Expert-Mediated Sequential Search☆

Meenal Chhabra[a,1], Sanmay Das[b,1,*], David Sarne[c,*]

[a] *Dept. of Computer Science, Virginia Tech, Blacksburg, VA, USA*
[b] *Dept. of Computer Science & Engineering, Washington University in St.Louis, St. Louis, MO, USA*
[c] *Dept. of Computer Science, Bar Ilan University, Ramat Gan, Israel*

Abstract

This paper studies markets, such as Internet marketplaces for used cars or mortgages, in which consumers engage in sequential search. In particular, we consider the impact of information-brokers (*experts*) who can, for a fee, provide better information on true values of opportunities. We characterize the optimal search strategy given a price and the terms of service set by the expert, and show how to use this characterization to solve the monopolist expert's service pricing problem. Our analysis enables the investigation of three common pricing schemes (pay-per-use, unlimited subscription, and package pricing) that can be used by the expert. We demonstrate that in settings characteristic of electronic marketplaces, namely those with lower search costs for consumers and lower costs of production of expert services, unlimited subscription schemes are favored. Finally, we show that the platform that connects consumers and experts can improve social welfare by subsidizing the purchase of expert services. The optimal level of subsidy forces the buyer to exactly fully internalize the marginal cost of provision of expert services. In electronic markets, this cost is minimal, so it may be worthwhile for the platform to make the expert freely available to consumers.

Keywords: E-commerce, Artificial intelligence, Sequential decision making, Expert-mediated search

1. Introduction

We study markets in which consumers search sequentially for a single item that they are interested in acquiring. A classic example is a consumer looking to buy a used car. She will

☆A preliminary version of some of these results appeared in the Proceedings of the Tenth International Conference on Autonomous Agents and MultiAgent Systems [12].

*Corresponding author

Email addresses: meenal@cs.vt.edu (Meenal Chhabra), sanmay@seas.wustl.edu (Sanmay Das), sarned@cs.biu.ac.il (David Sarne)

[1]Parts of this work were done when Chhabra and Das were at Rensselaer Polytechnic Institute and at Virginia Tech.

typically investigate cars one at a time until she decides on one she wants. Similar models apply in a range of settings beyond consumers purchasing goods: for example, sequential search theory has been applied to job-search, house search, technology R&D, and mate search [27, 35, 7, 40, 22, 42]. Analysis of search is becoming increasingly important in the context of modern electronic marketplaces because of the emergence of a plethora of online sellers and the ability to access them easily [3, 9, 20].

The main tradeoff in sequential search arises from the fact that there is a cost incurred in finding out the true value of any opportunity encountered [2, 24, 27]. For example, there is a cost to arranging a meeting to test drive a car you are considering purchasing, or to evaluating an offer received when shopping for a mortgage online. The searcher needs to trade off the potential benefit of continuing to search and seeing a possibly more valuable opportunity with the costs incurred in doing so [18]. An additional complexity is that searchers often only obtain a noisy signal of the true value. For example, the drivetrain of a used car may not be in good condition, even if the body of the car looks terrific. Similarly, when shopping for a loan online, a promising offer with a low interest rate might turn out to be a poor offer due to stringent terms and conditions, or excessive initiation fees and closing costs. The relaxation of the assumption of perfect values not only changes the optimal strategy for a searcher, it also leads to a niche in the marketplace for new information brokers. These information brokers, or *experts*, are service providers whose main role is to inform consumers or searchers about the values of opportunities. An expert offers the searcher the option to obtain a more precise estimate of the value of an opportunity in question, in exchange for the payment of a fee (which covers the cost of providing the service as well as the profit of the expert). In the used-car example, independent mechanics or agencies like Carfax that provide reports on car histories can serve as experts. The expert does not necessarily need to be an entity external to the platform the users use for reviewing opportunities. Instead, it can be a service offered by the platform itself.

Prior work on the problem of sequential search with noisy signals has treated the expert (or equivalent) as exogenous to the model, assuming that the searcher can purchase information on the true value of an opportunity as needed [29, 42]. In this work we investigate what happens when the expert and the platform connecting experts and searchers are both potentially strategic players with their own incentives.

This paper. We introduce a model of a search market with a self-interested expert that attempts to maximize her expected profit. She must decide how to set prices and other characteristics of the services she offers, choosing from a rich set of (possibly nonlinear) pricing models. The expert is taken to be a monopolist, since in many situations of interest the economies of scale involved in production of expert services favor the emergence of monopolies, and even when there are competitors in a market, the platform that connects consumers and experts may have a special relationship with one provider, or may provide expert services itself as a value-added product.

In order to study the monopolist's pricing problem, we first need to characterize the optimal response of a searcher to the expert's strategy. The most common pricing models used in today's markets are a la carte (pay-per-use) pricing, fixed-fee pricing (subscriptions to services) and non-linear package pricing (where the expert sells packages of a fixed number of uses of her services for a given price) (the first two pricing models are actually special cases of the third, but worthy of special consideration because of their ubiquity). For the general case of non-linear pricing we prove that, under a standard stochastic dominance assumption on the signal structure (that higher signals are "good news" [32]), the optimal search rule is characterized by a "double reservation value" strategy, whenever the expert can be used only upon purchasing a new package, and a "single reservation value" strategy whenever a purchased package has not been completely exhausted. With the double reservation value strategy, the searcher rejects all signals below a certain threshold, resuming search, and accepts all signals above another threshold, thus terminating search without querying the expert; the agent purchases a new service package and queries the expert for all signals that are in between the two thresholds. With the single reservation value the searcher rejects all signals below a certain threshold, resuming search, and queries the expert otherwise. This strategy has some similarities in structure to the optimal strategy in other partial information search models, but previous models do not analyze the general case of a service package.

We use the derivation of the optimal search strategy to find the monopolist expert's pricing strategy and then show how the expert's optimal strategy is affected by the fundamental features of the underlying environment, such as consumer search costs and the marginal cost of production of additional units of expert services. We find that fixed-fee or subscription models are likely to be preferred by the expert in settings with low search costs and low marginal costs of producing extra reports, which are both key features of online markets.

Finally, we use the combined characterization of optimal strategies for the searcher and the expert to study the problem from the perspective of the platform that brings them together. We find that the platform can increase social welfare by subsidizing access to expert services, and that social welfare is maximized when consumers pay exactly the marginal cost of production of expert services.

Contributions. To summarize, this paper makes four main contributions to the literature. (1) We characterize the structure of optimal search strategies when experts offer packages of (non-linearly priced) services, extending existing literature on a la carte service pricing. (2) We formulate and solve the monopolist expert's pricing problem, treating the expert as the first mover to whom searchers respond optimally. These two methodological contributions allow us to (3) formulate a novel model of noisy search where the cost of improved information acquisition is endogenous to the model, and derive results on prices as well as expert and searcher welfare in the model. Finally, (4) we derive implications for the platform that connects searchers and experts, and prove novel results about optimal subsidization in search markets. A key focus throughout our work is on understanding how important features of electronic markets, namely low search costs and low costs of delivery of expert services, affect outcomes in these expert-mediated search markets.

2. Related Work

This paper touches on several different literatures. The technical aspects of the problem from the searcher's perspective are grounded in the theory of sequential search [30, 27, 19]. The optimal stopping rule in sequential search has been widely studied, and is often a *reservation strategy*, where the searcher should terminate search once she encounters an opportunity which has a value above a certain reservation value or threshold [41, 31]. In some specific instances, the optimal strategy is a double reservation value one [29, 23].

Noisy signals have been introduced into optimal stopping problems and search models previously [43]. In particular, MacQueen [29] characterizes optimal search when disambiguating the uncertainty associated with a given opportunity incurs a fixed cost (the equivalent to the a la carte (pay-per-use) pricing method). However, MacQueen does not analyze more general information pricing schemes as we do in this paper. More importantly, MacQueen's work focuses solely on the searcher's problem, assuming fixed costs of information acquisition, and does not consider the strategic and systemic aspects of the problem from the perspective of the expert or the platform. Our model assumes that the disambiguation of noisy signals is performed by a strategic agent (the expert) who serves as an information broker and sets the cost to the decision-maker (searcher) of acquiring more information. We investigate the implications of this, focusing particularly on what it can tell us about strategies for the expert and for the platform that connects searchers with opportunities, and thus also with experts.

Other work on optimal stopping problems in the presence of noisy signals (e.g., the work of Monahan [34, 33]), is also typically limited to a la carte pricing, and significantly different from ours both in model formulation and technical detail. This also includes work on how much costly information it makes sense to acquire before making a decision [36] and on multi-attribute sequential search where additional attributes can be revealed at certain costs along the search path [26, 42]. In addition to the technical differences between these models and ours, as in the case of MacQueen, all these models also consider the cost of obtaining the information to be set exogenously (i.e., in a non-strategic manner) thus precluding any insights from the points of view of the expert and the platform.

Relaxation of the perfect signals assumption has also been explored in models of two-sided search [10], including marriage or dating markets [13] and markets with interviewing [25]. The standard assumption in these cases is usually that agents can get better signals through repeat interactions (e.g., cohabitation and repeated interviews [38]). Mediators in such models usually take the form of matchmakers rather than information brokers [8], and the focus of results is often on the stability of outcomes. The literature in this area has not, to this point, focused on the decision problems faced by self-interested knowledge brokers, or how their presence affects the market.

Our paper is focused on mathematical and computational modeling of search. There is also a literature that, while grounded in theory, takes econometric/statistical or experimental approaches to understanding search behavior. For example, Baryla et al. [5] use survey data to show that economic conditions are the dominant factor in search duration in the real-estate market, and Cason and Friedman [11] show, using lab experiments, that buyer behavior in posted price search markets can deviate from equilibrium predictions. Our model can provide theoretical grounding for future experimental and statistical studies of expert-mediated search markets.

Another major related strand of literature is on non-linear pricing. Sundararajan [39] analyzes optimal pricing for information goods (those with low or zero marginal costs of production) with a particular focus on explaining when and why subscription models (fixed-fee pricing) dominate usage-based or a la carte pricing from the seller's perspective. Starting from the empirical observation that unlimited subscription models are usually available in online services, where vanishing marginal costs and low search costs dominate, Sundararajan [39] explains why this would occur, and analyzes the optimal combination of fixed-fee and non-linear usage based pricing as a function of the characteristics of the market. Our work focuses on sequential search with informational experts, and as a specific case we can analyze environments with informational services with low marginal costs (for example, producing an additional report on a car's history). An interesting difference arises when we look at the impact of search costs on the optimality of fixed-fee pricing, because search costs are an intrinsic feature of the environment and directly impact the consumer rather than the expert. Our work is also somewhat related to recent literature on bundling of information goods [4, 17]. Much of this literature focuses on deriving general conditions under which bundling of complementary or substitute goods makes sense. While packages of expert *services* are, of course, a different case, we are able to extend the intuitions from the theory of product bundling of Geng et al. [17] to the search domain. For example, in product bundling of information goods, if consumers' values for future goods do not decrease too quickly, bundling can be optimal for the seller. Similarly, in search, the marginal utility of an additional query to the searcher turns out to be an important factor in the expert's decision as to whether to offer fixed-fee or usage-based pricing, and whether to package multiple queries into indivisible groups that must be purchased together (such schemes are common: for example, as of September 2013, Carfax offered a five-report package for $49.99, a single report for $39.99, and an unlimited package for one month for $54.99).

There is also related work on two-sided markets, i.e., platforms that bring together different sides of the market. In this literature, the usual problem that is studied is of how to improve welfare. For example, the impact of having the platform charging only one of the two participating sides and cases where consumers are in effect paid to use the platform were studied; the low price on one side not only attracts elastic consumers but also, as a result, leads to higher prices or more participation on the other side [1, 37]. Our work can be viewed in a similar vein, especially in the context of the platform subsidizing searchers' use of expert services, although the intuitions behind our results are quite different and grounded in the frictions of search. Baye and Morgan [6] study how a monopoly online platform can successfully attract a share of local trade, focusing on homogeneous product markets. A similar model in which firms are selling differentiated products was studied by Galeotti and Moraga-Gonzalez [16]. The main difference between this line of work and ours is that it assumes that if a buyer picks the platform alternative, then she sees all seller listings with no additional fee, hence there are no search considerations involved. Furthermore, observations are not noisy and the expert option is therefore irrelevant.

The literature on market interventions such as subsidies is, of course, enormous, so we limit our discussion to interventions in markets where search plays an important role. The two-sided search literature has considered the impact of search frictions on labor markets [35, 14]. One classic regulatory intervention in these models is the introduction of a minimum wage, which can be shown to be welfare increasing in many contexts [21]. We are unaware of any work on subsidizing information brokers, which is our main concern.

3. The Model

Our model considers an agent or *searcher* facing an infinite stream of opportunities from which she needs to choose one. The value v of each opportunity is a priori unknown, however the searcher can receive, at cost c_s, a noisy signal s, correlated with the true value according to a known probability density function $f_v(v|s)$. In addition, the searcher may query and obtain from a third party (the *expert*), for a fee, the true value v of an opportunity, for which signal s was received. It is assumed that rejected opportunities cannot be recalled and that the searcher is acquainted with the (stationary) probability density function from which signals are drawn, denoted $f_s(x)$ ($x \in \mathbb{R}$). The searcher's goal is to maximize her total utility, defined as the value of the opportunity eventually picked minus the overall cost incurred during the search.

The assumptions above are standard in the literature. In particular the above model is the same as the one used by MacQueen [29] and used as a basis by many others [42, 26]. When signals are not noisy, then the model is identical to the standard one used in the one-sided search literature in economics [31, 28].

Unlike prior work, we assume that the expert is a self-interested agent that sets its service terms strategically. Further, we do not limit the expert to offering her services based on a fixed per-query fee, instead allowing non-linearly priced service packages. In this case, a searcher can purchase a packet for a cost c_e^k. The searcher can then use the expert's service k times along her search path with no additional cost. The package service offer generalizes two other common service pricing schemes: a la carte (pay-per-use) pricing (setting $k = 1$, charging c_e for each requested query) and fixed-fee (subscriptions to services) pricing (setting $k \to \infty$ and charging cost c_e when the expert is first queried and none for each additional query). When $k > 1$, the searcher does not necessarily need to make use of all the k queries; if she does use them all she can purchase more packages as required. The searcher's strategy is a mapping from the signal received and the number of queries remaining to the set of actions $\{\texttt{terminate}, \texttt{resume}, \texttt{query}\}$, which have the following interpretations: `terminate` means taking the current opportunity and ending search; `resume` means incurring a cost c_s and moving on to the next opportunity; `query` means engaging the services of the expert, paying a cost c_e^k if no package has been purchased yet or if all k queries from the last purchased package have been exhausted. We assume that the provider of expert services only pays a marginal cost d_e per query and that her reports are truthful.

These additional assumptions are intuitively reasonable. The use of service packages is widespread both in conventional markets (e.g., packages of vouchers) and in expert-mediated markets (e.g., Carfax's use of packages as discussed above). It is justifiable to assume that the expert only pays a marginal cost per query whenever she has already performed the "startup work" necessary, which is usually the case with information services (since fixed costs are constant, they do not affect the price setting decision at the margin). Lastly, truthful reporting can be assumed for reputational or regulatory reasons.

The first question that arises is how to characterize the optimal strategy for the searcher, given the package terms (k, c_e^k) set by the expert (or the cost c_e set in the pay-per-use and subscription models). A second question is how the expert sets her package terms k and c_e^k in order to maximize her (per-searcher) expected profit, denoted π_e. In this paper we consider a monopolist provider of expert services (e.g., part of the platform, or approved by the platform). The searcher's optimal strategy is directly influenced by the service terms, and therefore implicitly determines the expected number of times the services of the expert are required, and thus the expert's revenue. The problem can be thought of as a Stackelberg game where the expert is the first mover, and wants to maximize her profits with respect to the service terms she sets.

When using a la carte pricing, the expert's per-searcher profit π_e is a function of the expected number of queries purchased by the searcher, denoted η_{c_e}, the cost of the service c_e and her cost of producing the service d_e. Formally, $\pi_e = (c_e - d_e)\eta_{c_e}$. Similarly, with subscription pricing $\pi_e = (c_e - d_e\eta_{c_e}^{\infty})P_{\text{purchase}}$ (where P_{purchase} is the probability a subscription is purchased and $\eta_{c_e}^{\infty}$ is the expected number of queries made upon purchasing a subscription). When the expert offers her services in packages, we need to distinguish between the expected number of packages purchased by the searcher, denoted η_b, and the expected number of queries used by the searcher (accounting for the fact that not all queries in a package may end up being used, since the searcher may terminate search), denoted $\eta_{c_e^k}$ ($\eta_{c_e^k} \leq \eta_b k$). The expert's expected profit in this case is given by: $\pi_e = c_e^k \eta_b - \eta_{c_e^k} d_e$.

4. Optimal Policies

In this section, we derive optimal strategies for searchers and the resulting per-searcher profit for monopolist experts. We first characterize the searcher's optimal strategy for the case of complete information and then generalize it to noisy signals. From this we can derive the expert's expected profit as a function of the price she sets, enabling maximization of the expert's revenue.

4.1. The searcher's optimal strategy

One-Sided Search with Complete Information. The optimal strategy if signals are fully informative is reservation-value based, as discussed in Section 2. A searcher reviews opportunities sequentially (in random order) and terminates the search once a value greater than a reservation value x^* is revealed, where the reservation value x^* satisfies $c_s = \int_{y=x^*}^{\infty}(y - x^*)f_v(y)dy$ [31].

Intuitively, x^* is the value where the searcher is precisely indifferent: the expected marginal utility from continuing search and obtaining the value of the next opportunity exactly equals the cost of obtaining that additional value. The reservation property of the optimal strategy is due to the stationarity of the problem – since the searcher is not limited by the number of opportunities she can explore, resuming search places her at the same position as at the beginning of the search [31]. Consequently, a searcher that follows a reservation value strategy will never decide to accept an opportunity she has once rejected, and the optimal search strategy is the same whether or not recall is permitted. The expected gain to the searcher from following the optimal strategy in this case is x^*: the searcher is indifferent between accepting an opportunity valued at x^* and resuming search (which has expected value x^*). The expected number of search iterations is simply the inverse of the success probability, $\frac{1}{1-F_v(x^*)}$, since this becomes a Bernoulli sampling process, as opportunities arise independently at each iteration.

One-Sided Search with Noisy Signals. When the searcher receives a noisy signal rather than the actual value of an opportunity, there is no guarantee that the optimal strategy is reservation-value based as it is in the case where values obtained are certain. The stationarity of the problem, however, still holds, and an opportunity that has been rejected will never be recalled. In the absence of any restriction over $f_s(s|v)$, the optimal strategy can be thought of as a set S of signal-value intervals for which the searcher terminates search. The expected utility in this case, denoted $V(S)$, is given by:

$$V(S) = -c_s + V(S)\int_{s\notin S} f_s(s)\,ds + \int_{s\in S} f_s(s)E[v|s]\,ds \tag{1}$$

The searcher incurs cost c_s for receiving a new signal. If the signal is not in S (with probability $\int_{s\notin S} f_s(s)\,ds$) the search resumes, and the expected utility is $V(S)$. For any signal $s \in S$, the search terminates, and the searcher obtains $E[v|s]$ in expectation. The fact that the optimal strategy may not be reservation-value based in this case is because there may be no systematic correlation between the signal and the true value of the opportunity. However, in most real-life cases, such correlation does exist. In particular, a fairly weak and commonly used restriction on the conditional distribution of the true value given the signal goes a long way towards allowing us to recapture a simple space of optimal strategies. This is the restriction that higher signal values are "good news" in the sense that when $s_1 > s_2$, the conditional distribution of v given s_1 first-order stochastically dominates that of v given s_2 [44, 32, 29]. Formally:

Definition 1. Higher signals are good news (HSGN) assumption: If $s_1 > s_2$, and the support of the value distribution is $[a, b]$, then, $\forall y \in (a, b), F_v(y|s_1) < F_v(y|s_2)$, where $F_v(y|s)$ is the cumulative distribution function (cdf) of values given the signal.

We note that for noisy environments satisfying the HSGN assumption, if $s_1 > s_2$, then it follows that $E[v|s_1] > E[v|s_2]$. This enables us to prove the following theorem.

Theorem 1 *For any probability density function $f_v(v|s)$ satisfying the HSGN assumption: (a) the optimal search strategy is a reservation-value rule, where the reservation value, t^*, satisfies:*

$$c_s = \int_{s=t^*}^{\infty} \Big(E[v|s] - E[v|t^*]\Big) f_s(s)\, ds \tag{2}$$

(b) the expected utility of using the optimal search strategy satisfies: $V(t^*) = E[v|t^*]$.

Proof: This proof uses standard techniques from dynamic programming and search theory; we defer the details to the supplementary material. Note that the expected number of search iterations is $\frac{1}{1-F_s(t^*)}$, since this is again a Bernoulli sampling process. $\square$

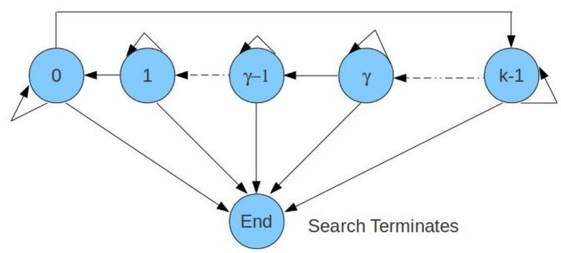

Figure 1: MDP representation of the searcher's problem when the expert offers packages of her services (state variable is the number of queries the searcher has remaining).

The Expert Option. The introduction of an expert extends the number of decision alternatives available to the searcher. Now, when receiving a noisy signal she can, in addition to rejecting or accepting, also query the expert in order to obtain the true value of an opportunity for an additional fee. Since the searcher cannot recall previous opportunities, her state depends only on the number of remaining pre-paid queries, denoted by γ. The system can thus be modeled as a Markov Decision Process with k search states ($\gamma = 0, 1, 2, \ldots, k-1$) and one termination state, as illustrated in Figure 1. Let V_γ denote the expected utility-to-go of following the optimal search strategy starting from state γ. The process begins when the searcher is in state $\gamma = 0$. In this state, upon receiving a signal s from the current opportunity, the searcher can either: (a) reject it and continue search, starting from the same state $\gamma = 0$; (b) accept it and terminate search; (c) query the expert to know the true value of the current opportunity, incurring a cost c_e^k, and, based on the value received, either accept it (terminating the search) or reject it and continue search from state $k-1$. When in state $\gamma > 0$, the searcher has the same options when receiving a signal s, except that when querying the expert the agent does not incur any cost, and if, based on the value received, she decides to continue the search, it is resumed from state $\gamma - 1$.

While terminating without querying the expert is a legitimate option in any states $\gamma > 0$, we can show that it is never preferred (see Lemma 4 in Appendix A). Intuitively, if the searcher believes that the opportunity is valuable based on the initial signal and there are

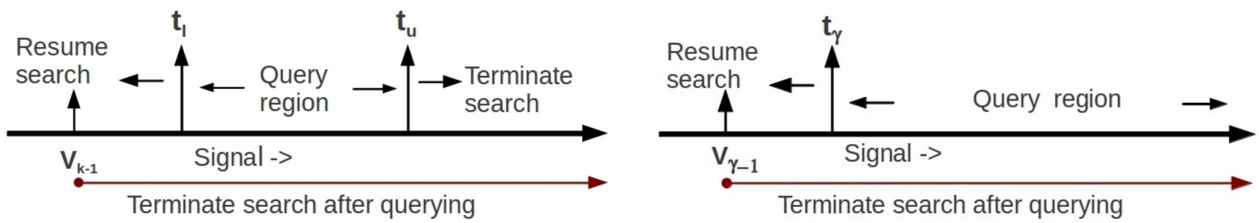

(a) For $\gamma = 0$ (b) For $\gamma > 0$

Figure 2: Characterization of the optimal strategy for noisy search with an expert offering packages.

still queries remaining from the purchased package, the queries would be worthless if the searcher terminates without using them. On the other hand, if the signals are very weak, the user prefers to resume search without exhausting a query. For any signal s received in state $\gamma > 0$, the expected utility if querying the expert, denoted by $M(s, V_{\gamma-1})$, is given by: $M(s, V_{\gamma-1}) = \int_{-\infty}^{\infty} \max(x, V_{\gamma-1}) f_v(x|s)\, dx$. Resuming search without querying the expert yields V_γ. Therefore, the optimal strategy is to query the expert if $M(s, V_{\gamma-1}) > V_\gamma$. The expected utility of using the optimal strategy, when starting from state $\gamma > 0$, is thus given by:

$$V_{\gamma>0} = -c_s + V_\gamma \int_{M(s,V_{\gamma-1})<V_\gamma} f_s(s)\, ds + \int_{M(s,V_{\gamma-1})>V_\gamma} f_s(s) M(s, V_{\gamma-1})\, ds \tag{3}$$

Similarly, when in state $\gamma = 0$, the expected utility when obtaining a signal s is: (a) $E[v|s]$ if terminating the search without querying; (b) V_0 if resuming the search without querying the expert; and (c) $M(s, V_{k-1}) - c_e^k$ on querying the expert. For any signal s, the choice which yields the maximum among the three should be made. Let $\Lambda = \max(E[v|s], V_0, M(s, V_{k-1}) - c_e^k)$ and let ζ_1, ζ_2 and ζ_3 define the sets of signal support such that Λ equals V_0, $E[v|s]$, and $M(s, V_{k-1}) - c_e^k$, respectively. The expected utility is then:

$$V_0 = -c_s + V_0 \int_{\zeta_1} f_s(s)\, ds + \int_{\zeta_2} E[v|s] f_s(s)\, ds + \int_{\zeta_3} f_s(s) \Big(M(s, V_{k-1}) - c_e^k \Big)\, ds \tag{4}$$

While the structure of the optimal strategy depends on $f_v(y|s)$, under the HSGN assumption the optimal strategy has a simple representation in the form of reservation values.

Theorem 2 *For $f_v(y|s)$ satisfying the HSGN assumption the optimal strategy for a searcher can be described as (see Figure 2):*
(1) a tuple (t_l, t_u, V_{k-1}), corresponding to state $\gamma = 0$, such that for any signal s: (a) the search should resume if $s \le t_l$; (b) the opportunity should be accepted if $s \ge t_u$; and (c) the expert should be queried if $t_l \le s \le t_u$ and the opportunity accepted (and search terminated)

if the value obtained from the expert is above the expected utility of resuming the search, V_{k-1}, otherwise search should resume; and
(2) a set of $(k-1)$ tuples $(t_\gamma, V_{\gamma-1})$ corresponding to states $\gamma \in 1, 2 \ldots, k-1$ such that: (a) the search should resume if $s \leq t_\gamma$; and (b) the expert should be queried if $s \geq t_\gamma$ and the opportunity accepted if the value obtained from the expert is above the expected utility of resuming the search, $V_{\gamma-1}$, otherwise search should resume.
The values V_0, t_l, t_u, t_γ and V_γ (for $1 \leq \gamma < k$) can be calculated from solving the set of Equations:

$$V_{\gamma>0} = V_\gamma F_s(t_\gamma) + \int_{s=t_\gamma}^{\infty} f_s(s)\Big(\int_{y=V_{\gamma-1}}^{\infty} y f_v(y|s)\,dy + V_{\gamma-1}F_v(V_{\gamma-1}|s)\Big)\,ds - c_s \tag{5}$$

$$V_0 = V_0 F_s(t_l) + \int_{s=t_u}^{\infty} f_s(s)E[v|s]\,ds + \int_{s=t_l}^{t_u} f_s(s)\Big(V_{k-1}F_v(V_{k-1}|s) + \int_{y=V_{k-1}}^{\infty} y f_v(y|s)\,dy - c_e^k\Big)\,ds - c_s \tag{6}$$

$$c_e^k = -V_0 + V_{k-1} + \int_{y=V_{k-1}}^{\infty} (y - V_{k-1}) f_v(y|t_l)\,dy \tag{7}$$

$$c_e^k = \int_{y=-\infty}^{V_{k-1}} (V_{k-1} - y) f_v(y|t_u)\,dy \tag{8}$$

$$V_\gamma = V_{\gamma-1}F_v(V_{\gamma-1}|t_\gamma) + \int_{y=V_{\gamma-1}}^{\infty} y f_v(y|t_\gamma)\,dy \tag{9}$$

Proof: See Appendix A. □

Equations 5 and 6 are the appropriate modifications of Equations 3 and 4 for the HSGN case (when the searcher's optimal decisions are made based on the thresholds specified in the theorem). Equations 7-9, that characterize the searcher's optimal strategy, can also be derived from the searcher's indifference conditions at signals t_l, t_u, and t_γ respectively. For example, t_l is the signal at which a searcher is indifferent between either resuming the search or querying the expert, i.e., $V_0 = \int_{y=V_{k-1}}^{\infty} y f_v(y|t_l)\,dy + V_{k-1}F_v(V_{k-1}|t_l) - c_e^k$, which transforms into Equation 7; alternatively, t_l can also be interpreted as a point where cost of purchasing the expert's service is equal to the expected increase in utility from consulting the expert when the searcher would otherwise reject and resume search.

There is also a reasonable degenerate case where $t_l = t_u (= t)$, i.e., a single threshold serves as the optimal strategy, as in the case with no expert. This happens when the cost of querying is so high that it never makes sense to engage the expert's services. In this case, a solution to the set of equations specified in Theorem 2 that satisfies $t_l < t_u$ does not exist.

Instead, we obtain a single reservation value to be used by the searcher using Theorem 1.[2]

Finally, we note that the above analysis is also applicable when the expert supplies a noisy signal rather than a fully informative one, as long as the expert's signal satisfies the HSGN assumption. This can be shown using a transformation due to MacQueen [29] (see details in the supplementary material).

A La Carte Pricing. As discussed above, the case where the searcher pays per-use is a specific case of package pricing, where at each time a "package" of size $k = 1$ is purchased for a cost c_e. Here, there is no transition between the different γ states. Instead, whenever choosing to query the expert and then deciding to resume search, the searcher always returns to the same state. The solution in this case complies with part (a) of Theorem 2, except that it uses V_0 rather than V_{k-1} as the threshold. The values t_l, t_u and V_0, for the HSGN case, can be calculated from solving the set of Equations 10-12:

$$V_0 = -c_s + V_0 F_s(t_l) + \int_{s=t_u}^{\infty} f_s(s) E[v|s]\, ds + \int_{s=t_l}^{t_u} f_s(s) \left(V_0 F_v(V_0|s) + \int_{y=V_0}^{\infty} y f_v(y|s)\, dy - c_e \right) ds \tag{10}$$

$$c_e = \int_{y=V_0}^{\infty} (y - V_0) f_v(y|t_l)\, dy \tag{11}$$

$$c_e = \int_{y=-\infty}^{V_0} (V_0 - y) f_v(y|t_u)\, dy \tag{12}$$

This latter result aligns with MacQueen's result for a similar a la carte model of noisy exploration where a more accurate signal can be obtained for a fixed fee, where both distributions of signals satisfy the HSGN assumption [29].

A Subscription Model: Fixed Fee Pricing. A specific case of importance is when the searcher can use as many queries as she would like upon payment of a fee to the expert (i.e., purchasing a package of size $k \to \infty$). In this case, the model reduces to two states. The searcher starts in state $\gamma = 0$ and continues in this state until she either terminates search or transitions to state $\gamma = \infty$. In the latter case the searcher can keep querying the expert for any reasonable opportunity until she finally finds a sufficiently good opportunity and terminates search (the existence of a search cost ensures that this querying process does not go on forever). Being in state $\gamma = \infty$ is equivalent to being in the world of perfect signals. The optimal reservation

[2] Another technical point worth noting is that Equation 8 is for the case when the support on signal s is unbounded. When there is an upper limit on s, i.e., $s \leq m$ for some m, once t_u reaches m (we never buy without querying), Equation 8 does not hold. Now the searcher rejects if the signal is below t_l or queries if it is above.

value when in state $\gamma = \infty$ can thus be found by solving the equation (e.g., [31]):

$$V_\infty = -c_s + V_\infty F_v(V_\infty) + \int_{y=V_\infty}^{\infty} y f_v(y)\, dy \tag{13}$$

For state $\gamma = 0$, we can use appropriate modifications of Equations 6-8, replacing V_{k-1} with V_∞ (realizing that the searcher transitions to state $\gamma = \infty$). The optimal strategy can be found by solving this set of four equations.

4.2. The Expert's Perspective

A monopolist expert moves first in the implied Stackelberg game by setting the size and price (k, c_e^k) of the package (or the price c_e in the other two pricing schemes). The searcher responds by following her optimal search strategy described above. Therefore, the expert can solve for the searcher's behavior, given knowledge of the size and price (k, c_e^k) of the package and the signal and value distributions. The expert's profit as a function of the size and price (k, c_e^k) of the package that she sets is given by $\pi_e = c_e^k \eta_b - \eta_{c_e^k} d_e$ (see Section 3). Given a way of computing η_b and $\eta_{c_e^k}$, the profit-maximizing package (or the profit-maximizing query/subscription price) characteristics can be determined. For expository purposes, we present here the calculation for the a la carte case and supply the details for the calculation with the general non-linear pricing model in the supplementary material. For the a la carte case, the search strategy (t_l, t_u, V) defines how many times the expert's services are used. Let A, B, C, D be the probabilities, respectively, that (1) the expert is queried and search then resumed, (2) the expert is not queried and search resumed, (3) the expert is not queried and search terminated, and (4) the expert is queried and search then terminated. These probabilities can be calculated using (note that $A + B + C + D = 1$):

$$A = \Pr(t_l \le s \le t_u \wedge v < V_0) = \int_{y=t_l}^{y=t_u} f_s(y) \int_{x=-\infty}^{x=V_0} f_v(x|y)\, dx\, dy \tag{14}$$

$$B = \Pr(s < t_l) = F_s(t_l) \tag{15}$$

$$C = \Pr(s > t_u) = 1 - F_s(t_u) \tag{16}$$

$$D = \Pr(t_l \le s \le t_u \wedge v \ge V_0) = \int_{y=t_l}^{y=t_u} f_s(y) \int_{x=V_0}^{x=\infty} f_v(x|y)\, dx\, dy \tag{17}$$

We use η_s to denote the expected number of searches executed by the searcher (i.e., the expected number of opportunities she runs into throughout the search). Since on each search iteration along the search the searcher queries the expert if the signal is between t_l and t_u, the expected overall number of queries made is given by:

$$\eta_{c_e} = \Pr(\text{Searcher queries an expert}) \times \eta_s$$

Now, $\eta_s = \frac{1}{C+D}$ (termination is a Bernoulli trial with probability $C+D$) and the probability that the searcher queries the expert at any search instance is $A+D$. Therefore: $\eta_{c_e} = \frac{A+D}{C+D}$. The expert can thus maximize her expected profit, $\pi_e = \mathbb{E}(\text{Profit}) = (c_e - d_e)\eta_{c_e}$, with respect to c_e (η_{c_e} decreases as c_e increases; see supplementary material for proof).

5. Strategic Implications For Experts

We can use the characterization of optimal search strategies to solve for the expert's optimal strategy and derive implications for how experts should price their services. Equilibrium in expert-mediated search derives from a complex set of dynamics. Many parameters affect the equilibrium, including the distribution of values, the correlation between signals and values, search frictions, the cost of querying the expert, and (in the case of non-linear pricing) package size. Uncovering phenomenological properties of the model is therefore difficult and restricted using a static analysis. Instead, we turn to an illustrative model that uses a particular, plausible distribution of signals and values. [3]

Suppose the signal is an upper bound on the true value. So the signal could be thought of as the searcher's optimistic estimate upon observing the opportunity (e.g., sellers and dealers offering cars for sale usually make cosmetic improvements to the cars in question, and proceed to advertise them in the most appealing manner possible, hiding defects using temporary fixes; mortgage lenders may advertise their most appealing features, such as a low introductory rate, while keeping troublesome terms and conditions hidden). Specifically, we assume signals s are uniformly distributed on $[0,1]$, and the conditional density of true values is linear on $[0,s]$. Thus

$$f_s(s) = \begin{cases} 1 & \text{for } 0 \le s \le 1 \\ 0 & \text{otherwise} \end{cases} \qquad f_v(y|s) = \begin{cases} \frac{2y}{s^2} & \text{for } 0 \le y \le s \\ 0 & \text{otherwise} \end{cases}$$

5.1. Searcher's Optimal Strategies and Expert's Optimal Pricing

We can use these distributions to solve the system of equations described in Section 4. Figure 3(a) illustrates how the reservation values t_l and t_u change as a function of c_s for $c_e = 0.05$ when using a la carte pricing. The vertical axis is the interval of signals. As can be

[3] We note here that, while particulars may vary, the form of our results is robust when we add a discount factor into the utility functions for agents. This is described in detail in the supplementary material.

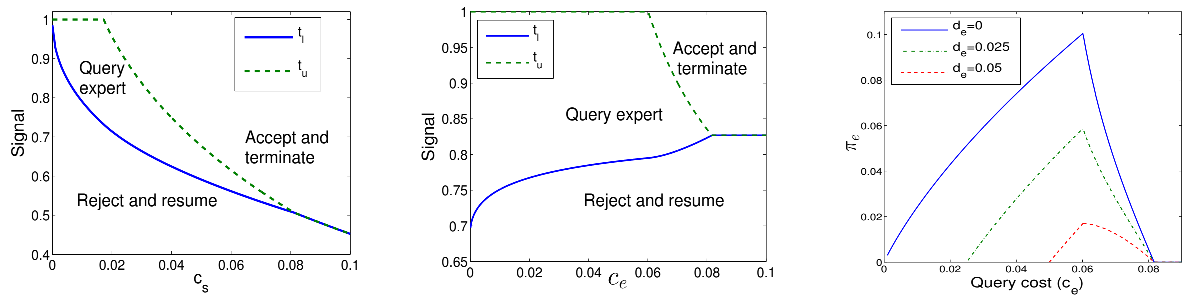

(a) (b) (c)

Figure 3: (a) Effect of c_s on the signal thresholds (t_l,t_u) for $c_e = 0.05$. (b) Effect of c_e on the signal thresholds (t_l,t_u) for $c_s = 0.01$. (c) Expert's profit as a function of c_e and d_e for $c_s = 0.01$. The model uses the a la carte pricing.

seen from the graph, in this specific case for very small search costs (c_s), the searcher never terminates search without querying the expert.[4] Due to the low search cost the searcher should only query the expert when she receives a high signal, because the cost of finding opportunities with high signals is relatively low. As the search cost c_s increases, there is some behavior that is not immediately intuitive. The reservation values t_l and t_u become closer to each other until coinciding at $c_s = 0.08$, at which point the expert is never queried anymore. The reason for this is that the overall utility of continuing search goes down significantly as c_s increases, therefore the cost of querying the expert becomes a more significant fraction of the total cost, making it comparatively less desirable.

A similar pattern holds when using packages of more than a single query: The dependency of t_l and t_u in c_s and c_e is similar to the one illustrated in Figure 3, and similarly, t_γ decreases as c_s increases and increases as c_e increases (see graphs in the supplementary material).

It is now easy to calculate numerically the value of c_e^* that maximizes π_e, trading off a decrease in the number of queries η_{c_e} in exchange for an increase in the revenue per query c_e. The expected number of queries can be found by substituting in the signal and noise distributions into Equations 14-17.

Figure 3(c) shows examples of the expected profit of the expert as a function of the expert's fee, c_e, for different values of d_e, the marginal cost to the expert of producing an extra report (fixing $c_s = 0.01$). Perhaps most interestingly, the shapes of the curves are the same and peak in the same region for significantly different values of d_e. It is also worth noting that the additional marginal costs are typically not being simply passed on

[4]When search costs are 0 the problem is ill-defined. The first point on the graph shows an extremely low, but non-zero search cost. In this case $t_u = 1$ and t_l is almost 1, but not exactly, and the expert is again always queried before search terminates.

to consumers – the expert suffers when the marginal cost of producing an additional report increases despite her monopoly position. This is explained, to some extent, by the decrease in searchers' tendency to use the expert as the cost of service increases (as illustrated in Figure 3(b)). An important implication of this phenomenon is that there may be welfare gains from additional efficiency in producing reports.

5.2. Heterogenous Customer Types

The above techniques can also be applied to solve for optimal expert strategies when there are multiple types of customers in the market. Assume there is a continuum of agent types distributed according to the pdf $f_w(w)$. Agents of different types differ in the values they assign to any subset of the parameters c_s, $f_s(s)$ and $f_v(y|s)$, such that every agent of type w is characterized by c_s^w, $f_s^w(s)$ and $f_v^w(y|s)$. Since each agent uses its best-response strategy based on the parameters c_s^w, $f_s^w(s)$,$f_v^w(y|s)$, k and c_e^k, it is not affected by the search strategy set by the other agents. Therefore, given (k, c_e^k), the strategy of agents of each agent type w can be calculated using Theorem 2, and consequently, the expected profit made by the expert if the agent type w is the only agent in the market, π_e^w. The expected profit of the expert in this case is given by $\pi_e = \int_w \pi_e^w f_w(w)dw$.

We illustrate this for the a la carte model through a case where the population consists of a proportion θ of "low" type agents, characterized by $f_v(y|s)(= 2y/s^2)$, as above. The remaining proportion $1-\theta$ of the population consists of "high" type agents, characterized by a different conditional distribution for valuations given signals, $f_v(y|s) = 3y^2/s^3(0 \le y \le s)$. Obviously the distribution of valuations of high-type agents stochastically dominates that of low-type agents.

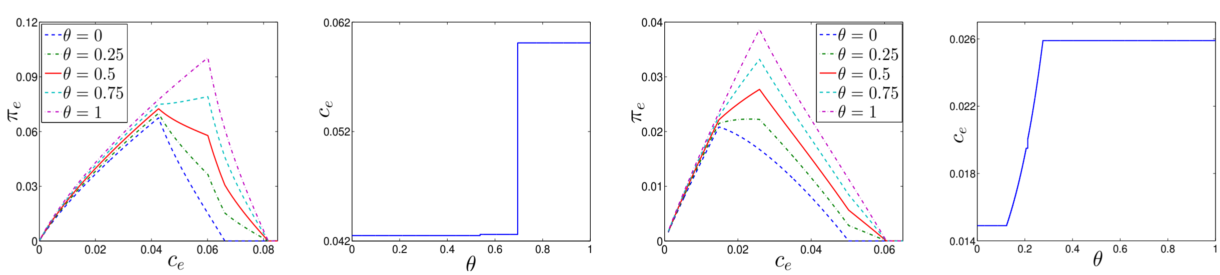

For $c_s = 0.01$ For $c_s = 0.05$

Figure 4: The expert's profit as a function of c_e for different proportions of the low type (θ) in the population, and the expert's profit-maximizing c_e as a function of θ, for two different costs of search c_s.

Figure 4 shows the profit of an expert for different values of θ vs. the query cost c_e, and the expert's profit-maximizing c_e as a function of θ, for two different costs of search c_s. In the

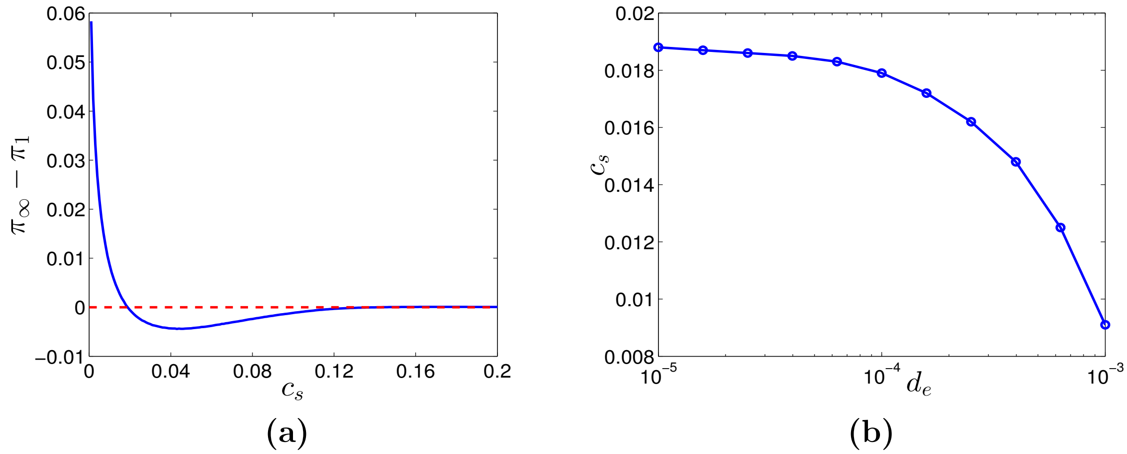

Figure 5: **(a)** Difference in expected profit of the expert for fixed-fee and per-use pricing for digital services ($d_e = 0$) vs. search cost. Fixed-fee pricing is preferred when search costs are low. **(b)** The search cost c_s at which the expert is indifferent between fixed-fee and per-use pricing, as a function of d_e, the marginal cost of producing an expert report. The expert prefers fixed-fee pricing if the search cost is lower than the one shown on the graph, and prefers per-use pricing if its higher.

expected profit graphs, each data point is a convex combination of the uppermost ($\theta = 1$) and lowest ($\theta = 0$) curves (weighted according to θ). We observe that the optimal query cost is a non-decreasing function of θ, the fraction of low type agents in the population. This is natural: low-type agents have more to gain from expert services since they receive less accurate signals. The expert charges a lower price for her services when the population is composed only of high-types as opposed to when it is composed only of low-types. Interestingly, as θ increases, the change in the expert's strategy from the strategy that is optimal for an entire population of high-types (a higher-volume, lower-margin strategy) to a strategy that is optimal for an entire population of low-types (a lower-volume, higher-margin strategy) is rather sharp (a phase-transition-like pattern). This is because of the asymmetric behavior of the expert's profit around the optimal price: it increases to the peak slowly and decreases rapidly once past the peak. This effect is more apparent for a lower value of c_s, which would typically characterize electronic markets.

5.3. Non-linear Pricing

We now turn our attention to another aspect of the expert's strategy, whether she should engage in non-linear pricing. First, we study conditions under which offering a fixed fee, unlimited use scheme makes sense for the expert. Figure 5(a) shows the difference in the expected profit of the expert between offering a fixed-fee plan as opposed to per-use pricing as a function of search cost c_s for the digital services case ($d_e = 0$). We see that fixed-fee pricing can yield significantly higher profits for the expert when search costs are low, but that per-use pricing is preferred when search costs are high (when they become too high consumers no longer use the expert and the profits of both expert strategies go to 0). In a world of high search costs, users do not expect to keep searching for more than a few opportunities, so they are unlikely to be willing to purchase a fixed-fee plan at a price that would make sense for the expert.

Figure 5(b) examines the threshold value of c_s for different marginal costs of producing expert reports. At search costs lower than the threshold shown in the graph, the expert expects higher profits from fixed-fee pricing, and at higher search costs, she expects higher profits from per-use pricing. We can see that the threshold decreases rapidly as a function of d_e (the horizontal axis is on a log scale), indicating that as the marginal cost of expert services increases, per-use pricing quickly becomes preferred by the expert.

Both these observations, that increasing either the marginal cost of producing expert services and/or the cost incurred by the buyer in searching lead to unit size or smaller package sizes being preferred, correspond well to the real world and to recent observations in the literature. Unlimited subscription models are usually available in online services, where vanishing marginal costs and low search costs dominate. Sundararajan [39] provides a model that explains why subscription services are often preferred in the world of information goods, focusing on the vanishing marginal costs mentioned above. He includes transaction costs for administering pricing as a key part of his model (and these costs are obviously more relevant when the costs of producing the goods are becoming smaller and smaller). Our result is a parallel one in the realm of sequential search rather than direct pricing of products: we show the effects of search costs, which are borne by the consumer, and how they affect the optimality of subscription pricing models for information services (rather than information goods).

In a different framework, that of bundling complementary information goods, Geng et al. [17], show that if consumers' values for future goods do not decrease too quickly, then bundling is (approximately) optimal for a monopolist, whereas if the values for future goods do decrease quickly bundling may not be optimal. In our framework, in the presence of high search costs, consumers are unlikely to want to use too many queries to the experts, so the marginal utility of an extra query is rapidly decreasing compared to the case where there are lower search costs.

Finite-size Packages. We can also solve numerically for finite k. Figure 6 depicts the expert's expected profit as a function of the package size for $c_s = 0.01$ and different d_e values.[5]

[5] For each package size, the appropriate optimal c_e^k is used.

One interesting observation is that when the marginal cost of producing an expert report is significantly higher than zero, the optimal number of queries to sell as a package tends to be small. The optimal package size decreases as the service production cost d_e increases. Surprisingly, the expert's increased profit from non-linearly priced packages of greater size does not come at the cost of the searcher – searcher utility also increases, so packaging is pareto-improving. The search overall is becoming more efficient, and the expert and the buyer can split the additional gain. When the marginal cost is zero (as in the "digital services" case: an extra Carfax report can be produced essentially for free), the expert continues to achieve high profits with very high k, corresponding with the results above for when fixed-fee pricing is preferred.

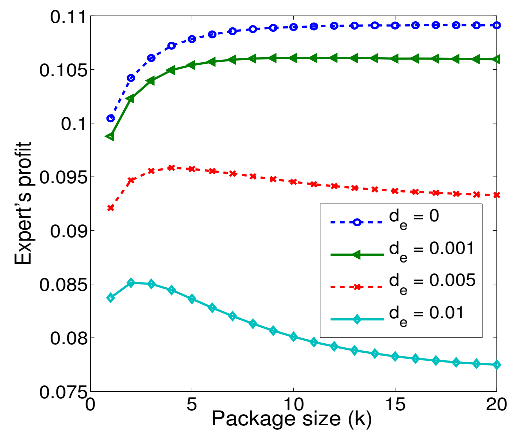

Figure 6: Expert profit as a function of the number of expert queries offered as a package for $c_s = 0.01$.

6. Platform Design

Consider the design of a large scale Internet website like `AutoTrader.com`. The listings for cars that users see are signals, and users may be unsure of a car's true value. `AutoTrader` (the platform) can partner with a provider of reports like Carfax (the expert), to make it easy for users to look up a car's value. As another example, an online mortgage broker may choose to provide, as a premium product, tools that allow customers to learn more about the good and bad features of an offered mortgage. In this case the platform itself would be serving the role of the expert. The platform wishes to attract customers by providing a high value shopping experience. The expert wishes to maximize its profits. Since the platform and the expert both have significant power, they are likely to come up with different relationship models. In this section, we focus on the strategic implications of the model of expert-mediated search developed in this paper for the platform that connects consumers and experts.

We focus here on one aspect of platform design, where the platform intervenes in the market by subsidizing expert queries. The platform can make payments to the expert in exchange for the expert providing services to users at a lower cost. A typical problem with subsidization is that it often decreases social welfare because the true cost of whatever is being subsidized is hidden from the consumer, leading to overconsumption of the resource. In this instance, however, the natural existence of monopolies in expert services, combined with the importance of search frictions, make it quite possible that subsidies will in fact increase social welfare.

We first define social welfare in our context and then describe the effects of subsidization on social welfare. For simplicity of exposition, and because it captures all the relevant intuition, we focus this discussion on the a la carte pricing model. It is also worth noting that this discussion is equally relevant to both private platforms (e.g., `AutoTrader.com`) as well as to regulators of large markets (e.g., government agencies).

6.1. Subsidization Model

Suppose a monopolist provider of expert services maximizes her profits by setting the query cost to c_e^*, yielding an expected profit $\pi_e = (c_e^* - d_e)\eta_{c_e^*}$. The platform can step in and negotiate a reduction of the fee c_e charged by the expert to a value c_e', for the benefit of the searchers. In return for the expert's agreement, the platform can offer a per-consumer payment β to the expert, which fully compensates the expert for the decreased revenue, leaving her total profit unchanged. Since $c_e^{'} < c_e^*$, $\eta_{c_e'} > \eta_{c_e^*}$ (the consumer queries more often because she has to pay less). The compensation for a requested decrease in the expert's fee from c_e^* to $c_e^{'}$ is thus $\beta = (c_e^* - d_e)\eta_{c_e^*} - (c_e^{'} - d_e)\eta_{c_e'}$. The overall welfare per agent in this case increases by $V_{c_e'} - V_{c_e^*}$, where $V_{c_e'}$ and $V_{c_e^*}$ are the expected utilities of searchers according to Equation 10-12, when the expert uses a fee $c_e^{'}$ and c_e^* respectively, at a cost β to the platform.

The social welfare is a function of utilities of all the parties involved. While this function can have any form, here we study utilitarian or additive social welfare, which is appropriate especially in the context of possible transfers from the platform that connects experts and searchers. We consider two representative agents, the searcher and the expert, and note that this generalizes to multiple searchers (each search process is independent, and social welfare scales up linearly in the number of searchers). Then: $W = V_{c_e^*} + \pi_e$. When the platform enters the picture by subsidizing expert queries, the social welfare must also take the subsidy into account. Since the expert is fully compensated for her loss due to the decrease in her fee, the change in the overall social welfare is $V_{c_e'} - V_{c_e^*} - \beta$. Under the new pricing scheme $c_e^{'}$, and given the subsidy β, the social welfare is given by $W^{'} = V_{c_e'} + \pi_e - \beta$.

6.2. Welfare-Maximizing Subsidies

The first important question that arises is what level the platform should set the subsidy at in order to maximize social welfare. Here we prove that, in the model of subsidization described above, social welfare is maximized at the point where the searcher pays exactly d_e per query, thus fully internalizing the cost to the expert of producing the service. If the searcher had to pay less, it would lead to inefficient overconsumption of expert services, whereas if she had to pay more, the expected decline in the utility she receives from participating in the search process would outweigh the savings to the platform from having to pay less subsidy.

Theorem 3 *Suppose the platform pays the expert a flat subsidy β (per-customer) in exchange for the expert reducing her per-query fee from c_e^* to $c_e^{'}$ such that $\beta = (c_e^* - d_e)\eta_{c_e^*} - (c_e^{'} - d_e)\eta_{c_e^{'}}$ then social welfare is maximized when β is set such that $c_e^{'} = d_e$.*

Proof: The searcher acts to maximize her utility. For convenience, we mildly abuse notation and assume that $\mathbb{E}[\omega()]$, $\eta_s()$ and $\eta_{c_e}()$ are functions defined by the searcher's optimal behavior from solving the Bellman equation. Here $\mathbb{E}[\omega()]$ is the expected value of the final opportunity that is taken (which is different from the expected utility of the searcher since it does not factor in costs), $\eta_s()$ is the expected number of opportunities examined by the searcher, analogous to η_s, and $\eta_{c_e}()$, analogous to η_{c_e}, the expected number of times the expert is queried. Then

$$W = \mathbb{E}[\omega(c_e)] - \eta_s(c_e)c_s - \eta_{c_e}c_e + \eta_{c_e}(c_e - d_e) = \mathbb{E}[\omega(c_e)] - \eta_s(c_e)c_s - \eta_{c_e}(c_e)d_e \qquad (18)$$

We claim that social welfare is maximized when $c_e = d_e$. We show this by appealing to the optimal solution of a single searcher's decision problem. The functions $\mathbb{E}[\omega()]$, $\eta_s()$ and $\eta_{c_e}()$ in this case are such that the expression in Equation 19 is maximized.

$$\mathbb{E}[\omega(d_e)] - \eta_s(d_e)c_s - \eta_{c_e}(d_e)d_e \qquad (19)$$

Suppose there exists some $c_{e'} \neq d_e$ that maximizes the social utility (Equation 18). Let us now define an alternative strategy for the single searcher above. The searcher pretends that the cost of querying the expert is $c_{e'}$ and follows the strategy given by that belief. Then the expected utility of the searcher can be found by assuming the searcher pays cost $c_{e'}$ but receives a "kickback" of $c_{e'} - d_e$ every time. The expected utility of the searcher is then $(\mathbb{E}[\omega(c_{e'})] - \eta_s(c_{e'})c_s - \eta_{c_e}(c_{e'})c_{e'}) + \eta_{c_e}(c_{e'})(c_{e'} - d_e)$. By the Bellman Optimality Principle, this cannot be greater than Equation 19. Therefore, such a $c_{e'} \neq d_e$ does not exist. In addition, note that this is actually just the right hand side of Equation 18. Therefore, it must be the case that Equation 18 is maximized at $c_e = d_e$. $\square$

An immediate implication of this theorem is for the "digital services" case, where $d_e = 0$ (producing an extra Carfax report, for example, typically has zero marginal cost). In this case, there is no societal cost to higher utilization of the expert's services, so subsidy is welfare improving right up to making the service free. These are the cases where it could make sense for the platform to take over offering the service itself (for example, the online mortgage broker providing mortgage comparison and evaluation tools), and making it free, potentially leveraging the increased welfare of consumers by attracting more consumers to their market, or by increasing fees for the use of platform.

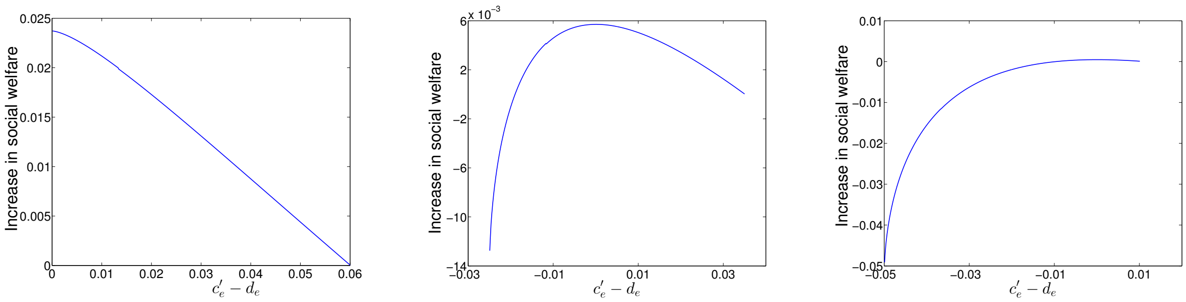

for $d_e = 0$ for $d_e = 0.025$ for $d_e = 0.05$

Figure 7: Increase in social welfare vs. the difference between the subsidized query cost ($c'_e < c^*_e$) due to subsidy and the marginal cost d_e. When $d_e = 0$, it is best for the platform to make the service available for free, but when there is some marginal cost involved, then the increase in social welfare is concave, peaking at zero, i.e., when the subsidized price is equal to the marginal cost.

	Without Subsidy						With Subsidy						
d_e	c_e	$\eta_{c_e} c_e$	$c_s \eta_s$	$\mathbb{E}(\omega)$	V	W	β	c_e	$\eta_{c_e} c_e$	$c_s \eta_s$	$\mathbb{E}(\omega)$	V	W
0	0.06	0.10	0.08	0.7472	0.5653	0.6656	0.1003	0	0	0.1036	0.7928	0.6893	0.6893
0.025	0.06	0.10	0.08	0.7472	0.5653	0.6238	0.0585	0.025	0.05	0.09	0.7712	0.6295	0.6295
0.05	0.061	0.09	0.08	0.7365	0.5637	0.5806	0.0169	0.05	0.09	0.08	0.7536	0.5824	0.5824

Table 1: The different components of social welfare with and without subsidy for $c_s = 0.01$. $\mathbb{E}(\omega)$ is the expected value of the opportunity eventually picked. Initially the decrease in query cost contributes more to the increase in social welfare, but as d_e increases, this contribution becomes less significant. Note that the first two rows in the case without subsidy are similar because the profit-maximizing c_e is the same and the searcher's cost depends only on the value of the selected c_e, not d_e.

6.3. Numerical Simulations

Simulations using the example distributions in Section 5 serve to both confirm our theoretical results and provide insight into where the improvements come from. Figure 7 shows the improvement in social welfare as a function of the difference between the subsidized query cost c'_e and the marginal cost d_e for various d_e values (where $c_s = 0.01$). From the graphs, we do indeed find that subsidization can lead to substantial increases in social welfare, even when there is a significant marginal cost of producing an expert report. While this could be from a reduction in search and query costs or an increase in the expected value of the opportunity finally taken, the data in Table 1 indicates that the latter explanation is the dominant factor in this case. Also, in accordance with Theorem 3, social welfare is maximized at the point where the searcher pays exactly d_e per query, thus fully internalizing the cost to the expert of producing the extra report.

7. Discussion

There are many online platforms for which our model is relevant. For example, when using `Yet2.com`, a global technology online marketplace, "buyers" can search the inventions listed by industrial firms, entrepreneurial ventures, research universities and individual inventors, and use experts (e.g., professionals with expertise in areas that complement the buyer's main business) to obtain a better valuation of any listing. There are also websites like `alibaba.com`, `made-in-china.com` and `gobizkorea.com` that serve as platforms for connecting buyers and sellers. The products listed come from different sellers; the quality of the product and the reputation of the seller can be unclear, leaving room for reputation systems or specialists to serve as experts that are capable of determining their genuine value. However, we note that our model is general and can apply to both electronic and offline markets. Two parameters, the marginal cost of production of expert services and the search cost, play key roles in understanding the differences between online and offline markets. Both of these costs tend to be lower in online markets. We find that experts may prefer fixed-fee pricing over usage-based pricing in electronic markets. This result supports that of Sundararajan [39], who finds that fixed-fee pricing may be preferred in direct selling of digital or information goods. His model is not in the setting of sequential search, and is relevant to producers of information/digital goods rather than information services; further, his model and results focus on the cost to the producer of administering the pricing scheme, while our focus is on the cost to the consumer of continuing search.

We also analyze the issues faced by platforms that bring searchers and experts together. In many search domains, the expert has a special relationship with the platform (e.g., an auto-trading website contracts with Carfax) or may even be embedded within the platform (e.g., an online mortgage broker could develop and offer a tool for understanding terms of different mortgages better). We show that the platform can create surplus by subsidizing expert queries to the point that searchers pay exactly the cost faced by the expert. The platform can capture some of the surplus through commissions or advertising, or through increased market share if it provides a better experience to consumers. In fact, if the platform can quantify its indirect benefit from the increase in the users social welfare, it can use the analysis given in this paper to determine the optimal subsidy, considering the tradeoff between the payment to the expert and the additional benefit deriving from the resulting increase in users' welfare. The same holds for a platform that is also offering the expert-like premium service — the tradeoff between the expected loss of direct income (i.e., query fees) and the expected gain in indirect benefits dictates the pricing of the service. In the important case of digital services, where the marginal cost of producing an extra unit of services is zero, our analysis suggests that expert services should be free to consumers, and therefore the platform may benefit particularly from a special relationship with, or ownership of, the expert.

8. Conclusions and Future Work

This paper studies markets with informational intermediaries in the context of sequential search theory. We investigate the interplay between costly sequential search and costly information acquisition, and are able to gain new insights into how a monopolist information broker (an *expert*) should price her services, and into how the platform that connects experts and consumers can maximize welfare. In doing so, we also characterize the optimal strategy for searchers to follow; this state-dependent "double reservation value"/ "single reservation value" strategy is a novel result for the type of information signals (non-additive) we consider here.

Extending this work from the monopolistic setting to understand the dynamics of competition between experts (who could, for example, compete not just in price but also in quality of information provided, as in the third-party certification literature [15]) is an important future direction. Our analysis still holds if the monopolist expert provides a noisy signal of true value (and the expected noise level could be a measure of quality), as long as the HSGN assumption holds, as discussed in the supplementary material; however, the extension to oligopolistic markets is non-trivial.

Another worthwhile extension would be to incorporate sellers' responses to the market dynamics as part of the equilibrium analysis, instead of modeling opportunities as exogenous. Finally, in some cases the platform may be in a position to *charge* experts for the privilege of connecting them with buyers; this poses some intriguing questions that would entail analysis of the dynamics of multi-platform/multi-expert competition. The model presented in this paper should provide a good foundation for these extensions.

Acknowledgements

This work was supported by a US-Israel Binational Science Foundation grant (2008404), a US National Science Foundation CAREER award (IIS-0952918/1303350), and an Israel Science Foundation grant (1083/13).

References

[1] Armstrong, M., 2006. Competition in two-sided markets. RAND Journal of Economics 37 (3), 668–691.

[2] Atakan, A. E., 2006. Assortative matching with explicit search costs. Econometrica 74 (3), 667–680.

[3] Bakos, J., 1997. Reducing buyer search costs: Implications for electronic marketplaces. Management Science 43 (12), 1676–1692.

[4] Bakos, Y., Brynjolfsson, E., 1999. Bundling information goods: Pricing, profits, and efficiency. Management Science 45 (12), 1613–1630.

[5] Baryla, E. A., Zumpano, L. V., Elder, H. W., 2000. An investigation of buyer search in the residential real estate market under different market conditions. Journal of Real Estate Research 20 (1), 75–91.

[6] Baye, M. R., Morgan, J., 2001. Information gatekeepers on the internet and the competitiveness of homogeneous product markets. American Economic Review 91 (3), 454–474.

[7] Bental, B., Peled, D., 2002. Quantitative growth effects of subsidies in a search theoretic R&D model. Journal of Evolutionary Economics 12 (4), 397–423.

[8] Bloch, F., Ryder, H., 2000. Two-sided search, marriages, and matchmakers. International Economic Review 41 (1), 93–115.

[9] Brynjolfsson, E., Smith, M., 2000. Frictionless commerce? a comparison of Internet and conventional retailers. Management Science 46 (4), 563–585.

[10] Burdett, K., Wright, R., 1998. Two-sided search with nontransferable utility. Review of Economic Dynamics 1, 220–245.

[11] Cason, T. N., Friedman, D., 2003. Buyer search and price dispersion: a laboratory study. Journal of Economic Theory 112 (2), 232–260.

[12] Chhabra, M., Das, S., Sarne, D., 2011. Expert-mediated search. In: Proceedings of the International Joint Conference on Autonomous Agents and Multi-Agent Systems. pp. 415–422.

[13] Das, S., Kamenica, E., August 2005. Two-sided bandits and the dating market. In: Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence. Edinburgh, UK, pp. 947–952.

[14] Diamond, P., 1982. Aggregate demand management in search equilibrium. The Journal of Political Economy 90 (5), 881–894.

[15] Dranove, D., Jin, G., 2010. Quality disclosure and certification: Theory and practice. Journal of Economic Literature 48 (4), 935–63.

[16] Galeotti, A., Moraga-Gonzalez, J. L., 2009. Platform intermediation in a market for differentiated products. European Economic Review 53 (4), 417–428.

[17] Geng, X., Stinchcombe, M., Whinston, A., 2005. Bundling information goods of decreasing value. Management Science 51 (4), 662–667.

[18] Grosfeld-Nir, A., Sarne, D., Spieglier, I., 2009. Modeling the search for the least costly opportunity. European Journal of Operational Research 197 (2), 667–674.

[19] Hak Chun, Y., 1996. Selecting the best choice in the weighted secretary problem. European Journal of Operational Research 92 (1), 135–147.

[20] Haynes, M., Thompson, S., March 2008. Price, price dispersion and number of sellers at a low entry cost shopbot. International Journal of Industrial Organization 26 (2), 459–472.

[21] Ho, L., 2000. Wage subsidies as a labour market policy tool. Policy Sciences 33 (1), 89–100.

[22] Janetos, A., 1980. Strategies of female mate choice: A theoretical analysis. Behavioral Ecology and Sociobiology 7 (2), 107–112.

[23] Kang, B., 2005. Optimal stopping problem with double reservation value property. European Journal of Operational Research 165 (3), 765–785.

[24] Kephart, J. O., Hanson, J. E., Greenwald, A. R., 2000. Dynamic pricing by software agents. Computer Networks 32, 731–752.

[25] Lee, R., Schwarz, M., 2009. Interviewing in two-sided matching markets, NBER Working Paper.

[26] Lim, C., Bearden, J., Smith, J., 2006. Sequential search with multiattribute options. Decision Analysis 3 (1), 3–15.

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

[28] Lippman, S., Mcoall, J., 1976. The economics of job search: A survey. Economic Inquiry 14, 155–189.

[29] MacQueen, J. B., 1964. Optimal policies for a class of search and evaluation problems. Management Science 10 (4), 746–759.

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

[31] McMillan, J., Rothschild, M., 1994. Search. In: Aumann, R., Hart, S. (Eds.), Handbook of Game Theory with Economic Applications. North-Holland, pp. 905–927.

[32] Milgrom, P., 1981. Good news and bad news: Representation theorems and applications. The Bell Journal of Economics 12, 380–391.

[33] Monahan, G., 1980. Optimal stopping in a partially observable Markov process with costly information. Operations Research 28 (6), 1319–1334.

[34] Monahan, G., 1982. Optimal stopping in a partially observable binary-valued Markov chain with costly perfect information. Journal of Applied Probability 19 (1), 72–81.

[35] Mortensen, D., Pissarides, C., 1999. New developments in models of search in the labor market. Handbook of Labor Economics 3, 2567–2627.

[36] Moscarini, G., Smith, L., 2003. The optimal level of experimentation. Econometrica 69 (6), 1629–1644.

[37] Rysman, M., 2009. The economics of two-sided markets. Journal of Economic Perspectives 23 (3), 125–43.

[38] Sahib, P., Gu, X., 2002. To tie the knot or not: Cohabitation, marriage and individuals' attitudes to risk. Demographic Research 6 (13), 355–382.

[39] Sundararajan, A., 2004. Nonlinear pricing of information goods. Management Science 50 (12), 1660–1673.

[40] Terwiesch, C., Loch, C. H., 2004. Collaborative prototyping and the pricing of custom-designed products. Management Science 50 (2), 145–158.

[41] Weitzman, M., 1979. Optimal search for the best alternative. Econometrica 47 (3), 641–654.

[42] Wiegmann, D., Weinersmith, K., Seubert, S., 2010. Multi-attribute mate choice decisions and uncertainty in the decision process: A generalized sequential search strategy. Journal of Mathematical Biology 60 (4), 543–572.

[43] Wilson, K. E., Szechtman, R., Atkinson, M. P., 2011. A sequential perspective on searching for static targets. European Journal of Operational Research 215 (1), 218 – 226.

[44] Wright, R., 1986. Job search and cyclical unemployment. The Journal of Political Economy 94 (1), 38–55.

Appendix A. Proofs

Lemma 4 *In state $\gamma > 0$, querying the expert dominates terminating search without querying the expert*

Proof: For any strategy that terminates search upon obtaining a signal s when in state $\gamma > 0$, consider instead a modification of that strategy which queries the expert, and terminates only if the true value v, revealed by the expert, is greater than $V_{\gamma-1}$. This new strategy dominates the one it is based on , since if the true value is less than $V_{\gamma-1}$ the searcher is better off resuming search, and she pays no marginal cost to obtain the expert's service in this instance. $\square$

Proof of Theorem 2: For Part (1), we first show that if, optimally, the searcher should resume her search given a signal s, then she must also do so given any signal $s' < s$. Then, we show that if, optimally, the searcher should terminate her search given a signal s, then she must also do so given any signal $s'' > s$.

If the optimal strategy given signal s is to resume search then the following two inequalities hold, describing the superiority of resuming search over terminating search (A.1) and querying the expert (A.2):

$$V_0 > E[v|s] \qquad \text{(A.1)} \qquad\qquad V_0 > M(s, V_{k-1}) - c_e^k \quad \text{(A.2)}$$

Given the HSGN assumption and since $s' < s$, (A.1) holds also for s'. Similarly, notice that:

$$\begin{aligned} M(s, V_{k-1}) &= \int_{y=V_{k-1}}^{\infty} y f_v(y|s)\, dy + V_{k-1} F_v(V_{k-1}|s) \\ &= \Big[y F_v(y|s) \Big]_{y=V_{k-1}}^{\infty} - \int_{y=V_{k-1}}^{\infty} F_v(y|s)\, dy + V_{k-1} F_v(V_{k-1}|s) \\ &= \int_{y=V_{k-1}}^{\infty} (1 - F_v(y|s))\, dy + V_{k-1} > M(s', V_{k-1}) \quad \forall s' < s, \quad \text{due to HSGN.} \end{aligned} \tag{A.3}$$

Consequently (A.2) holds also for s', proving 1(a).

The proof for $s'' > s$ is exactly equivalent: the expected utility of accepting the current opportunity can be shown to dominate both resuming the search and querying the expert, proving 1(b) and thus (together with 1(a)) also 1(c). The optimal strategy when in state $\gamma = 0$ can thus be described by the tuple (t_l, t_u, V_{k-1}) as stated in the theorem and consequently calculated as the Bellman Equation captured by Equation 6 for the case of $\gamma = 0$. Taking the derivative of Equation 6 w.r.t. t_l and equating to zero, we obtain a unique t_l which maximizes the expected utility.

$$\begin{aligned} \frac{\partial V_0}{\partial t_l} =& V_0 f_s(t_l) + F_s(t_l) \frac{\partial V_0}{\partial t_l} - f_s(t_l) V_{k-1} F_v(V_{k-1}|t_l) + \int_{t_l}^{t_u} f_s(s) \frac{\partial (V_{k-1} F_v(V_{k-1}|s))}{\partial t_l} \\ & - f_s(t_l) \int_{V_{k-1}}^{\infty} y f_v(y|t_l)\, dy - \int_{t_l}^{t_u} f_s(s) V_{k-1} f_v(V_{k-1}|s) \frac{\partial V_{k-1}}{\partial t_l}\, ds + c_e^k f_s(t_l) \end{aligned}$$

where $F_s(s)$ is the cumulative distribution function of the signal s. After some algebraic manipulations, we get:

$$\frac{\partial V_0}{\partial t_l}(1 - F_s(t_l)) = f_s(t_l) \left(V_0 - V_{k-1} - \int_{V_{k-1}}^{\infty} (1 - F_v(y|t_l))\, dy + c_e^k \right) + \frac{\partial V_{k-1}}{\partial t_l} \int_{t_l}^{t_u} f_s(s) F_v(V_{k-1}|s)\, ds \tag{A.4}$$

Substituting $\frac{\partial V_0}{\partial t_l} = 0$ and $\frac{\partial V_{k-1}}{\partial t_l} = 0$[6] in (A.4), we obtain Equation 7: $c_e^k = \int_{V_{k-1}}^{\infty} (1 - F_v(y|t_l))\, dy + V_{k-1} - V_0 = V_{k-1} - V_0 + \int_{V_{k-1}}^{\infty} (y - V_{k-1}) f_v(y|t_l)\, dy$.

[6] This holds whenever $\frac{\partial V_0}{\partial t_l} = 0$. It can be seen by subsituting $k-1$ for γ in the section on calculation of derivatives in supplementary material.

To confirm that this is a maximum, we compute the second derivative of V_0 w.r.t t_l:

$$\frac{\partial^2 V_0}{\partial {t_l}^2}(1 - F_s(t_l)) - f_s(t_l)\frac{\partial V_0}{\partial t_l} = f_s^{'}(t_l)\left(V_0 - V_{k-1} - \int_{V_{k-1}}^{\infty}(1 - F_v(y|t_l))\,dy + c_e^k\right)$$
$$+ f_s(t_l)\left(\frac{\partial V_0}{\partial t_l} - \frac{\partial V_{k-1}}{\partial t_l} + (1 - F_v(V_{k-1}|t_l))\frac{\partial V_{k-1}}{\partial t_l} + \int_{V_{k-1}}^{\infty} dy \frac{\partial F_v(y|t_l)}{\partial t_l}\right)$$
$$+ \frac{\partial^2 V_{k-1}}{\partial t_l^2}\int_{t_l}^{t_u} f_s(s)F_v(V_{k-1}|s)\,ds + \frac{\partial V_{k-1}}{\partial t_l}\frac{\partial}{\partial t_l}\left(\int_{t_l}^{t_u} f_s(s)F_v(V_{k-1}|s)\,ds\right) \tag{A.5}$$

Substituting $\frac{\partial V_0}{\partial t_l} = 0$, $\frac{\partial V_{k-1}}{\partial t_l} = 0$, $c_e^k = \int_{V_{k-1}}^{\infty}(1 - F_v(y|t_l))\,dy + V_{k-1} - V_0$, and $\frac{\partial^2 V_{k-1}}{\partial t_l^2} = \prod_{i=1}^{k-1}\Pr(v \leq V_{i-1}|s \geq t_i)\frac{\partial^2 V_0}{\partial t_l^2} + \frac{\partial}{\partial t_l}\left(\prod_{i=1}^{k-1}\Pr(v \leq V_{i-1}|s \geq t_i)\right)\frac{\partial V_0}{\partial t_l}$ (we can get this by substituting $k-1$ in place of γ in the expression for derivatives given in supplementary material) obtains:

$$\frac{\partial^2 V_0}{\partial {t_l}^2}\left(1 - F_s(t_l) - \left(\prod_{i=1}^{k-1}\Pr(v \leq V_{i-1}|s \geq t_i)\right)\Pr(v \leq V_{k-1} \wedge t_l \leq s \leq t_u)\right)$$
$$= f_s(t_l)\left(\int_{V_{k-1}}^{\infty} dy \frac{\partial F_v(y|t_l)}{\partial t_l}\right)$$

The left hand side is positive, since:

$$1 - F_s(t_l) - \left(\prod_{i=1}^{k-1}\Pr(v \leq V_{i-1}|s \geq t_i)\right)\Pr(v \leq V_{k-1} \wedge t_l \leq s \leq t_u)$$
$$= 1 - F_s(t_l) - \left(\prod_{i=1}^{k-1}\Pr(v \leq V_{i-1}|s \geq t_i)\right)\Pr(v \leq V_{k-1}|t_l \leq s \leq t_u)\Pr(t_l \leq s \leq t_u)$$
$$\geq 1 - F_s(t_l) - \Pr(t_l \leq s \leq t_u) = 1 - F_s(t_l) - (F_s(t_u) - F_s(t_l)) = 1 - F_s(t_u) \geq 0$$

By the HSGN assumption, $F_V(y|t_l)$ is a decreasing function of t_l, therefore, the derivative is negative, and the term inside the integral is negative, meaning the above term is negative, which completes the proof, showing that Equation 7 must hold when using the optimal search strategy. The proof for Equation 8 is similar.

For Part (2), given Lemma 4, we only need to prove that if, optimally, the searcher should resume search given a signal s, then she must also do so given signal $s' < s$. The proof is similar to the above, ignoring the option to terminate without querying the expert and the query cost, and replacing V_{k-1} by $V_{\gamma-1}$ whenever applicable. If the optimal strategy given signal s is to resume search then the following equality (A.6) holds, describing the superiority of resuming search over querying the expert:

$$V_{\gamma>0} > M(s, V_{\gamma-1}) \tag{A.6}$$

Since $M(s, V_{\gamma-1}) > M(s', V_{\gamma-1})$ (substitute γ for k in (A.3)), we obtain that (A.6) holds also for s'. This, together with Lemma 4 suggests that if, optimally, the searcher should terminate her search given a signal s, then she must also do so given any signal $s'' > s$. Taking the derivative of Equation 5 (the Bellman Equation for state $\gamma > 0$) w.r.t. t_γ and equating to zero (similar to the way given above), we obtain a unique t_γ which maximizes the expected utility. We can use Equations 6, 7, 8, and the set of $2k-2$ instances of Equations 5 and 9 (one for each γ value, $1 \leq \gamma < k$, i.e., a total of $2k+1$ equations), to solve for $2k+1$ unknowns: V_0, t_l, t_u, and for each $\gamma = 1, 2, \ldots, k-1$ the appropriate V_γ and t_γ. $\square$

We note that all these proofs generalize easily to cases where the value distribution has support $[a, b]$. Furthermore, the above analysis is applicable also for the case where the expert supplies a noisy signal rather than a definite value, whenever the expert's signal complies with HSGN. This is achieved using a simple transformation that can be found in MacQueen [29] (see details in the supplementary material).

S.1. Supplementary Material for "Expert-Mediated Sequential Search"

S.1.1. Proof of Theorem 1

We show that, optimally, if the searcher should resume her search given a signal s, then she must necessarily also do so given any signal $s' < s$. Let V denote the expected utility to the searcher upon resuming the search if signal s is obtained. Since the optimal strategy given signal s is to resume search, we know $V > E[v|s]$. From the HSGN assumption, $E[v|s] \geq E[v|s']$ for $s' < s$. Therefore, $V > E[v|s']$, proving that the optimal strategy is reservation-value and search should be resumed for any $s' < s$. Then, the expected utility when using reservation signal t is given by:

$$V(t) = -c_s + V(t)\int_{s=-\infty}^{t} f_s(s)\,ds + \int_{s=t}^{\infty} E[v|s]f_s(s)\,ds = -c_s + V(t)F_s(t) + \int_{s=t}^{\infty} E[v|s]f_s(s)\,ds \tag{S.1}$$

where $F_s(s)$ is the cumulative distribution function of the signal s. Taking the first derivative with respect to t on both sides and then setting it to 0 at t^*, we get:

$$\frac{dV(t)}{dt} = \frac{dV(t)}{dt}F_s(t) + V(t)f_s(t) - E[v|t]f_s(t) = 0$$

And since $\frac{dV(t)}{dt} = 0$ for $t = t^*$ we obtain:

$$V(t^*) = E[v|t^*] \qquad f_s(t^*) \neq 0 \tag{S.2}$$

$V(t^*) = E[v|t^*]$ implies that the reservation value t^* is the signal for which the searcher's utility of resuming search is equal to the expected value of the opportunity associated with that signal.

To verify that $V(t)$ reaches its maximum for t^*, we calculate the second derivative.

$$\frac{d^2V}{dt^2} = F_s(t)\frac{d^2V(t)}{dt^2} + f_s(t)\frac{dV(t)}{dt} + \frac{d(V(t)f_s(t))}{dt} - \frac{d(E[v|t]f_s(t))}{dt}$$
$$\frac{d^2V}{dt^2}(1 - F_s(t)) = f_s(t)\frac{dV(t)}{dt} + \frac{d((V(t) - E[v|t])f_s(t))}{dt}$$
$$\frac{d^2V}{dt^2}(1 - F_s(t)) = f_s(t)\frac{dV(t)}{dt} + \frac{df_s(t)}{dt}(V(t) - E[v|t]) + f_s(t)\frac{d(V(t) - E[v|t])}{dt}$$

At $t = t^*$, $\frac{d(V(t)}{dt}|t^* = 0$ and $V(t^*) = E(v|t^*)$, therefore substituting these values in the above equation, we get:

$$\frac{d^2V}{dt^2}(1-F_s(t^*)) = 0+0+0-f_s(t^*)\frac{dE[v|t])}{dt}|t^* = \frac{d^2V}{dt^2} = -\frac{f_s(t^*)}{1-F_s(t^*)}\left(\frac{dE[v|t]}{dt}|t^*\right)$$

Given the HSGN assumption, $E[v|t]$ is an increasing function in t. Therefore $\frac{dE[v|t]}{dt} > 0$, hence the value of the second derivative evaluated at t^* is less than zero, which confirms that t^* indeed yields a maximum.

Finally substituting $V(t^*) = E[V|t^*]$ in (S.1) we obtain Equation 2 as follows:

$$E[v|t^*] = -c_s + E[v|t^*]\int_{s=-\infty}^{t^*} f_s(s)\,ds + \int_{s=t^*}^{\infty} E[v|s]f_s(s)\,ds \Rightarrow c_s = \int_{s=t^*}^{\infty}(E[v|s]-E[v|t^*])f_s(s)\,ds$$

The value t^* can be calculated using the above equation. □

S.1.2. Noisy Expert Signal

If the expert is producing a noisy signal s_e rather than a definite (true) value then the analysis provided in the paper, and in particular the optimal policy, is applicable conditioned that the expert's signal also complies with HSGN. This is facilitated by using the transformation $Z = \mathbb{E}(v|s, s_e)$ suggested in [29] for the specific case of a la carte pricing. In this case we can show that the optimal policy is two-thresholds-based when the searcher is in state $\gamma = 0$, i.e, buy the package if the signal s is below t_l, accept without querying if the signal s is greater than t_u, and if the signal s lies in between t_l and t_u, query the expert. On querying the expert and receiving the signal s_e, if $Z = \mathbb{E}_v(y|s, s_e) > V_{k-1}$, then accept the opportunity, otherwise resume. In any other state $\gamma > 0$, the optimal policy is to query the expert if the signal s is greater than t_γ, otherwise reject; on querying the expert and receiving the signal s_e, accept the opportunity if $Z = \mathbb{E}_v(y|s, s_e) > V_{\gamma-1}$, otherwise reject.

This is achieved by replacing $f_v(y|s)$ with $f_z(z|s)$, where $f_z(z|s)$ is the conditional pdf of Z given s (the searcher's signal), and similarly replacing $F_v(y|s)$ with $F_z(z|s)$, the conditional cdf of Z given s, in Equations 5-9. The resulting equations in this case are:

$$V_{\gamma>0} = V_\gamma F_s(t_\gamma) + \int_{s=t_\gamma}^{\infty} f_s(s)\Big(\int_{y=V_{\gamma-1}}^{\infty} yf_z(y|s)\,dy + V_{\gamma-1}F_z(V_{\gamma-1}|s)\Big)\,ds - c_s \tag{S.3}$$

$$V_0 = V_0F_s(t_l) + \int_{s=t_u}^{\infty} f_s(s)E[v|s]\,ds + \int_{s=t_l}^{t_u} f_s(s)\Big(V_{k-1}F_z(V_{k-1}|s) + \int_{y=V_{k-1}}^{\infty} yf_z(y|s)\,dy - c_e^k\Big)\,ds - c_s \tag{S.4}$$

$$c_e^k = -V_0 + V_{k-1} + \int_{y=V_{k-1}}^{\infty} (1 - F_z(y|t_l))\, dy \tag{S.5}$$

$$c_e^k = \int_{y=-\infty}^{V_{k-1}} (V_{k-1} - y) f_z(y|t_u)\, dy \tag{S.6}$$

$$V_\gamma = V_{\gamma-1} F_z(V_{\gamma-1}|t_\gamma) + \int_{y=V_{\gamma-1}}^{\infty} y f_z(y|t_\gamma)\, dy \tag{S.7}$$

We now show how to calculate $F_z(z|s)$ and $f_z(z|s)$, the conditional cdf and pdf of Z given the searcher's signal s respectively

$$\begin{aligned} F_z(z|s) &= \Pr(Z \leq z|s) \\ &= \Pr(\mathbb{E}_v(y|s, s_e) \leq z|s) \\ &= \Pr(\int_{-\infty}^{\infty} y f_v(y|s, s_e)\, dy \leq z|s) \end{aligned}$$

and consequently:

$$f_z(z|s) = \frac{dF_z(z|s)}{dz}$$

This assumes that the joint distribution $f(v, s, s_e)$ is given, therefore, in the above set of equations $f_v(y|s, s_e)$ can be calculated using the joint distribution $f(v, s, s_e)$, as $f_v(y|s, s_e) = \frac{f(v,s,s_e)}{\int_{-\infty}^{\infty} f(v,s,s_e)\, dv}$.

S.1.3. Number of queries, η_{c_e}, decreases as c_e increases

Intuitively, given that the expert is providing similar quality at two different query cost, the searcher who is paying the cost will not want to query more often at a higher price than the number of times she queries at a lower price if she is playing optimally.

Lemma 5 *The optimal number of times the searcher queries the expert decreases as the price charged by the expert increases.*

Proof: We use notation similar to that in the proof of Theorem 3, assuming that $\mathbb{E}[\omega()]$, $\eta_s()$, and $\eta_{c_e}()$ are functions defined by the searcher's optimal behavior from solving the Bellman equation. Here $\mathbb{E}[\omega(c_e)]$ is the expected worth of the final opportunity that is accepted, $\eta_s(c_e)$ (analogous to η_s), is the expected number of opportunities examined by the searcher, and $\eta_{c_e}(c_e)$, analogous to η_{c_e}, the expected number of times the expert is queries when the searcher follows optimal strategy given the query price is c_e. Note, that these functions also depend on the search cost c_s but we are dropping it from the notation because we will keep it constant.

Using the notation discussed above for any c_e, the searcher's utility $V()$ can be formulated as: $V(c_e) = \mathbb{E}[\omega(c_e)] - c_s\eta_s(c_e) - c_e\eta_{c_e}(c_e)$. At any c_e this is the highest utility the searcher can achieve. Therefore:

$$\begin{aligned} V(c_e) &= \mathbb{E}[\omega(c_e)] - c_s\eta_s(c_e) - c_e\eta_{c_e}(c_e) \\ &\geq \mathbb{E}[\omega(c'_e)] - c_s\eta_s(c'_e) - c_e\eta_{c_e}(c'_e) \quad \forall \quad c'_e \end{aligned} \tag{S.8}$$

In order to show that the number of queries requested, when using the optimal strategy, decreases with an increase in query cost, we consider two query costs $c_{e_{\text{low}}}$ and $c_{e_{\text{high}}}$, such that, WLOG, $c_{e_{\text{low}}} < c_{e_{\text{high}}}$. Using Equation S.8 we get the following conditions:

$$\begin{aligned} V(c_{e_{\text{low}}}) &= \mathbb{E}[\omega(c_{e_{\text{low}}})] - c_s\eta_s(c_{e_{\text{low}}}) - c_{e_{\text{low}}}\eta_{c_e}(c_{e_{\text{low}}}) \\ &\geq \mathbb{E}[\omega(c_{e_{\text{high}}})] - c_s\eta_s(c_{e_{\text{high}}}) - c_{e_{\text{low}}}\eta_{c_{e_{\text{high}}}}(c_{e_{\text{high}}}) \\ \mathbb{E}[\omega(c_{e_{\text{low}}})] - c_s\eta_s(c_{e_{\text{low}}}) - c_{e_{\text{low}}}\eta_{c_e}(c_{e_{\text{low}}}) + c_{e_{\text{low}}}\eta_{c_{e_{\text{high}}}}(c_{e_{\text{high}}}) &\geq \mathbb{E}[\omega(c_{e_{\text{high}}})] - c_s\eta_s(c_{e_{\text{high}}}) \end{aligned} \tag{S.9}$$

$$\begin{aligned} V(c_{e_{\text{high}}}) &= \mathbb{E}[\omega(c_{e_{\text{high}}})] - c_s\eta_s(c_{e_{\text{high}}}) - c_{e_{\text{high}}}\eta_{c_e}(c_{e_{\text{high}}}) \\ &\geq \mathbb{E}[\omega(c_{e_{\text{low}}})] - c_s\eta_s(c_{e_{\text{low}}}) - c_{e_{\text{high}}}\eta_{c_{e_{\text{low}}}}(c_{e_{\text{low}}}) \\ \mathbb{E}[\omega(c_{e_{\text{high}}})] - c_s\eta_s(c_{e_{\text{high}}}) &\geq \mathbb{E}[\omega(c_{e_{\text{low}}})] - c_s\eta_s(c_{e_{\text{low}}}) - c_{e_{\text{high}}}\eta_{c_{e_{\text{low}}}}(c_{e_{\text{low}}}) + c_{e_{\text{high}}}\eta_{c_e}(c_{e_{\text{high}}}) \end{aligned} \tag{S.10}$$

Using Equations S.9 and S.10, we obtain:

$$\begin{aligned} \mathbb{E}[\omega(c_{e_{\text{low}}})] - c_s\eta_s(c_{e_{\text{low}}}) - c_{e_{\text{low}}}\eta_{c_e}(c_{e_{\text{low}}}) + c_{e_{\text{low}}}\eta_{c_{e_{\text{high}}}}(c_{e_{\text{high}}}) &\geq \mathbb{E}[\omega(c_{e_{\text{high}}})] - c_s\eta_s(c_{e_{\text{high}}}) \\ &\geq \mathbb{E}[\omega(c_{e_{\text{low}}})] - c_s\eta_s(c_{e_{\text{low}}}) - c_{e_{\text{high}}}\eta_{c_{e_{\text{low}}}}(c_{e_{\text{low}}}) + c_{e_{\text{high}}}\eta_{c_e}(c_{e_{\text{high}}}) \\ \Rightarrow 0 &\geq (\eta_{c_e}(c_{e_{\text{high}}}) - \eta_{c_e}(c_{e_{\text{low}}}))(c_{e_{\text{high}}} - c_{e_{\text{low}}}) \end{aligned} \tag{S.11}$$

Since $c_{e_{\text{high}}} > c_{e_{\text{low}}}$, Equation S.11 can hold only if $\eta_{c_e}(c_{e_{\text{high}}}) \leq \eta_{c_e}(c_{e_{\text{low}}})$. Therefore, with a higher query price, the optimal number of queries requested by the searcher is never greater than with a lower query price. $\square$

S.1.4. Computations For Non-Linear Package Pricing

Expected Number of Packages Purchased (η_b) and Expected Number of Queries used ($\eta_{c_e^k}$). In order to calculate η_b and $\eta_{c_e^k}$, we define and compute the following probabilities:

	Represents	General formulation **HSGN formulation**
$P_{\gamma\to\gamma-1}$	Pr (querying and resuming, moving from state $\gamma>0$ to $\gamma-1$)	$\int_{M(s,V_{\gamma-1})>V_\gamma} f_s(s)F_v(V_{\gamma-1}\|s)\,ds$ $\Pr(\mathbf{s}>\mathbf{t}_\gamma \text{ and } \mathbf{v}<\mathbf{V}_{\gamma-\mathbf{1}})$
$P_{\gamma\to\gamma}$	Pr (resuming without querying, staying in state $\gamma>0$)	$\int_{M(s,V_{\gamma-1})<V_\gamma} f_s(s)\,ds$ $\Pr(\mathbf{s}<\mathbf{t}_\gamma)$
$P_{\gamma\to ter}$	Pr (querying and then terminating from state $\gamma>0$)	$\int_{M(s,V_{\gamma-1})>V_\gamma} f_s(s)(1-F_v(V_{\gamma-1}\|s))\,ds$ $\Pr(\mathbf{s}\geq\mathbf{t}_\gamma \text{ and } \mathbf{v}\geq\mathbf{V}_{\gamma-\mathbf{1}})$
$P_{0\to k-1}$	Pr (querying and resuming, moving from state $\gamma=0$ to $k-1$)	$\int_{\zeta_3} f_s(s)F_v(V_{k-1}\|s)\,ds$ $\Pr(\mathbf{t_l}\leq\mathbf{s}\leq\mathbf{t_u} \text{ and } \mathbf{v}<\mathbf{V}_{\gamma-\mathbf{1}})$
$P_{0\to 0}$	Pr (resuming without querying when $\gamma=0$)	$\int_{\zeta_1} f_s(s)\,ds$ $\Pr(\mathbf{s}<\mathbf{t_l})$
$P_{0\to ter}^{\neg query}$	Pr (terminating without querying when $\gamma=0$)	$\int_{\zeta_2} f_s(s)\,ds$ $\Pr(\mathbf{s}>\mathbf{t_u})$
$P_{0\to ter}^{query}$	Pr (querying and terminating when $\gamma=0$)	$\int_{\zeta_3} f_s(s)(1-F_v(V_{\gamma-1}\|s))\,ds$ $\Pr(\mathbf{t_l}\leq\mathbf{s}\leq\mathbf{t_u} \text{ and } \mathbf{v}\geq\mathbf{V}_{\gamma-\mathbf{1}})$

Notice that $P_{\gamma\to\gamma-1}+P_{\gamma\to\gamma}+P_{\gamma\to ter}=1$, and similarly $P_{0\to k-1}+P_{0\to 0}+P_{0\to ter}^{\neg query}+P_{0\to ter}^{query}=1$. Let $P_\gamma(T)$ denote the probability of eventually transitioning, when in state γ, to state $\gamma-1$ (for $\gamma=0$ it is the probability of eventually transitioning to state $k-1$). Let P_{cycle} be the probability of starting at a given state and getting back to it after going through all other states (excluding search termination state). The values of $P_\gamma(T)$ and P_{cycle} can be calculated as $P_\gamma(T)=\sum_{j=0}^{\infty}(P_{\gamma\to\gamma})^j P_{\gamma\to\gamma-1}=\frac{P_{\gamma\to\gamma-1}}{1-P_{\gamma\to\gamma}}$; $P_{\text{cycle}}=\prod_{i=0}^{k-1}P_i(T)$. Now let $P_\gamma(\text{Term})$ denote the probability of terminating the search in current state γ without transitioning to state $\gamma-1$, and $P_0(\text{Buy}\wedge\text{Term})$ denote the probability of eventually purchasing the package and then terminating the search when starting from state $\gamma=0$. Then:

$$P_\gamma(\text{Term})=1-P_\gamma(T)=\frac{P_{\gamma\to ter}}{1-P_{\gamma\to\gamma}};\quad P_0(\text{Buy}\wedge\text{Term})=\frac{P_{0\to ter}^{query}}{1-P_{0\to 0}};$$

$$P_0(\neg\ \text{Buy}\wedge\text{Term})=\frac{P_{0\to ter}^{\neg query}}{1-P_{0\to 0}}$$

Expected number of Packages purchased (η_b). In order for the searcher to purchase the jth package, she first needs to purchase $j-1$ packages and use them fully (with probability

$(P_{\text{cycle}})^{j-1}$), then purchase exactly one additional package. There are three ways for this to occur: (1) the searcher purchases the package and terminates the search in state $\gamma = 0$ (with probability $P_0(\text{Buy} \wedge \text{Term})$); (2) the searcher purchases the package and terminates the search in some state $\gamma > 0$ (with probability $P_0(T) - P_{\text{cycle}}$); and (3) the searcher purchases the package, fully consumes it (i.e., returns to state $\gamma = 0$) and then terminates search without further package purchases (with probability $P_{\text{cycle}} P_0(\neg\ \text{Buy} \wedge \text{Term})$). Therefore,

$$\begin{aligned}\Pr(\text{purchase exactly } j \text{ packages}) = (P_{\text{cycle}})^{j-1}\Big(&P_0(\text{Buy} \wedge \text{Term}) + P_0(T) - P_{\text{cycle}} \\ &+ P_{\text{cycle}} P_0(\neg\ \text{Buy} \wedge \text{Term})\Big)\end{aligned}$$

$P_0(T) = 1 - P_0(\neg\ \text{Buy} \wedge \text{Term}) + P_0(\text{Buy} \wedge \text{Term})$. Therefore, the above expression becomes $P_{\text{cycle}}^{j-1}(P_0(\text{Buy} \wedge \text{Term}) + P_0(T))(1 - P_{\text{cycle}})$. Then the expected number of packages purchased, η_b, is given by:

$$\begin{aligned}\eta_b &= \sum_{j=1}^{\infty} j \Pr(\text{purchase exactly } j \text{ packages}) \\ &= \sum_{j=1}^{\infty} j (P_{\text{cycle}})^{j-1}(P_0(\text{Buy} \wedge \text{Term}) + P_0(T))(1 - P_{\text{cycle}})\end{aligned}$$

Expected number of queries used ($\eta_{c_e^k}$). Let $P_m(Q)$ be the probability that exactly m queries are used. If $m < k$ then the only way to exhaust m queries is by transitioning from state $\gamma = 0$ to state $k - m + 1$ and eventually terminating the search without transitioning to the next state. $m = k$ implies that the searcher transitions to state $\gamma = 1$ and either terminates without transitioning to $\gamma = 0$ or terminates after transitioning to $\gamma = 0$ without purchasing a new package. Therefore:

$$P_m(Q) = \begin{cases} P_0(\text{Buy} \wedge \text{Term}) & m = 1 \\ P_0(T)\Big(\prod_{j=k-m+2}^{j=k-1} P_j(T)\Big) P_{k-m+1}(\text{Term}) & m < k \\ P_0(T)\Big(\prod_{j=2}^{j=k-1} P_j(T)\Big) P_1(\text{Term}) + P_{\text{cycle}} P_0(\neg\text{Buy} \wedge \text{Term}) & m = k \end{cases}$$

For $m > k$, we represent m as $jk + i$ where $i = m\%k$ and $j = \lfloor \frac{m}{k} \rfloor$. Here, j represents the number of full cycles completed (when all queries in a package are used) and i represents the number of queries used prior to terminating before finishing the $(j+1)$th round. This cyclic nature gives us the following recurrence:

$$P_{m=jk+i}(Q) = P_{\text{cycle}} P_{((j-1)k+i)}(Q) = \cdots = (P_{\text{cycle}})^j P_i(Q)$$

Therefore the expected number of queries is given by:

$$\eta_{c_e^k} = \sum_{m=0}^{\infty} m P_m(Q) = \sum_{j=0}^{\infty}\sum_{i=1}^{k} (jk+i)(P_{\text{cycle}})^j P_i(Q) = \sum_{j=0}^{\infty}\sum_{i=1}^{k} (jk+i) \left(\prod_{l=0}^{k-1} P_l(T)\right)^j P_i(Q) \tag{S.12}$$

Expected value of an opportunity received. Let ω be the random variable representing the value of an opportunity received, then $E_i(\omega)$ represents the expected value of the opportunity if the search terminates in state i and $P_i(\omega)$ represents the probability of terminating at that state.

$$E_i(\omega) = \begin{cases} \mathbb{E}[v|v > V_{i-1}, s > t_i] & i > 0 \\ \frac{\mathbb{E}[v|v>V_{k-1}, t_l<s<t_u]\Pr(v>V_{k-1}\wedge t_l<s<t_u)+\mathbb{E}[v|s>t_u]\Pr(s>t_u)}{\Pr(v>V_{k-1}\wedge t_l<s<t_u)+\Pr(s>t_u)} & i = 0 \end{cases}$$

$$P_i(\omega) = \begin{cases} \sum_{j=0}^{j=\infty} P_{\text{cycle}}^j P_0(T)\left(\prod_{h=i+1}^{h=k-1} P_h(T)\right) P_i(\text{Term}) & i > 0 \\ \sum_{j=0}^{j=\infty} P_{\text{cycle}}^j (1 - P_0(T)) & i = 0 \end{cases}$$

$$\mathbb{E}(\omega) = \mathbb{E}(\text{Expected value of opportunity}) = \sum_{i=0}^{i=k-1} E_i(\omega) P_i(\omega)$$

Expected number of opportunities examined. By definition V_0 is the value of the opportunity accepted minus the total search and query cost incurred ($V_0 = \mathbb{E}(\omega) - \eta_s c_s - \eta_b c_e^k$). Therefore, $\eta_s = \frac{\mathbb{E}(\omega) - \eta_b c_e^k - V_0}{c_s}$.

S.1.5. The dependency of t_l, t_u and t_γ in c_s and c_e

A pattern similar to the one reported in Section 5.1 holds when using packages of more than a single query: The dependency of t_l and t_u in c_s and c_e is similar to the one illustrated in Figure 3, and similarly, t_γ decreases as c_s increases and increases as c_e increases (see Figure S.1).

S.1.6. Robustness: Adding Discounting

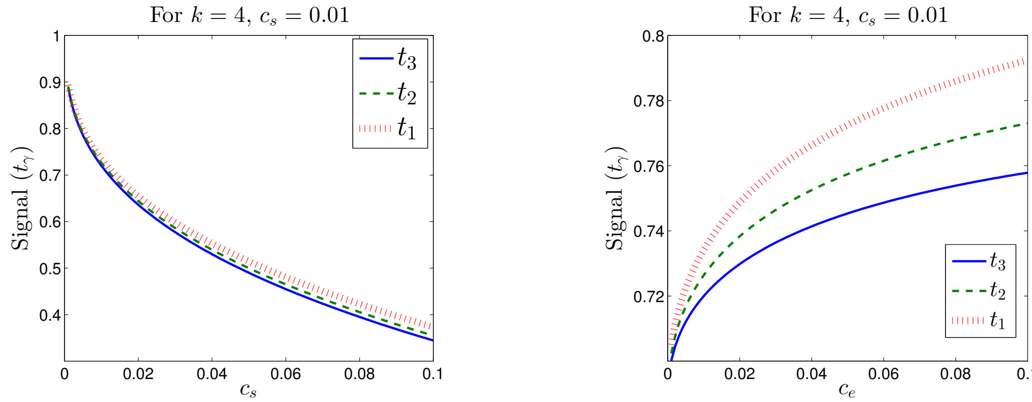

(a) (b)

Figure S.1: (a) Effect of c_s on the signal thresholds (t_γ) for $\gamma = 1, 2, 3$, $c_e^k = 0.01$ and package size $k = 4$. (b) Effect of c_e^k on the signal thresholds (t_γ) for states $\gamma = 1, 2, 3$, package size $k = 4$ and $c_s = 0.01$.

We note here that, while particulars may vary, the form of our results is robust when we add a discount factor δ into the utility functions.[1] We illustrate with an appropriate modification for the a la carte model variant. The changes for Equations 10-12 are easy; Equation 10 is updated to Equation S.13 and Equations 11-12 remain the same.

$$V = -c_s + \delta \left(V F_s(t_l) + \int_{s=t_u}^{\infty} f_s(s) E[v|s]\, ds + \int_{s=t_l}^{t_u} f_s(s) \left(V F_v(V|s) + \int_{y=V}^{\infty} y f_v(y|s)\, dy - c_e \right) ds \right) \tag{S.13}$$

The only minor technical change in the remaining analysis is that one cannot use the simple formulation for expert's strategy that we outline in Section 4.2 because simply multiplying the expected number of queries by the cost per query is insufficient – the timing of queries matters as well. Let P_s and P_q be the probabilities that the searcher resumes the search (given by $A + B$, as in Section 4.2) and that the searcher queries the expert (given

[1]Indeed, in many sequential search environments, discounting can be neglected if considering a short-term search where search costs dominate, as in most online markets. Still, discounting of gains and costs is important in more traditional search settings.

by $A + D$, as in Section 4.2), respectively. The expert's profit π_e is given by:

$$\pi_e = (c_e - d_e) \sum_{t=0}^{\infty} \delta^t \Pr(\text{Search is not terminated by time } t-1) \times \Pr(\text{Searcher queries at time t})$$
$$= (c_e - d_e) \sum_{t=0}^{\infty} \delta^t P_s^t P_q = \frac{(c_e - d_e)P_q}{1 - \delta P_s} = \frac{(c_e - d_e)(A + D)}{1 - \delta(A + B)} \quad \text{(S.14)}$$

Figure S.2 shows a comparison of the expert's profit as a function of query cost for the same environment as in Section 5 for a discounted case (with discount factor of 0.95) and the non-discounted case. We find that the introduction of discounting does not significantly affect the nature of our results, so we do not consider it in this paper.

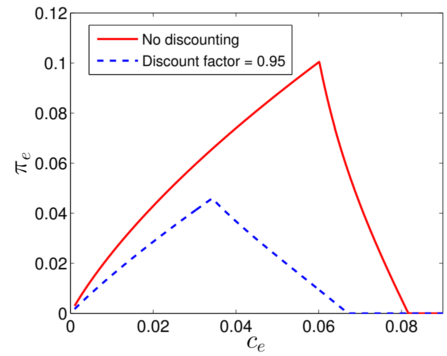

Figure S.2: Expert's profit with and without discount factor for $c_s = 0.01$.

S.1.7. Calculation of $\frac{\partial V_\gamma}{\partial t_l}$ and $\frac{\partial^2 V_\gamma}{\partial t_l^2}$

From Equation 5:

$$V_{\gamma>0} = V_\gamma F_s(t_\gamma) + \int_{s=t_\gamma}^{\infty} f_s(s)\Big(\int_{y=V_{\gamma-1}}^{\infty} y f_v(y|s)\,dy + V_{\gamma-1}F_v(V_{\gamma-1}|s)\Big)\,ds - c_s$$

$$\frac{\partial V_\gamma}{\partial t_l}(1 - F_s(t_\gamma)) = \frac{\partial V_{\gamma-1}}{\partial t_l}\int_{t_\gamma}^{\infty} f_s(s)F_v(V_{\gamma-1}|s)\,ds$$

$$= \frac{1}{1 - F_s(t_\gamma)}\frac{\partial V_{\gamma-1}}{\partial t_l}\int_{t_\gamma}^{\infty} f_s(s)F_v(V_{\gamma-1}|s)\,ds$$

$$= \frac{\Pr(v \le V_{\gamma-1} \wedge s \ge t_\gamma)}{\Pr(s \ge t_\gamma)}\frac{\partial V_{\gamma-1}}{\partial t_l}$$

$$= \Pr(v \le V_{\gamma-1}|s \ge t_\gamma)\frac{\partial V_{\gamma-1}}{\partial t_l}$$

$$= \prod_{i=1}^{\gamma}\Pr(v \le V_{i-1}|s \ge t_i)\frac{\partial V_0}{\partial t_l}$$

$$\frac{\partial^2 V_\gamma}{\partial t_l^2} = \prod_{i=1}^{\gamma}\Pr(v \le V_{i-1}|s \ge t_i)\frac{\partial^2 V_0}{\partial t_l^2} + \frac{\partial}{\partial t_l}\left(\prod_{i=1}^{\gamma}\Pr(v \le V_{i-1}|s \ge t_i)\right)\frac{\partial V_0}{\partial t_l}$$