

A Two-Agent Model of Sequential Search and Choice*

Eeva Mauring[†]

This draft: January 27, 2016

First draft: March 20, 2015

Abstract

This paper extends the standard sequential search model by allowing the agent who compiles the choice set via search (the “searcher”) to differ from the agent who chooses from the set (the “chooser”). I show for a general joint distribution of the agents’ preferences that the searcher’s optimal policy is a threshold rule. In contrast to the standard model, the threshold is weakly decreasing in time (i.e., exhibits the “discouragement effect”), although the search horizon is infinite and the search environment stationary. I characterise the threshold and discuss the testable implications of the discouragement effect. The characteristics of my model differ from two single-agent search models that feature a time-varying threshold (convex search costs or deadline). In particular, my model features a threshold that decreases endogenously over time and never generates return to an item rejected earlier, in contrast to the other models.

JEL classification: D83, D79, D99.

Keywords: *sequential search, multi-agent search, threshold policy, discouragement effect.*

*I thank my advisors Ran Spiegel and Martin Cripps for support and advice, and the editor, two anonymous referees, Jan Eeckhout, Andrea Pirrone, Nikita Roketskiy, Ludo Visschers, Lukas Wenner, and audiences at UCL, RES Junior Symposium, SaM Annual Conference, and RES Women’s Meeting for comments. Financial support from ERC grant no. 230251; ESRC; Kristjan Jaak Scholarship funded by the Ministry of Education and Research, Estonia, and the Archimedes Foundation; and Krediidipank is gratefully acknowledged. Earlier title: “Two-Person Search”. doi:10.1016/j.jebo.2015.12.010

[†]Department of Economics, University College London, Gower Street, London, WC1E 6BT, United Kingdom; eevamauring@gmail.com

1 Introduction

Standard sequential search models with recall build on the assumption that the search and choice stages comprise an undivided whole: the person who searches can stop and choose an item from the accumulated choice set at any time during the search process. This is an innocuous assumption if the preferences of the person are stable over time. In this paper, I extend the standard search model by allowing the preferences according to which the final choice is made to differ from the preferences according to which search is conducted. The set-up has two natural interpretations. First, the preferences belong to different parties: a “searcher” compiles a choice set via sequential search and a “chooser” chooses from the collected choice set. Second, the preferences belong to one individual, but change between the search and choice stages. I show that the searcher’s optimal policy is a threshold rule and characterise the threshold.

Examples of such search problems are an HR manager collecting applications for a boss who wants to hire a new worker and a real estate agent collecting offers for a client interested in buying a flat. An example involving a person and a set of individuals is a spouse looking for a job that determines the living place of the couple. A person who is interested in the return while searching for an investment opportunity, but later tempted to invest in an option that involves the least paperwork is a “multi-selves” example. More generally, many household decisions, organisational decisions involving different phases and multiple agents, choice processes partially outsourced to external partners, and political decisions involving advisers feature one party compiling a choice set for another party via search.

In this paper I analyse the optimal policy of a searcher who compiles a choice set for a chooser. I describe the model as a two-agent search problem. The searcher (he) and the chooser (she) have preferences over all items in some grand set of alternatives and the preferences are distributed according to a general full-support distribution function. The searcher has access to an arrival process. In each period, one item arrives and the searcher discovers how much utility both he and the chooser receive from the item if it is chosen. The searcher decides in each period whether to stop or continue the search process. If he stops the process, all the items that have arrived are presented to the chooser. The chooser then chooses the best item in the choice set according to her preferences, unless all the items in the set yield her less utility than her exogenous outside option. Utilities are realised when an item or the outside option is chosen. The process ends after the chooser moves. The searcher’s problem is to choose an optimal policy, knowing the chooser’s choice rule.

First, I derive and characterise the searcher’s optimal policy. His optimal policy is

a threshold rule and the threshold depends on the item that would be chosen by the chooser were the searcher to stop immediately, x_m . The searcher’s threshold is the lower the higher the chooser values this item because it acts as a restriction for the searcher: if he is unsatisfied with the utility he would receive from x_m , a new item is chosen only if the chooser’s utility from it exceeds her utility from x_m . This has two implications. First, if the searcher finds an item that has a very high value for the chooser, he optimally stops searching regardless of the value that the item yields him. Second, the observed threshold that the searcher uses is weakly decreasing in time, although the search horizon is infinite and the search environment stationary. I call this the “discouragement effect”. This is in contrast with the standard single-agent search model where a stationary environment translates into a stationary threshold. The searcher’s threshold in my model is defined implicitly by a differential equation. I use a specific joint distribution where the utilities’ correlation is captured by a single parameter to numerically show that an increase in the correlation parameter unambiguously increases the searcher’s threshold, in line with intuition.

Second, I compare the optimal threshold of the searcher in the main model with imperfectly correlated preferences to the benchmark where the agents’ preferences are perfectly aligned. I first show that in the main model the searcher’s threshold is always lower than in the benchmark: the searcher is “less picky”. The reason is that the chooser chooses according to her preferences not the searcher’s, which lowers the latter’s continuation value, thus, his threshold. I then provide an example where, as a result of a mean-preserving spread, the searcher’s threshold decreases in the main model, while it always increases in the benchmark: the searcher is “more conservative” in the main model.¹ A mean-preserving spread increases the probability that an item arrives that yields very high utility to the chooser, which restricts the searcher and lowers his continuation value.

Third, I explain how the model’s characteristics differ from those of two single-agent search models that feature a threshold that varies in time. The first model has convex search costs and the second, a deadline. Both models result in an optimal threshold that decreases over time (for a fixed searcher-preferred item in the choice set) because they assume non-stationarity of the environment. My model features a time-decreasing threshold in a stationary environment. Also, in those models returning to an item found earlier is possible, while search always stops with the item found last in my model. The searcher returns to an item found earlier in the models with convex search costs or a deadline because his threshold decreases exogenously over time. In my model, the

¹The terms “less picky” and “more conservative” are borrowed from Albrecht et al. (2010).

decrease is endogenous: it happens only if there is a change in the item that would be chosen if the searcher stopped. I suggest three tests on data that allow us to reject one or more of the three models.

Finally, I extend the model to allow the searcher to hide items. The searcher can hide an item only upon its arrival and succeeds with some given probability that is strictly less than one. He observes whether he succeeded before making the decision whether to stop or continue. As in the main model, I find that the searcher's threshold is unambiguously decreasing in the value that the chooser receives from the chooser-preferred element in the choice set. For independent uniform utilities, I show that the threshold is strictly increasing in the probability that the searcher succeeds in hiding. The constraint of having to account for the chooser's preferences becomes the less restrictive the likelier that the searcher can ignore those preferences.

Related literature. This paper is closely related in spirit to other papers on multi-agent search. In “committee” search problems the committee has a common arrival process and must agree on when to stop.² In “couple” search problems each person has his own arrival process, but they pool income.³ In these papers some part of the entire search process is joint, while distinct parties are engaged in distinct stages of the process in my paper. I borrow the terms “less picky” and “more conservative” in their specific meaning from Albrecht et al. (2010) (AAV henceforth). AAV analyse a committee search problem, where M members of an N -member committee must agree in order to stop the search process. They find that a committee is both less picky and more conservative than a single searching agent. I find that the searcher is both less picky and more conservative if his and the chooser's preferences are misaligned as opposed to when they are perfectly aligned. These results echo those of AAV, but the reasons behind the results are somewhat different, as I explain in Section 5.

The paper is also related to the literature on delegated choice in a principal-agent set-up.⁴ In these papers an agent makes the final choice and the principal either restricts the set of items that the agent can choose from or designs a contract. The available set of items is assumed to be given so the problem is static, unlike in this paper.

Finally, the paper is related to a recent literature on delegated search, where an agent conducts search on behalf of the principal in a principal-agent set-up, by Postl (2004), Armstrong and Vickers (2008), Lewis (2012), Ulbricht (2016), and Kováč et al.

²For example, see Albrecht et al. (2010), Compte and Jehiel (2010), Bergemann and Välimäki (2011), Kamada and Muto (2015), and Moldovanu and Shi (2013).

³For example, see Dey and Flinn (2008), Ek and Holmlund (2010), Flabbi and Mabli (2012), and Guler et al. (2012).

⁴For example, see Holmström (1977), Armstrong (1995), Alonso and Matouschek (2008), Armstrong and Vickers (2010), Amador and Bagwell (2013), and Kováč and Krähmer (2015).

(2014). These models study how the principal can direct the agent to conduct search in the best possible manner for her, whereas here I focus on the “agent’s” optimal policy in the case when the “principal” cannot affect the search process directly. In Postl (2004), Lewis (2012), and Ulbricht (2016) utility is transferable, the agent does not receive direct utility from a chosen item, and the focus is on the principal’s optimal contracts. In the other two papers, which are the most closely related to mine, utility is non-transferable. I point out how the modelling choices in Armstrong and Vickers (2008) (AV henceforth) and Kováč et al. (2014) (KKT henceforth) result in a simpler optimal policy for the agent in AV and KKT than in my paper. In AV, the working paper version of Armstrong and Vickers (2010), the agent makes the final choice and the principal permits the choice among a subset of all items. In an extension to the main model of AV, the agent collects the items via costly search. The most important aspect of my paper that sets it apart from AV is that the final choice is made by the “principal”, whereas in AV it is made by the agent. In AV the agent’s threshold is constant in time. In KKT the agent’s preferences differ cardinally, but not ordinally, from the principal’s and the final choice is made by the principal. They solve for the principal’s optimal mechanism. The most important aspect of my paper that sets it apart from KKT is that the agent and principal do not necessarily agree on which item is the best, whereas they do in KKT. The latter assumption together with assuming that the agent wants to stop with any item trivialises the unrestricted agent’s optimal stopping rule in KKT, unlike in this paper.

The next section contains the details of the model. Section 3 introduces the benchmarks: the searcher’s optimal policy when the chooser’s preferences are either perfectly aligned or opposed to his. Section 4 solves for the searcher’s optimal policy when the agents’ preferences are arbitrarily imperfectly correlated. Section 5 characterises the solution. The final section concludes. All omitted proofs are in the Appendix.

2 Model

A “chooser” (“she”) has to choose an item. The chooser makes the final choice from a choice set, but the choice set is compiled by someone else, a “searcher” (“he”). The final choice determines the agents’ payoffs. The searcher compiles the choice set over time via sequential search. He chooses when to stop and take the accumulated choice set to the chooser.

The arrival process. Time, t , is discrete and $t = 1, 2, \dots, \infty$. The searcher uncovers one new item in each period. He cannot affect how frequently or which items arrive.

Search costs $c > 0$ per period for the searcher.⁵

Preferences. The t 'th item that arrives, $x_t = (u_t, v_t)$, is worth u_t to the searcher and v_t to the chooser. The x_t are independent draws from a time-invariant distribution $H(u, v)$ with support $[0, 1]^2$. The marginal distributions of u and v are $F_u(\cdot)$ and $F_v(\cdot)$ respectively. The conditional distribution of u is $G(u|v)$. The associated pdfs are h , f_u , f_v and g , with $f_u(u) > 0$ for all $u \in [0, 1]$. I assume that the cost of search is relatively low, $c < \mathbb{E}[u]$, so that at least one period of search is desirable for the searcher. In the main part of the paper, Section 4, the joint distribution H has full support. In the benchmarks of Section 3, the preferences are either perfectly aligned (i.e., $P(v = u) = 1$ and $u \sim F_u$) or opposed (i.e., $P(v = 1 - u) = 1$ and $u \sim F_u$). The chooser has an outside option $\bar{v} \in [0, 1]$.

Actions. In period t the searcher chooses an action $a_t \in \{S, C\}$. If the searcher continues ($a_t = C$), the chooser does not get to act in period t . If the searcher stops ($a_t = S$), he takes the accumulated choice set to the chooser. He cannot hide or lie about items.⁶ Let the v -maximal element in a choice set be denoted $x_m = (u_m, v_m)$. From any non-empty choice set brought to the chooser, she chooses x_m if $v_m \geq \bar{v}$, which yields utility v_m to the chooser and u_m to the searcher. If $v_m < \bar{v}$, the chooser chooses her outside option, which yields utility \bar{v} to her and zero utility to the searcher. The entire process ends after the choice of an item or the outside option.

Timing. In each period, first, the searcher uncovers an item and pays the search cost. Second, he chooses whether to stop or continue to search. If the searcher stops searching, the chooser chooses an item from the choice set or her outside option, after which utilities are realised and the entire process ends. If the searcher continues search, the process moves to the next period.

Problem. The searcher's problem is to maximise his expected utility from the search process, taking as given the chooser's outside option and her choice rule. The searcher refuses to start the search process if his expected utility is negative.

3 Two benchmarks

In this section, I derive the searcher's optimal policy in two benchmark versions of the model: in the first benchmark, the chooser's preferences are perfectly aligned with the searcher's, and in the second, her preferences are perfectly opposed to the searcher's.

⁵The main results of the paper are not sensitive to the assumption of an additive search cost as opposed to discounting. I will comment on the result that does change in Section 4.

⁶The assumption of no hiding is relaxed in Section 6.

3.1 Perfectly aligned preferences

Suppose that the chooser's preferences are perfectly aligned with the searcher's, i.e., $v = u$ with probability one and $u \sim F_u$. If the chooser's outside option is zero, the searcher's problem is equivalent to the standard single-agent sequential search problem (see McCall (1970) for the seminal contribution). In the standard search model, the searcher's optimal policy is a threshold rule and the optimal threshold is equal to the searcher's expected value from starting the search process (if the value is positive).

The presence of the outside option changes the analysis, but not considerably. The searcher's optimal policy is a threshold rule because of the standard argument: if it is optimal for the searcher to stop with some item which yields him u , then it is optimal for him to stop with any item that yields him $u' > u$. If the u -value of an item that he finds exceeds the threshold, he stops and the chooser chooses this item; the searcher continues otherwise. The threshold does not depend on time as the problem is stationary. Since the chooser has an outside option worth \bar{v} , the searcher must wait for an item that exceeds the outside option to receive non-zero utility from the search process. The searcher uses one of two possible thresholds depending on the size of the chooser's outside option. First, if the outside option is low, the threshold equals \tilde{u} , the optimal threshold in the standard single-agent sequential search problem. This threshold is optimal if $\bar{v} \leq \tilde{u}$. McCall (1970) shows that \tilde{u} solves

$$\int_{\tilde{u}}^1 u - \tilde{u} dF_u(u) = c.$$

The threshold decreases in the cost of search, in line with intuition.

Second, if the chooser's outside option \bar{v} is higher than \tilde{u} , the searcher is restricted by \bar{v} in the sense that in the absence of it he would stop with items with $u \in [\tilde{u}, \bar{v})$, but has to continue in the presence of \bar{v} . Then the searcher optimally searches until he finds the first item that exceeds the chooser's outside option (if his expected payoff from search is positive). His expected payoff from the search process, $U_a(\cdot)$, is

$$U_a(\bar{v} > \tilde{u}) = P(u \geq \bar{v})\mathbb{E}[u|u \geq \bar{v}] + P(u < \bar{v})U_a(\bar{v} > \tilde{u}) - c. \quad (1)$$

The first term on the right-hand side (RHS) accounts for the possibility that the first item's value exceeds the outside option, in which case the searcher stops and receives u . If the value of the first item is lower than the outside option, he continues and his expected continuation value is the same as at the start of today. The searcher also pays the search cost c . The expected payoff decreases in the cost of search and the chooser's outside option. If the outside option is very high, i.e., $\bar{v} > v_a^*$ where v_a^*

solves $U_a(v_a^* > \tilde{u}) = 0$, the searcher refuses to search because his expected payoff from searching is negative.

3.2 Perfectly opposed preferences

The chooser's preferences are perfectly opposed with the searcher's if $v = 1 - u$ with probability one and $u \sim F_u$. I argue that the searcher's optimal threshold is zero. That is, he stops with any first item acceptable to the chooser, if he searches at all.

The searcher knows that the chooser chooses according to v . As a result, he optimally stops after uncovering the first item which satisfies $v \geq \bar{v}$ if he searches at all. The reason is as follows. Without loss of generality, suppose that $v_1 \geq \bar{v}$. Then for any realisation of u_1, x_2 satisfies one of the following. Either $v_2 \geq v_1$ (so that $u_2 \leq u_1$) or $v_2 < v_1$ (so that $u_2 > u_1$). If the searcher stopped after observing the second item, in the former case the chooser would choose x_2 and in the latter case x_1 . Since $u_2 \leq u_1$ in the former case and the searcher would end up getting x_1 at an extra search cost in the latter case, he prefers stopping after x_1 for any realisation of x_2 .

The searcher's expected payoff from stopping with the first item which satisfies $v \geq \bar{v}$, $U_o(\cdot)$, is

$$U_o(\bar{v}) = P(v \geq \bar{v})\mathbb{E}[1 - v | v \geq \bar{v}] + P(v < \bar{v})U_o(\bar{v}) - c.$$

The equation is interpreted analogously to equation (1), except that here the searcher receives $u = 1 - v$ when he stops. The searcher's expected payoff is decreasing in the cost of search and the chooser's outside option. The expected payoff from searching is positive if $\bar{v} < v_o^*$ where v_o^* solves $U_o(v_o^*) = 0$. If the expected payoff from searching is negative, the searcher refuses to search and receives zero.

The results of the first two sections are summarised in the following:

Lemma 1. (i) *If the agents' preferences are perfectly aligned, the searcher's optimal policy is a threshold rule and the threshold is $\max\{\tilde{u}, \bar{v}\}$, where \tilde{u} solves*

$$\int_{\tilde{u}}^1 u - \tilde{u} dF_u(u) = c.$$

(ii) *If the agents' preferences are perfectly opposed, the searcher's optimal policy is a threshold rule and the threshold is zero.*

4 Imperfectly correlated preferences

In this section the searcher's preferences, u , are imperfectly correlated with the chooser's preferences, v . The only assumption I make is that their joint distribution $H(u, v)$ has full support. I derive the first main result of this paper: the searcher's optimal policy is a threshold rule, the threshold is $\bar{u} = \max\{0, \bar{u}(v_m)\}$, and decreases in v_m .

The searcher's optimal policy is a threshold rule because of the standard argument: if it is optimal for the searcher to stop with some item $x = (u, v)$ (which yields him u), then it is optimal for him to stop with any $x' = (u', v)$ with $u' > u$. The searcher stops if the utility that he receives from the item that would be chosen by the chooser if the searcher was to stop immediately (i.e., the v -maximal item in the choice set) exceeds the threshold, and continues otherwise. The optimal threshold $\bar{u}(\cdot)$ is equal to the searcher's value from continuing, $U(\cdot)$, if the value is positive as in the standard search problem (and zero if the value is negative).

The threshold differs from the standard search problem's in several ways. Recall that $x_m = (u_m, v_m)$ is the v -maximal item in the choice set. First, the threshold is defined only if the choice set contains at least one item with $v \geq \bar{v}$, i.e., if $v_m \geq \bar{v}$. If no such item exists in the choice set, the searcher optimally continues for all u_m (if it was optimal for him to start the search process). Second, the positive part of the threshold $\bar{u}(\cdot)$, or, equivalently, the searcher's continuation value $U(\cdot)$, depends on v_m if $v_m > \bar{v}$, but not on u_m . Suppose that the v -maximal item in the choice set is $x_m = (u', v')$ with $v' = v_m > \bar{v}$. If the searcher stops, x_m is chosen and he receives u' . If he is not satisfied with u' , i.e., if $u' < \bar{u}(\cdot)$, he continues. I argue that his continuation value depends on v' , but not on u' . The first part is simple: the continuation value depends on v' because a new item $x'' = (u'', v'')$ is chosen only if $v'' > v'$. The second part requires considering two scenarios. Suppose a new item (u'', v'') arrives. If $v'' > v'$, then v_m changes to v'' : the new item would be chosen if the searcher stopped. Hence, u' is irrelevant. If $v'' \leq v'$, then v_m stays v' . But then nothing has changed as compared to the previous period (when the searcher continued) so that $u' < \bar{u}(\cdot)$ must still hold. In neither of the cases does the searcher's continuation value depend on u' . Hence, the expected value of continuing depends only on v_m so I write $U(v_m)$ and $\bar{u}(v_m)$. Third, the threshold is weakly positive for $v_m \in [\bar{v}, v^*]$. The lower bound \bar{v} is explained above. The upper bound v^* is present because v_m acts as a restriction: if v_m is very high, finding an item with $v > v_m$ becomes so unlikely that the searcher's continuation value becomes negative and he is better off accepting any positive u_m . The upper bound is defined as $\bar{u}(v_m = v^*) = 0$ and $v^* \in [0, 1]$ is guaranteed to exist for all $c > 0$ because of the full support assumption.

I solve for the searcher's continuation value for various values that v_m can take. Suppose first that $v_m \leq \bar{v}$. Then the searcher has no item in the choice set that is acceptable to the chooser. The searcher's expected value from continuing (equivalently, starting) the search process is

$$U(\bar{v}) = \int_{\bar{v}}^1 \int_0^1 \max\{u, U(v)\} h(u, v) \, du \, dv + \int_0^{\bar{v}} U(\bar{v}) f_v(v) \, dv - c, \quad (2)$$

if the next item that he finds is (u, v) . The first term on the RHS accounts for the possibility that v exceeds the chooser's outside option: the item would be chosen if the searcher stopped. The searcher chooses optimally whether to stop or continue, where the continuation value is now a function of v . The second term on the RHS accounts for the possibility that v is below the chooser's outside option: the searcher continues and the continuation value is the same as at the start of today. The last term accounts for the search cost. As $U(\bar{v})$ is also the searcher's expected payoff from the entire search process, he optimally starts search if and only if $U(\bar{v}) \geq 0$.⁷

Suppose now that $v_m > \bar{v}$. The searcher's expected value from continuing is

$$U(v_m) = \int_{v_m}^1 \int_0^1 \max\{u, U(v)\} h(u, v) \, du \, dv + \int_0^{v_m} U(v_m) f_v(v) \, dv - c, \quad (3)$$

which is interpreted analogously to (2), except that \bar{v} in (2) is replaced by v_m in (3).

As argued above, the searcher's optimal threshold $\bar{u}(\cdot)$ is equal to his continuation value $U(\cdot)$ if the latter is positive, which is true for $v_m \in [\bar{v}, v^*]$ for some $v^* \in [0, 1]$ that satisfies $U(v^*) = 0$. The optimal threshold is zero if the continuation value is negative, i.e., if $v_m > v^*$. Thus, by substituting $U(v_m) = \bar{u}(v_m)$ into equation (3) and using the fact that the agent continues for $x = (u, v)$ s.t. $v > v_m$ and $u < \bar{u}(v)$, I can rewrite (3) for $v_m \in [\bar{v}, v^*]$ as

$$\bar{u}(v_m) = [1 - F_v(v_m)]^{-1} \int_{v_m}^{v^*} \left[\int_0^{\bar{u}(v)} \bar{u}(v) h(u, v) \, du + \int_{\bar{u}(v)}^1 u h(u, v) \, du \right] \, dv, \quad (4)$$

where v^* satisfies $\int_0^1 \int_{v^*}^1 u h(u, v) \, dv \, du = c$. As search becomes more costly, i.e., c increases, v^* decreases. Intuitively, if search becomes more costly, it becomes unprofitable for the searcher to continue searching at lower values of v_m . Note that the function $\bar{u}(v_m)$ does not depend on \bar{v} : the chooser's outside option affects the searcher's expected payoff from the entire search process, but does not affect his optimal policy for $v_m > \bar{v}$.

⁷If I assume discounting instead of an additive search cost, then the searcher optimally starts search for all $\bar{v} < 1$ as $U(\bar{v}) > 0$ for all positive discount factors. This is the only result that qualitatively changes in a model with discounting.

The above integral equation can be converted into an ODE by differentiation:

$$\frac{\partial \bar{u}(v_m)}{\partial v_m} = -\frac{f_v(v_m)}{1 - F_v(v_m)} \int_{\bar{u}(v_m)}^1 u - \bar{u}(v_m) dG(u|v_m). \quad (5)$$

On the RHS of the equation, the first term is the hazard rate, i.e., the probability that an item with $v = v_m$ arrives, given that such an item has not arrived earlier. The integral term is the utility that the searcher expects to get from the new item, u , in excess of the continuation value $\bar{u}(v_m)$, which is realised only if u exceeds the continuation value. The slope of the threshold depends on the search cost c indirectly via the level of the threshold $\bar{u}(v_m)$, because the level depends on the search cost via v^* (see equation (4)).

Equation (5) does not, in general, have a closed form solution (I provide a closed form solution for an example below; further examples can be provided under the assumption that u and v are independent). However, equation (5) together with the initial condition $\bar{u}(v^*) = 0$ fully pin down the function $\bar{u}(v_m)$, which proves the first main result of this paper.

Proposition 1. *If the agents' preferences are imperfectly correlated, the searcher's optimal policy is a threshold rule. The threshold is $\bar{u} = \max\{0, \bar{u}(v_m)\}$, where $\bar{u}(v_m)$ solves*

$$\frac{\partial \bar{u}(v_m)}{\partial v_m} = -\frac{f_v(v_m)}{1 - F_v(v_m)} \int_{\bar{u}(v_m)}^1 u - \bar{u}(v_m) dG(u|v_m),$$

for $v_m \in [\bar{v}, v^*]$, where v^* satisfies $\bar{u}(v^*) = 0$.

The threshold $\bar{u}(v_m)$ is clearly decreasing in v_m for any joint distribution of u and v with full support. The chooser's value from the v -maximal item in the choice set, v_m , acts as a restriction on the searcher's problem because a new item $x' = (u', v')$ is chosen only if $v' > v_m$. The stricter the restriction, the lower the searcher's expected payoff from the process. I provide sufficient conditions for the threshold $\bar{u}(v_m)$ to be concave in the Appendix.

The fact that the threshold $\bar{u}(v_m)$ decreases in v_m means that the searcher does not start searching for high enough outside options for the chooser, \bar{v} . The chooser's outside option acts as a similar restriction on the searcher's problem as v_m . If the outside option is very high, i.e., $\bar{v} > v^*$, the searcher prefers to receive payoff zero to starting the process and making a loss in expectation. The critical outside option above which the searcher prefers not to start searching is lower if the agents' preferences are misaligned as opposed when they are perfectly aligned ($v^* < v_a^*$), in line with intuition.

Example (Analytic solution). Suppose the utilities are independent and h uniform on $[0, 1]^2$. The searcher's expected value from continuing, equation (4), simplifies to

$$\bar{u}(v_m) = \frac{1}{2}(1 - v_m)^{-1} \int_{v_m}^{v^*} 1 + \bar{u}(v)^2 \, dv,$$

and the associated ODE (5) to

$$\frac{\partial \bar{u}(v_m)}{\partial v_m} = -\frac{1}{2}(1 - v_m)^{-1}(1 - \bar{u}(v_m))^2. \quad (6)$$

In the Appendix I solve equation (6) in detail using a standard method and the initial condition $\bar{u}(v^*) = 0$. The explicit solution is

$$\bar{u}(v_m) = 1 - \left[\frac{1}{2} \ln \left(\frac{1 - v_m}{1 - v^*} \right) + 1 \right]^{-1}, \quad (7)$$

where the value v^* has a simple closed form: $v^* = 1 - 2c$. The conditions that guarantee that $\bar{u}(v_m)$ is concave in v_m are satisfied. An example of the threshold as a function of v_m is depicted in Figure 1.

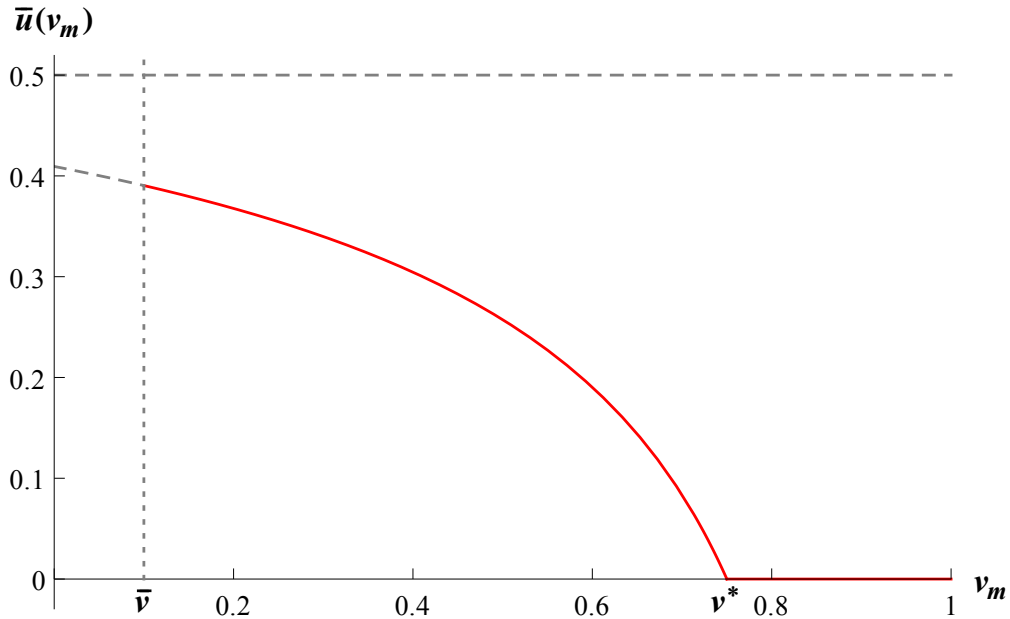


Figure 1: Searcher's threshold as a function of v_m if utilities are independent and uniform for $\bar{v} = \frac{1}{10}$ and $c = \frac{1}{8}$. Dashed line: \tilde{u} .

5 Characterisation

In this section, I first prove that the threshold policy exhibits the “discouragement effect” and derive testable implications of my model (Sections 5.1 and 5.2). I then show that, as compared to when the agents’ preferences are perfectly aligned, if the preferences are misaligned, the searcher is “less picky” and “more conservative” (Sections 5.3 and 5.4).⁸ These two results echo the results of AAV, but the mechanisms are slightly different as I explain below. I then provide a numerical example where a higher correlation of preferences results in an unambiguously higher threshold. Finally, I show how the implications of my model differ from two models that extend the standard single-agent search model and generate a threshold that varies over time. Omitted details are in the Appendix.

5.1 The threshold exhibits the discouragement effect

I present the second main result of the paper, that the searcher’s threshold exhibits the discouragement effect.

Definition (Discouragement effect). *A threshold policy exhibits the discouragement effect if the threshold weakly decreases in time.*

I call a time-decreasing threshold the discouragement effect because the longer the searcher searches, the lower-utility items he is willing to accept and the more likely he is to stop: he is discouraged from searching longer and for better items as time goes on.

Proposition 2. *The threshold policy $\bar{u}(v_m)$ exhibits the discouragement effect.*

Proof. The threshold $\bar{u}(v_m)$ is weakly decreasing in time if and only if v_m is weakly increasing in time because $\frac{\partial \bar{u}(v_m)}{\partial v_m} < 0$ (by Proposition 1). Since v_m is the maximum utility that the chooser gets from a choice set collected up to some time t , v_m is formally the t th (or largest) order statistic of the choice set at time t : $v_m = \max\{v_1, v_2, \dots, v_t\}$. But the t th order statistic must weakly increase in t : for any v_m at date t , either $v_{t+1} \leq v_m$, in which case v_m is left unchanged, or $v_{t+1} > v_m$, in which case v_m takes on the new, higher value v_{t+1} . \square

Proposition 2 states that the searcher’s threshold stochastically weakly decreases over time, despite the search environment being stationary. This is in contrast to the standard sequential search model where a stationary environment translates into a stationary threshold. The source of the time-decreasing threshold in my model is

⁸The terms “less picky” and “more conservative” are borrowed from AAV.

the endogenous restriction on the searcher’s problem that becomes more stringent over time, i.e., the v -value of the v -maximal element in the choice set, v_m , increases in time. If search has gone on for long enough, v_m is high and the searcher optimally accepts a newly arriving item (u, v) with $v > v_m$ regardless of how low u is. Suppose that the HR manager in the example in the Introduction cares about the amiability of a future colleague and the boss about the new worker’s qualifications. The discouragement effect means that if the HR manager receives the application of a highly qualified worker who does not seem like an amiable colleague, he stops looking for other applicants since he anticipates that the boss hires the highly qualified applicant.

5.2 Testable implications of the discouragement effect

I describe the testable implications of my model and especially the discouragement effect. I assume that we have data on multiple search instances involving the same (or a group of representative) searcher(s) and the same chooser. Each search instance is identified with an observation i . I assume that each observation contains information on the duration of search (denoted D_i), the identity of item that is the final choice (denoted M_i), and on the searcher’s and chooser’s utility from the finally chosen item $x_{M_i} = (u_{M_i}, v_{M_i})$. Data on utilities is unlikely to occur in field settings, but can be generated in a laboratory experiment. The model has three testable implications.

1. u_{M_i} and v_{M_i} are negatively correlated across i . This is a direct implication of the negatively-sloped threshold. If at any point during the search process v_m is low, the searcher only stops if u_m is high. Conversely, if v_m is high, then a low u_m is sufficient for the searcher to stop. Across many instances, the utilities from the finally chosen item should, thus, be negatively correlated.
2. D_i and u_{M_i} (v_{M_i}) are negatively (positively) correlated across i . The longer the searcher searches, the more likely is v_m to be high. But when v_m is high, the searcher accepts items with lower u -values as compared to when v_m is low. Across many instances, the search duration and the searcher’s utility from the finally chosen item should, thus, be negatively correlated.
3. $D_i = 1$ for some observations i . The searcher stops after any first item x_1 if $v_1 > v^*$. In this case, the utility that he gets from the item, u_1 , is irrelevant for his stopping problem and he stops immediately.

I explain in Section 5.6 how my model’s implications differ from the implications of two other models with a time-varying threshold.

5.3 The searcher is less picky if the preferences are misaligned

I borrow the terminology from AAV and say that a searcher with a lower threshold is “less picky” than a searcher with a higher threshold.

Proposition 3. *The searcher is less picky if the agents’ preferences are misaligned as compared to when they are perfectly aligned, i.e., $\max\{\tilde{u}, \bar{v}\} > \bar{u}(v_m)$ for all v_m .*

Proof. In the Appendix. □

The reason why the searcher whose preferences differ from the chooser’s is willing to accept a lower utility item is that his search process is restricted by the chooser’s preferences: an item that the searcher “likes” is chosen only if the chooser “likes” it, too. If their preferences are not perfectly aligned, the agents “like” the same item with a probability less than one. This reduces the searcher’s value from continuing with the search process, hence, his optimal stopping threshold. In terms of the hiring example, if amiable people have higher qualifications, then the HR manager is optimally satisfied with a more amiable new colleague as opposed to when amiability and qualifications are not perfectly correlated, despite the HR manager’s preferences not changing.

In AAV, a committee is less picky than a single searcher because the committee members need to compromise. In my model, the searcher is less picky if his and the chooser’s preferences are misaligned because choice is made according to the chooser’s not the searcher’s preferences.

5.4 The searcher is more conservative if the preferences are misaligned

I say that the searcher is “more conservative” if his threshold may decrease as a result of a mean-preserving spread (MPS) to the distribution of utilities, following AAV. In order to make the comparison between the benchmark and full model and give an unambiguous meaning to a MPS, I assume in this section that the marginal distributions of u and v equal: $F_u = F_v =: F$. I show by example that the searcher is more conservative if his and the chooser’s preferences are misaligned as compared to when they are perfectly aligned.

Proposition 4. *If $F_u = F_v$ and the joint distribution of the utilities is subjected to a mean-preserving spread, the searcher’s threshold always increases when the preferences are perfectly aligned, but for a certain set of parameter values the threshold decreases when the agents’ preferences are misaligned.*

Proof. (a) If the joint distribution of the utilities is subjected to a mean-preserving spread, the searcher's threshold \tilde{u} always increases when the preferences are perfectly aligned. For proof, see e.g., AAV.

(b) If the joint distribution of the utilities is subjected to a mean-preserving spread, the searcher's threshold may decrease when the agents' preferences are misaligned. By example: I provide an example where a MPS leads to a decrease in $\bar{u}(v_m)$. Assume that u and v are independent and have the same uniform marginal $\mathcal{U}[a, b]$ for $0 < a < b$. An example of a MPS for the distribution is $\mathcal{U}[a', b']$ such that $0 < a' < a$, $b' > b$ and $a + b = a' + b'$.

The threshold of the searcher satisfies the ODE

$$\frac{\partial \bar{u}(v_m)}{\partial v_m} = -\frac{(b - \bar{u}(v_m))^2}{2(b - a)(b - v_m)}.$$

The searcher prefers continuing to accepting any item if his value from continuing is higher, i.e., if $U(v_m) \geq a$. Define v^* by $U(v^*) = a$, or $\bar{u}(v^*) = a$. The closed form is $v^* = b - 2c$. Then I can solve the ODE by using similar methods as in the proof of Claim 1 (in the Appendix). The result is

$$\bar{u}(v_m) = \frac{2(b + a) - 2b - b \ln\left[\frac{b - v^*}{b - v_m}\right]}{2 - \ln\left[\frac{b - v^*}{b - v_m}\right]}.$$

Differentiating the expression with respect to b while keeping $a + b$ constant gives

$$\frac{\partial \bar{u}(v_m)}{\partial b} = 4v_m - 2(a + b) + (b - v_m) \left(\ln \left[\frac{b - v^*}{b - v_m} \right] \right)^2.$$

This is negative, for example, if $b = 3$, $a = 1$, $v_m = 1.5$ and $c = \frac{1}{4}$.

□

The reason why a MPS leads to an unambiguous increase in the searcher's optimal threshold in a single-agent search problem is that a MPS increases the option value of searching by making really high (and really low) draws possible. The gain from the really high draws outweighs the loss from the really low draws because the latter are not accepted.

Conversely, if the searcher's and chooser's preferences are misaligned, the searcher's optimal threshold may decrease under a MPS: the searcher behaves more conservatively under more risk by accepting lower-utility items. Two counteracting effects lie behind the result. On the one hand, a MPS of the u -value distribution benefits the searcher

through the same mechanism as in the single-agent setup: the option value of searching increases and the threshold rises. On the other, a MPS increases the probability of items with high v -value occurring. A high v -value acts as a stricter constraint for the searcher, thus decreasing his value of continuing and his threshold. In the example in the proof of Proposition 4, the negative effect outweighs the positive.

In AAV, a committee is more conservative than a single searcher. The result emanates from one committee member exerting a negative externality on another, e.g., in a unanimity committee a member can veto stopping in a case where everyone but the vetoing member would receive a high utility from the last item. The externality can become more severe under a MPS. Here, the result emanates from the fact that under a MPS, the searcher is more likely to find an item with a v -value that restricts him more.

5.5 The searcher prefers more aligned preferences

I use a simple parametric family of joint distributions to demonstrate numerically that the searcher is unambiguously better off as the agents' preferences become more aligned. Analytic results are not available even for this simple family. Let the marginals of u and v be uniform on $[0, 1]$ throughout. The family of distributions deals with positive and negative correlation separately. For positive correlation, let the correlation be governed by parameter $q \in [0, 1]$. The conditional distribution of u given v is

$$u|v = \begin{cases} v & \text{with probability } q, \\ \sim \mathcal{U}[0, 1], u \perp\!\!\!\perp v & \text{---} \\ & 1 - q. \end{cases} \quad (8)$$

The searcher's optimal threshold is the solution to a system of two differential equation that have the form

$$\frac{\partial \bar{u}(v_m)}{\partial v_m} = -\frac{1}{2}(1 - v_m)^{-1}[A(1 - \bar{u}(v_m))^2 + B(v_m - \bar{u}(v_m))],$$

where A and B depend only on q . I provide the exact differential equations in the Appendix, but omit them here as they do not possess closed form solutions that could be interpreted.

For negative correlation, let the correlation be governed by parameter $r \in [0, 1]$. The conditional distribution of u given v is

$$u|v = \begin{cases} 1 - v & \text{with probability } r, \\ \sim \mathcal{U}[0, 1], u \perp\!\!\!\perp v & \text{---} \\ & 1 - r. \end{cases} \quad (9)$$

The searcher's optimal threshold is the solution to a system of at most two differential

equations, with a general form given by

$$\frac{\partial \bar{u}(v_m)}{\partial v_m} = -\frac{1}{2}(1 - v_m)^{-1}[A(1 - \bar{u}(v_m))^2 + B(1 - v_m - \bar{u}(v_m))],$$

where A and B depend only on r . Again, I provide the exact differential equations in the Appendix.

The positive and negative correlation parameters q and r can be comprised in a single parameter ρ : $\rho = q$ for $q \geq 0$ and $r = 0$, and $\rho = -r$ for $r \geq 0$ and $q = 0$. I show numerically that $\bar{u}(v_m)$ increases in ρ . The result is illustrated in Figure 2 (a lower curve corresponds to a lower level of correlation). The dashed grey line corresponds to $\tilde{u} = 1 - \sqrt{2c}$. The intuition behind the result is simple: if the utilities become more correlated, then the searcher is better off as his continuation value, thus, his threshold, increases. An increase in ρ affects the threshold via several channels. First, a higher ρ directly increases the searcher's value from stopping because a high v is more likely to be accompanied by a high u . Second, a higher ρ indirectly increases the searcher's continuation value both through the decreased likelihood that a high v is restrictive and through the increased future value of stopping.

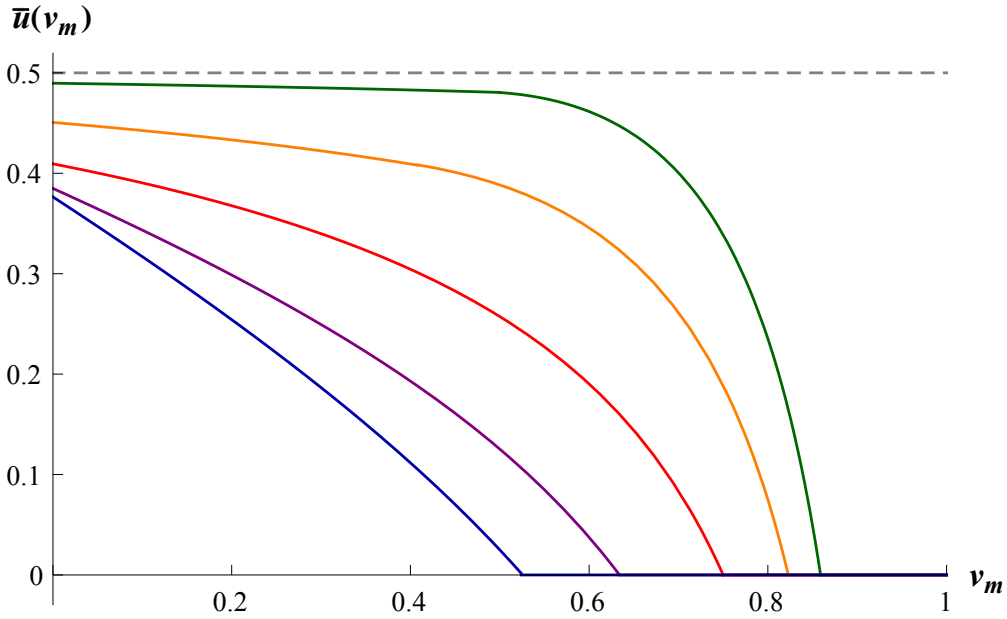


Figure 2: Searcher's threshold $\bar{u}(v_m)$ at correlation levels $\rho = -0.9, -0.5, 0, 0.5, 0.9$ (from lowest to highest curve respectively; dashed line: \tilde{u}), for $\bar{v} = 0$, $c = \frac{1}{8}$.

5.6 Differences with other models with time-varying thresholds

Here, I first explain how the characteristics of my model differ from two single-agent search models that feature a time-varying threshold. I then explain how data allows us to test between the three models. Recall that in my model the searcher's threshold decreases (weakly) over time because v_m increases (weakly) over time. Extensions to the standard single-agent search model that produce a time-varying threshold are, for example, a finite horizon, i.e., a deadline, (see Gronau (1971) and Lippman and McCall (1976)) or convex search costs (see Stiglitz (1987)). For a fixed u -value of the u -maximal item in a choice set (denoted u_w), the threshold in these models is decreasing over time.

First, I compare the reasons behind the decreasing thresholds. The reason behind a decreasing threshold (for a fixed u_w) in all the models is that the searcher's value of continuing decreases over time. However, the deeper reasons differ. In the case of a deadline, the decrease is due to the exogenous end of search possibilities beyond the deadline. In the case of convex search costs, the decrease is due to search becoming exogenously more expensive over time. In fact, in these models the environment is non-stationary to start with (so that a non-stationary optimal policy is expected), contrary to my model.

Second, for a fixed u_w , the time path of the threshold in a model with a deadline or convex search costs is strictly concave in time (for infinitesimally short time periods) whereas it is not concave in time in my model. In a model with a deadline, the threshold is concave in time because for a fixed time increment, the loss of future search opportunities for the searcher becomes larger the closer the deadline. In a model with convex search costs, the threshold is concave in the number of items (equivalently, in the number of time periods). In my model, the threshold is not concave in time because with positive probability, a new item's v -value does not change v_m , in which case the threshold is constant in time, thus, not concave. The time path of the threshold cannot be constant in time in a model with either a deadline or convex search costs.

Third, the models generate different outcomes. In my model, the item that is finally chosen is always the one uncovered last. This is because the threshold decreases only if the v -maximal item changes. If a new item becomes the v -maximal item in the choice set, it changes the threshold $\bar{u}(v_m)$, and the new item may be attractive for the searcher to stop with (if $u_m \geq \bar{u}(v_m)$). Alternatively, the new item does not become the v -maximal item in the choice set and the u -value of the v -maximal item still falls below the threshold. In contrast, there may be return to an item uncovered earlier in a model with either a deadline or convex search costs. This is because in those models

the threshold can decrease independently of the changes in the choice set. An item that yielded too little utility to warrant stopping in the past may exceed the decreased threshold. In sum, in my model the final choice is always the item uncovered last despite the searcher's decreasing threshold.

Finally, I explain how data allows to test between the three models. For this exercise, let us first assume that each observation i in the data corresponds to a search instance and contains information on the duration of the search process, D_i , and the identity of the item that is the final choice, M_i . Suppose we have data on $i \geq 1$ of such search instances. Then the following predictions are made.

1. If in any observation the finally chosen item is not the last item, i.e., $M_i \neq D_i$ for an i , we can reject my model. This prediction is a straightforward implication of the last of the three differences between the three models that I discussed above.
2. If the deadline in the deadline model is some known number T_i for each observation and the duration of search exceeds the deadline in any observation, i.e., $D_i > T_i$ for some i , we can reject the deadline model.

If each observation would additionally contain information about the utilities of all the items that the searcher has observed, a further prediction is made.

3. If the finally chosen item does not have the highest utility for the searcher among all the observed items, i.e., if for some i , $u_{M_i} \neq \operatorname{argmax}_{x \in C_{D_i}} u(x)$, where C_{D_i} is the choice set collected until the process ends, we can reject the deadline and the convex search cost models. In my model, the final choice is made according to the chooser's preferences, v , but is made according to the searcher's preferences, u , in the two other models. Thus, in the two models the final choice must be the u -maximal item, but not in my model.

6 An extension: hiding

Suppose the searcher can hide the items that he wants to at the arrival of the items, but is not always successful.⁹ The searcher takes two actions in any period: $a_{1t} \in \{H, D\}$ and $a_{2t} \in \{S, C\}$, where H stands for trying to hide, D for not trying to hide, S for stopping, and C for continuing. In particular, after an item x_t arrives and the searcher has found out its utilities (u_t, v_t) , he can attempt to hide the item x_t . He succeeds with probability $p \in (0, 1)$. After taking the hiding action and observing its outcome, he chooses whether to stop or continue.

⁹I thank Ludo Visschers for proposing this particular hiding technology.

I derive the searcher's optimal hiding and stopping policy. In sum, the searcher (weakly) prefers hiding all items that he does not want to stop with. He prefers stopping to continuing if his value from stopping exceeds the value from continuing. Let us consider all the possible cases that the searcher may encounter. Suppose that the v -maximal item found until today is $x_m = (u_m, v_m)$, the searcher decided to continue yesterday (so that $u_m < U(v_m)$), and the item found today is $x = (u, v)$. I describe the searcher's optimal policy for all possible values of x_m and x . If $v < v_m$, the searcher continues (with or without attempting to hide x) because x_m is chosen if the searcher stopped (so that his continuation value is definitely $U(v_m)$) and $u_m < U(v_m)$ still holds. If $v \geq v_m$, it must be that $U(v) \leq U(v_m)$ because a higher v_m is a greater restriction for the searcher (I verify later that this property holds). Then the maximum continuation value that the searcher can achieve is, after successfully hiding x , $U(v_m)$. Hence, if $u > U(v_m)$ it is optimal for the searcher to not to hide the item and stop. Suppose that $u < U(v_m)$: the searcher would like to continue if he could guarantee himself the continuation value $U(v_m)$. Hence, the searcher attempts hiding, $a_{1t} = H$, and if he succeeds, continues as $u_m < U(v_m)$. If he fails, he compares u to his continuation value $U(v)$: if $u \geq U(v)$, he stops, and continues otherwise. In sum, the optimal sequence of actions for the agent after receiving $x = (u, v)$ with $v \geq v_m$ is

$$(a_{1t}, a_{2t}) = \begin{cases} (D, S) & \text{if } u \geq U(v_m), \\ (H, C) & \text{if } u < U(v_m) \text{ and hiding succeeds,} \\ (H, S) & \text{if } u \in [U(v), U(v_m)) \text{ and hiding fails,} \\ (H, C) & \text{if } u < U(v) \text{ and hiding fails.} \end{cases} \quad (10)$$

The optimal sequence of actions for the agent after receiving $x = (u, v)$ with $v < v_m$ is

$$(a_{1t}, a_{2t}) = (\{H, D\}, C). \quad (11)$$

Thus, the optimal policy of the searcher can no longer be characterised by a threshold only, but the buyer's continuation value $U(v_m)$ is sufficient to describe his optimal policy.

Note that if v_m is very high, then the searcher's continuation value is negative for the same reason as in the main model. For any $p < 1$, if v_m is very high, the probability that an item arrives with $v > v_m$ (and $u > u_m$) is very low, so that the searcher's benefit from continuing is less than the cost c . As before, let the smallest v_m s.t. $U(v_m) \leq 0$ be v^* .

Formally, the searcher's continuation value when the v -maximal item is (u_m, v_m)

and the newly arrived item is (u, v) satisfies

$$\begin{aligned}
U(v_m) = & \int_0^{v_m} U(v_m) f_v(v) dv + \int_{v_m}^{v^*} \int_0^{U(v)} [pU(v_m) + (1-p)U(v)] h(u, v) du + \quad (12) \\
& + \int_{U(v)}^{U(v_m)} [pU(v_m) + (1-p)u] h(u, v) du + \int_{U(v_m)}^1 uh(u, v) du dv + \\
& + \int_{v^*}^1 \int_0^{U(v_m)} [pU(v_m) + (1-p)u] h(u, v) du + \int_{U(v_m)}^1 uh(u, v) du dv - c,
\end{aligned}$$

where v^* satisfies $U(v^*) = 0$. Each of the terms corresponds to the optimal sequence of actions for the searcher as described in equations (10) and (11). For example, the first term reads that if the new item's v -value v falls below the v -value of the v -maximal item found so far, then the searcher optimally continues to search (as described in equation (11)). The last double integral reads that if $v > v^*$ (so we know that the searcher's continuation value is negative if he fails to hide x), the searcher stops for sure if u exceeds his best possible continuation value $U(v_m)$ (last term). If u falls short of this continuation value, the searcher tries to hide x . If he succeeds, he continues. If he fails, he stops.

In a similar fashion as in the main model, I derive the differential equation that the continuation value satisfies by differentiating (12) with respect to v_m and obtain the ODE

$$\frac{\partial U(v_m)}{\partial v_m} \int_{v_m}^1 \int_0^1 h(u, v) du - p \int_0^{U(v_m)} h(u, v) du dv = - \int_{U(v_m)}^1 h(u, v_m)(u - U(v_m)) du,$$

with terminal condition $U(v^*) = 0$, or $\int_{v^*}^1 \int_0^1 uh(u, v) du dv = c$. Since the RHS of the ODE is negative and the multiplier after $\frac{\partial U(v_m)}{\partial v_m}$ on the LHS positive, the continuation value is unambiguously decreasing in v_m . I explain why v^* is independent of p . The value v^* is defined as the smallest v_m such that $U(v_m = v^*) \leq 0$ so that the searcher wants to stop with any u_m when $v_m = v^*$. Suppose that $v_m = v^*$ and the searcher continues (so that $u_m = 0$). Then the best continuation value that he can hope for is $U(v^*) = 0$. If an item (u, v) arrives with $v > v^*$, the searcher optimally accepts any u because $U(v) < 0$ for all $v > v^*$. But the expected value of u (given that $v > v^*$) is independent of p , hence, v^* is independent of p .

The comparative static derivative of $U(v_m)$ with respect to p cannot be determined without solving for the function $U(v_m)$. I show for an example below that, in line with intuition, the searcher's continuation value is strictly increasing in p for all $v_m < v^*$.

Example (Analytic solution, continued). *Suppose that the utilities are independent and*

h uniform on $[0, 1]^2$. The searcher's expected value from continuing simplifies to

$$U(v_m)(1 - v_m) = \int_{v^*}^1 \int_0^{U(v_m)} [pU(v_m) + (1 - p)u] du + \int_{U(v_m)}^1 u du dv - c + \quad (13)$$

$$+ \int_{v_m}^{v^*} \int_0^{U(v)} [pU(v_m) + (1 - p)U(v)] du + \int_{U(v)}^{U(v_m)} [pU(v_m) + (1 - p)u] du + \int_{U(v_m)}^1 u du dv.$$

The derivative of the above is

$$\frac{\partial U(v_m)}{\partial v_m} = -\frac{1}{2}(1 - v_m)^{-1}(1 - pU(v_m))^{-1}(1 - U(v_m))^2,$$

with terminal condition $U(v^*) = 0$, which yields $v^* = 1 - 2c$. Note that the equation collapses to equation (6) if $p = 0$.

I use the same method to solve the ODE as before, using the initial condition $U(1 - 2c) = 0$. The implicit solution for $U(v_m)$ is

$$2(1 - p)\frac{U(v_m)}{1 - U(v_m)} - 2p \ln(1 - U(v_m)) = \ln\left(\frac{1 - v_m}{1 - v^*}\right).$$

The RHS is constant in $U(\cdot)$ and p so the derivative $\frac{dU(\cdot)}{dp} = -\frac{LHS_p}{LHS_{U(\cdot)}} > 0$ as $LHS_{U(\cdot)} > 0$ and $LHS_p < 0$ (straightforward to verify from above). In line with intuition, the searcher becomes unambiguously better off as his probability of successfully hiding the items that he wants to hide increases: he is less likely to have to account for the chooser's preferences as p increases. The result is illustrated on Figure 3: higher curves stand for higher p .

Allowing the searcher to hide items acts in a similar manner as making the agents' preferences more aligned in the sense that the searcher's expected continuation value increases. However, the parametric examples illustrate that there are differences (see Figure 2 and Figure 3). For a given correlation structure, an increase in p shifts the searcher's continuation value up, without changing the value of v_m above which the continuation value is zero (v^*). Conversely, for a given p , an increase in the correlation parameter shifts both the searcher's continuation value and v^* up. The reason behind the difference is as follows. If the utilities become more positively correlated, then a high v_m is not as great a restriction on the searcher's problem: the probability that an item arrives tomorrow with a v -value that exceeds v_m does not change, but it is more likely that the high v -value is accompanied by a high u -value. Thus, the searcher's continuation value at any given v_m goes up, including at $v_m = v^*$. If, instead, the probability that the searcher successfully hides an item increases, then if the searcher

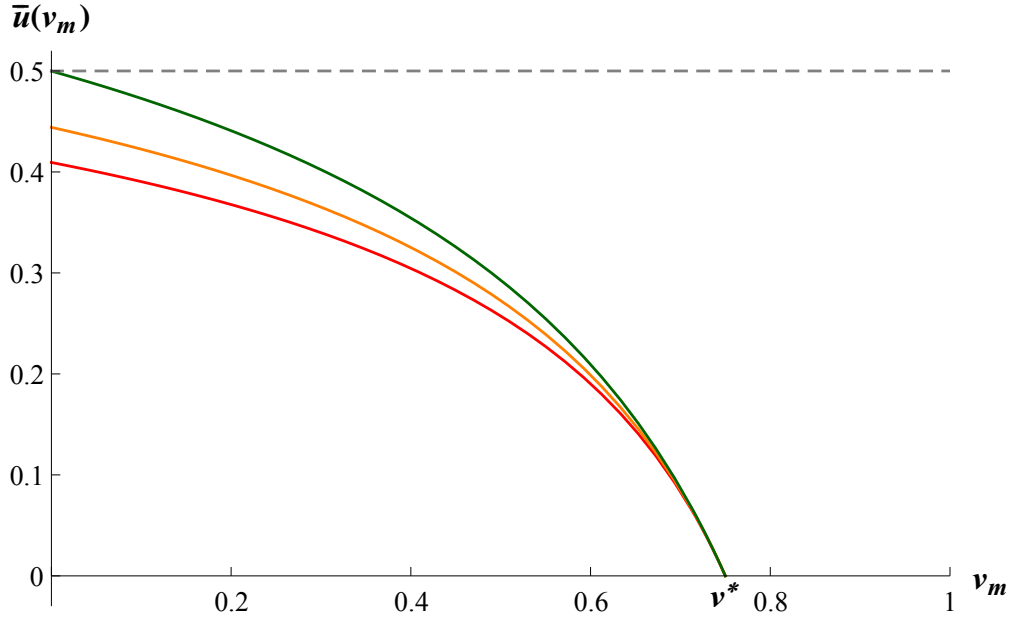


Figure 3: Searcher’s threshold $\bar{u}(v_m)$ at hiding probabilities $p = 0, 0.5, 0.999$ (from lowest to highest curve respectively), for $\bar{v} = 0$, $c = \frac{1}{8}$. Dashed line: \tilde{u} .

fails to hide an item with a high v_m , the high v_m is exactly as great a restriction on the searcher’s continuation problem as in the original model (i.e., where $p = 0$).

7 Conclusion

I study a sequential search problem where the preferences according to which the final item is chosen differ from the preferences according to which search is conducted. A natural interpretation of this set-up is that the preferences belong to separate parties: a searcher and a chooser. Alternatively, the preferences of an individual change between the search and choice stages. I show that the optimal policy of the searcher is a threshold rule and that the threshold depends on the items that the searcher has found so far. Due to this dependence, the search behaviour features the discouragement effect: the threshold decreases weakly in time. The threshold is characterised in detail in Section 5. The characteristics of my model differ from two single-agent search models that feature a time-varying threshold (convex search costs or deadline). In particular, my model features a threshold that decreases endogenously over time and never generates return to an item rejected earlier, in contrast to the other models.

I interpret some of the results in the context of the multi-selves example presented in Introduction. The threshold decreasing in v_m means that a person who has an investment opportunity readily available that involves minimal paperwork (i.e., if the outside option for the chooser is high) optimally does not even attempt to look for an oppor-

tunity with a higher return as he knows he will choose the minimal-paperwork option when the time to invest arrives. A person who is intimidated by paperwork optimally stops searching at an opportunity that offers a lower return (but little paperwork) than a person who likes paperwork: the former is less picky. He is less picky not because he cares less about returns, but because he anticipates that otherwise he chooses a minimal paperwork option when investing.

The model can be extended in several ways. First, if the model is limited to the principal-agent set-up it is reasonable to think that the chooser has a direct influence on the searcher's search process. The possibilities of extending the model are rich due to the many possible assumptions that can be made on the action space and commitment power of the principal. For example, the principal's optimal restrictions on the length of search would be interesting to study because these are prevalent in real life.¹⁰ Second, the model in this paper is very general with respect to the joint distribution of utilities that is considered. If the model is restricted to some application, an application-motivated simplifying restriction on the joint distribution would allow for a more detailed analysis of various aspects of the optimal threshold. Third, my model can be enriched by deriving the chooser's outside option from the model by considering many searchers competing for the chooser. I conjecture that a robust equilibrium is such where each searcher reports to the chooser as soon as he finds the first item. Fourth, my model provides a natural framework in which to think about issues related to naivete and sophistication in a search framework. This forms part of my planned future work.

¹⁰I have derived partial results on time restrictions, specifically, the minimal search duration restriction optimal for the principal. Full results are complex to derive because a maximal search duration restriction (a deadline) means that the agent's threshold depends directly on the amount of time available before the deadline. In fact, it is unclear that the optimal restriction in the time dimension is a connected interval. In addition, I have considered the principal's optimal restriction in the v -dimension. Details of both extensions are available upon request.

References

- ALBRECHT, J., A. ANDERSON, AND S. VROMAN (2010): “Search by Committee,” *Journal of Economic Theory*, 145, 1386–1407.
- ALONSO, R. AND N. MATOUSCHEK (2008): “Optimal Delegation,” *The Review of Economic Studies*, 75, 259–293.
- AMADOR, M. AND K. BAGWELL (2013): “The Theory of Optimal Delegation With an Application to Tariff Caps,” *Econometrica*, 81, 1541–1599.
- ARMSTRONG, M. (1995): “Delegation and Discretion,” *MPRA Paper, University of Munich*, 17069.
- ARMSTRONG, M. AND J. VICKERS (2008): “A Model of Delegated Project Choice,” *MPRA Paper, University of Munich*, 8963.
- (2010): “A Model of Delegated Project Choice,” *Econometrica*, 78, 213–244.
- BERGEMANN, D. AND J. VÄLIMÄKI (2011): “Efficient Search by Committee,” *Discussion Paper, Cowles Foundation*, 1782.
- COMPTE, O. AND P. JEHIÉL (2010): “Bargaining and Majority Rules: A Collective Search Perspective,” *Journal of Political Economy*, 118, 189–221.
- DEY, M. AND C. FLINN (2008): “Household search and health insurance coverage,” *Journal of Econometrics*, 145, 43–63.
- EK, S. AND B. HOLMLUND (2010): “Family Job Search, Wage Bargaining, and Optimal Unemployment Insurance,” *The B.E. Journal of Economic Analysis & Policy*, 10, 1–27.
- FLABBI, L. AND J. MABLI (2012): “Household Search or Individual Search: Does It Matter? Evidence from Lifetime Inequality Estimates,” *IZA Discussion Paper*, 6908.
- GRONAU, R. (1971): “Information and Frictional Unemployment,” *The American Economic Review*, 61, 290–301.
- GULER, B., F. GUVENEN, AND G. L. VIOLANTE (2012): “Joint-Search Theory: New Opportunities and New Frictions,” *Journal of Monetary Economics*, 59, 352–369.
- HOLMSTRÖM, B. (1977): “Delegation and Discretion,” Ph.D. dissertation, Graduate School of Business, Stanford University.

- KAMADA, Y. AND N. MUTO (2015): “Multi-Agent Search with Deadline,” *mimeo*.
- KOVÁČ, E. AND D. KRÄHMER (2015): “Optimal Sequential Delegation,” *mimeo*.
- KOVÁČ, E., D. KRÄHMER, AND T. TATUR (2014): “Optimal Stopping in a Principal-Agent Model with Hidden Information and No Monetary Transfers,” *mimeo*.
- LEWIS, T. R. (2012): “A Theory of Delegated Search for the Best Alternative,” *The RAND Journal of Economics*, 43, 391–416.
- LIPPMAN, S. A. AND J. J. MCCALL (1976): “The Economics of Job Search: A Survey,” *Economic Inquiry*, 14, 155–189.
- MCCALL, J. J. (1970): “Economics of Information and Job Search,” *The Quarterly Journal of Economics*, 84, 113–26.
- MOLDOVANU, B. AND X. SHI (2013): “Specialization and Partisanship in Committee Search,” *Theoretical Economics*, 8, 751–774.
- POSTL, P. (2004): “Delegated Search: Procedure Matters,” *Discussion Paper, University of Birmingham*, 04-17.
- STIGLITZ, J. E. (1987): “Competition and the Number of Firms in a Market: Are Duopolies More Competitive than Atomistic Markets?” *Journal of Political Economy*, 95, 1041–1061.
- ULBRICHT, R. (2016): “Optimal Delegated Search with Adverse Selection and Moral Hazard,” *Theoretical Economics*, *mimeo*.

A Appendix

A.1 Imperfectly correlated preferences

The following provides sufficient conditions for the searcher’s threshold $\bar{u}(v_m)$ to be concave and is a corollary to Proposition 1.

Corollary 1. *Sufficient conditions for the threshold $\bar{u}(v_m)$ to be concave are: (a) monotone increasing hazard rate for f_v : $\frac{\partial}{\partial v_m} \frac{f_v(v_m)}{1-F_v(v_m)} \geq 0$, and (b) $\frac{\partial g(u|v_m)}{\partial v_m} \geq 0$ for all $u \in [\bar{u}(v_m), 1]$.*

Proof. I derive the sufficient conditions for $\bar{u}(v_m)$ to have a negative second derivative. Differentiating equation (5) gives

$$\begin{aligned} \frac{\partial^2 \bar{u}(v_m)}{\partial v_m^2} = & - \left(\frac{\partial}{\partial v_m} \frac{f_v(v_m)}{1 - F_v(v_m)} \right) \int_{\bar{u}(v_m)}^1 g(u|v_m)(u - \bar{u}(v_m)) \, du \\ & - \frac{f_v(v_m)}{1 - F_v(v_m)} \left[\int_{\bar{u}(v_m)}^1 \frac{\partial g(u|v_m)}{\partial v_m} (u - \bar{u}(v_m)) - g(u|v_m) \frac{\partial \bar{u}(v_m)}{\partial v_m} \, du \right], \end{aligned}$$

which is negative for sure if (a) $\frac{\partial}{\partial v_m} \frac{f_v(v_m)}{1 - F_v(v_m)} \geq 0$ and (b) $\frac{\partial g(u|v_m)}{\partial v_m} \geq 0$ for $u \in [\bar{u}(v_m), 1]$ as $\frac{\partial \bar{u}(v_m)}{\partial v_m} < 0$. For example, log-concave distributions satisfy (a) and independent u and v satisfy (b). \square

Claim 1. For u independent of v and h uniform on $[0, 1]^2$, the searcher's threshold is given by equation (7).

Proof. I solve the differential equation for $\bar{u}(v_m)$ as follows. Rewrite (6) as

$$\frac{dy}{dx} = -\frac{1}{2}(1-x)(1-y)^2,$$

and rearrange to separate the variables:

$$(1-y)^{-2} \, dy = -\frac{1}{2}(1-x) \, dx.$$

I use the standard method, integrating and rearranging, to solve the above equation. Integrating both sides of the equation yields

$$(1-y)^{-1} = \frac{1}{2} \ln(1-x) + k,$$

where k is a constant and which can be rearranged to yield

$$y = 1 - \left\{ \frac{1}{2} \ln(1-x) + k \right\}^{-1}. \quad (14)$$

The constant k is pinned down by the initial condition $\bar{u}(v^*) = \bar{u}(1 - 2c) = 0$:

$$k = 1 - \frac{1}{2} \ln(1 - v^*).$$

Inserting k to (14) and reverting to the original notation yields the result. \square

A.2 Characterisation

This section contains the omitted details of Section 5.

A.2.1 The searcher is less picky if the preferences are misaligned

Proposition 3. *The searcher is less picky if the agents' preferences are misaligned as compared to when they are perfectly aligned, i.e., $\max\{\tilde{u}, \bar{v}\} > \bar{u}(v_m)$ for all v_m .*

Proof. I show that $\tilde{u} > \bar{u}(v_m)$ for all v_m . Recall that \tilde{u} is the optimal threshold if the searcher could choose himself and equals the searcher's value from the search process. Recall that $\bar{u}(v_m)$ is the optimal threshold if the chooser chooses instead of the searcher and equals the searcher's value from the search process. The proof is by noting that the searcher can do as least as well in the absence of a chooser as in the presence of her. Let us call a searcher who can choose himself A and call a searcher who searches for a chooser B . Consider a sequence of items drawn from the distribution H with misaligned preferences $(u_1, v_1), (u_2, v_2), (u_3, v_3)$, etc., that is relevant for B and the equivalent sequence $(u_1, u_1), (u_2, u_2), (u_3, u_3)$, etc., relevant for A . A can mimic the behaviour of B by using the threshold $\bar{u}(v_1)$ in the first period, threshold $\bar{u}(\max\{v_1, v_2\})$ in the second, and so on. However, A can do strictly better in expectation as he can ignore the v -value of the items. For example, if $v_1 > v^*$, then $\bar{u}(v_1) = 0$ so that B would accept (u_1, v_1) even if $u_1 \leq \epsilon$ for $\epsilon > 0$ very small. A can do better by continuing to search. The probability of an item with $u_1 < \epsilon$ and $v_1 \geq v^*$ occurring is positive under H because of the full support assumption. Since A 's value from searching equals \tilde{u} , B 's value equals $\bar{u}(\bar{v})$, and $\bar{u}(v_m)$ decreases in v_m , it must be that $\tilde{u} \geq \bar{u}(v_m)$ for all v_m . \square

A.2.2 The searcher prefers more aligned preferences

For positive correlation parametrised by q , the differential equations that pin down the solution are, for $v_m < \hat{v}$,

$$\frac{\partial \bar{u}(v_m)}{\partial v_m} = -\frac{1}{2}(1 - v_m)^{-1}(1 - q)(1 - \bar{u}(v_m))^2,$$

with initial condition $\bar{u}(\hat{v}) = \hat{v}$, and for $v_m \in (\hat{v}, v^*)$,

$$\frac{\partial \bar{u}(v_m)}{\partial v_m} = -\frac{1}{2}(1 - v_m)^{-1}[(1 - q)(1 - \bar{u}(v_m))^2 + 2q(v_m - \bar{u}(v_m))],$$

with initial condition $\bar{u}(v^*) = 0$. The differential equations do not have manageable closed form solutions (i.e., the solutions are nonlinear functions involving Bessel and gamma functions) and are thus omitted.

For negative correlation parametrised by r , the relevant system of equations depends on the size of c (or, more precisely, on whether the curve $\bar{u}(v_m)$ intersects with the line $1 - v_m$). For c large enough, the curve and line do not intersect and the differential

equation that pins down the solution is for any $v_m < v^*$,

$$\frac{\partial \bar{u}(v_m)}{\partial v_m} = -\frac{1}{2}(1 - v_m)^{-1}[(1 - r)(1 - \bar{u}(v_m))^2 + 2r(1 - v_m - \bar{u}(v_m))],$$

with initial condition $\bar{u}(v^*) = 0$. Again, the differential equation does not have a manageable closed form solution and is omitted.

For small search cost, the curve $\bar{u}(v_m)$ and line $1 - v_m$ intersect. Let $\bar{u}(\hat{v}_i) = 1 - \hat{v}_i$ for $i = 1, 2$ such that $\hat{v}_1 < \hat{v}_2$. The differential equation that pins down the solution is for $v_m < \hat{v}_1$ and $v_m \in (\hat{v}_2, v^*)$

$$\frac{\partial \bar{u}(v_m)}{\partial v_m} = -\frac{1}{2}(1 - v_m)^{-1}[(1 - r)(1 - \bar{u}(v_m))^2 + 2r(1 - v_m - \bar{u}(v_m))],$$

with initial conditions $\bar{u}(\hat{v}_1) = 1 - \hat{v}_1$ and $\bar{u}(v^*) = 0$ for $v_m < \hat{v}_1$ and $v_m \in (\hat{v}_2, v^*)$ respectively, and for $v_m \in (\hat{v}_1, \hat{v}_2)$

$$\frac{\partial \bar{u}(v_m)}{\partial v_m} = -\frac{1}{2}(1 - v_m)^{-1}(1 - r)(1 - \bar{u}(v_m))^2,$$

with initial condition $\bar{u}(\hat{v}_2) = 1 - \hat{v}_2$. The solutions to these equations are not informative and thus omitted.