Optimal delegated search with learning and no monetary transfers
Preliminary and incomplete

Elnaz Bajoori* Julia Wirtz†

August 2022

Abstract

A principal delegates the search for the cheapest price to an agent with private information about the price distribution. We do not allow for any monetary transfers to incentivize the agent. The optimal pooling search rule features strictly increasing thresholds, which reflect the principal's updated belief about the price distribution. We show that the pooling rule can be improved upon by a separating menu of search rules with fixed thresholds and a minimum number of offers. Then, we find conditions under which either rule is preferred. Finally, we characterize the optimal separating search rule with a minimum number of offers.

JEL Codes. D83, D82, D86.

Keywords. Delegated search, Asymmetric information, Bandit problems, Dynamic contracts.

*Department of Economics, University of Bath, e.bajoori@bath.ac.uk.

†School of Economics, University of Bristol, j.wirtz@bristol.ac.uk.

1 Introduction

Search is often delegated to experts who are most qualified to find a suitable alternative. The expert has superior knowledge of the market to their client, who may not know precisely what level of outcome they can expect. Examples are estate agents searching for a house or buyer for a client, a hiring committee looking for job candidates for an organisation, engineers developing a new product for a business, or researchers coming up with a new project for an institute. In this paper, we use the example of a procurement team tasked to source an input at the lowest price.

In a standard search problem (McCall; 1970) the optimal policy consists in a fixed threshold, e.g. a maximum acceptable price. However, when the client is uncertain about the price distribution, they do not know what the optimal threshold should be. One possibility is to adjust the thresholds as they learn about the distribution during the search. If the expert fails to meet a threshold, the client becomes more convinced that prices are high and could increase the threshold accordingly. This would be equivalent to a situation where the client is searching himself from an unknown distribution, without the help of an expert. But with the help of an expert who knows the price distribution the client may be able to do better than this. We study how the client can design a simple search rule to incentivize the expert to reveal their knowledge. This is not trivial, because, if the client sets different thresholds depending on the market conditions, the expert has an incentive to report the condition with the easiest threshold. We show that the client can get the expert to reveal the truth by requiring a minimum number of offers in challenging conditions. In procurement rules requiring a minimum number of offers for expensive acquisitions are common. This paper provides a new rationale for such search rules.[1]

We study a principal-agent model in which an object needs to be procured that is paid for by the principal (she), while the search for the best price is delegated to the agent (he). The principal is impatient and incurs a search cost until he receives the object, while the agent bears a cost for each search he performs. While the principal wants to buy supplies as cheaply as possible, the agent is not intrinsically interested in the search outcome, i.e. the purchase price. Thus he would like to finish the search as soon as possible. The principal imposes a search rule that specifies under which conditions the agent may terminate the search.

The employee is more knowledgeable about the market conditions and knows more about

[1] Another reason to require a minimum number of offers is to prevent collusion with a specific supplier. This is less relevant when the transaction is not significant enough for the supplier to make side-payments to the expert worthwhile.

the distribution of prices. For simplicity we assume that the price distribution is either high (state H) or low (state L). When there is no asymmetry of information, i.e. the principal knows the state of the world, it is known that the optimal search rule prescribes a threshold which depends on the state. This first best threshold is strictly higher in state H than in state L. When the principal does not know the state of the world, the first best threshold rules violate incentive compatibility: No matter the true state, the agent prefers to report state H in order to obtain a higher threshold and search less. In order to overcome this problem, we propose two simple search rules: The principal can either set the same search rule for both types (pooling rule) or require a minimum number of offers in the high state in order to incentivize the agent to tell the truth in the low state (separating rule).

With a pooling rule the problem becomes equivalent to a single-agent search with an unknown price distribution. When the principal sets a single rule for both states, she learns about the price distribution over time. The optimal sequence of pooling thresholds can be derived in advance. The optimal threshold level for each number of offers only depends on the fact that the agent has not found an acceptable price up until this point.

To achieve incentive compatibility for a separating rule, the principal needs to discourage the agent in state L from misreporting the state as H. We show that incentive compatibility can be achieved with a rule that uses the first best thresholds, but requires a minimum number of offers when state H is reported. By offering different rules for states H and L, she elicits the agent's private information before the start of the search. While the principal updates her beliefs about the state of the world over time with a pooling rule, she learns the state immediately with a separating rule.

Next, we compare both rules and analyze when this separating rule outperforms the pooling equilibrium.[2] Finally, we characterize the optimal second best separating rule featuring a minimum number of offers, where thresholds are not fixed to first best levels. In the optimal rule, the threshold for state H is lower than the first best threshold, while the threshold for state L is higher.

Our search rule is simple and robust as it does not rely on monetary transfers and minimizes the monitoring burden of the principal. A mechanism using monetary transfers requires very precise information in this setting. In theory, the principal could align incentives by precisely compensating the agent for his search cost. However, if the agent is paid slightly too little, he still wants to stop as soon as possible. If he is paid slightly too much, he does not want to stop at all and has an incentive to manipulate offers to make them look worse.

[2]Example 1 in the Appendix is a numerical example illustrating this result.

Moreover, there are many situations where it is not possible or practical to make monetary transfers conditional on the search outcome. This could be the case because the quality of the outcome is difficult to verify externally, such as in the case of a house or a job candidate. Another reason could be that the agent is a an employee who performs many different searches as part of his job and writing a contract based on all outcomes would be excessively complex. Therefore, we assume that the employee is on a fixed salary, i.e. the agent's pay is independent of the number of searches performed or the result of the search.

Moreover, we restrict our attention to threshold rules that only depend on the number of offers and lowest price found. In theory, the search rule could depend on the exact value of all the offers presented to the principal. Using these, the principal could more precisely update on the probability of the state ex post and punish the agent if he is likely to have lied. However, this would require a substantially more complex search rule. Moreover, the agent could be tempted to manipulate or hide offers. The principal would need to commit to costly monitoring and verification of offers to prevent this. Our search rule creates no temptation for manipulation and does not require monitoring. It only requires that the number of offers is verifiable or that it is costly for the agent to fabricate fake offers.

The remainder of the paper is organized as follows. Section 2 discusses related literature. Section 3 introduces the model and section 4 presents the first best benchmark, i.e. the case that the principal observes the state of the world. Section 5 analyses pooling and separating rules in the asymmetric information case. Section 6 concludes.

2 Related Literature

Our paper is related to the literature on delegated search. In Postl (2004), Armstrong and Vickers (2010), Lewis and Ottaviani (2008), Lewis (2012), and Ulbricht (2016) a principal delegates search to an agent. The closest work to ours in this literature is Ulbricht (2016). We share the assumption that the agent has ex-ante information about the outcome distribution and no stake in the search outcome. Differently from us, the principal cannot observe the number of searches performed, but can use transfers to incentivize the agent.

In contrast, Kovác, Krähmer and Tatur (2013) study delegated search with hidden information where monetary transfers are infeasible, similarly to us. However, their setting differs from ours in that the agent has no ex ante informational advantage. Instead, they privately observe the search outcomes, which can only take two values and cannot be verified by the principal. The agent has different preferences about stopping the search and needs to be incentivized to truthfully reveal the outcomes.

Situated between the literature on delegated search and delegated choice, Mauring (2016) studies a model where one agent searches and a different agent chooses the preferred option, where preferences may diverge. However, she focuses on the optimal stopping policy of the first agent, given that the second agent will choose their preferred option among all that have been revealed. Krähmer and Kováč (2016) study a model of sequential delegated choice. Similarly to us, transfers are not possible and the agent has private ex ante information, which the principal may want to elicit. The agent has private information about the distribution of the state which they learn over time. The principal then delegates the choice to the agent within a restricted choice set.

Finally, our paper is related to a small literature on single agent search from an unknown distribution. Rothschild (1978) studies an agent searching for the lowest price, but does not know the price distribution. For a Dirichlet distribution, he shows that the optimal search rule is equivalent to the case of a known distribution. Bikhchandani and Sharma (1996) extends the result to a more general setting. We are able to show that these results can be adopted for the optimal pooling contract with threshold rules in the principal agent setting.

3 Model

We study a principal-agent model where the principal delegates the search for the cheapest price to an agent.

Let $F_a(p)$ denote the cumulative distribution function of prices $p > 0$ in state $a \in \{H, L\}$, which has full support, no atoms and is twice differentiable. We assume that prices tend to be higher in the high state. Specifically, the high price distribution first order stochastically dominates the low price distribution: $F_H(p) \leq F_L(p)$ for all p. While the agent knows the state, the principal initially believes that the state is H with a probability $\rho_0 \in (0, 1)$ and L with probability $(1 - \rho_0)$.

The agent is paid a fixed wage w and needs to exert effort to search for an offer, i.e. he bears a cost s for every search. A search consists in a price drawn from the distribution $F_a(p)$. The agent's payoff with t offers is thus given by:

$$U = w - st.$$

The principal assigns a value V to the object.[3] Differently from the agent, she has to cover the price p of the object. She is impatient and bears a waiting cost c for every search. Thus,

[3]The results are identical if the value depends on the state.

the principal's payoff from receiving the object at price p after t searches is given by:

$$W = V - ct - p.$$

Writing $\mathbb{E}_a[\cdot]$ for the expectation in state a, we assume $c \leq V - \mathbb{E}_a[p]$ for $a \in \{H, L\}$, such that search is efficient for the principal in both states in the absence of information asymmetries.[4] The principal is risk neutral and maximizes her expected payoff. She fixes a menu of search rules, which consist in a sequence of thresholds: The agent may stop the search when he finds a price below the threshold $y(t)$ after t searches.

The timing of the model is as follows: The principal specifies a menu of search rules, the agent chooses a rule and starts searching.[5] When the requirements of the chosen search rule are met, the agent stops the search, presents the offers and the principal buys the object at the minimum price.

4 Symmetric information

Suppose that the principal can observe the state. The principal's problem then reduces to the standard search problem of a single agent.[6] The optimal search rule takes the form of a threshold: In state a the agent is required to search until he finds a price below a given threshold y_a^*, which is constant over time. The optimal stopping rule is myopic: It is set as if time ended after the next search. At the optimal threshold, the cost of searching once more equals the expected savings from finding a lower price:

$$c = \int_0^{y_a^*} (y_a^* - p)\, dF_a(p).$$

Lemma 1 *If F_H first order stochastically dominates F_L, then $y_H^* \geq y_L^*$.*

Proof. See Appendix.

As soon as he finds a price below the threshold, the agent may stop the search. The expected search duration in state a given a threshold y is given by $\mathbb{E}_a(t|y) = \dfrac{1}{F_a(y)}$. Clearly, the agent expects to search less, the higher the threshold.

[4]Otherwise the principal would prefer to shut down search in the other state and use the first best rule for the other state.

[5]Later, this will be relaxed to allow the agent to start searching before choosing a rule.

[6]For derivations see, e.g., McCall (1970).

5 Asymmetric information

In this section we consider the more interesting case in which the principal does not observe the state. Under asymmetric information the principal cannot simply propose the first best thresholds y^*_H, y^*_L for each state, because incentive compatibility would be violated. The agent's best response would be to announce the state H with the higher threshold in order to minimize the expected search duration.

We restrict our attention to threshold rules, which only depend on the minimum price. The search rule could depend on the whole vector of reported prices $\boldsymbol{p}$. However, such a rule would potentially be sensitive to manipulation by the agent. In practice, the principal usually only verifies the minimum reported price at which she buys the object. When the principal does not have the capacity to verify any further prices, a search rule that depends on these would not be robust. Assume that, for a given minimum price, there is some report $\boldsymbol{p}$ such that the principal does not buy the object and a different report $\hat{\boldsymbol{p}}'$ for which she does. The agent could then always report $\hat{\boldsymbol{p}}'$ without being found out. Moreover, it is possible that the agent colludes with the suppliers to inflate prices. In this case, a robust rule must be a threshold rule such that the principal buys the object if the minimum price is below some threshold. Assume instead a search rule which is not a threshold rule, i.e. there are two prices p, p' with $p < p'$ such that the principal accepts p', but not p. However, then the agent could inflate the price from p to p' and get the principal to buy the object. In conclusion, a search rule which is robust to manipulating non-verified offers takes the form of a sequence of thresholds $\overline{y} = y_1, y_2, \ldots$. Therefore, we restrict attention to robust threshold rules of the following form: Let y_t be the threshold with $t \in \mathbb{R}^+$ offers.[7] The principal buys the object if the minimum reported price with t offers is below y_t.

In section 5.1 we find the optimal pooling rule, which specifies a common search rule for both states and takes into account that the principal updates her belief about the state over time. In section 5.2 we find a separating search rule which achieves incentive compatibility by imposing a minimum number of offers for state H. In section 5.3 we investigate under which conditions the principal prefers to propose a separating menu of search rules and when the principal prefers a pooling rule. In section 5.4 we find the optimal separating menu of search rules featuring a minimum number of offers.

[7] When t is a non-integer number, one can use a random mechanism that requires $\lfloor t \rfloor$ offers with probability one, and an additional search with probability $t - \lfloor t \rfloor$, independently of the outcome found.

5.1 Optimal pooling search rule

In this section we derive the optimal pooling rule, with a common threshold for both states. With a pooling rule the principal is not able to distinguish the states before the agent begins the search. However, the principal will update her belief about the state depending on the search duration the agent needs to find an acceptable price. We give a condition such that the optimal pooling search rule prescribes monotonically increasing thresholds, i.e. $y_t \leq y_{t+1}$. We denote by $\rho_t \equiv \rho(t, y_t)$ the principal's posterior belief that the state is H, given that the agent has not found a price below threshold y_t with $t \in \mathbb{N}$ offers. Using Bayesian updating, the posterior is given by:

$$\rho_t = \frac{\rho_0 \left[1 - F_H(y_t)\right]^t}{\rho_0 \left[1 - F_H(y_t)\right]^t + (1 - \rho_0) \left[1 - F_L(y_t)\right]^t}$$

We would like to find the optimal sequence of pooling thresholds $\{y_t^P\}$ that maximizes the principal's expected payoff, given by

$$W\left(\{y_t^P\}\right) = \rho_0 \left[V_H - c\mathbb{E}_H(t|\{y_t^P\}) - \mathbb{E}_H\left(p|\{y_t^P\}\right)\right] + (1-\rho_0)\left[V_L - c\mathbb{E}_L(t|\{y_t^P\}) - \mathbb{E}_L\left(p|\{y_t^P\}\right)\right].$$

Proposition 1 *The optimal threshold with t offers y_t^P fulfills:*

$$c = \int_0^{y_t^P} \left(y_t^P - p\right) \left[\rho_t F_H(p) + (1 - \rho_t) F_L(p)\right] dp.$$

Proof. With a pooling rule, the principal's problem is equivalent to that of a single agent searching from an unknown distribution without the agency problem. We can thus make use of results from Bikhchandani and Sharma (1996) concerning this setting. For details see the Appendix.

Proposition 1 states that it is optimal to stop at the first price such that the expected saving from one additional search is smaller than the cost. Thus, the optimal threshold y_t^P is myopic. Since the optimal threshold y_t^P with t offers only depends on the previous threshold y_{t-1}^P and the parameters of the model, the optimal sequence $\{y_t^P\}$ can be determined in advance, before any offers have been received. Moreover, the optimal threshold with t offers with uncertainty about the state is identical to the optimal (constant) threshold with a known price distribution, $\hat{F}_t(p)$, which is equal to the expected posterior distribution given that the agent does not find a price below the threshold y_t^P in t searches:

$$\hat{F}_t(p) = \rho_t F_H(p) + (1 - \rho_t) F_L(p).$$

Therefore, every element of the sequence of optimal thresholds must lie between the first best thresholds for state H and L: $y_L^* \leq y_t^P \leq y_H^*$.

Next we consider how ρ_t develops over time. Under the following assumption and with non-decreasing thresholds, the principal becomes increasingly sure that the state is high.

$$\frac{1 - F_H(p_1)}{1 - F_L(p_1)} \leq \frac{1 - F_H(p_2)}{1 - F_L(p_2)} \quad y_L^* \leq p_1 \leq p_2 \leq y_H^*. \tag{A1}$$

Assumption A1 requires the failure probability ratio in the high and low state to be increasing with the threshold. This assumption holds for pairs of many common distributions such as normal distributions, exponential distributions, uniform distributions etc. that display FOSD.

Lemma 2 *ρ_t monotonically increases in t with an increasing pooling threshold y_t if Condition (A1) holds.*

Proof. See Appendix.

It is more likely that the agent fails to find a price below any given threshold when the state is H than if the state is L. If Condition (A1) holds, a failure with a higher threshold is more informative about state H. In this case the probability that the state is H increases over time with increasing thresholds.

Lemma 2 shows that ρ_t increases for increasing thresholds if condition (A1) holds. The optimal threshold increases over time, as long as ρ_t is increasing. Therefore, the optimal threshold increases over time if condition (A1) holds. As ρ_t is approaching 1, y_t^P approaches y_H^*.

5.2 Optimal separating search rule with first best thresholds

In this section we show how the principal can elicit the state from the agent by offering a menu of two different search rules designed for either state. The search rules need to be incentive compatible: Given the realized state, the agent prefers the corresponding search rule and, thus, is willing to report the state truthfully. This can be achieved by imposing a minimum number of offers in the high state. The separating rule then allows the principal to distinguish the states before the agent begins the search.

We consider the search rule that uses the first best thresholds for both states y_H^*, y_L^*, combined with a minimum number of offers $k_H, k_L > 0$. Formally, the thresholds for state a are:

$$y_{a,t} = \begin{cases} 0 & \text{for } t < k_a \\ y_a^* & \text{for } t \geq k_a. \end{cases}$$

Effectively, the agent is asked to acquire at least k_a offers. If he can find prices below the threshold y_a^* among the first k_a offers, the lowest of these is taken. Otherwise, he has to continue the search for a price below the threshold y_a^*.

We would like to find the minimum numbers of offers k_H, k_L that maximize the principal's expected payoff subject to incentive compatibility.

The Separating Problem

$$\max_{k_H, k_L} W(k_H, y_H, k_L, y_L) = \rho_0 W_H(k_H, y_H) + (1 - \rho_0) W_L(k_L, y_L), \tag{1}$$

$$\text{subject to} \quad \mathbb{E}_L(t|k_L, y_L) \leq \mathbb{E}_L(t|k_H, y_H) \tag{ICL}$$

$$\mathbb{E}_H(t|k_H, y_H) \leq \mathbb{E}_H(t|k_L, y_L), \tag{ICH}$$

where $W_H = V_H - c\mathbb{E}_H(t|k_H, y_H) - \mathbb{E}_H(p|k_H, y_H)$, $W_L = V_L - c\mathbb{E}_L(t|k_L, y_L) - \mathbb{E}_L(p|k_L, y_L)$ and all thresholds are set to the first best level: $y_H = y_H^*, y_L = y_L^*$.

The incentive constraints ensure that the agent prefers to follow the appropriate search rule for the state. Recall that, as long as the principal purchases the object, the agent only cares about minimizing the search duration. ICL states that in state L the expected search duration with the search rule (k_L, y_L^*) should be lower than with the rule (k_H, y_H^*). ICH states that in state H the expected search duration with search rule (k_H, y_H^*) should be lower than with the rule (k_L, y_L^*).

If the following assumption holds, a separating equilibrium exists:

$$\frac{1}{F_H(p_1)} - \frac{1}{F_L(p_1)} \geq \frac{1}{F_H(p_2)} - \frac{1}{F_L(p_2)} \quad y_L^* \leq p_1 \leq p_2 \leq y_H^* \tag{A2}$$

Assumption A2 says that the difference between the expected search duration in state H and L is decreasing with the threshold. An example where this assumption holds can be seen in Figure 1. This assumption guarantees that imposing a minimum number of offers is not too costly when the true state is H. Otherwise, it may not be possible to fulfil ICL without violating ICH: The agent could prefer the search rule designed for state L in either state in order to avoid having to provide a minimum number of offers. As with A1, this assumption holds for pairs of many common distributions such as normal distributions,

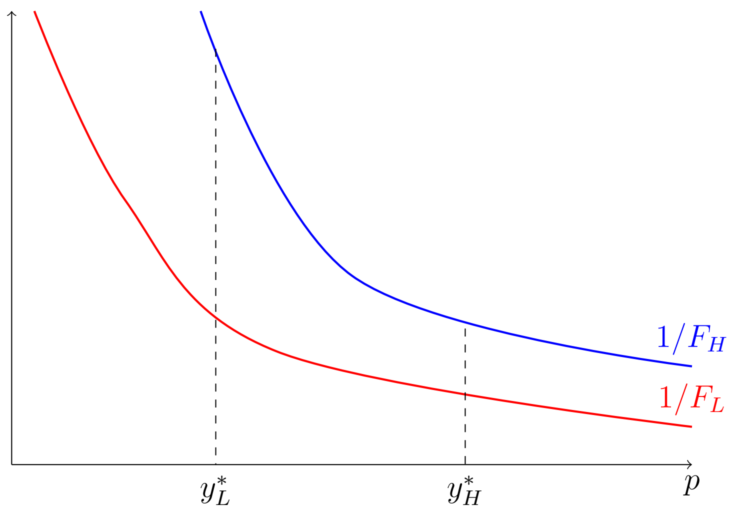

Figure 1: Assumption A2.

exponential distributions, uniform distributions etc. that display FOSD.

Proposition 2 *Given the thresholds y_L^*, y_H^*, the principal's expected payoff is maximized by setting the minimum number of offers for the low state to $k_L^* = 0$ and for the high state to k_H^* which solves*

$$k_H^* = \frac{1}{F_L(y_L^*)} - \frac{[1 - F_L(y_H^*)]^{k_H^*}}{F_L(y_H^*)}.$$

Proof. See Appendix.

5.3 Comparison of separating and pooling search rules

This section compares the separating rule described above and the optimal pooling rule and gives conditions under which the principal receives a higher payoff from one or the other.

Note that the principal's expected payoffs in state L and H are maximized at y_L^* and y_H^*, respectively. Therefore, the separating menu $R^* = ((0, y_L^*), (k_H^*, y_H^*))$ is clearly optimal for the principal conditional on the state being low, since she is setting the optimal threshold y_L^* with no minimum number of offers: $W_L(y_L^*) \geq W_L\left(\{y_t^P\}\right)$. Therefore, the separating rule is clearly superior conditional on the state being L. Similarly, once the agent has obtained k_H^* offers, the separating menu is also optimal for the principal if the state is high: she is setting the optimal threshold y_H^* and there are no further obligatory offers to be obtained: $W_H(y_H^*) \geq W_H\left(\{y_t^P\}\right)$. The only constellation in which the pooling rule has an advantage over the separating rule is for the high state, before the agent has obtained k_H^* offers. With the pooling rule, the agent may stop searching if he finds a sufficiently low price. Instead, with the separating rule (k_H^*, y_H^*) there is no possibility of stopping before k_H^* offers have been obtained.

Let $f(k)$ be the difference in expected search duration under the search rule (k, y_H^*) and the first best rule $(0, y_H^*)$. We have

$$f(k) = k - \frac{1 - [1 - F_H(y_H^*)]^k}{F_H(y_H^*)}.$$

Proposition 3 then gives a sufficient condition such that the separating rule generates a higher payoff for the principal than the optimal separating menu:

Proposition 3 *The optimal separating menu R^* is preferred to the optimal pooling threshold $\{y_t^P\}$ if $cf(k_H^*) \leq W_H(y_H^*) - W_H(\{y_t^P\})$*

Proof. See Appendix.

The principal receives a higher payoff from the separating menu conditional on the state being L. The principal must prefer the separating menu overall if she also receives a higher payoff from the separating menu in state H. This is guaranteed if the condition in Proposition 3 is fulfilled.

Proposition 4 *The optimal pooling threshold $\{y_t^P\}$ is preferred to the optimal separating menu R^* iff the prior ρ_0 is sufficiently high and $W_H(\{y_t^P\}) > W_H((k_H^*, y_H^*))$.*

Proof. See Appendix.

The principal receives a higher payoff from the separating menu in state L. However, it is possible that the principal receives a higher payoff from the pooling threshold in state H. In order for the principal to prefer the pooling threshold overall, this must be the case and, in addition, the probability of state H must be high enough.

Covert searching While either the separating rule or the pooling threshold could be preferred by the principal, there is a caveat to the implementability of the separating rule. If agents are able to start searching in secret before announcing the state, a separating rule is infeasible. When offered a menu of the form $((0, y_L), (k_H, y_H))$ with $y_L < y_H$ and $k_H > 0$, the agent would start searching covertly before announcing the state. If he finds a price below y_L before obtaining k_H offers, he announces the state to be L in either state. If he does not, he announces the state to be H and proceeds to search with the higher threshold

y_H. Effectively, the agent faces the pooling threshold $\{y'\}$ with:

$$y'_t = \begin{cases} y_L & \text{for } t < k_H \\ y_H & \text{for } t \geq k_H. \end{cases}$$

However, $\{y'\}$ must result in a lower payoff to the principal than the optimal pooling threshold $\{y_t^P\}$ as characterized in Section 5.1. When the agent, instead, cannot search covertly before announcing the state, either the separating rule or the pooling threshold could be preferred by the principal.

5.4 Optimal separating search rule

In Section 5.2 we fixed the thresholds of the separating search rule to the first best levels, which allowed for a simple comparison to the optimal pooling rule. However, without this restriction, it is possible to improve the principal's payoff from a separating rule. In this section we thus characterize the optimal menu of separating search rules in the class of rules with a minimum number of offers. We denote this class of search rules as follows:

$$\mathcal{R} = \{(k_L, y_L), (k_H, y_H) \mid k_a \geq 0, y_a \geq 0, a \in \{H, L\}\} \tag{2}$$

In order to characterize the optimal incentive compatible menu within the class $\mathcal{R}$, we would like to find the menu $((k_L, y_L), (k_H, y_H))$ that solves the Separating Problem (1). Proposition 5 characterizes the optimal menu within the class $\mathcal{R}$:

Proposition 5 *The optimal menu for the principal within the class $\mathcal{R}$ is $\hat{R} = \{(0, \hat{y}_L), (\hat{k}_H, \hat{y}_H)\}$, in which $\hat{y}_L > y_L^*$, $\hat{y}_H < y_H^*$, and $\hat{k}_H > 0$ solves*

$$\hat{k}_H = \frac{1}{F_L(\hat{y}_L)} - \frac{[1 - F_L(\hat{y}_H)]^{\hat{k}_H}}{F_L(\hat{y}_H)}.$$

Proof. See Appendix.

For state L, the optimal separating menu sets a threshold that is greater than the first best threshold y_L^*. Moreover, the agent does not need to provide a minimum number of offers ($\hat{k}_L = 0$). For state H, the optimal menu sets a threshold that is smaller than the first best threshold y_H^*. At the same time, the optimal minimum number of offers $\hat{k}_H$ is smaller than in the separating rule with first best thresholds.

6 Conclusion

This paper proposes simple search rules for search delegated to an expert with superior knowledge of the underlying distribution. With a known distribution, it is optimal to search until a fixed threshold is met. However, when the distribution is unknown, the principal can improve their payoff with a threshold that changes with the number of searches performed. The optimal pooling rule prescribes strictly increasing thresholds as the principal successively learns about the state: She becomes increasingly convinced that the state of the world is unfavourable when the agent fails to find a price below a certain threshold. Alternatively, the principal can extract the expert's knowledge ex ante with a separating contract which uses first best thresholds, but imposes a minimum number of offers when the distribution is unfavourable. This provides a theoretical explanation to the prevalence of such rules in practice. We find conditions under which this separating rule outperforms the optimal pooling rule. Last but not the least, we characterize the optimal separating rule with a minimum number of offers. We find that under this rule, the optimal threshold in the low state is above the first best and the optimal threshold for the high state is below the first best.

A Appendix

Proof of Lemma 1. Let $\phi_a(y) \equiv \int_0^y (y-p)dF_a(p)$. First we show that for every $y > 0$, $\phi_L(y) \geq \phi_H(y)$. Using integration by parts we have

$$\phi_a(y) = (y-p)F_a(p)\Big|_0^y + \int_0^y F_a(p)\, dp = \int_0^y F_a(p)\, dp$$

Since F_H first order stochastically dominates F_L, for every $y > 0$ we have

$$\int_0^y F_H(p)\, dp \leq \int_0^y F_L(p)\, dp$$

Therefore,

$$\phi_L(y) \geq \phi_H(y) \quad \forall y.$$

We know that at the optimal thresholds y_H^* and y_L^* we have $\phi_H(y_H^*) = \phi_L(y_L^*) = c$. As $\phi_L(y) \geq \phi_H(y)$ for every $y > 0$ and $\phi_a(y)$ is increasing in y, we can conclude that $y_H^* \geq y_L^*$. $\square$

Proof of Proposition 1. Let $\mathbb{E}[F](p) \equiv \rho_0 F_H(p) + (1-\rho_0)F_L(p)$ denote the expected price distribution from the perspective of the principal, given her prior belief ρ_0. The posterior price distribution with t offers with thresholds $y_1, \cdots, y_t$ is denoted by $\mathbb{E}[F|y_1, \cdots, y_t](p) \equiv \rho_t F_H(p) + (1-\rho_t)F_L(p)$.

Bikhchandani and Sharma (1996) show that the optimal stopping rule is myopic when the following assumption is fulfilled for all t and all observations $y_1, \cdots, y_t$[8]:

$$\mathbb{E}\left[F|y_1, \ldots, y_t\right](p) \leq \mathbb{E}\left[F|y_1, \ldots, y_{t-1}\right](p) \leq \cdots \leq \mathbb{E}\left[F\right](p) \quad \forall p < \min(y_1, \ldots, y_t). \tag{3}$$

Since we have $\rho_t < \rho_{t-1}$ and $F_H(p) < F_L(p)$ $\forall p$ we can show that Condition (3) holds in our setting. For all k and for all thresholds $y_1, \cdots, y_t$ we have:

$$\begin{aligned}
&\mathbb{E}\left[F|y_1, \cdots, y_t\right](p) = \rho_t F_H(p) + (1-\rho_t)F_L(p) \\
\leq &\mathbb{E}\left[F|y_1, \cdots, y_{t-1}\right](p) = \rho_{t-1} F_H(p) + (1-\rho_{t-1})F_L(p) \\
\leq &\cdots \\
\leq &\mathbb{E}\left[F\right](p) = \rho_0 F_H(p) + (1-\rho_0)F_L(p) \quad \forall p
\end{aligned}$$

Therefore, the optimal stopping rule is myopic. Specifically, Bikhchandani and Sharma (1996) show that, when condition (3), holds it is optimal to stop at the first price such that the expected saving from a further search is smaller than c. The optimal threshold y_t^P then fulfills the following condition:

$$\begin{aligned}
c &= \int_0^{y_t^P} \left(y_t^P - p\right) \mathbb{E}\left[F|y_t^P, t\right](p)\, dp \\
&= \int_0^{y_t^P} \left(y_t^P - p\right) \left[\rho_t F_H(p) + (1-\rho_t)F_L(p)\right] dp.
\end{aligned}$$

□

Proof of Lemma 2.

We want to show that Assumption A1 implies $\rho_t \leq \rho_{t+1}$ for all $t \geq 1$. We have

$$\rho_t = \frac{\rho_0 \left[1 - F_H(y_t)\right]^t}{\rho_0 \left[1 - F_H(y_t)\right]^t + (1-\rho_0)\left[1 - F_L(y_t)\right]^t}$$

[8]Bikhchandani and Sharma (1996) study optimal search from an unknown distribution. This is equivalent to our setting, since the agent is willing to search in both states, given the thresholds $y_1, \cdots, y_t$. Moreover, they include the possibility of right-censored observations, i.e., it can only be observed that the price is greater than a particular level. This would be equivalent to our setting, where the principal can only infer with certainty that the minimum price the agent has found with t offers is higher than the threshold y_t.

and

$$\rho_{t+1} = \frac{\rho_0 \left[1 - F_H(y_{t+1})\right]^{t+1}}{\rho_0 \left[1 - F_H(y_{t+1})\right]^{t+1} + (1-\rho_0)\left[1 - F_L(y_{t+1})\right]^{t+1}}.$$

Thus we can write $\rho_{t+1} = h_t \rho_t$.

$$h_t = \frac{\left[1 - F_H(y_{t+1})\right]^{t+1}}{\left[1 - F_H(y_t)\right]^t} \frac{\rho_0 \left[1 - F_H(y_t)\right]^t + (1-\rho_0)\left[1 - F_L(y_t)\right]^t}{\rho_0 \left[1 - F_H(y_{t+1})\right]^{t+1} + (1-\rho_0)\left[1 - F_L(y_{t+1})\right]^{t+1}}.$$

If $h_t \geq 1$ then $\rho_t \leq \rho_{t+1}$ for all $t \geq 1$. Finally, we show that $\frac{1-F_H(y_{t+1})}{1-F_L(y_{t+1})} \geq \frac{1-F_H(y_t)}{1-F_L(y_t)}$ implies $h_t \geq 1$ for all $t \geq 1$:

$$\begin{aligned} &\frac{1 - F_H(y_{t+1})}{1 - F_L(y_{t+1})} \geq \frac{1 - F_H(y_t)}{1 - F_L(y_t)} \\ \Leftrightarrow &\frac{\left[1 - F_L(y_{t+1})\right]^t}{\left[1 - F_H(y_{t+1})\right]^t} \leq \frac{\left[1 - F_L(y_t)\right]^t}{\left[1 - F_H(y_t)\right]^t} \end{aligned}$$

Since we have $\frac{1-F_L(y_{t+1})}{1-F_H(y_{t+1})} \leq 1$, this implies:

$$\begin{aligned} \Rightarrow &\frac{\left[1 - F_L(y_{t+1})\right]^{t+1}}{\left[1 - F_H(y_{t+1})\right]^{t+1}} \leq \frac{\left[1 - F_L(y_t)\right]^t}{\left[1 - F_H(y_t)\right]^t} \\ \Leftrightarrow &(1-\rho_0)\frac{\left[1 - F_L(y_{t+1})\right]^{t+1}}{\left[1 - F_H(y_{t+1})\right]^{t+1}} \leq (1-\rho_0)\frac{\left[1 - F_L(y_t)\right]^t}{\left[1 - F_H(y_t)\right]^t} \\ \Leftrightarrow &\rho_0 + (1-\rho_0)\frac{\left[1 - F_L(y_{t+1})\right]^{t+1}}{\left[1 - F_H(y_{t+1})\right]^{t+1}} \leq \rho_0 + (1-\rho_0)\frac{\left[1 - F_L(y_t)\right]^t}{\left[1 - F_H(y_t)\right]^t} \\ \Leftrightarrow &\rho_0 \frac{\left[1 - F_H(y_{t+1})\right]^{t+1}}{\left[1 - F_H(y_{t+1})\right]^{t+1}} + (1-\rho_0)\frac{\left[1 - F_L(y_{t+1})\right]^{t+1}}{\left[1 - F_H(y_{t+1})\right]^{t+1}} \\ &\quad \leq \rho_0 \frac{\left[1 - F_H(y_t)\right]^t}{\left[1 - F_H(y_t)\right]^t} + (1-\rho_0)\frac{\left[1 - F_L(y_t)\right]^t}{\left[1 - F_H(y_t)\right]^t} \\ \Leftrightarrow &\frac{\rho_0 \left[1 - F_H(y_{t+1})\right]^{t+1} + (1-\rho_0)\left[1 - F_L(y_{t+1})\right]^{t+1}}{\left[1 - F_H(y_{t+1})\right]^{t+1}} \\ &\quad \leq \frac{\rho_0 \left[1 - F_H(y_t)\right]^t + (1-\rho_0)\left[1 - F_L(y_t)\right]^t}{\left[1 - F_H(y_t)\right]^t} \\ \Leftrightarrow &\frac{\left[1 - F_H(y_{t+1})\right]^{t+1}}{\left[1 - F_H(y_t)\right]^t} \frac{\rho_0 \left[1 - F_H(y_t)\right]^t + (1-\rho_0)\left[1 - F_L(y_t)\right]^t}{\rho_0 \left[1 - F_H(y_{t+1})\right]^{t+1} + (1-\rho_0)\left[1 - F_L(y_{t+1})\right]^{t+1}} \geq 1 \\ \Leftrightarrow &h_t \geq 1. \end{aligned}$$

$\square$

Proof of Proposition 2.

First, we find the minimum number of offers k_H^*, k_L^* that maximize the principal's expected payoff subject to the incentive constraint (ICL). Then we show that, for the resulting search rules (k_H^*, y_H^*) and (k_L^*, y_L^*) assumption A2 ensures that the incentive constraint ICH is also fulfilled.

$$\max_{k_H, k_L} W(k_H, y_H^*, k_L, y_L^*) =$$
$$\rho_0 \left(V_H - c\mathbb{E}_H(t|k_H, y_H^*) - \mathbb{E}_H(p|k_H, y_H^*)\right) + (1-\rho_0)\left(V_L - c\mathbb{E}_L(t|y_L^*) - \mathbb{E}_L(p|y_L^*)\right)$$
$$\text{s.t. } \mathbb{E}_L(t|k_L, y_L^*) \leq \mathbb{E}_L(t|k_H, y_H^*) \qquad \text{ICL}.$$

As $W_a(k_a, y_a^*)$ is decreasing in k_a it is optimal to set k_a as small as the incentive constraint allows.

The expected search duration with rule (k_b, y_b) when the state is a is given by:

$$\mathbb{E}_a(t|k_b, y_b) = k_b + \frac{[1 - F_a(y_b)]^{k_b}}{F_a(y_b)}.$$

ICL then becomes

$$k_L + \frac{[1 - F_L(y_L)]^{k_L}}{F_L(y_L)} \leq k_H + \frac{[1 - F_L(y_H)]^{k_H}}{F_L(y_H)}.$$

As the LHS increases in k_L, ICL becomes easier to fulfil the smaller k_L. Therefore it is optimal to set $k_L^* = 0$. Thus, ICL becomes:

$$\frac{1}{F_L(y_L)} \leq k_H + \frac{[1 - F_L(y_H)]^{k_H}}{F_L(y_H)}. \tag{4}$$

Note that, clearly $k_H = 0$ does not satisfy ICL, so we must have $k_H > 0$. As the RHS increases in k_H, it is optimal to set k_H^* such that Equation ICL is satisfied with equality, given thresholds y_H^* and y_L^*:

$$k_H^* = \frac{1}{F_L(y_L^*)} - \frac{[1 - F_L(y_H^*)]^{k_H^*}}{F_L(y_H^*)}.$$

Finally, we show that for the resulting search rules (k_H^*, y_H^*) and $(0, y_L^*)$ assumption A2 implies that the incentive constraint ICH is fulfilled. According to assumption A2:

$$\frac{1}{F_H(y_L^*)} - \frac{1}{F_L(y_L^*)} \geq \frac{1}{F_H(y_H^*)} - \frac{1}{F_L(y_H^*)}.$$

Moreover, we know that

$$F_L(y) \leq F_H(y)$$

$$\begin{aligned}
&\frac{1}{F_H(y_L^*)} - \frac{1}{F_L(y_L^*)} \geq \frac{1}{F_H(y_H^*)} - \frac{1}{F_L(y_H^*)} \\
\Rightarrow &\frac{1}{F_H(y_L^*)} - \frac{1}{F_L(y_L^*)} \geq \frac{1}{F_H(y_H^*)} - \frac{1}{F_L(y_H^*)} \geq \frac{[1-F_H(y_H^*)]^{k_H}}{F_H(y_H^*)} - \frac{[1-F_L(y_H^*)]^{k_H}}{F_L(y_H^*)} \\
\Rightarrow &\frac{1}{F_L(y_L^*)} - \frac{[1-F_L(y_H^*)]^{k_H}}{F_L(y_H^*)} \leq \frac{1}{F_H(y_L^*)} - \frac{[1-F_H(y_H^*)]^{k_H}}{F_H(y_H^*)} \\
\Rightarrow &\frac{1}{F_L(y_L^*)} - \frac{[1-F_L(y_H^*)]^{k_H^*}}{F_L(y_H^*)} + \frac{[1-F_H(y_H^*)]^{k_H}}{F_H(y_H^*)} \leq \frac{1}{F_H(y_L^*)} \\
\Rightarrow &k_H + \frac{[1-F_H(y_H^*)]^{k_H}}{F_H(y_H^*)} \leq \frac{1}{F_H(y_L^*)} \\
\Rightarrow &\mathbb{E}_H(t|k_H, y_H^*) \leq \mathbb{E}_H(t|0, y_L^*) \qquad \forall k_H > 0.
\end{aligned}$$

□

Proof of Proposition 3. The principal's expected payoff under the common threshold rule $\{y_t^P\}$ is given by

$$W\left(\{y_t^P\}\right) = \rho W_H\left(\{y_t^P\}\right) + (1-\rho) W_L\left(\{y_t^P\}\right).$$

Let R^* stand for the separating menu $((0, y_L^*), (k_H^*, y_H^*))$. The principal's expected payoff given the menu R^* is then given by

$$W(R^*) = \rho W_H(k_H^*, y_H^*) + (1-\rho) W_L(y_L^*).$$

Clearly, $W_L(y_L^*) \geq W_L\left(\{y_t^P\}\right)$. We also know that the expected price in state H with search rule (k_H^*, y_H^*) is lower than with the first-best threshold y_H^* without a minimum number of offers.[9]

[9] This is because with (k_H^*, y_H^*) there is a possibility that the agent finds more than one price below y_H^* before completing the mandatory search duration k_H^*.

Now, assume $cf\left(k_H^*\right) \leq W_H\left(y_H^*\right) - W_H\left(\left\{y_t^P\right\}\right)$, then

$$\begin{aligned}
&c\left[\mathbb{E}_H(k|\left(k_H^*, y_H^*\right)) - \mathbb{E}_H(k|y_H^*)\right] \leq W_H\left(y_H^*\right) - W_H\left(\left\{y_t^P\right\}\right) \\
\Leftrightarrow &c\mathbb{E}_H(k|\left(k_H^*, y_H^*\right)) - c\mathbb{E}_H(k|y_H^*) \leq V_H - c\mathbb{E}_H(k|y_H^*) - \mathbb{E}_H\left(p|y_H^*\right) - W_H\left(\left\{y_t^P\right\}\right)
\end{aligned}$$

Since we have $\mathbb{E}_H\left(p|y_H^*\right) \geq \mathbb{E}_H\left(p|\left(k_H^*, y_H^*\right)\right)$ this implies:

$$\begin{aligned}
\Rightarrow W_H\left(\left\{y_t^P\right\}\right) &\leq V_H - \mathbb{E}_H\left(p|y_H^*\right) - c\mathbb{E}_H(k|\left(k_H^*, y_H^*\right)) \\
&\leq V_H - \mathbb{E}_H\left(p|\left(k_H^*, y_H^*\right)\right) - c\mathbb{E}_H(k|\left(k_H^*, y_H^*\right)) \\
&= W_H(R^*)
\end{aligned}$$

This shows that $cf\left(k_H^*\right) \leq W_H\left(y_H^*\right) - W_H\left(\left\{y_t^P\right\}\right)$ is a sufficient condition such that the principal's expected profit under the separating menu R^* is higher than under the pooling threshold $\{y_t^P\}$. □

Proof of Proposition 4. The optimal pooling threshold $\{y_t^P\}$ is preferred to the optimal separating menu R if the following holds:

$$\begin{aligned}
&W\left(\left\{y_t^P\right\}\right) \geq W(R^*) \\
\Leftrightarrow &\rho_0 W_H\left(\left\{y_t^P\right\}\right) + (1-\rho_0)W_L\left(\left\{y_t^P\right\}\right) \geq \rho_0 W_H(R^*) + (1-\rho_0)W_L(R^*) \\
\Leftrightarrow &\rho_0\left[W_H\left(\left\{y_t^P\right\}\right) - W_H(R^*)\right] \geq (1-\rho_0)\left[W_L(R^*) - W_L\left(\left\{y_t^P\right\}\right)\right]
\end{aligned}$$

If $\left[W_H\left(\left\{y_t^P\right\}\right) - W_H(R^*)\right] > 0$, this is equivalent to:

$$\rho_0 \geq \frac{W_L(R^*) - W_L\left(\left\{y_t^P\right\}\right)}{W_L(R^*) - W_L\left(\left\{y_t^P\right\}\right) + W_H\left(\left\{y_t^P\right\}\right) - W_H(R^*)}$$

□

Lemma 3 *Suppose there is only one state and the principal offers a search rule (k, y), with $k > 0$ exogenously fixed. Then it is optimal for the principal to offer the rule (k, y^*) in which y^* is the first best threshold.*

Proof. Suppose the agent has already searched k times. Further suppose the minimum price he has found is higher than y^*. The principal can then either buy the good at that price or ask the agent to continue to search for a price lower than y^*. Clearly, the principal's expected payoff is higher in the latter case. Suppose now that the agent has found a price $\tilde{p} \leq y^*$. The principal's payoff if she buys the good at price $\tilde{p}$ is higher than the expected payoff if the agent continues to search. Therefore the optimal price threshold is equal to y^*.

□

Proof of Proposition 5. We would like to find the menu $\{(\hat{k}_L, \hat{y}_L), (\hat{k}_H, \hat{y}_H)\}$ that maximizes the principal's expected payoff subject to incentive constraints ICL and ICH. We know that when there are no incentive compatibility constraints (ICCs), it is optimal to perform an extra search if and only if the expected saving in price is larger than the search cost. Since the use of a threshold is always more efficient than the use of a minimum number of offers, $\hat{k}_L$ and $\hat{k}_H$ should be set as small as the constraints allow. Thus it is optimal to set $k_L = 0$ for any given expected search duration that needs to be achieved. Thus ICL becomes:

$$\frac{1}{F_L(y_L)} \leq k_H + \frac{[1 - F_L(y_H)]^{k_H}}{F_L(y_H)} \tag{5}$$

We first solve the Separating Problem (1), subject only to the constraint ICL. At the end of the proof we then show that Assumption A2 implies that ICH is satisfied for $k_L = 0$, $k_H = \hat{k}_H$, $y_L = \hat{y}_L$, and $y_H = \hat{y}_H$.

$$\begin{aligned} &\max_{k_H, y_H, y_L} W(k_H, y_H, y_L) = \\ &\rho_0 \left(V_H - c\mathbb{E}_H(t|k_H, y_H) - \mathbb{E}_H(p|k_H, y_H)\right) + (1 - \rho_0) \left(V_L - c\mathbb{E}_L(t|y_L) - \mathbb{E}_L(p|y_L)\right) \\ &\text{s.t. } \mathbb{E}_L(t|y_L) \leq \mathbb{E}_L(t|k_H, y_H). \end{aligned}$$

The Lagrangian is given by

$$\begin{aligned} \mathcal{L} =&\rho_0 \left(V_H - c\mathbb{E}_H(t|k_H, y_H) - \mathbb{E}_H \left(p|k_H, y_H\right)\right) + (1 - \rho_0) \left(V_L - c\mathbb{E}_L(t|y_L) - \mathbb{E}_L \left(p|y_L\right)\right) \\ &- \lambda \left(\mathbb{E}_L(t|y_L) - \mathbb{E}_L(t|k_H, y_H)\right) \end{aligned}$$

The complementary slackness conditions are: $\lambda \left(\mathbb{E}_L(t|y_L) - \mathbb{E}_L(t|k_H, y_H)\right) = 0$ and $\lambda \geq 0$. We know that ICL will bind, thus $\mathbb{E}_L(t|y_L) = \mathbb{E}_L(t|k_H, y_H)$ and $\lambda > 0$.

The first-order condition for y_L is:

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial y_L} =& - (1 - \rho_0)c\frac{\partial \mathbb{E}_L(t|\hat{y_L})}{\partial y_L} - (1 - \rho_0)\frac{\partial \mathbb{E}_L(p|\hat{y_L})}{\partial y_L} - \lambda\frac{\partial \mathbb{E}_L(t|\hat{y_L})}{\partial y_L} = 0 \\ \Rightarrow& - c\frac{\partial \mathbb{E}_L(t|\hat{y_L})}{\partial y_L} - \frac{\partial \mathbb{E}_L(p|\hat{y_L})}{\partial y_L} = \frac{\lambda}{1 - \rho_0}\frac{\partial \mathbb{E}_L(t|\hat{y_L})}{\partial y_L}. \end{aligned}$$

We have $c\dfrac{\partial \mathbb{E}_L(t|y_L)}{\partial y_L} - \dfrac{\partial \mathbb{E}_L(p|y_L)}{\partial y_L} = \dfrac{\partial W_L(y_L)}{\partial y_L}$. We know $\dfrac{\partial \mathbb{E}_L(t|y_L)}{\partial y_L} < 0$, $\dfrac{\partial \mathbb{E}_L(p|y_L)}{\partial y_L} > 0$.

Moreover, we know $\lambda > 0$ and $(1-\rho_0) > 0$. This implies $\frac{\partial W_L}{\partial y_L}(\hat{y}_L) < 0$. Since $W_L(y_L)$ is maximized at $y_L = y_L^*$, W_L increases for $y_L < y_L^*$ and decreases for $y_L > y_L^*$. Therefore, it must be that $\hat{y}_L > y_L^*$.

The optimal threshold in state L is higher than the first best threshold. This means that it is optimal to stop searching when the expected saving in price is still greater than the cost of a search. The decrease in search cost is smaller than the increase in price as the threshold increases.

The first-order condition for y_H is:

$$\begin{aligned}
\frac{\partial \mathcal{L}}{\partial y_H} =& -\rho_0 c \frac{\partial \mathbb{E}_H(t|\hat{k}_H, \hat{y}_H)}{\partial y_H} - \rho_0 \frac{\partial \mathbb{E}_H(p|\hat{k}_H, \hat{y}_H)}{\partial y_H} + \lambda \frac{\partial \mathbb{E}_L(t|\hat{k}_H, \hat{y}_H)}{\partial y_H} = 0 \\
\Rightarrow & -c \frac{\partial \mathbb{E}_H(t|\hat{k}_H, \hat{y}_H)}{\partial y_H} - \frac{\partial \mathbb{E}_H(p|\hat{k}_H, \hat{y}_H)}{\partial y_H} = -\frac{\lambda}{\rho_0} \frac{\partial \mathbb{E}_L(t|\hat{k}_H, \hat{y}_H)}{\partial y_H}.
\end{aligned}$$

We have $-c\frac{\partial \mathbb{E}_H(t|k_H, y_H)}{\partial y_H} - \frac{\partial \mathbb{E}_H(p|k_H, y_H)}{\partial y_H} = \frac{\partial W_H(k_H, y_H)}{\partial y_H}$. We know $\frac{\partial \mathbb{E}_H(t|k_H, y_H)}{\partial y_H} < 0$, $\frac{\partial \mathbb{E}_H(p|k_H, y_H)}{\partial y_H} > 0$. Moreover, we know $\lambda > 0$ and $\rho_0 > 0$. This implies $\frac{\partial W_H}{\partial y_H}(\hat{k}_H, \hat{y}_H) > 0$. Furthermore, we know that $\frac{\partial W_H(k_H, y_H)}{\partial y_H}$ is maximized at y_H^* for $k_H = 0$. Using lemma 3, we can infer that the same holds for any $k_H > 0$. Therefore, W_H increases for $y < y_H^*$ and decreases for $y > y_H^*$. Thus, it must be that $\hat{y}_H < y_H^*$.

The optimal threshold in state H is lower than the first best threshold. This means that it is optimal to keep searching, even when the cost of a search is greater than the expected saving in price. The decrease in search cost is larger than the increase in price as the threshold increases.

Finally, we show that Assumption A2 implies that ICH is satisfied for the menu $\left\{(0, \hat{y}_L), \left(\hat{k}_H, \hat{y}_H\right)\right\}$. Clearly, $= \hat{y}_L < \hat{y}_H$, otherwise a common threshold without a minimum number of offers would be optimal for the principal.

From Assumption A2 we get the following result:

$$
\begin{aligned}
&\frac{1}{F_H(\hat{y}_L)}-\frac{1}{F_L(\hat{y}_L)}\geq\frac{1}{F_H(\hat{y}_H)}-\frac{1}{F_L(\hat{y}_H)}\\
\Rightarrow&\frac{1}{F_H(\hat{y}_L)}-\frac{1}{F_L(\hat{y}_L)}\geq\frac{1}{F_H(\hat{y}_H)}-\frac{1}{F_L(\hat{y}_H)}\geq\frac{\left[1-F_H\left(\hat{y}_H\right)\right]^{\hat{k}_H}}{F_H\left(\hat{y}_H\right)}-\frac{\left[1-F_L\left(\hat{y}_H\right)\right]^{\hat{k}_H}}{F_L\left(\hat{y}_H\right)}\\
\Rightarrow&\frac{1}{F_L\left(\hat{y}_L\right)}-\frac{\left[1-F_L\left(\hat{y}_H\right)\right]^{\hat{k}_H}}{F_L\left(\hat{y}_H\right)}\leq\frac{1}{F_H\left(\hat{y}_L\right)}-\frac{\left[1-F_H\left(\hat{y}_H\right)\right]^{\hat{k}_H}}{F_H\left(\hat{y}_H\right)}\\
\Rightarrow&\hat{k}_H\leq\frac{1}{F_H\left(y_L\right)}-\frac{\left[1-F_H\left(\hat{y}_H\right)\right]^{\hat{k}_H}}{F_H\left(\hat{y}_H\right)}\\
\Rightarrow&\hat{k}_H+\frac{\left[1-F_H(\hat{y}_H)\right]^{\hat{k}_H}}{F_H(\hat{y}_H)}\leq\frac{1}{F_H(\hat{y}_L)}\\
\Rightarrow&\mathbb{E}_H(t|\hat{k}_H,\hat{y}_H)\leq\mathbb{E}_H(t|0,\hat{y}_L).
\end{aligned}
$$

□

Example 1: Assume that in state L the set of prices is $L=\{5,10,20\}$, with probability distribution $(\frac{1}{3},\frac{1}{3},\frac{1}{3})$ and in state H the price set is $H=\{20,40,60\}$ with probability distribution $(\frac{1}{10},\frac{1}{10},\frac{8}{10})$. The value of the object is 25 in state L and and 100 in state H. The search cost is $c=2$. Also, assume that the state is L with probability $\frac{2}{3}$ and H with probability $\frac{1}{3}$. It is easy to check that the first best thresholds are $y_L^*=10$ and $y_H^*=20$. In state L, the resulting expected search duration is 1.5 and the expected price is 7.5. In state H, the resulting expected search duration in state H is 10 and the expected price is 20. The principal's expected first best payoff is 29.666.

When the principal does not observe the state, the first best thresholds $y_L^*=10$ and $y_H^*=20$ violate the incentive compatibility condition. This is because in state L, the agent would prefer to announce state H to get a threshold of 20, which reduces the search duration he has to perform to 1. The principal's optimal pooling rule sets a threshold of 10 for the first search and 20 after that, i.e. $\{y_t\}=(10,20,20,\ldots)$. The expected search duration in state $a\in\{H,L\}$ is

$$1+\frac{1-F_a(10)}{F_a(20)}.$$

This implies that expected search duration is $\dfrac{4}{3}$ in state L and 11 in state H. The expected price in state H is clearly 20. In state L it is derived as follows: If the agent finds a price below 10 in the first search the price is either 5 or 10 (each with probability $\dfrac{1}{2}$), otherwise the price will be either 5, 10, or 20 (each with probability $\frac{1}{3}$), therefore we have

$$F_L(10)\left(\frac{1}{2}\cdot 5+\frac{1}{2}\cdot 10\right)+(1-F_L(10))\left(\frac{1}{3}\cdot 5+\frac{1}{3}\cdot 10+\frac{1}{3}\cdot 20\right)=\frac{5}{3}+\frac{10}{3}+\frac{5}{9}+\frac{10}{9}+\frac{20}{9}=\frac{80}{9}.$$

The resulting payoff is

$$\frac{2}{3}(25-\frac{4}{3}(2)-\frac{80}{9})+\frac{1}{3}(100-11(2)-20)=(\frac{2}{3})(\frac{121}{9})+(\frac{1}{3})(58)=\frac{242+522}{27}\approx 28.286.$$

Now, consider the separating menu $\hat{R}=(10,(2,20))$. The rule for state L sets a threshold of $y_L^*=10$ and the rule for state H sets a minimum of 2 offers and a threshold of $y_H^*=20$. This menu is clearly incentive compatible. The expected search duration in state H is $2+(1-\frac{1}{10})^2(10)=10.1$.

The principal's expected payoff is

$$\frac{2}{3}(25-(1.5)(2)-\frac{15}{2})+\frac{1}{3}(100-(10.1)2-20)=(\frac{2}{3})(\frac{29}{2})+(\frac{1}{3})(59.8)=\frac{88.8}{3}=29.6$$

This is higher than the expected payoff of the principal with the optimal pooling threshold.

References

Armstrong, M. and Vickers, J. (2010). A model of delegated project choice, *Econometrica* **78**(1): 213–244.

Bikhchandani, S. and Sharma, S. (1996). Optimal search with learning, *Journal of Economic Dynamics and Control* **20**(1-3): 333–359.

Kovác, E., Krähmer, D. and Tatur, T. (2013). Optimal stopping in a principal-agent model with hidden information and no monetary transfers. mimeo.

Krähmer, D. and Kováč, E. (2016). Optimal sequential delegation, *Journal of Economic Theory* **163**: 849–888.

Lewis, T. R. (2012). A theory of delegated search for the best alternative, *The RAND Journal of Economics* **43**(3): 391–416.

Lewis, T. R. and Ottaviani, M. (2008). Search agency. mimeo.

Mauring, E. (2016). A two-agent model of sequential search and choice, *Journal of Economic Behavior & Organization* **123**: 122–137.

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

Postl, P. (2004). Delegated search: Procedure matters. Department of Economics Discussion Paper, University of Birmingham.

Rothschild, M. (1978). Searching for the lowest price when the distribution of prices is unknown, *Uncertainty in Economics*, Elsevier, pp. 425–454.

Ulbricht, R. (2016). Optimal delegated search with adverse selection and moral hazard, *Theoretical Economics* **11**(1): 253–278.