangle, \{b_0, s, \dots, sn_e\}; \tilde{\pi})] - z - c_e \\
&= \mathbb{E}_{\mathbf{X}} \left[b - \int_z^b \prod_{k=1}^{n_e} \frac{\partial \Theta(\Omega_t, \{s_k\}, w)}{\partial w} dw \right] - z - c_e
\end{aligned}$$

Note that $\Theta(\Omega_t, \{s_k\}, w) = w + \max\{0, -c_s + \mathbb{E}_Y[\max\{x_k + Y - w, 0\}]\} = w + \max\{0, Q_s(x_k, c_s, w)\}$. Since $Q_s(x_k, c_s, w)$ is weakly decreasing in w and $Q_s(x_k, c_s, w) = 0$ for $w = z_k^s$, it follows that

$$\frac{\partial \Theta(\Omega_t, \{s_k\}, w)}{\partial w} = \begin{cases} 1 & \forall w \geq z_k^s \\ 1 + \frac{\partial Q_s(x_k, c_s, w)}{\partial w} = F(w - x_k) & \forall w < z_k^s \end{cases} \quad (9)$$

Let $\tilde{\mathbf{X}} \equiv [X_{(1)}, \dots, X_{(n_e)}]$ where $X_{(1)}, \dots, X_{(n_e)}$ are order statistics of X (with cdf $G(x)$), such that the corresponding search values (which are non-linear transformations of X) satisfy

$Z_{(1)}^s \leq Z_{(2)}^s \cdots \leq Z_{(n_e)}^s$.³⁶ Using the above and defining $Z_{(0)}^s \equiv z$ then yields:³⁷

$$\begin{aligned}
Q_e(c_e, c_s, z) &= \mathbb{E}_{\tilde{\mathbf{X}}} \left[b - \int_{\max\{Z_{(n_e)}^s, z\}}^b 1 dz - \int_{\max\{Z_{(n_e-1)}^s, z\}}^{\max\{Z_{(n_e)}^s, z\}} F(w - X_{(n_e)}) dw - \cdots \right. \\
&\quad \left. \cdots - \int_z^{\max\{Z_{(1)}^s, z\}} \prod_{k=0}^{n_e-1} F(w - X_{(n_e-k)}) dw \right] - z - c_e \\
Q_e(c_e, c_s, z) &= \mathbb{E}_{\tilde{\mathbf{X}}} \left[\max\{Z_{(n_e)}^s, z\} - \int_{\max\{Z_{(n_e-1)}^s, z\}}^{\max\{Z_{(n_e)}^s, z\}} F(w - X_{(n_e)}) dw - \cdots \right. \\
&\quad \left. \cdots - \int_z^{\max\{Z_{(1)}^s, z\}} \prod_{k=0}^{n_e-1} F(w - X_{(n_e-k)}) dw \right] - z - c_e \\
&= \mathbb{E}_{\tilde{\mathbf{X}}} \left[\int_{\max\{Z_{(n_e-1)}^s, z\}}^{\max\{Z_{(n_e)}^s, z\}} [1 - F(w - X_{(n_e)})] dw - \cdots \right. \\
&\quad \left. \cdots - \int_z^{\max\{Z_{(1)}^s, z\}} \left[1 - \prod_{k=0}^{n_e-1} F(w - X_{(n_e-k)}) \right] dw \right] - c_e \\
Q_e(c_e, c_s, z) &= \mathbb{E}_{\tilde{\mathbf{X}}} \left[\sum_{q=0}^{n_e-1} \int_{\max\{Z_{(n_e-q-1)}^s, z\}}^{\max\{Z_{(n_e-q)}^s, z\}} \left[1 - \prod_{k=0}^q F(w - X_{(n_e-k)}) \right] dw \right] - c_e \quad (10)
\end{aligned}$$

In the case when each expansion only reveals a single product ($n_e = 1$), the above simplifies to

$$\begin{aligned}
Q_e(c_e, c_s, z) &= \mathbb{E}_X \left[\int_z^{\max\{Z^s, z\}} [1 - F(w - X)] dw \right] - c_e \\
&= \mathbb{E}_X \left[\max \left\{ 0, \int_z^{Z^s} [1 - F(w - X)] dw \right\} \right] - c_e \\
&= \mathbb{E}_X \left[\max \left\{ 0, \int_z^\infty [1 - F(w - X)] dw - \int_{Z^s}^\infty [1 - F(w - X)] dw \right\} \right] - c_e \\
&= \mathbb{E}_X \left[\max \left\{ 0, \int_z^\infty [1 - F(w - X)] dw - c_s \right\} \right] - c_e \\
&= \mathbb{E}_X [\max \{0, Q_s(X, c_s, z)\}] - c_e
\end{aligned}$$

To derive $\frac{\partial Q_e(c_e, c_s, z)}{\partial z}$ from (10) and show that $\frac{\partial Q_e(c_e, c_s, z^e)}{\partial z} < 0$ (required for Proposition 2), I present the case for $n_e = 2$ and with Y being independent of X , thus simplifying exposition. More general cases can be derived in a similar fashion. The steps are shown below, where I define $X_{(0)} \equiv z - \xi$.

³⁶The order in $X_{(1)}, \dots, X_{(n_e)}$ can differ if the conditional distributions vary such that $\xi_{(k)} < \xi_{(k-1)}$ in (2).

³⁷The second step follows from expanding with $-\max\{Z_{(n_e)}^s, z\} + \max\{Z_{(n_e-1)}^s, z\} + \int_{\max\{Z_{(n_e-1)}^s, z\}}^{\max\{Z_{(n_e)}^s, z\}} 1 dw - \max\{Z_{(n_e-1)}^s, z\} + \max\{Z_{(n_e-2)}^s, z\} + \int_{\max\{Z_{(n_e-2)}^s, z\}}^{\max\{Z_{(n_e-1)}^s, z\}} 1 dw \cdots - \max\{Z_{(1)}^s, z\} + z + \int_z^{\max\{Z_{(1)}^s, z\}} 1 dw$.

$$\begin{aligned}
Q_e(c_e, c_s, z) &= \mathbb{E}_{\tilde{\mathbf{X}}} \left[\sum_{q=0}^{n_e-1} \int_{\max\{X_{(n_e-q)}+\xi, z\}}^{\max\{X_{(n_e-q)}+\xi, z\}} \left[1 - \prod_{k=0}^q F(w - X_{(n_e-k)}) \right] dw \right] - c_e \\
&= -c_e + \int_{z-\xi}^{\infty} \int_{-\infty}^{z-\xi} \cdots \int_{-\infty}^{z-\xi} \left[\int_z^{x_{(n_e)}+\xi} 1 - F(w - x_{(n_e)}) dw \right] d\tilde{G}(\tilde{\mathbf{x}}) \\
&\quad + \int_{z-\xi}^{\infty} \int_{z-\xi}^{\infty} \int_{-\infty}^{z-\xi} \cdots \int_{-\infty}^{z-\xi} \left[\int_{x_{(n_e-1)}+\xi}^{x_{(n_e)}+\xi} 1 - F(w - x_{(n_e)}) dw \right. \\
&\quad \quad \left. + \int_z^{x_{(n_e-1)}+\xi} 1 - F(w - x_{(n_e)}) F(w - x_{(n_e-1)}) dw \right] d\tilde{G}(\tilde{\mathbf{x}}) \\
&\quad + \cdots \\
&\quad + \int_{z-\xi}^{\infty} \cdots \int_{z-\xi}^{\infty} \left[\sum_{q=0}^{n_e-1} \int_{x_{(n_e-q-1)}+\xi}^{x_{(n_e-q)}+\xi} 1 - \prod_{k=0}^q F(w - x_{(n_e-k)}) dw \right] d\tilde{G}(\tilde{\mathbf{x}}) \quad (11)
\end{aligned}$$

For $n_e = 2$, this reduces to

$$\begin{aligned}
Q_e(c_e, c_s, z) &= -c_e + \int_{z-\xi}^{\infty} \int_{-\infty}^{z-\xi} \left[\int_z^{x_{(2)}+\xi} 1 - F(w - x_{(2)}) dw \right] d\tilde{G}(\tilde{\mathbf{x}}) \\
&\quad + \int_{z-\xi}^{\infty} \int_{z-\xi}^{\infty} \left[\int_{x_{(1)}+\xi}^{x_{(2)}+\xi} 1 - F(w - x_{(2)}) dw \right. \\
&\quad \quad \left. + \int_z^{x_{(1)}+\xi} 1 - F(w - x_{(2)}) F(w - x_{(1)}) dw \right] d\tilde{G}(\tilde{\mathbf{x}})
\end{aligned}$$

Differentiating with respect to z then yields

$$\begin{aligned}
\frac{\partial Q_e(c_e, c_s, z)}{\partial z} &= - \int_{-\infty}^{z-\xi} \left[\int_z^z 1 - F(w - x_{(2)}) dw \right] d\tilde{G}(x_{(1)}|x_{(2)} = z - \xi) \\
&\quad + \int_{z-\xi}^{\infty} \left[\int_z^{x_{(2)}+\xi} 1 - F(w - x_{(2)}) dw \right] d\tilde{G}(x_{(2)}|x_{(1)} = z - \xi) \\
&\quad - \int_{z-\xi}^{\infty} \int_{-\infty}^{z-\xi} [1 - F(z - x_{(2)}) dw] d\tilde{G}(\tilde{\mathbf{x}}) \\
&\quad - \int_{z-\xi}^{\infty} [\cdots] d\tilde{G}(x_{(1)}|x_{(2)} = z - \xi) \\
&\quad - \int_{z-\xi}^{\infty} \left[\int_z^{x_{(2)}+\xi} 1 - F(w - x_{(2)}) dw \right. \\
&\quad \quad \left. + \int_z^z 1 - F(w - x_{(2)}) F(w - x_{(1)}) dw \right] dG(x_{(2)}|x_{(1)} = z - \xi) \\
&\quad - \int_{z-\xi}^{\infty} \int_{z-\xi}^{\infty} [1 - F(z - x_{(2)}) F(z - x_{(1)}) dw] dG(\tilde{\mathbf{x}}) \\
&= - \int_{z-\xi}^{\infty} \int_{-\infty}^{z-\xi} [1 - F(z - x_{(2)}) dw] dG(\tilde{\mathbf{x}}) \\
&\quad - \int_{z-\xi}^{\infty} \int_{z-\xi}^{\infty} [1 - F(z - x_{(2)}) F(z - x_{(1)}) dw] dG(\tilde{\mathbf{x}})
\end{aligned}$$

Note that $-\int_{z-\xi}^{\infty} [\cdots] d\tilde{G}(x_{(1)}|x_{(2)} = z - \xi) = 0$, as $x_{(2)} = z - \xi$ implies that $x_{(1)} \geq z - \xi$ occurs with probability zero.

For $\frac{\partial Q_e(c_e, c_s, z^e)}{\partial z} = 0$ to hold, the above shows that the following two conditions need to hold:

$$\begin{aligned} \int_{z^e - \xi}^{\infty} \int_{-\infty}^{z^e - \xi} [1 - F(z^e - x_{(2)})] dw \, dG(\tilde{\mathbf{x}}) &= 0 \\ \int_{z^e - \xi}^{\infty} \int_{z^e - \xi}^{\infty} [1 - F(z^e - x_{(2)}) F(z^e - x_{(1)})] dw \, dG(\tilde{\mathbf{x}}) &= 0 \end{aligned}$$

Given that $F(\cdot)$ is a cumulative density,³⁸ these two conditions imply the following:

$$\begin{aligned} A &\equiv \int_{z^e - \xi}^{\infty} \int_{-\infty}^{z^e - \xi} \left[\int_{z^e}^{x_{(2)} + \xi} 1 - F(w - x_{(2)}) \, dw \right] d\tilde{G}(\tilde{\mathbf{x}}) \\ &\leq \int_{z^e - \xi}^{\infty} \int_{-\infty}^{z^e - \xi} [1 - F(z^e - x_{(2)})] [x_{(2)} + \xi - z^e] \, d\tilde{G}(\tilde{\mathbf{x}}) \\ &\leq 0 \\ B &\equiv \int_{z^e - \xi}^{\infty} \int_{z^e - \xi}^{\infty} \left[\int_{x_{(1)} + \xi}^{x_{(2)} + \xi} 1 - F(w - x_{(2)}) \, dw \right] d\tilde{G}(\tilde{\mathbf{x}}) \\ &\leq \int_{z^e - \xi}^{\infty} \int_{z^e - \xi}^{\infty} [1 - F(z^e - x_{(2)})] [x_{(2)} - x_{(1)}] \, d\tilde{G}(\tilde{\mathbf{x}}) \\ &\leq \int_{z^e - \xi}^{\infty} \int_{z^e - \xi}^{\infty} [1 - F(z^e - x_{(2)}) F(z^e - x_{(1)})] [x_{(2)} - x_{(1)}] \, d\tilde{G}(\tilde{\mathbf{x}}) \\ &\leq 0 \\ C &\equiv \int_{z^e - \xi}^{\infty} \int_{z^e - \xi}^{\infty} \left[\int_{z^e}^{x_{(1)} + \xi} 1 - F(w - x_{(2)}) F(w - x_{(1)}) \, dw \right] d\tilde{G}(\tilde{\mathbf{x}}) \\ &\leq \int_{z^e - \xi}^{\infty} \int_{z^e - \xi}^{\infty} [1 - F(z^e - x_{(2)}) F(z^e - x_{(1)})] [x_{(1)} + \xi - z^e] \, d\tilde{G}(\tilde{\mathbf{x}}) \\ &\leq 0 \end{aligned}$$

As $Q_e(c_e, c_s, z) = A + B + C - c_e$, it follows that $\frac{\partial Q_e(c_e, c_s, z^e)}{\partial z} = 0$ implies $Q_e(c_e, c_s, z^e) \leq -c_e < 0$.

To show that z^e is linear in the mean of X given independence of Y , a change in variables is applied to (11). In particular, substituting $\bar{x}_{(k)} = x_{(k)} - \mu_X \forall k \in \{1, \dots, n_e\}$ (again denoting $\bar{x}_{(0)} \equiv z - \xi$) and denoting $\bar{\mathbf{x}} = [\bar{x}_{(1)}, \dots, \bar{x}_{(n_e)}]$ in (11) yields

$$\begin{aligned} Q_e(c_e, c_s, z) &= -c_e + \int_{z - \mu_X - \xi}^{\infty} \int_{-\infty}^{z - \mu_X - \xi} \dots \int_{-\infty}^{z - \mu_X - \xi} \left[\int_z^{\bar{x}_{(n_e)} + \xi} 1 - F(w - \bar{x}_{(n_e)}) \, dw \right] d\tilde{G}(\bar{\mathbf{x}}) \\ &\quad + \int_{z - \mu_X - \xi}^{\infty} \int_{z - \mu_X - \xi}^{\infty} \int_{-\infty}^{z - \mu_X - \xi} \dots \int_{-\infty}^{z - \mu_X - \xi} \left[\int_{\bar{x}_{(n_e-1)} + \xi}^{\bar{x}_{(n_e)} + \xi} 1 - F(w - \bar{x}_{(n_e)}) \, dw \right. \\ &\quad \left. + \int_z^{\bar{x}_{(n_e-1)} + \xi} 1 - F(w - \bar{x}_{(n_e)}) F(w - \bar{x}_{(n_e-1)}) \, dw \right] d\tilde{G}(\bar{\mathbf{x}}) \\ &\quad + \dots \\ &\quad + \int_{z - \mu_X - \xi}^{\infty} \dots \int_{z - \mu_X - \xi}^{\infty} \left[\sum_{q=0}^{n_e-1} \int_{\bar{x}_{(n_e-q-1)} + \xi}^{\bar{x}_{(n_e-q)} + \xi} 1 - \prod_{k=0}^q F(w - \bar{x}_{(n_e-k)}) \, dw \right] d\tilde{G}(\bar{\mathbf{x}}) \end{aligned}$$

³⁸Being a cumulative density implies $F(y) \in [0, 1] \forall y \in \mathbb{R}$ and $\frac{\partial F(y)}{\partial y} \geq 0$.

Let $\Xi \equiv z - \mu_X$, such that

$$Q_e(c_e, c_s, \Xi + \mu_X) = \mathbb{E}_{\tilde{\mathbf{X}}} \left[\sum_{q=0}^{n_e-1} \int_{\max\{\bar{X}_{(n_e-q)} + \xi, \Xi\}}^{\max\{\bar{X}_{(n_e-q)} + \xi, \Xi\}} \left[1 - \prod_{k=0}^q F(w - \bar{X}_{(n_e-k)}) \right] dw \right]$$

where the expectations operator now works on the order statistics $\tilde{\mathbf{X}} = [\bar{X}_{(1)}, \dots, \bar{X}_{(n_e)}]$. When Ξ solves $Q_e(c_e, c_s, \Xi + \mu_X) = 0$, $z = \Xi + \mu_X$ solves $Q_e(c_e, c_s, z) = 0$, and thus $z^e = \mu_X + \Xi$ when Y is independent of X .

C MONOTONICITY AND EXTENSIONS

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

$$\begin{aligned} 0 \leq & \mathbb{E}_{\mathbf{X}, Y, n_e, J, t} \left[\Theta(\Omega_{t+1}, \tilde{A}_{t+1}, g_t^e) \right] \\ & - \mathbb{E}_{\mathbf{X}, Y, n_e, J, t+1} \left[\Theta(\Omega_{t+2}, \tilde{A}_{t+2}, g_{t+1}^e) \right] \end{aligned} \quad (12)$$

where g_t^e is the Gittins index of expanding (defined by (8)), and $\tilde{A}_{t+1} \equiv \{e, s1, \dots, sn_e\}$ is the set of actions available in $t+1$ containing the newly revealed products and (if available) the possible future expansion. The expectation operator $\mathbb{E}_{\mathbf{X}, Y, n_e, J, t} [\cdot]$ integrates over the following random realizations, where the respective joint distribution may be time-dependent: (i) Partial valuations drawn from $\mathbf{X} = [X_1, \dots, X_{n_e}]$; (ii) conditional distributions $F_{Y|X=x}(y)$; (iii) the number of revealed alternatives (n_e); (iv) whether an additional expansion is available, 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 4 provides several specifications that are of interest and for which (12) holds (see also Section 3).

Proposition 4. (12) holds for the below deviations from the baseline model:

- i) Y is independent of X . The revealed partial valuations in $\mathbf{X}$ are i.i.d. with time-dependent cumulative density $G_t(x)$ such that $G_t(x) \geq G_{t+1}(x) \forall x \geq z^e - \xi$.
- ii) The consumer does not know how many alternatives are revealed when expanding awareness. Instead, he has beliefs such that with each expansion, at most the same number of alternatives is revealed as in the last expansion ($n_{e,t+1} \leq n_{e,t}$).

Proof. Each part is proven using slightly different arguments.

- i) Let $\tilde{x} \equiv \max_{k \in \{1, \dots, n_e\}} x_k$. If $\tilde{z}^s = \tilde{x} + \xi \leq z^e$, $\Theta(\Omega_{t+1}, \tilde{A}_{t+1}, z^e) = 1$, whereas for $\tilde{x} > z^e - \xi$, $\frac{\partial \Theta(\Omega_{t+1}, \{e, s1, \dots, sn_e\}, z^e)}{\partial \tilde{x}} \geq 0$ (which follows from (9)). Independence implies that the cumulative density of the maximum $\tilde{x}$ is $\tilde{G}_t(x) = G_t(x)^{n_e}$. Consequently, whenever the distribution of X shifts such that $G_t(x) \geq G_{t+1}(x) \forall x \geq z^e - \xi$, larger values of $\Theta(\Omega_{t+1}, \tilde{A}_{t+1}, g_t^e)$ become less likely in $t+1$, and hence (12) holds.
- ii) Since $\frac{\partial \Theta(\Omega_{t+1}, \{s_k\}, w)}{\partial w} \leq 1$, we have $\frac{\partial \Theta(\Omega_{t+1}, \tilde{A}_{t+1}, g_t^e)}{\partial n_e} \geq 0$. Hence (12) holds given $n_{e,t+1} \leq n_{e,t}$.

□

Based on this monotonicity condition, Proposition 5 generalizes Theorem 1. It implies that whenever (12) holds, the expansion value can be calculated based on the expected myopic net gain of expanding over immediately taking the hypothetical outside option. Hence, whenever (12) holds, the optimal policy continues to be fully characterized by reservation values that can be obtained without having to consider many future periods.

Proposition 5. *Whenever (12) is satisfied, Theorem 1 continues to hold (with appropriate adjustment of the expansion value's time-dependence).*

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

D COSTLY RECALL EXAMPLE

The following simple example provides some intuition on why the reservation value policy is not optimal when there is costly recall.

Suppose only 2 alternatives are available. The consumer initially is aware only of the first alternative with partial valuation $x_1 = 0.5$. X and Y follow discrete uniform distributions with respective support $x_2 \in \{0, 0.5, 1\}$ and $y_j \in \{-0.5, 0, 0.5\}$ for $j = 1, 2$. Search and expansion costs are given by $c_s = 0.2$ and $c_e = 0.1$ respectively. Moreover, recall is costly such that the consumer incurs costs $c_r = 2c_s$ if he searches alternative 1 only after having expanded awareness. In this environment, the policy of Theorem 1 prescribes to first expand to discover the second alternative.³⁹ Following this policy yields an expected total payoff of

$$\begin{aligned} U_o &= \frac{1}{3} \left[-c_e - c_r + \frac{1}{2} \right] \\ &\quad + \frac{1}{3} \left[-c_e - c_s + \frac{1}{3} \left(\frac{1}{2} + 1 \right) + \frac{1}{3} \left(-c_r + \frac{1}{3} \left(\frac{1}{2} + 1 \right) \right) \right] \\ &\quad + \frac{1}{3} [-c_e - c_s + 1] \\ &\approx 0.31 \end{aligned}$$

Now consider an alternative policy that prescribes to first search alternative 1, and then continue based on the reservation value policy (i.e. expand only if $x_1 + y_1 \leq z^e$). Following this policy yields an expected total payoff of

$$\begin{aligned} U_a &= \frac{1}{3} [-c_s + 1] + \frac{1}{3} \left[-c_s + \frac{1}{2} \right] \\ &\quad + \frac{1}{3} \left[-c_s - c_e + \frac{1}{3} \left(-c_s + \frac{1}{2} \right) + \frac{1}{3} (-c_s + 1) \right] \\ &\approx 0.39 \end{aligned}$$

Following this alternative policy thus yields a larger expected total payoff, showing that the reservation value policy is not optimal.

This results from the expansion value z^e being based only on the net benefits of finding a value $x_2 > x_1$. However, due to the recall costs, an expansion not only reveals a possibly better product, but also makes searching product 1 more costly. By not accounting for these added costs, z^e overstates the expected benefits of expanding. As a result, comparing z_1^s with z^e does not appropriately reflect a comparison of the expected total future payoffs of these actions, and thus may fail to prescribe the optimal action to take.

³⁹In particular, we have $z^e = 0.5$ and $z_1^s = 0.45$. Furthermore, at cost $c_r = 2c_s$, we have $z_1^{sf} = 0.15$.

REFERENCES

- Adam, K. (2001). Learning while searching for the best alternative. *Journal of Economic Theory* 101(1), 252–280.
- 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(PA), 293–307.
- 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.
- Bikhchandani, S. and S. Sharma (1996, January). Optimal search with learning. *Journal of Economic Dynamics and Control* 20(1), 333–359.
- Chade, H., J. Eeckhout, and L. Smith (2017). Sorting through search and matching models in economics. *Journal of Economic Literature* 55(2), 493–544.
- 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.
- Choi, M., A. Y. Dai, and K. Kim (2018). Consumer search and price competition. *Econometrica* 86(4), 1257–1281.
- De Los Santos, B., A. Hortaçsu, M. R. Wildenbeest, 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.
- Doval, L. (2018). Whether or not to open Pandora’s box. *Journal of Economic Theory* 175, 127–158.
- Gabaix, X., D. Laibson, G. Moloche, and S. Weinberg (2006). Costly information acquisition: Experimental analysis of a boundedly rational model. *The American Economic Review* 96(4), 1043–1068.
- 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.
- Goeree, M. S. (2008). Limited information and advertising in the u.s. personal computer industry. *Econometrica* 76(5), 1017–1074.
- Honka, E. (2014). Quantifying search and switching costs in the US auto insurance industry. *The RAND Journal of Economics* 45(4), 847–884.
- 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.
- 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.
- Kleinberg, R., B. Waggoner, and E. G. Weyl (2017). Descending price optimally coordinates search.
- Kohn, M. G. and S. Shavell (1974). The theory of search. *Journal of Economic Theory* 9(2), 93–123.
- Koulayev, S. (2014). Search for differentiated products: identification and estimation. *The RAND Journal of Economics* 45(3), 553–575.
- 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.
- Morozov, I. (2019). Measuring benefits from new products in markets with information frictions.
- Olszewski, W. and R. Weber (2015). A more general Pandora rule? *Journal of Economic Theory* 160, 429–437.
- Rosenfield, D. B. and R. D. Shapiro (1981). Optimal adaptive price search. *Journal of Economic Theory* 25(1), 1–20.
- Rothschild, M. (1974). Searching for the lowest price when the distribution of prices is unknown. *Journal of Political Economy* 82(4), 689–711.
- 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.