The Loser's Curse in the Search for Advice

Pak Hung Au*

December 5, 2016

Abstract

An agent searches sequentially for advice from experts concerning the payoff of taking an operation. Consulting each expert incurs a positive search cost. There are infinitely many experts, each has access to an identically and conditionally independent signal structure about the payoff, and makes a recommendation after observing the signal realization. The interests of the experts and the agent are partially aligned. We show that a loser's curse effect is present, hampering the quality of information gathered. Conditions that ensure information aggregation as the search cost vanishes are identified. The more surprising finding is that both the agent's payoff and the total efficiency can be strictly lower than the alternative scenario in which the agent can consult only a single expert.

Keywords Search, Expert Advice, Information Transmission, Information Aggregation

JEL Codes D83

1 Introduction

A patient with some medical symptoms is unsure whether a surgery is an appropriate treatment. For advice, he consults a doctor, who diagnoses the case, and makes a recommendation. Suppose the interests of the doctor and the patient are aligned, so the patient's only concern is that the doctor's diagnosis is not accurate (rather than the incentives for lying). He may consult another doctor for a second opinion. If the recommendations of the two doctors match, then he is more confident about the appropriate course of treatment. Nonetheless, he may continue the process of consulting other doctors for more accurate information. If he is eventually confident enough that the surgery is necessary, then

*Assistant Professor, Nanyang Technological University, email: phau@ntu.edu.sg.

he undergoes the surgery with one of the doctors that recommend him to do so. It is natural to expect that by consulting more doctors, the patient could gather better information about the appropriate treatment. A similar scenario arises in the process of a customer looking for a repair service, and an entrepreneur looking for financial investment of a venture capitalist.

In this paper, we analyze a model in which an agent sequentially consults experts for advice on whether to undergo an operation or not. His payoff of having the operation is uncertain, taking a positive value if the operation is suitable for him (an event denoted by the state $\omega = 1$), and a negative value if the operation is unsuitable for him (an event denoted by the state $\omega = 0$). There are infinitely many potential experts for the agent to consult, and each expert has access to an identical and conditionally independent signal structure for learning about the agent's state ω. For each visit and consultation of an expert, the agent has to incur a positive search cost. An expert consulted makes a recommendation after privately learning her own signal realization, but not those of previously consulted experts. If she recommends for the operation, the agent can decide whether to undergo the operation with her, or seek more advice from other experts. If she recommends against the operation, then the agent cannot have the operation with her (but he may still consult another expert). Eventually, the agent may undergo the operation with one of the experts that recommend him to do so, or quit the process of consulting experts without taking the operation.

The payoffs of the agent and the experts depend on the state and whether the operation is performed. We assume the interests of the experts and the agent are partially aligned. More specifically, if the agent undergoes the operation with an expert, then the sign of their payoffs are equal. On the other hand, if the agent does not undergo the operation with her, then she gets a zero payoff. For simplicity, the payoff structure is exogenously fixed, so we have abstracted away from considerations such as the pricing of the operation, and bargaining over the division of liability of failed operation (i.e., the operation is carried out with the agent's state being $\omega = 0$).

Our objective is to investigate the equilibrium outcome and welfare consequence of allowing the agent to sequentially search for experts' advice. In particular, we would like to compare the agent's welfare and the social welfare under two scenarios: (i) a benchmark setting in which there is only one expert available for consultation; and (ii) a setting in which there are infinitely many experts available for consultation at an infinitesimal search cost. At first sight, it seems immediate that the agent would prefer the second scenario, as he would be able to learn the state almost costlessly. Our analysis shows that whether this conjecture holds or not depends crucially on the experts' signal structure. We find that there are informative signal structures under which the agent would not be able to learn the true state in scenario (ii) above, despite having access almost costlessly to infinitely many experts. More strikingly, with some payoff and informative signal structures, both the agent's payoff and the social welfare are strictly lower in scenario (ii) than in scenario (i).

The key driving force behind our results is a *loser's curse* effect: in an equilibrium in which the agent searches for advice, each expert understands that the agent decides to undergo the operation if and only if he has received sufficiently favorable information from other experts, and that her recommendation affects her payoff if and only if she is pivotal. As a result, an *endogenous conflict of interest* can arise between the agent and an expert, even if their interests are perfectly aligned in a one-shot game without search. In particular, the loser's curse effect leads each expert to adopt a recommendation rule that is more lenient than that would be adopted were she base her recommendation only on her privately observed signal. This in turn hurts the agent because it worsens the quality of the information gathered by consulting multiple experts. We find that as the search cost vanishes, the loser's curse effect becomes extremely severe: every expert almost always discards her private signal and recommends for the operation.

It is natural that the comparison of agent's welfare and efficiency mentioned above depends crucially on whether information is perfectly aggregated, i.e., whether the agent is able to learn the true state almost surely or not. We find that perfect information aggregation arises as an limit equilibrium outcome only if there exists a signal that completely reveals that the agent is not suitable for operation (i.e., $\omega = 0$). The intuition is as follows. Because of the loser's curse effect, the experts very often recommend for the operation. A recommendation against the operation is thus a very strong and clear signal that the agent is indeed unsuitable, under a signal structure with a fully revealing signal that $\omega = 0$. With a negligibly small search cost, an agent with $\omega = 0$ would eventually be able to find a negative recommendation from a large sample of consulted experts. With such a signal structure, having (almost) free access to multiple experts always benefits the agent; in fact, both the agent's payoff and the social welfare can almost achieve the highest feasible level.

On the other hand, the agent may take the ex-post incorrect action with positive probability if the expert's signal structure does not contain a signal that fully reveals $\omega = 0$. Furthermore, if the relative loss of the expert in the case of a failed operation is lower than that of the agent, there are signal structures under which both the agent's equilibrium payoff and the social welfare are strictly lower than the alternative scenario of a single expert. The intuition is that the loser's curse is particularly strong if the expert's liability is small, making the quality of information gathered through search very poor.

Our result therefore suggests that if the agent has access to multiple experts, it is possible that both his own payoff and the social welfare can be improved if he could commit to consult a single expert only. Note that this result does not arise from the costs involved in search and diagnosis, as they are assumed to be vanishingly small. Instead, it arises because the loser's curse effect mentioned above becomes so severe in the limit that the information transmission becomes highly ineffective.

1.1 Related Literature

Our model has some flavour of the credence goods market: after performing a diagnosis, the expert is more informed about the type and quality of the service the customer needs.[1] Inefficiency in the competitive market for credence goods has been studied in Wolinsky (1993), Pesendorfer and Wolinsky (2003), Wolinsky (2005), and Alger and Salanié (2006). Wolinsky (1993) shows that competition can mitigate experts' incentives to prescribe overtreatment (i.e., providing unnecessarily expensive treatment) if there are firms that specialize in providing low-cost repair. Alger and Salanié (2006) show that price competition in the low-cost repair induces overtreatment. In contrast, our model does not feature any price competition, and overtreatment is not a problem here. Pesendorfer and Wolinsky (2003) analyze a setting in which experts need to exert effort to learn about the treatment needed, and show that price competition leads to inefficient effort exertion. Wolinsky (2005) analyzes a setting in which experts exert effort to devise an appropriate plan for the customer. Inefficiency arises because the customer does not internalize the effort cost of the experts sampled. Our model, on the other hand, does not have any moral hazard.

Our model also bears similarities with a common-value auction, as the state of the agent is common to all experts. Milgrom (1979) identifies necessary and sufficient conditions for full information aggregation as the number of bidders grow to infinity. Our condition has a similar flavor, though a tighter condition is necessary to ensure the uniqueness of equilibrium outcome. Pesendorfer and Swinkels (1997) analyze information aggregation in a common-value auction with a large supply. They show that the winner's curse and the loser's curse balance each other, leading to perfect information aggregation. In the context of our model, it is reasonable to fix the supply at one, as each agent's problem is independent. Several studies found that in a common-value auction, an increase in the number of bidders may decrease the expected winning bid, a result due to the effect of the winner's curse. Bulow and Klemperer (2002) illustrate this possibility with asymmetric bidders. Hong and Shum (2002) illustrate this with a small number of bidders. In contrast, we allow the number of symmetric experts to grow to infinity. Moreover,

[1] See Dulleck and Kerschbamer (2006) for a survey.

we show that not only the agent's payoff (which corresponds to the seller's revenue) goes down with a lower search cost (which corresponds to having more bidders), but also the total efficiency.

The outline of the paper is as follows. The model is set up in Section 2. We analyze two benchmark cases in Section 3. The first benchmark assumes there is only one available expert. The second benchmark assumes experts mistakenly believes that they are the only expert consulted. In the main analysis of the model in Section 4, we first establish equilibrium existence and provide some characterizations. We then consider the limit equilibrium in which the search cost vanishes. Section 5 discusses a few alternative settings. Lengthy proofs are relegated to the appendix.

2 Model

A (male) agent can either undergo an operation (denoted by $a = 1$) or not (denoted by $a = 0$). His payoff of undergoing the operation depends on a binary state of the world $\omega \in \{0, 1\}$. If the state is $\omega = 1$ ($\omega = 0$), then the operation is suitable (unsuitable) for the agent. His prior belief about the state is denoted by $\pi \equiv \Pr(\omega = 1) \in (0, 1)$. The operation must be carried out by an (female) expert. There are infinitely many ex-ante identical experts. In each of the infinitely many periods, the agent can visit one expert. Upon a visit, each expert conducts a test which generates an informative signal about the state ω. After privately observing the signal, she then makes a recommendation to the agent. The experts have a common payoff function, as well as a common information acquisition technology. The agent has free access to one expert. For each additional visit and consultation of other experts, the agent has to incur a search cost of $c > 0$. Each of the infinitely many experts is drawn (without recall) with equal probability in every period.

Each expert can run a test costlessly to obtain a signal $s \in [\underline{s}, \bar{s}] \subset [0, \infty]$ about ω. The signal of each expert is distributed identically and independently (conditional on the state ω). Specifically, the signal is generated according to conditional density function $f(s|\omega)$, with corresponding conditional distribution function $F(s|\omega)$. It is without loss to label the signals as their corresponding likelihood ratios, i.e., $s \equiv \frac{f(s|1)}{f(s|0)}$. With this labelling, a high signal is more indicative of $\omega = 1$. Moreover, the signal structure is informative if and only if $\underline{s} \in [0, 1)$ and $\bar{s} \in (1, \infty]$. In addition to admitting a conditional density function, we further assume that $f(s|\omega)$ has full support on $[\underline{s}, \bar{s}]$, for $\omega = 0, 1$.

The signal realization of the test is unverifiable and observed privately by the consulted expert. Moreover, we assume that the signal realization is so complicated that it is infeasible to communicate its full content to the agent. Instead, each expert makes a binary recommendation of having the operation or not. Denote the set of recommendations by $\{Y, N\}$, where Y stands for a non-binding recommendation for the operation, and N stands for a recommendation against the operation. If the expert recommends against the operation (i.e. recommendation N), it means she refuses to carry out the operation for the agent, who must then part with the expert. On the other hand, if the expert suggests to have the operation (i.e. recommendation Y), it means she is willing to perform the operation for the agent, who can choose to whether or not to have the operation with the expert. In other words, a recommendation N is a rejection by the expert; whereas a recommendation Y means the expert provides an option for the agent to undergo the operation with her. If the expert makes recommendation Y and the agent agrees to have the operation, then they collect their respective payoffs to be described below, and the game is over. If the agent chooses not to have the operation, then he parts with the expert. He can then either consult another expert, or stop the search for advice altogether.

The payoff received by the agent and the expert for different scenarios are tabulated below:

		Action	
		$a = 1$	$a = 0$
State	$\omega = 1$	$1, 1$	$0, 0$
	$\omega = 0$	$-L, -l$	$0, 0$

Here, L and l are positive numbers. The agent receives a positive payoff normalized to one if the state is $\omega = 1$ and the operation is carried out. He suffers a loss L if the state is $\omega = 0$ and the operation is carried out. His payoff of not undergoing the operation is normalized to 0. On the other hand, the expert who perform the operation for the agent has a partially-aligned payoff function: a positive payoff normalized to one if the agent is indeed suitable for the operation; a negative payoff equals $-l$ if the agent is not suitable. Finally, her payoff is normalized to zero if she does not carry out the operation for the agent. If $L = l$, then the interest of the expert and the agents are perfectly aligned.

We assume the agent cannot communicate his history of recommendations received from previously consulted experts. One justification for this assumption is that if the message about this history is cheap talk, there is always an equilibrium in which such communication is babbling. Moreover, in an alternative model specification in which each expert has to incur a positive effort cost in running the test and diagnosis, the only equilibrium outcome is babbling, regardless of how small the effort cost is. The reason is that an expert would exert effort if and only if she believes that she is sufficiently likely to be pivotal. Therefore, the agent has incentives to lie about the history to increase an expert's belief that she is pivotal, rendering such communication completely ineffective.[2] Furthermore, we assume the

[2] In our analysis, we abstract away effort cost in running test and diagnosis. Introducing such cost would further reduce

search process is without recall. This is an innocuous assumption as all experts are identical replica of each other. If there is a small positive cost of returning to a previously visited expert, the agent strictly prefers to carry out the operation with the current expert.

To summarize, the timeline of game is as follows. Events unfold in each period in the following order:

1. The agent decides whether to consult an expert (that he has not visited before) or not. If he chooses not to, nothing else happens in this period. If he decides to consult an expert, he has to incur a search cost c.

2. The expert consulted privately observes a signal s about the state ω. She then makes a recommendation for or against the operation, denoted as recommendations Y and N respectively. If she recommends N, then she refuses to carry out the operation for the agent, and nothing else happens in this period.

3. If she recommends Y, then the agent chooses whether or not to have the operation with the expert. If he does, they collect their respective state-dependent payoffs. Otherwise, the expert collects a zero payoff and her role in the game is over.

In our model, a pure strategy of an expert is the set of signals, denoted by A, under which she recommends the operation (i.e., choosing Y). A history of the agent after consulting n experts is a sequence of recommendations $\{Y, N\}^n$. Denote the set of all possible recommendation histories by $H = \cup_{n \in \mathbb{N} \cup \{0\}} \{Y, N\}^n$. At the beginning of a period, the agent decides whether to visit an expert that he has not consulted before. At the end of a period, if he has consulted an expert and the expert recommends Y, then he decides whether to have the operation with the current expert or not. A pure strategy of the agent, denoted by $\beta = (\beta_0, \beta_1)$, consists of two components, both of which are mappings from H to $\{0, 1\}$: $\beta_0(h) = 0$ ($\beta_0(h) = 1$) stands for the decision to search (not to search) at the beginning of a period, provided that the current history is h; whereas $\beta_1(h) = 1$ ($\beta_1(h) = 0$) stands for the decision of having the operation (not having the operation) at the end of a period after being recommended Y by the current expert.[3] Denote the set of all pure strategies of the agent by Λ, and the set of all mixed strategies by $\triangle\Lambda$.

the quality of information transmission in the search for advice through a mechanism similar to Pesendorfer and Wolinsky (2003).

[3] In short, for $i = 0, 1$, $\beta_i(h) = 1$ stands for stopping the search; and $\beta_i(h) = 0$ stands for continuing the search.

We focus on a sequential equilibrium in which all experts play an identical pure strategy. The requirement of sequential equilibrium rules out the trivial weak perfect Bayesian equilibrium that all experts make a negative recommendation, anticipating that the agent chooses not to have the operation.

3 Benchmark Models

We consider two benchmark models, which are useful for later comparison of results. First, we consider the case in which the agent can consult only one expert. Conditions under which information is transmitted are identified. We also compute the highest agent welfare and social efficiency among the class of feasible signal structures. Second, we consider the scenario in which the agent has access to infinitely many experts, but each expert adopts a fixed (not necessarily optimal) recommendation strategy. We show in this case, as the search cost vanishes, the agent necessarily learns the true state ω with probability arbitrarily close to one.

3.1 Benchmark I: Single Expert

Suppose there is only a single expert. The expert, after receiving signal s, updates her belief on ω according to Bayes' rule:

$$\Pr(\omega = 1|s) = \frac{\pi f(s|1)}{\pi f(s|1) + (1-\pi) f(s|0)} = \frac{1}{1 + \left(\frac{1-\pi}{\pi}\right)\frac{1}{s}}.$$

It is optimal for the expert to recommend the operation if

$$\Pr(\omega = 1|s) - l \Pr(\omega = 0|s) \geq 0 \Leftrightarrow s \geq \frac{1-\pi}{\pi} l \equiv \tilde{s}. \tag{1}$$

If the agent follows the expert's recommendation with positive probability, then it is strictly optimal to recommend the operation whenever $s > \tilde{s}$. Therefore, in each sequential equilibrium, the expert recommends the operation if and only if $s \geq \tilde{s}$. As a result, the expert's recommendation is (partially) informative if and only if $\tilde{s} \in (\underline{s}, \bar{s})$, or equivalently, $\pi \in \left(\frac{l}{\bar{s}+l}, \frac{l}{\underline{s}+l}\right)$.

Conditional on a positive recommendation, the agent's payoff of taking the operation is

$$\Pr(\omega = 1|s \geq \tilde{s}) + (-L) \Pr(\omega = 0|s \geq \tilde{s}) = -L + \frac{1}{1 + \frac{1-\pi}{\pi}\frac{1-F(\tilde{s}|0)}{1-F(\tilde{s}|1)}}(1+L).$$

The agent follows the recommendation if the payoff above is nonnegative, or equivalently,

$$\frac{1-\pi}{\pi}\frac{1-F(\tilde{s}|0)}{1-F(\tilde{s}|1)} \leq \frac{1}{L}. \tag{2}$$

Consequently, there exists an equilibrium in which the agent always follows the expert's recommendation if and only if (2) holds. If (2) fails, then the only equilibrium outcome is that the agent never takes the operation. The observations above imply that the agent's payoff, denoted by $U(\pi)$, is given by

$$U(\pi) = \begin{cases} \pi\left(1 - F\left(\frac{1-\pi}{\pi}l|1\right)\right) - L(1-\pi)\left(1 - F\left(\frac{1-\pi}{\pi}l|0\right)\right) & \text{if (2) holds and } \pi \in \left(\frac{l}{\bar{s}+l}, \frac{l}{\underline{s}+l}\right); \\ \pi - L(1-\pi) & \text{if } \pi > \max\left\{\frac{l}{l+\underline{s}}, \frac{L}{1+L}\right\}; \\ 0 & \text{otherwise.} \end{cases} \tag{3}$$

Similarly, the total payoff of the expert and the agent, denoted by $T(\pi)$ is given by

$$T(\pi) = \begin{cases} 2\pi\left(1 - F\left(\frac{1-\pi}{\pi}l|1\right)\right) - (L+l)(1-\pi)\left(1 - F\left(\frac{1-\pi}{\pi}l|0\right)\right) & \text{if (2) holds and } \pi \in \left(\frac{l}{\bar{s}+l}, \frac{l}{\underline{s}+l}\right); \\ 2\pi - (L+l)(1-\pi) & \text{if } \pi > \max\left\{\frac{l}{l+\underline{s}}, \frac{L}{1+L}\right\}; \\ 0 & \text{otherwise.} \end{cases} \tag{4}$$

3.1.1 Bounds on Payoffs

For later comparison, we derive tight upper bounds on $U(\pi)$ and $T(\pi)$ **among all permissible signal structures** $(f(s|1), f(s|0))$**, fixing the lower bound of the signal space** $\underline{s} \in [0,1)$**.** If $\underline{s} \geq \tilde{s}$, or equivalently, $\pi \geq \frac{l}{l+\underline{s}}$, the expert always recommends the operation. Therefore, the payoff functions $U(\pi)$ and $T(\pi)$ are independent of the details of the signal structure, and given by (3) and (4) above. Below we focus on the case $\pi < \frac{l}{l+\underline{s}}$.

Suppose for the time being that a discrete signal structure is permitted, i.e., the signal space can be finite. The most informative signal structure, for a fixed lower bound of the signal space $\underline{s}$, is denoted by $F_{\underline{s}}$. It has a binary support: $\{\underline{s}, \bar{s}\}$, where $\bar{s}$ fully reveals the state being $\omega = 1$. That is, $F_{\underline{s}}(\underline{s}|0) = 1$, and

$$F_{\underline{s}}(s|1) = \begin{cases} 0 & \text{if } s < \underline{s} \\ \underline{s} & \text{if } s \in [\underline{s}, \bar{s}] \\ 1 & \text{if } s = \bar{s} \end{cases}.$$

Lemma 1 *Fix the lower bound of the expert's signal space at* $\underline{s}$*, and suppose a discrete signal structure is permitted.*

(i) If either (a) $l \leq L$*, or (b)* $l > L$ *and* $\pi \leq \frac{L}{\underline{s}+L}$*, then* $U(\pi)$ *is maximized if the expert's signal structure is the most informative, i.e.,* $F_{\underline{s}}$*.*

(ii) If $l > L$ *and* $\pi \in \left(\max\left\{\frac{L}{\underline{s}+L}, \frac{l}{1+l}\right\}, \frac{l}{l+\underline{s}}\right)$*, then* $U(\pi)$ *is maximized if the expert's signal structure is completely uninformative.*

(iii) If $l > L$, $\frac{L}{L+\underline{s}} < \frac{l}{l+1}$, *and* $\pi \in \left(\frac{L}{L+\underline{s}}, \frac{l}{l+1}\right]$, *then* $U(\pi)$ *is maximized if the expert's signal structure is partially informative with a binary support* $\left\{\underline{s}, \frac{1-\pi}{\pi} l\right\}$.

An upper bound on $U(\pi)$ *for the benchmark game is given by* $\bar{U}_{\underline{s}} : [0,1] \to [0,1]$ *defined as follows.*

$$\bar{U}_{\underline{s}}(\pi) \equiv \begin{cases} \pi(1-\underline{s}) & \text{if } l \leq L, \text{ or if } l > L \text{ and } \pi \leq \frac{L}{\underline{s}+L} \\ \pi(1-\pi)(1-\underline{s})\frac{l-L}{l-\pi(l+\underline{s})} & \text{if } \frac{L}{L+\underline{s}} < \frac{l}{l+1} \text{ and } \frac{L}{L+\underline{s}} < \pi \leq \frac{l}{l+1} \\ -L+\pi(1+L) & \text{if } \max\left\{\frac{l}{l+1}, \frac{L}{L+\underline{s}}\right\} < \pi \end{cases}.$$

The signal structure that maximizes $U(\pi)$ can be computed by considering the following auxiliary game. Take the same setting as the first benchmark model, but allow the agent to choose the signal structure of the expert; and mandate the expert to report truthfully the learned signal along with her recommendation. As the agent in this auxiliary game can never be worse off than his counterpart in the benchmark model, solving the optimal signal structure of this auxiliary game gives an upper bound to $U(\pi)$. The fact that the upper bound is achievable follows from the simplicity of the optimal signal structure: the agent does not need to know the realized signal to implement the optimal decision rule.

Lemma 1 is quite intuitive. As the interests of the agent and the expert are partially aligned, it is not surprising that the agent would prefer an informative expert for a wide range of parameters, as shown in case (i). In cases (ii) and (iii), as $l > L$, the expert is overly cautious from the agent's perspective. The agent would therefore prefer a signal structure that maximizes the probability that the operation is recommended.

It is not difficult to see that the signal structures in Lemma 1 can be approximated arbitrarily well using signal structures with a connected support $[\underline{s}, \infty]$, and well-defined conditional density functions. The lemma below explicitly constructs the required approximation.

Lemma 2 *Fix a* $\underline{s} \in [0,1)$, *a prior* $\pi < \frac{l}{\underline{s}+l}$, *and a positive real number* $\varepsilon > 0$. *There exists a pair of conditional density functions* $f(\cdot|1)$ *and* $f(\cdot|0)$ *with support* $[\underline{s}, \infty]$ *such that the agent's payoff in the benchmark game exceeds* $\bar{U}_{\underline{s}}(\pi) - \varepsilon$.

The two lemmas above therefore imply that $\bar{U}_{\underline{s}}(\pi)$ is indeed a tight upper bound for the agent's payoff $U(\pi)$.

Proposition 1 *Suppose* $\pi < \frac{l}{l+\underline{s}}$ *and the lower bound on the signal space is* $\underline{s} \in [0,1)$. *For each* $\varepsilon > 0$, *there exists a signal structure (with a connected support and conditional density functions) under which the agent's equilibrium payoff in the benchmark game is no less than* $\bar{U}_{\underline{s}}(\pi) - \varepsilon$.

Following a similar analysis, the tight upper bound on the total payoff of the agent and the expert can be obtained.

Proposition 2 *Suppose $\pi < \frac{l}{\underline{s}+l}$ and the lower bound on the signal space is $\underline{s} \in [0,1)$. Suppose further that a discrete signal structure is permitted.*

(i) If $l \leq L$, or if $l > L$ and $\pi \leq \frac{L+l}{L+l+2\underline{s}}$, then $T(\pi)$ is maximized by the most informative signal structures $F_{\underline{s}}$.

(ii) If $l > L$, $\frac{L+l}{2\underline{s}+L+l} < \frac{l}{1+l}$ and $\frac{L+l}{2\underline{s}+L+l} < \pi < \frac{l}{1+l}$, then $T(\pi)$ is maximized by a signal structure with a binary support $\left\{\underline{s}, \frac{1-\pi}{\pi} l\right\}$.

(iii) If $l > L$ and $\max\left\{\frac{L+l}{2\underline{s}+L+l}, \frac{l}{1+l}\right\} \leq \pi < \frac{l}{\underline{s}+l}$, then $T(\pi)$ is maximized by a completely uninformative signal structure.

An upper bound on $T(\pi)$ for the benchmark game is given by $\bar{T}_{\underline{s}} : [0,1] \to [0,1]$ defined as follows.

$$\bar{T}_{\underline{s}}(\pi) = \begin{cases} 2\pi(1-\underline{s}) & \text{if } l \leq L, \text{ or if } l > L \text{ and } \pi \leq \frac{L+l}{L+l+2\underline{s}} \\ \pi(1-\pi)L(2l+1)\frac{1-\underline{s}}{(1-\pi)l-\pi\underline{s}} & \text{if } l > L,\ \frac{L+l}{2\underline{s}+L+l} < \frac{l}{1+l} \text{ and } \frac{L+l}{2\underline{s}+L+l} < \pi < \frac{l}{1+l} \\ -(l+L)+\pi(2+l+L) & \text{if } l > L \text{ and } \max\left\{\frac{L+l}{2\underline{s}+L+l}, \frac{l}{1+l}\right\} \leq \pi < \frac{l}{\underline{s}+l} \end{cases}.$$

Furthermore, for each $\varepsilon > 0$, there exists a signal structure (with a connected support and conditional density functions) under which the total payoff in the benchmark game is no less than $\bar{T}_{\underline{s}}(\pi) - \varepsilon$.

3.2 Benchmark II: Non-strategic Experts

In this subsection, we consider an alternative benchmark in which the agent can sequentially consult multiple experts for advice, but each expert (mistakenly) believes that she is the only expert consulted by the agent. The main purpose is to illustrate that as the search cost c vanishes, the agent learns the true state ω with probability arbitrarily close to one, so enjoys an ex-ante expected payoff arbitrarily close to π. Moreover, the optimal search strategy of the agent characterized in this benchmark plays a key role in the subsequent analysis.

We begin by noting that the optimal search strategy of the agent has a very simple structure:

Lemma 3 *Denote by p the agent's belief that $\omega = 1$, given the history of recommendations received. The optimal search strategy of the agent in this benchmark game is characterized by a pair of posteriors $p_0, p_1 \in [0,1]$, with $p_0 \leq p_1$ such that the agent quits searching if $p < p_0$; takes the operation with the current expert if $p > p_1$; and continues searching if $p \in (p_0, p_1)$. At $p = p_0$, the agent is indifferent between quitting and continuing search. At $p = p_1$, the agent is indifferent between having the operation and continuing search.*

According to the lemma, the agent's optimal search strategy has the Markov property, with the state being the agent's current belief. This is very intuitive, as each expert's recommendation strategy is invariant of her position in the agent's order of search.

A sharper characterization of the agent's optimal search strategy can be obtained by computing his continuation value function. Recall from (1) that, if the expert believes that she is the only expert consulted, then she finds it optimal to recommend the operation (i.e., making recommendation Y) if and only if her observed signal is no less than $\tilde{s}$. Therefore, if the agent approaches the expert with a prior belief p, then conditional on a positive recommendation Y, his posterior belief becomes $\left(1+\frac{1-p}{p}\frac{1-F(\tilde{s}|0)}{1-F(\tilde{s}|1)}\right)^{-1}$; whereas conditional on a negative recommendation N, his posterior belief becomes $\left(1+\frac{1-p}{p}\frac{F(\tilde{s}|0)}{F(\tilde{s}|1)}\right)^{-1}$. Denote by $V:[0,1]\to\mathbb{R}$ the agent's beginning-of-period continuation value, as a function of his current belief $p\in[0,1]$. Using observations above, $V(p)$ can be recursively defined by

$$V(p)=\max\left\{\begin{array}{c} 0, -c+\left[p\left(1-F\left(\tilde{s}|1\right)\right)+\left(1-p\right)\left(1-F\left(\tilde{s}|0\right)\right)\right] \\ \times\max\left\{V\left(\frac{1}{1+\frac{1-p}{p}\frac{1-F(\tilde{s}|0)}{1-F(\tilde{s}|1)}}\right), -L+\frac{1}{1+\frac{1-p}{p}\frac{1-F(\tilde{s}|0)}{1-F(\tilde{s}|1)}}\left(1+L\right)\right\} \\ +\left[pF\left(\tilde{s}|1\right)+\left(1-p\right)F\left(\tilde{s}|0\right)\right]V\left(\frac{1}{1+\frac{1-p}{p}\frac{F(\tilde{s}|0)}{F(\tilde{s}|1)}}\right) \end{array}\right\}. \tag{5}$$

To understand expression (5), note that after paying a search cost c, the agent may receive a recommendation Y, at which point, he can either (i) leave the current expert without undergoing the operation; or (ii) agree to have the operation with the current expert. If he receives a recommendation N, then he must leave the current expert without undergoing the operation.

Given the value function V, the agent's optimal search strategy can be easily computed by solving for p_0 and p_1: $p_0=\max\{p\in[0,1]:V(p)=0\}$; and p_1 is the solution to $V(p)=-L+p(1+L)$. It remains to identify the function V.

Proposition 3 *There exists a unique value function $V:[0,1]\to\mathbb{R}$ that satisfies (5). Moreover, V is nondecreasing and weakly convex.*

The proof of Proposition 3 provides an algorithm for computing the value function V recursively. Specifically, denote by V_n the value function if the agent can consult at most n experts. Using observations following (5), the sequence of functions $\{V_n(p)\}$ can be iteratively defined. The proof of Proposition 3 shows that the sequence of functions $\{V_n\}$ converges uniformly to V.

For later comparison, we compute the agent's payoff in this scenario if the search cost is vanishingly small. It is straightforward to verify that $V(p)=p$ is the unique solution to (5) if $c=0$. Therefore, the agent's ex-ante payoff is just his prior belief π. Moreover, as $c\to 0$, we have $p_1\to 1$ and $p_0\to 0$: the agent continues searching for advice until he is almost certain of his state ω. In other words, information is fully aggregated in the limit. The intuition of these results is simple: if the search cost is extremely low, the agent can sample a large number of recommendations by incurring a negligible total search cost. As the recommendations are (conditionally) independently and identically distributed, the agent can learn the true state with a probability close to one, because of the law of large numbers.

4 Multiple Strategic Experts

We analyze the main model in which the agents can consult multiple experts, who respond to the agent's search strategy optimally. We show that an equilibrium exists and characterize it using a value function similar to (5). Then we consider the issues of information aggregation and welfare in the limiting case of vanishing search cost. We provide necessary and sufficient condition for full information aggregation in equilibrium. Finally, we show that if information is not fully aggregated, agent's welfare, as well as the social welfare, may be strictly lower than the first benchmark model in which the agent commits to consult only one expert.

4.1 Equilibrium Existence and Characterization

The analysis here is similar to the second benchmark case above except that we need to take into account the optimal response of the experts. Recall a pure strategy of the expert is the set of signals A under which she recommends Y. A pure strategy of the agent β is a mapping from the set of recommendation histories H to the decision of whether to search for advice or not at the beginning of a period, and whether to have the operation or not at the end of a period (after being recommended Y). Note that the posterior belief associated with history $h \in H$ depends only on the experts' strategy A, so we can write $p(h; A)$ to stand for the induced posterior belief. Also, denote by $q_\omega(h; A, \beta)$ the ex-ante probability that history $h \in H$ arises, given a strategy profile (A, β) and the state being $\omega \in \{0, 1\}$.

Consider the problem of the experts. First, note that the expert's recommendation matters (to her payoff) if and only if she is pivotal. Specifically, a recommendation N against the operation gives the expert a sure payoff of zero. A recommendation Y for the operation gives the expert a non-zero payoff if and only if the agent follows her recommendation and takes the operation with her. Therefore, when deciding her recommendation, the expert should compare her payoffs, conditioning on the event that the agent would follow her Y recommendation and takes the operation with her. Moreover, each expert has two other pieces of information when making her recommendation: (i) she is being consulted, and

(ii) she observes a private signal s. As a result, facing a strategy profile (A, β), an individual expert makes a Y recommendation if and only if the private signal $s \in [\underline{s}, \bar{s}]$ is such that

$$\Pr(\omega = 1 | \text{consulted}, s, \{h \in H : \beta_1(h, Y) = 1\}, A, \beta) \geq \frac{l}{1+l}. \tag{6}$$

Note that under the assumption that there are infinitely many experts, each of whom is drawn with equal probability, the conditioning event that the expert is being consulted has a zero probability. To get an idea of how the probability in (6) should be, suppose momentarily that the number of available experts is finite at n. The probability can be written as follows:

$$\begin{aligned}
& \Pr(\omega = 1 | \text{consulted}, s, \{h \in H : \beta_1(h, Y) = 1\}, A, \beta) \\
= & \left(1 + \frac{\Pr(\text{consulted}, s, \{h \in H : \beta_1(h, Y) = 1\} | \omega = 0, A, \beta)}{\Pr(\text{consulted}, s, \{h \in H : \beta_1(h, Y) = 1\} | \omega = 1, A, \beta)} \frac{1 - \Pr(\omega = 1 | A, \beta)}{\Pr(\omega = 1 | A, \beta)}\right)^{-1} \\
= & \left(1 + \frac{\Pr(\text{consulted}, \{h \in H : \beta_1(h, Y) = 1\} | \omega = 0, A, \beta) f(s|0)}{\Pr(\text{consulted}, \{h \in H : \beta_1(h, Y) = 1\} | \omega = 1, A, \beta) f(s|1)} \frac{1 - \pi}{\pi}\right)^{-1} \\
= & \left(1 + \frac{1}{s} \frac{1 - \pi}{\pi} \frac{\sum_{h \in H} \Pr(\text{consulted}, h | \omega = 0, A, \beta) \beta_1(h, Y)}{\sum_{h \in H} \Pr(\text{consulted}, h | \omega = 1, A, \beta) \beta_1(h, Y)}\right)^{-1} \\
= & \left(1 + \frac{1}{s} \frac{1 - \pi}{\pi} \frac{\sum_{h \in H} q_0(h; A, \beta) \beta_1(h, Y)}{\sum_{h \in H} q_1(h; A, \beta) \beta_1(h, Y)}\right)^{-1}.
\end{aligned} \tag{7}$$

The first equality is Bayes' rule. The second equality makes use of the fact that conditional on ω, the private signal s is independent of other events of the game. The third equality makes use of definitions of β and s.[4] The fourth equality is a result of the observation that $\Pr(\text{consulted}, h | \omega, A, \beta)$ is proportional to $q_\omega(h; A, \beta)$. To see this, denote the length of history h by $|h|$. Then we have

$$\Pr(\text{consulted}, h | \omega, A, \beta) = \frac{n-1}{n} \frac{n-2}{n-1} \cdots \frac{n-|h|-2}{n-|h|-1} \frac{1}{n-|h|} q_\omega(h; A, \beta) = \frac{1}{n} q_\omega(h; A, \beta).$$

The probability in (7) therefore is independent of the number of experts n. In the subsequent discussion, we impose that (7) holds even if there are infinitely many experts.

If the agent adopts a mixed strategy $\alpha \in \triangle\Lambda$, a distribution over the set of pure strategies, then an expert recommends Y if and only if

$$\begin{aligned}
& \int \Pr(\omega = 1 | \text{consulted}, s, \{h \in H : \beta(h, Y) = 1\}, A, \beta) \, d\alpha(\beta) \geq \frac{l}{1+l} \\
\Leftrightarrow & \int \left(1 + \frac{\sum_{h \in H} q_0(h; A, \beta) \beta_1(h, Y)}{\sum_{h \in H} q_1(h; A, \beta) \beta_1(h, Y)} \frac{1}{s} \frac{1 - \pi}{\pi}\right)^{-1} d\alpha(\beta) \geq \frac{l}{1+l}.
\end{aligned} \tag{8}$$

It is clear that the left hand side is increasing in s. Thus, the experts necessarily adopt a cutoff strategy in equilibrium: recommend the operation if and only if $s \geq s^*$, for some $s^* \in [\underline{s}, \bar{s}]$. Without loss of

[4] Note that the set of histories H is countable.

generality, instead of A, we simply use the cutoff s^* to stand for the experts' strategy. In sum, the optimal strategy of an individual expert $\rho : [\underline{s}, \bar{s}] \times \triangle\Lambda \to [\underline{s}, \bar{s}]$, given the strategy s^* of other experts and a mixed strategy α of the agent, is as follows.

$$\rho(s^*, \alpha) \equiv \begin{cases} \underline{s} & \text{if } \int \left(1 + \frac{\sum_{h \in H} \beta_1(h,Y) q_0(h; s^*, \beta)}{\sum_{h \in H} \beta_1(h,Y) q_1(h; s^*, \beta)} \frac{1}{\underline{s}} \frac{1-\pi}{\pi}\right)^{-1} d\alpha(\beta) \geq \frac{l}{1+l} \\ s & \text{if} \int \left(1 + \frac{\sum_{h \in H} \beta_1(h,Y) q_0(h; s^*, \beta)}{\sum_{h \in H} \beta_1(h,Y) q_1(h; s^*, \beta)} \frac{1}{s} \frac{1-\pi}{\pi}\right)^{-1} d\alpha(\beta) = \frac{l}{1+l} \text{ for some } s \in [\underline{s}, \bar{s}] \\ \bar{s} & \text{otherwise} \end{cases} . \tag{9}$$

Next consider the agent's problem. Given the experts' strategy s^*, the best response of the agent can be characterized in a similar way to that in the second benchmark model. Setting $\tilde{s} = s^*$ in (5) gives the continuation value function $V(p; s^*)$. The agent's best response is characterized by cutoffs $p_0(s^*)$ and $p_1(s^*)$ such that $p_0(s^*) \equiv \max\{p \in [0,1] : V(p) = 0\}$ and $V(p_1(s^*); s^*) = -L + p_1(s^*)(1+L)$.

Using the observations above, we define an equilibrium as follows.

Definition 1 *An equilibrium is characterized by a signal cutoff $s^* \in [\underline{s}, \bar{s}]$, and an agent search (possibly mixed) strategy α, such that*

(i) $s^* = \rho(s^*, \alpha)$; and

(ii) for all pure search strategy β on the support of α, it is necessary that

$$\beta_0(h) = \begin{cases} 0 & \text{if } p(h; s^*) \in (p_0(s^*), 1] \\ 1 & \text{if } p(h; s^*) \in [0, p_0(s^*)) \end{cases} ; \text{ and}$$

$$\beta_1(h) = \begin{cases} 0 & \text{if } p(h; s^*) \in [0, p_1(s^*)) \\ 1 & \text{if } p(h; s^*) \in (p_1(s^*), 1] \end{cases} .$$

Recall the definition of the optimal expert strategy in (9) relies on the assumption that the experts compute conditional probability according to (7). Proposition 7 in the appendix justifies this definition of equilibrium by showing that every sequence of equilibria of games with finitely many experts can only converge, as the number of experts grows to infinity, to an equilibrium defined above.

The proposition below establishes the existence of an equilibrium.

Proposition 4 *An equilibrium exists.*

Below we discuss some properties of the equilibrium. Recall the pivotal nature of the experts' problem: her recommendation matters (to her payoff) if and only if the agent would take the operation immediately following her recommendation for it. This observation leads to two implications. First, an *endogenous conflict of interest* between the agent and the experts arises. To see this, suppose $l = L$, so that if the game were one-shot, the interests of the expert and the agent are perfectly aligned, and the agent's preferred cutoff coincide with the expert's optimal cutoff, $\tilde{s} = \frac{1-\pi}{\pi} L$ (recall (1)). Now return to the game of multiple experts. Suppose the equilibrium expert cutoff is $s^* \in (\underline{s}, \bar{s})$. An expert is pivotal if and only if the agent's belief prior to consulting her is sufficiently close to $p_1(s^*)$: denote by $p_1'(s^*)$ the minimum belief for the consulted expert being pivotal. That is, if the beginning-of-period belief is $p_1'(s^*)$, after receiving a recommendation Y from the current expert, the agent's posterior would jump up to $p_1(s^*)$. Thus, $p_1(s^*)$ and $p_1'(s^*)$ are related by

$$p_1(s^*) = \frac{1}{1 + \frac{1-p_1'(s^*)}{p_1'(s^*)} \frac{1-F(s^*|0)}{1-F(s^*|1)}} \Leftrightarrow p_1'(s^*) = \frac{1}{\left(\frac{1}{p_1(s^*)} - 1\right) \frac{1-F(s^*|1)}{1-F(s^*|0)} + 1}. \tag{10}$$

As $p(h; s^*) \geq p_1'(s^*)$ for all $h \in H$ such that $\beta_1(h, Y) = 1$ and β is a best response to s^*, we have

$$\Pr(\omega = 1 | \text{consulted}, s^*, \{h \in H : \beta_1(h, Y) = 1\}, A, \beta) \geq \left(1 + \frac{1}{s^*} \frac{1 - p_1'(s^*)}{p_1'(s^*)}\right)^{-1}.$$

Combining with (9), we have

$$\frac{1}{1 + \frac{1-p_1'(s^*)}{p_1'(s^*)} \frac{1}{s^*}} \leq \frac{l}{1+l} \Leftrightarrow s^* \leq \frac{1 - p_1'(s^*)}{p_1'(s^*)} l = \frac{1 - p_1'(s^*)}{p_1'(s^*)} L.$$

Thus, whenever $p_1'(s^*) > \pi$, it is necessary that $s^* < \tilde{s}$. That is, unlike the one-shot game, the recommendation rule adopted by experts are no longer perfectly aligned with the agent's ideal rule. In the next subsection, we show that the condition $p_1'(s^*) > \pi$ necessarily arises in equilibrium when the search cost is sufficiently small and π is not too extreme.

The second implication is that the experts face a *loser's curse* if they are not strategic. When making her recommendation decision, the expert should take into account that the agent would agree to have the operation if and only if he has received sufficiently favorable recommendations from other experts. That is, his posterior belief is at least $p_1'(s^*)$, which may be well above the prior π. Failure to take this into account would lead to an excessively tough recommendation rule. Moreover, it is clear that the higher the value of $p_1'(s^*)$, the more lenient the expert is in recommending the operation.

4.2 Vanishing Search Cost

In this subsection, we consider the scenario in which the agent's search cost c is vanishingly small. Specifically, take an arbitrary sequence $\{c_n\}$ such that $c_n > 0$ for all $n \in \mathbb{N}$, and $\lim_{n\to\infty} c_n = 0$. This gives a sequence of equilibria, each of which belongs to the game with the corresponding search cost. We are interested in the information and welfare properties of the limit equilibrium, and compare them against the first benchmark model with a single available expert (considered in Section 3.1).

Note that for the agent to seek second opinions, the experts' recommendations must be informative, and a necessary condition is that $\pi \in \left(\frac{l}{l+\bar{s}}, \frac{l}{l+\underline{s}}\right)$. If $\pi \in \left[0, \frac{l}{l+\bar{s}}\right] \cup \left[\frac{l}{l+\underline{s}}, 1\right]$, the experts' recommendation is independent of the signals received, and there always exists an equilibrium in which the agent stops at the first expert. Without any informative recommendation, the agent takes the operation if and only if $\pi \geq \frac{L}{1+L}$. Moreover, if $\pi < \frac{l}{l+\underline{s}} \leq \frac{L}{1+L}$, there is an equilibrium in which experts recommend Y regardless of their signal, as the experts know that the agent would not follow their recommendation. Consequently there is no search in this case. As our main results concern the comparison of the setting with multiple experts against that with a single expert, we rule out these trivial equilibria by focusing on the case $\pi \in \left(\max\left\{\frac{l}{l+\bar{s}}, \frac{L}{1+L}\right\}, \frac{l}{l+\underline{s}}\right)$.

The reason why the aforementioned comparison is not trivial (even though the search cost is vanishingly small) is that the loser's curse effect lowers the quality of information transmission, and the effect becomes extremely severe in the limit as the search cost vanishes. Specifically, Lemma 4 below states that as $c_n \to 0$, the experts' recommendation in every equilibrium becomes completely uninformative. Denote by $s^*(c, \underline{s})$ the experts' cutoff in an equilibrium of the game in which the search cost is c and the lower bound of the expert's signal space is $\underline{s} \in [0, 1)$.[5] With $\pi \in \left(\frac{l}{l+\bar{s}}, \frac{l}{l+\underline{s}}\right)$, the equilibrium cutoff $s^*(c, \underline{s}) \notin \{\underline{s}, \bar{s}\}$.[6]

Lemma 4 *For each $\underline{s} \in [0, 1)$, if $\pi \in \left(\max\left\{\frac{l}{l+\bar{s}}, \frac{L}{1+L}\right\}, \frac{l}{l+\underline{s}}\right)$, then $\lim_{n\to\infty} s^*(c_n, \underline{s}) = \underline{s}$.*

The intuition of the lemma is as follows. Suppose $\lim_{n\to\infty} s^*(c_n, \underline{s}) > \underline{s}$, so that the experts recommend N for a positive measure of unfavorable signals. In this case, even though the search cost is vanishingly small, the expert's advice remains somewhat informative. Consequently, the agent would sample advice until he is almost sure that $\omega = 1$ before taking the operation. This in turn implies

[5] If the game has more than one equilibrium, the function $s^*(c, \underline{s})$ selects an arbitrary equilibrium cutoff.

[6] Suppose $s^*(c, \underline{s}) \in \{\underline{s}, \bar{s}\}$. Then the experts' recommendation are uninformative, and the agent would not search beyond the first (free) expert. Therefore, upon being consulted, the expert believes that she is the only one. As $\pi \in \left(\frac{l}{l+\bar{s}}, \frac{l}{l+\underline{s}}\right)$, she should make recommendation based on the received private signal, contradicting $s^*(c, \underline{s}) \in \{\underline{s}, \bar{s}\}$.

that regardless of the private signal s, as long as the agent takes the operation following a recommendation Y, the expert's belief that the agent has $\omega = 1$ is very close to one. This contradicts that $\lim_{n\to\infty} s^*(c_n, \underline{s}) > \underline{s}$.

Because of Lemma 4, the agent's search strategy in the limiting equilibrium is nontrivial. On one hand, sampling an additional advice is almost costless; on the other hand, the advice is almost completely uninformative. The lemma below characterizes the limit search strategy of the agent. Recall the agent's strategy in equilibrium is characterized by two cutoffs p_0 and p_1 (see condition (ii) of Definition 1).

Lemma 5 *For each* $\underline{s} \in [0,1)$ *and* $\pi \in \left(\max\left\{\frac{l}{l+\bar{s}}, \frac{L}{1+L}\right\}, \frac{l}{l+\underline{s}}\right)$,

(i) $\lim_{n\to\infty} p_1(s^*(c_n, \underline{s})) = \frac{l}{\underline{s}+l}$*;*

(ii) $\lim_{n\to\infty} p_0(s^*(c_n, 0)) = 0$*; and*

(iii) If $\underline{s} > 0$ *and* $l < L$*, then* $\liminf_{n\to\infty} p_0(s^*(c_n, \underline{s})) > 0$.

The intuition for part (i) of the lemma above is as follows. By Lemma 4, as search cost vanishes, $s^*(c_n, \underline{s})$ is arbitrarily close to $\underline{s}$, i.e., the experts are willing to recommend Y despite receiving an unfavorable signal. This can be optimal to an individual expert if and only if conditional on the agent's taking the operation following a Y recommendation, the expert's belief prior to learning the private signal is sufficiently high. More specifically, suppose $s^*(c_n, \underline{s}) = \underline{s} + \varepsilon$, for some $\varepsilon > 0$, then the expert's prior belief p (that $\omega = 1$) must satisfy

$$\frac{1}{1 + \frac{1-p}{p}\frac{1}{\underline{s}+\varepsilon}} \geq \frac{l}{1+l} \Leftrightarrow p \geq \frac{l}{\underline{s}+\varepsilon+l}.$$

On the other hand, p cannot strictly exceed $\frac{l}{\underline{s}+l}$, for otherwise, all experts recommend Y regardless of signals, and there is no value in searching for advice regardless of c. Therefore, p approaches $\frac{l}{\underline{s}+l}$ in the limit. Furthermore, as $s^*(c_n, \underline{s})$ is arbitrarily close to $\underline{s}$, the increase in the agent's posterior after learning each Y recommendation is almost zero. This implies that conditional on being a pivotal consultation, the distribution of the agent's belief prior to learning the expert's recommendation must be almost degenerate at $p_1(s^*(c_n, \underline{s}))$. As the expert's belief about this distribution is correct in equilibrium, it is necessary that $p_1(s^*(c_n, \underline{s}))$ approaches $\frac{l}{\underline{s}+l}$ in the limit.

In the proof of part (ii) of Lemma 5, we compute the limiting payoff of the following simple (but necessarily suboptimal) search strategy of the agent: sample a fixed number of experts and have the operation in the end if and only if all of them recommend Y. It is shown that by choosing the fixed number of experts appropriately, the agent can attain an ex-ante payoff arbitrarily close to the highest possible level of π as search cost vanishes. This means that in the limit, the agent necessarily learns the true state with a negligible total search cost. As a result, the agent almost always finds it profitable to search for additional advice, even if the current belief that the operation is suitable is very low, i.e., $p_0(s^*(c,0))$ approaches 0 in the limit.

The intuition for part (iii) of the lemma is as follows. The condition $l < L$ means that the expert suffers relatively less in the case of a failed operation (i.e., performing the operation at $\omega = 0$). Thus, the expert is too lenient in recommending the operation even in the one-shot game. Compounded with the loser's curse effect, as the search cost vanishes, $s^*(c_n, \underline{s})$ approaches $\underline{s} > 0$ so quickly that the agent does not find it profitable to search if his current belief is sufficiently low.

We say information is perfectly aggregated if the agent takes the ex-post correct action with probability one. Information is almost perfectly aggregated if and only if $p_0(s^*(c, \underline{s}))$ and $p_1(s^*(c, \underline{s}))$ are close to 0 and 1 respectively. The previous lemma implies that information is perfectly aggregated in the limit as search cost vanishes only if $\underline{s} = 0$.

Proposition 5 *Perfect information aggregation arises as the unique limit equilibrium outcome (i.e.,* $\lim_{n\to\infty} p_1(s^*(c_n, \underline{s})) = 1$ *and* $\lim_{n\to\infty} p_0(s^*(c_n, \underline{s})) = 0$*) if and only if* $\underline{s} = 0$ *and* $\pi > \max\left\{\frac{l}{l+\bar{s}}, \frac{L}{1+L}\right\}$.

The condition on perfect information aggregation identified in the proposition above depends crucially on the lower bound of the signal space $\underline{s}$. This echoes the condition for information aggregation in the literature of common-value auction (see Milgrom (1979)). In an auction for an object with common value V, a random variable with state space $\{v_1, v_2, ...\}$, perfect information aggregation (in the limit as the number of bidders grow to infinity) requires that for each k, there exists a signal that distinguishes $\{V = v_k\}$ from $\{V < v_k\}$. As bidders concern about the winner's curse in an auction, whereas experts in our model concern about the loser's curse, it is natural to expect that the analogous condition for information aggregation is the existence of a signal that allows the expert to distinguish $\{\omega = 0\}$ from $\{\omega > 0\}$, which is exactly the requirement that $\underline{s} = 0$. However, to ensure that the unique equilibrium outcome has information fully aggregated, an extra condition on the signal structure is needed: the upper bound of the signal space $\bar{s}$ must exceed $l\frac{1-\pi}{\pi}$. If $\bar{s} \leq l\frac{1-\pi}{\pi}$, there always exists an equilibrium in which the experts recommend N regardless of signals.

Finally, we consider welfare in the limit equilibrium. In the case $\underline{s} = 0$, as information is perfectly aggregated in the limit equilibrium, it is clear that the agent takes the correct action with probability one, so his resulting payoff is the highest possible level of π. Moreover, full efficiency necessarily arises in this case, i.e., the total payoffs of all players achieves the highest possible level of 2π. On the other hand, if $\underline{s} > 0$, then information is imperfectly aggregated in the limit equilibrium. As the agent may take the ex-post incorrect action with positive probability, his equilibrium payoff is bounded away from π, even if the search cost is infinitesimal. More strikingly, as the lemma below shows, it is possible that the agent becomes worse off than the first benchmark model analyzed in Section 3.1 in which he can consult only one expert. If this happens, the total welfare is also lower.

Proposition 6 *Suppose* $\underline{s} \in [0,1)$ *and* $\pi \in \left(\max\left\{\frac{l}{l+\bar{s}}, \frac{L}{1+L}\right\}, \frac{l}{l+\underline{s}}\right)$. *Denote by* $U(c, \underline{s})$ *the agent's expected payoff in an equilibrium of the game in which the search cost is* c *and the lower bound of the experts' signal space is* $\underline{s}$. *Denote by* $T(c, \underline{s})$ *the total expected payoff of the agent and all experts in the corresponding game.*[7]

(i) $\lim_{n\to\infty} U(c_n, 0) = \pi$, *and* $\lim_{n\to\infty} T(c_n, 0) = 2\pi$.

(ii) If $\underline{s} > 0$, *then* $\limsup_{n\to\infty} U(c_n, \underline{s}) \leq \max\left\{\pi\left(1 - \frac{L}{l}\underline{s}\right), 0\right\}$, *and* $\limsup_{n\to\infty} T(c_n, \underline{s}) \leq \pi\left(1 - \frac{L}{l}\underline{s}\right) + \pi(1 - \underline{s})$.

The most interesting implication of the proposition above arises in the case $\underline{s} > 0$ and $L > l$. In this case, the agent's equilibrium payoff with an infinitesimal search cost and access to infinitely many experts is no more than $\max\left\{\pi\left(1 - \frac{L}{l}\underline{s}\right), 0\right\}$ **for all signal structures** of experts. On the other hand, if he has access to only one expert as in the first benchmark model analyzed in Section 3.1, his expected payoff, denoted by $U(\pi)$,[8] is approximately $\pi(1 - \underline{s})$ if the signal structure is close to $F_{\underline{s}}$ (recall Proposition 1). As a result, if the signal structure of the experts is sufficiently close to $F_{\underline{s}}$, then the agent *is strictly worse off with the access to infinitely many experts at vanishingly small cost*, compared to the alternative scenario in which he has only one expert for consultation. Furthermore, in this case, the social welfare is also lower with infinitely many experts and negligibly search cost: the total payoff of all players is no more than $\pi\left(1 - \frac{L}{l}\underline{s}\right) + \pi(1 - \underline{s})$; whereas the total payoff in the case of a single expert is arbitrarily close to $2\pi(1 - \underline{s})$ (recall Proposition 2).

Corollary 1 *Suppose* $\underline{s} > 0$, $L > l$, *and* $\pi \in \left(\frac{L}{1+L}, \frac{l}{l+\underline{s}}\right)$. *If the experts' signal structure is close to* $F_{\underline{s}}$, *then*

$$\limsup_{n\to\infty} U(c_n, \underline{s}) < U(\pi)\text{; and } \limsup_{n\to\infty} T(c_n, \underline{s}) < T(\pi).$$

The intuition of the result above is as follows. Recall from Proposition 5 that if $\underline{s} > 0$, the agent may not learn the true state ω eventually, so his equilibrium payoff can be bounded away from the highest

[7] If the game has more than one equilibrium, the functions $U(c, \underline{s})$ and $T(c, \underline{s})$ select the payoffs of an arbitrary equilibrium.

[8] Recall the definition of $U(\pi)$ in (3).

possible level of π. The reason for the low quality of information gathered through the search process is the presence of the loser's curse. Its effect is particularly strong if l is small, as the cost of a mistake to the expert is low. As a result, in this case, the experts' recommendation become uninformative very quickly as the search cost vanishes.

5 Concluding Remarks

For simplicity and tractability, our model of searching for advice has abstracted away from some realistic considerations. Below we briefly discuss a few possible variations and extensions.

Alternative expert preference In the main model, we assume that the expert derives a non-zero payoff from the interaction with the agent if and only if the agent takes the operation with her. Alternatively, each expert may care about the agent's eventual payoff provided that she has recommended the operation to him. Specifically, suppose the expert's payoff is given by the table below:

	$\omega = 1$	$\omega = 0$
Recommends Y and agent has one eventually (not necessarily with her)	1	$-l$
Recommend N, or agent does not have operation in the end	0	0

Then the experts essentially face the same structure of incentives when making recommendation, so all of our results remain valid, except for those concerning the total efficiency as we have changed the payoffs of the experts.

Cheap talk message In the main model, we impose a restrictive binary message space for the experts. If we allow for a general message space, more information can be transmitted in some equilibrium, though the equilibrium we consider and characterize still exists. For instance, along with the recommendation decision, the expert can send a cheap talk message. Now if an expert does not recommend the operation, she may just as well fully reveal the signal learned through her message, as her payoff is constant at zero. However, as Lemma 4 shows, the experts almost always recommends the operation as the search cost vanishes. This use of the cheap talk message therefore does not affect our limit results.

Nonrandom search order In the main model, we assume that the agent samples experts randomly. Suppose, in contrast, there are a few prominent experts whom are consulted first. The prominent and non-prominent experts compute the conditional probability in (6) differently. It is possible that the prominent expert would adopt a more informative recommendation policy, knowing that she can be pivotal even if the agent has not had much previous consultations with other experts. It is therefore conjectured that the existence of prominent experts can improve the quality of information gathered through search.

Appendix

Proof of Lemma 1: The upper bound on $U(\pi)$ can be computed by considering an auxiliary game that makes the following two modifications to the benchmark game. First, the agent chooses the signal structure of the expert; and second, the expert is forced to report truthfully the learned signal, along with her recommendation. It is clear that the agent's payoff in this auxiliary game is no less than what he can obtain in the benchmark model considered here. The agent's payoff in this auxiliary game can be computed using the technique developed by Kamenica and Gentzkow (2011): the optimal signal structure can be found by looking for the concave closure of the payoff as a function of realized posterior.

As the lower bound of the signal space $\underline{s}$ is fixed, the lowest possible posterior belief after learning the signal is $\underline{p} \equiv \left(1 + \frac{1-\pi}{\pi}\frac{1}{\underline{s}}\right)^{-1}$. Denote by $\tilde{U}_{\underline{s}} : \left[\underline{p}, 1\right] \to \mathbb{R}$ the agent's payoff as a function of the expert's posterior after learning the signal:

$$\tilde{U}_{\underline{s}}(p) = \begin{cases} 0 & \text{if } p < \frac{l}{1+l} \\ \max\{0, -L + p(1+L)\} & \text{if } p \geq \frac{l}{1+l} \end{cases}.$$

To understand the payoff function, note that if the expert has a posterior belief $p < \frac{l}{1+l}$, then she recommends no operation, and the agent gets a zero payoff. If the expert has a posterior belief $p > \frac{l}{1+l}$, then she recommends the operation, and the agent's expected payoff is $\max\{0, -L + p(1+L)\}$. The lemma below follows immediately from Corollary 2 of Kamenica and Gentzkow (2011).

Lemma 6 *Fix the lower bound of the signal space $\underline{s} \in [0,1)$. The agent's payoff in the auxiliary game is given by the concave closure of $\tilde{U}_{\underline{s}}(\cdot)$, denoted by con $\left[\tilde{U}_{\underline{s}}\right](\cdot)$, evaluated at the prior π. The concave closure is defined by*

$$con\left[\tilde{U}_{\underline{s}}\right](\cdot) \equiv sup\left\{z \in \mathbb{R} : (p,z) \in co(\tilde{U}_{\underline{s}}(p))\right\},$$

where co $\left(\tilde{U}_{\underline{s}}(p)\right)$ is the convex hull of the graph of $\tilde{U}_{\underline{s}}(p)$.

The lemma allows us to compute the agent's payoff in the auxiliary game by constructing the concave closure of $\tilde{U}_{\underline{s}}$.

(i) For case (a) $l \leq L$, a plot of the function $\tilde{U}_{\underline{s}}(p)$ is given below.

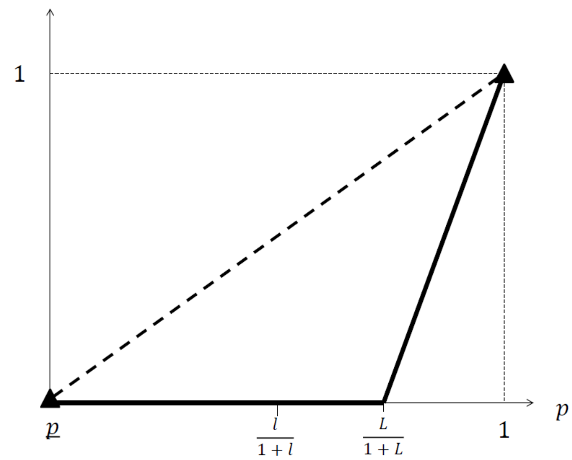

Figure 1: Case $l \leq L$

In the figure, the solid curve is the function $\tilde{U}_{\underline{s}}(p)$, and its concave closure is represented by the dotted line. It is clear that the dotted line is obtained by connecting the two points represented by the triangles. As a result, the optimal signal structure is exactly $F_{\underline{s}}$. Consequently, in this case, we have $con\left[\tilde{U}_{\underline{s}}\right](\pi) = \pi(1-\underline{s})$.

For case (b), a plot of the function $\tilde{U}_{\underline{s}}(p)$ is given below.

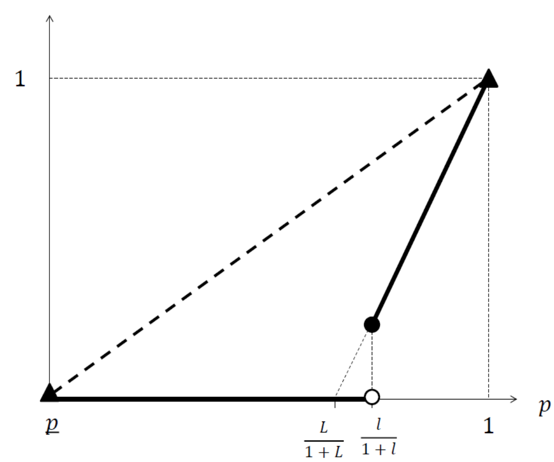

Figure 2: Case $l > L$ and $\pi < \frac{L}{L+\underline{s}}$.

Note that the condition $\pi \leq \frac{L}{L+\underline{s}}$ ensures that $\frac{L}{1+L} \geq \underline{p}$. It is clear from Figure 2 that the concave closure of $\tilde{U}_{\underline{s}}$ in this case is obtained by connecting the two points represented by the triangles. That is, the most informative signal structure $F_{\underline{s}}$ is optimal. Consequently, in this case, we have $con\left[\tilde{U}_{\underline{s}}\right](\pi) = \pi(1-\underline{s})$.

(ii) A plot of the function $\tilde{U}_{\underline{s}}(p)$ for the case $l > L$ and $\pi \in \left(\max\left\{\frac{l}{1+l}, \frac{L}{\underline{s}+L}\right\}, \frac{l}{l+\underline{s}}\right)$ is given below.

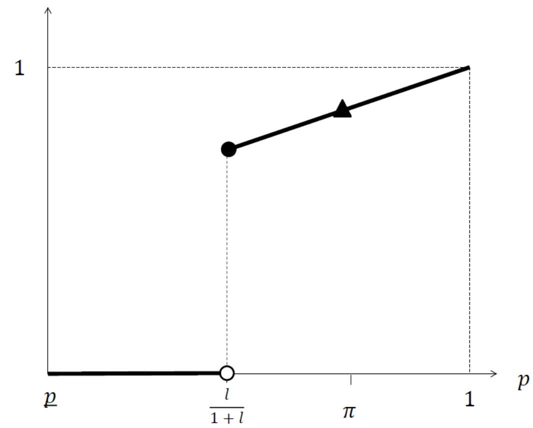

Figure 3: Case $l > L$ and
$\pi \in \left(\max\left\{\frac{l}{1+l}, \frac{L}{\underline{s}+L}\right\}, \frac{l}{l+\underline{s}}\right)$.

Note that the condition $\pi > \frac{L}{\underline{s}+L}$ implies that maximal informativeness is not optimal. If, in addition, $\pi > \frac{l}{1+l}$, then it is clear from Figure 3 that the concave closure of $\tilde{U}_{\underline{s}}$ coincides with $\tilde{U}_{\underline{s}}$ itself at $p = \pi$. Thus, an uninformative signal structure is optimal for the agent. Consequently, in this case, we have $con\left[\tilde{U}_{\underline{s}}\right](\pi) = -L + \pi(1 + L)$.

(iii) Suppose $\frac{L}{\underline{s}+L} < \frac{l}{1+l}$. A plot of the function $\tilde{U}_{\underline{s}}(p)$ for the case $l > L$ and $\pi \in \left(\frac{L}{\underline{s}+L}, \frac{l}{1+l}\right]$ is given below.

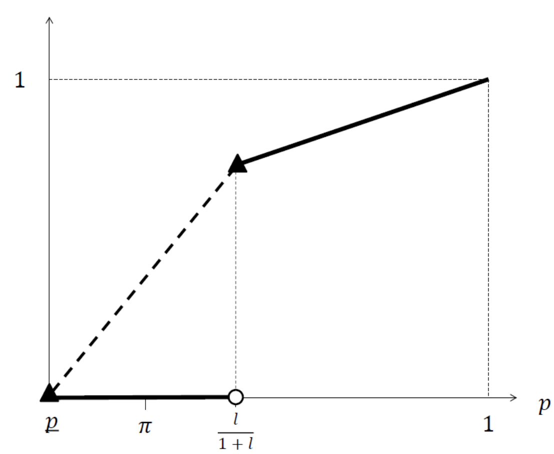

Figure 4: Case $l > L$ and $\pi \in \left(\frac{L}{\underline{s}+L}, \frac{l}{1+l}\right]$

Again, the condition $\pi > \frac{L}{\underline{s}+L}$ implies that maximal informativeness is not optimal. The condition $\pi \leq \frac{l}{1+l}$ implies that the concave closure of $\tilde{U}_{\underline{s}}$ around π is obtained by connecting points $(\underline{p}, 0)$ and

$\left(\frac{l}{1+l}, \tilde{U}_{\underline{s}}\left(\frac{l}{1+l}\right)\right)$. Thus, the optimal information signal is one that induces a posterior distribution with a binary support $\left\{\underline{p}, \frac{l}{1+l}\right\}$. Consequently, in this case, we have $con\left[\tilde{U}_{\underline{s}}\right](\pi) = \pi(1-\pi)(1-\underline{s})\frac{l-L}{l-\pi(l+\underline{s})}$.

Finally, we explain why the upper bound on the agent's payoff $con\left[\tilde{U}_{\underline{s}}\right](\pi)$ identified above is (almost) achievable in the benchmark model in which the signal realization is privately observed by the expert. This is clear for case (ii), as it calls for a completely uninformative signal. For cases (i) and (iii), the optimal signal structure has a binary support: the lower signal leads to recommendation N, thus a zero payoff to the agent; the higher signal leads to recommendation Y, thus a positive expected payoff to the agent. Consequently, the agent's expected payoff remains unchanged even if he does not know the expert's signal realization. Q.E.D.

Proof of Lemma 2: The construction for case (iii) of Lemma 1 is analogous to that of case (i), so we illustrate only case (i). Below we explicitly construct a sequence of conditional distribution functions that converge pointwise to $F_{\underline{s}}$. For each $n \in \mathbb{N}$, define a density function over posterior belief $f_n : \left[\underline{p}, 1\right] \to \mathbb{R}_+$ as follows:

$$f_n(p) \equiv \begin{cases} n\left(1-\pi(1-\underline{s})-\frac{m_n^1}{n}\right) & \text{if } p \in \left[\underline{p}, \underline{p}+\frac{1}{n}\right] \\ \frac{1}{1-\frac{2}{n}-\underline{p}}\left(\frac{m_n^1+m_n^2}{n}\right) & \text{if } p \in \left[\underline{p}+\frac{1}{n}, 1-\frac{1}{n}\right] \\ n\left(\pi(1-\underline{s})-\frac{m_n^2}{n}\right) & \text{if } p \in \left[1-\frac{1}{n}, 1\right] \end{cases},$$

where

$$\begin{aligned} m_n^1 - m_n^2 &\equiv \frac{1}{\frac{1-\pi}{1-\pi+\pi\underline{s}}+\frac{1}{n}}\left(-2\pi(1-\underline{s})+1\right), \\ m_n^1, m_n^2 &> 0, \text{ and } \frac{m_n^1}{n}, \frac{m_n^2}{n} \to 0. \end{aligned}$$

It is straightforward to verify that (i) there exists a pair of sequences $\{m_n^1\}$ and $\{m_n^2\}$ that satisfies the conditions above, (ii) f_n is a well-defined density function; (ii) the expected value of the posterior distribution is equal to π; (iii) for all $p \in \left(\underline{p}, 1\right)$, $f_n(p)$ converges to 0 as $n \to \infty$.

Next, define $f_n(s|0)$, $f_n(s|1) : [\underline{s}, \infty] \to \mathbb{R}_+$ by

$$f_n(s|0) \equiv \frac{f_n\left(p\right)}{\pi s + (1-\pi)}, \text{ and } f_n(s|1) = s f_n(s|0).$$

Now suppose $\pi < \frac{l}{\underline{s}+l}$. The expert recommends operation if and only if $s \in \left[\frac{1-\pi}{\pi}l, \infty\right]$, an event that happens with probability $\int_{\frac{l}{1+l}}^{1} f_n(p)\,dp$, which converges to $\pi(1-\underline{s})$ as $n \to \infty$. It is straightforward that conditional on a recommendation for the operation, the agent's belief that $\omega = 1$ converges to 1, because $\frac{1-F_n(\tilde{s}(l)|0)}{1-F_n(\tilde{s}(l)|1)} = \frac{\int_{\frac{1-\pi}{\pi}l}^{\infty} f_n(s|0)ds}{\int_{\frac{1-\pi}{\pi}l}^{\infty} f_n(s|1)ds} \to 0$ as $n \to \infty$. This immediately implies that (2) holds for n sufficiently large. Therefore, the agent's payoff converges to $\pi(1-\underline{s})$ as $n \to \infty$. Q.E.D.

Proof of Proposition 2: Consider an auxiliary game that modifies the benchmark game as follows: (i) a third-party designer, with the objective of maximizing the joint payoff of the agent and the expert, chooses the signal structure of the expert; (ii) the realized signal is shown to the expert, who can then choose recommendation Y or N; and (iii) upon a recommendation Y, the designer decides on behalf of the agent whether to take the operation or not. It is clear that the total payoff of the agent and the expert in this auxiliary game is no less than that in the benchmark game, thus serving as an upper bound of the latter.

In the auxiliary game, the operation is carried out if and only if (a) the expert is willing to carry it out after learning the signal; and (b) the total payoff is nonnegative. Denote by $\tilde{T}_{\underline{s}} : [\underline{p}, 1] \to \mathbb{R}$ the total payoff of the expert and the agent, as a function of the posterior learned by the expert:

$$\tilde{T}_{\underline{s}}(p) = \begin{cases} 0 & \text{if } p < \max\left\{\frac{l}{1+l}, \frac{l+L}{2+l+L}\right\} \\ -(L+l) + p(2+L+l) & \text{if } p \geq \max\left\{\frac{l}{1+l}, \frac{l+L}{2+l+L}\right\} \end{cases}.$$

The efficient signal structure in the auxiliary game can be obtained by finding the concave closure of $\tilde{T}_{\underline{s}}(\cdot)$. If $l \leq L$, then $\frac{l}{1+l} \leq \frac{l+L}{2+l+L}$. Moreover, $\pi < \frac{l}{l+\underline{s}}$ implies that $\underline{p} \leq \frac{l+L}{2+l+L}$. A plot of $\tilde{T}_{\underline{s}}$ in this case is given by the figure below:

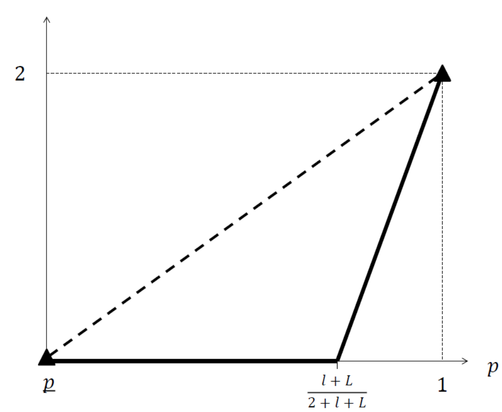

Figure 5: Case $l \leq L$ and $\underline{p} \leq \frac{l+L}{2+l+L}$

It is therefore clear that the concave closure is obtained by connecting the two triangles in Figure 5, and $con\left[\tilde{T}_{\underline{s}}\right](\pi) = 2\pi(1-\underline{s})$. Thus, if $l \leq L$, the most informative signal structure $F_{\underline{s}}$ is optimal. It is easy to see that with $F_{\underline{s}}$, the designer does not need to observe the signal realization to determine the socially optimal action: he/she can simply follow the expert's recommendation.

On the other hand, if $l > L$, then $\frac{l}{1+l} > \frac{l+L}{2+l+L}$. Suppose in addition, $2+L+l \geq \frac{2}{1-\underline{p}}$, then $\tilde{T}_{\underline{s}}$ takes the following form:

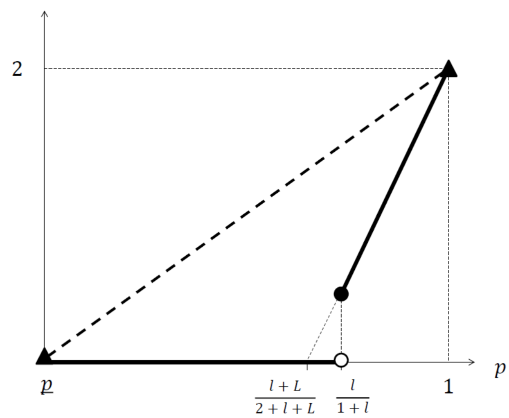

Figure 6: Case $l > L$ and
$2 + L + l \geq \frac{2}{1-\underline{p}}$

Note that in the figure above, the condition $2 + L + l \geq \frac{2}{1-\underline{p}}$ ensures that the slope of the solid line (the latter portion) is no less than that of the dotted line. It is clear that maximal informativeness is efficient.

Next, suppose $l > L$ and $2 + L + l < \frac{2}{1-\underline{p}}$. Then $\tilde{T}_{\underline{s}}$ takes the following form:

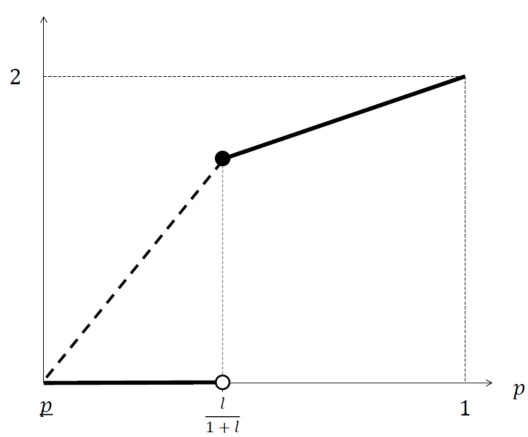

Figure 7: Case $l > L$ and
$2 + L + l < \frac{2}{1-\underline{p}}$

The condition $2 + L + l \geq \frac{2}{1-\underline{p}}$ implies that maximal information is not efficient. Moreover, the efficient information structure depends on whether π is above or below $\frac{l}{1+l}$: if $\pi < \frac{l}{1+l}$, then a posterior distribution with binary support $\left\{\underline{p}, \frac{l}{1+l}\right\}$ is efficient; if $\pi \geq \frac{l}{1+l}$, then a degenerate posterior distribution at π is efficient.

Summarizing the discussion above, in the case $l > L$ and $\pi < \frac{l}{l+\underline{s}}$,

$$
\begin{aligned}
con\left[\tilde{T}_{\underline{s}}\right](\pi) &= \begin{cases} 2\pi\left(1-\underline{s}\right) & \text{if } 2+L+l \geq \frac{2}{1-\underline{p}} \\ \left(\left(-l+L\right)+\frac{l}{1+l}\left(1+l+L\right)\right)\left(\frac{\pi-\underline{p}}{\frac{l}{1+l}-\underline{p}}\right) & \text{if } \pi < \frac{l}{1+l} \text{ and } 2+L+l < \frac{2}{1-\underline{p}} \\ -\left(l+L\right)+\pi\left(2+l+L\right) & \text{if } \pi \geq \frac{l}{1+l} \text{ and } 2+L+l < \frac{2}{1-\underline{p}} \end{cases} \\
&= \begin{cases} 2\pi\left(1-\underline{s}\right) & \text{if } \pi \leq \frac{L+l}{2\underline{s}+L+l} \\ \pi\left(1-\pi\right)L\left(2l+1\right)\frac{1-\underline{s}}{(1-\pi)l-\pi\underline{s}} & \text{if } \frac{L+l}{2\underline{s}+L+l} < \frac{l}{1+l} \text{ and } \frac{L+l}{2\underline{s}+L+l} < \pi < \frac{l}{1+l} \\ -\left(l+L\right)+\pi\left(2+l+L\right) & \text{if } \max\left\{\frac{L+l}{2\underline{s}+L+l}, \frac{l}{1+l}\right\} \leq \pi < \frac{l}{\underline{s}+l} \end{cases},
\end{aligned}
$$

where the second equality follows from the definition of $\underline{p}$ and re-arranging. The first case above calls for maximal informativeness (i.e., $F_{\underline{s}}$); the second case calls for partial informativeness (with binary support $\left\{\underline{s}, \frac{1-\pi}{\pi}l\right\}$); and the final case calls for no information.

Again, with these signal structures, the designer can simply follow the expert's recommendation to determine the socially optimal decision, so direct observation of the signal realization is not necessary. Therefore, $\bar{T}_{\underline{s}}$ given in the proposition indeed gives an upper bound on the total payoff in the benchmark game.

Finally, it follows from an argument similar to that in Lemma 2 that the total payoff of the benchmark game T can be made arbitrarily close to $\bar{T}_{\underline{s}}$ with a signal structure with a connected support and conditional density function. Q.E.D.

Proof of Lemma 3: First, we show that there exists a $p_0 \in [0,1]$ such that the agent searches if the beginning-of-period belief $p > p_0$. Suppose the agent's beginning-of period belief is $p' \in [0,1)$. Suppose also that by searching with a continuation strategy β, the agent can achieve a nonnegative payoff. Now if the agent's belief is $p'' \in (p',1)$ instead of p', he can still search with the same strategy β and achieve a weakly higher payoff. Consequently, an agent with belief p'' would find it optimal to continue searching. Therefore, the set of beginning-of-period beliefs under which the agent finds it strictly optimal to search takes the form $(p_0, 1)$. Moreover, at belief p_0, the agent must be indifferent between stopping and continuing searching. This follows because fixing a search strategy, the agent's continuation payoff is continuous in her current belief p.

Next, we show that there exists a $p_1 \in [0,1]$ such that the agent undergoes the operation if the end-of-period belief $p > p_1$. Suppose the agent is recommended the operation and his end-of period belief is $p' \in [0,1)$. Suppose the agent finds it optimal to have the operation (instead of quitting and continuing the search). We show that if the agent's belief is $p'' > p'$ instead, then he would find it strictly optimal to have the operation. Suppose not, i.e., there exists a continuation search strategy β such that the agent with belief p'' can achieve a payoff no less than $-L+p''(1+L)$. The search strategy β prescribes the set of histories, denoted by H_β, under which the agent undergoes the operation. The agent's continuation payoff $u(p'';\beta)$, under belief p'' and searching with continuation strategy β, can be expressed as

$$u(p'';\beta) = -C(\beta) + \sum_{h\in H_\beta} (p''[\Pr(h|\omega=1) + \Pr(h|\omega=0)L] - \Pr(h|\omega=0)L),$$

where $\Pr(h|\omega)$ is the probability that history h realizes conditional on the state being ω, and $C(\beta)$ is the expected search cost implied by strategy σ. By hypothesis, we have $u(p'';\beta) \geq -L + p''(1+L)$. Therefore,

$$\begin{aligned}
& u(p';\beta) - (-L + p'(1+L)) \\
= & \left\{ u(p'';\beta) + (p' - p'') \sum_{h\in H_\beta} [\Pr(h|\omega=1) + \Pr(h|\omega=0)L] \right\} \\
& - \{(-L + p''(1+L)) + (p' - p'')(1+L)\} \\
= & u(p'';\beta) - (-L + p''(1+L)) + (p'' - p')\left[(1+L) - \sum_{h\in H_\beta} [\Pr(h|\omega=1) + \Pr(h|\omega=0)L]\right] \\
> & 0
\end{aligned}$$

This is in contradiction to the assumption that the agent finds it optimal to have the operation at belief p'.

Finally, the indifference at $p = p_1$ follows from the continuity of the agent's payoff with respect to p, holding the search strategy fixed. Q.E.D.

Proof of Proposition 3: Consider an auxiliary game with exactly the same structure as benchmark model II, except that the agent can consult at most $M \in \mathbb{N}$ experts. Denote by $V_t : \{0, 1, ..., M\} \times [0, 1] \to \mathbb{R}$ the beginning-of-period value function given a current belief $p \in [0, 1]$, provided that there remains $t \in \{0, 1, ..., M\}$ unconsulted experts. Denote by $W_t : \{1, ..., M\} \times [0, 1] \to \mathbb{R}$ the beginning-of-period value function given a current belief p, provided that there remains t unconsulted experts and the agent **decides to search this period**.

The value functions V_t and W_t can be recursively constructed. If $t = 0$, there is no more expert available, and we define $V_0(p) = \max\{-L + p(1+L), 0\}$. Clearly, V_0 is nondecreasing and weakly convex. Next, if $t \geq 1$, there are still available experts and is defined recursively as follows: for

$t \in \{1, ..., M\}$,

$$V_t(p) = \max\{0, W_t(p)\}; \text{ and} \tag{11}$$

$$\begin{aligned} W_t(p) &= -c + [p(1 - F(\tilde{s}|1)) + (1-p)(1 - F(\tilde{s}|0))] \\ &\quad \times \max\left\{ V_{t-1}\left(\frac{1}{1 + \frac{1-p}{p}\frac{1-F(\tilde{s}|0)}{1-F(\tilde{s}|1)}}\right), -L + \frac{1}{1 + \frac{1-p}{p}\frac{1-F(\tilde{s}|0)}{1-F(\tilde{s}|1)}}(1+L) \right\} \\ &\quad + [pF(\tilde{s}|1) + (1-p)F(\tilde{s}|0)] V_{t-1}\left(\frac{1}{1 + \frac{1-p}{p}\frac{F(\tilde{s}|0)}{F(\tilde{s}|1)}}\right) \end{aligned} \tag{12}$$

We show below that V_t is nondecreasing and weakly convex for all $t \geq 0$. Suppose V_{t-1} is nondecreasing and weakly convex. To see V_t is nondecreasing, it suffices to note that (i) W_t is nondecreasing in p, and (ii) $\frac{1}{1+\frac{1-p}{p}\frac{1-F(\tilde{s}|0)}{1-F(\tilde{s}|1)}} \geq \frac{1}{1+\frac{1-p}{p}\frac{F(\tilde{s}|0)}{F(\tilde{s}|1)}}$.

We now show V_t is weakly convex. Take a pair $p', p'' \in [0,1]$ with $p' < p''$ and an $\alpha \in (0,1)$. Denote $\hat{p} \equiv \alpha p' + (1-\alpha) p''$.

$$\begin{aligned} &\alpha W_t(p') + (1-\alpha) W_t(p'') \\ = &\ -c + \left[\begin{array}{c} \alpha [p'(1 - F(\tilde{s}|1)) + (1-p')(1 - F(\tilde{s}|0))] \\ \times \max\left\{ -L + \frac{1+L}{1+\frac{1-p'}{p'}\frac{1-F(\tilde{s}|0)}{1-F(\tilde{s}|1)}}, V_{t-1}\left(\left(1 + \frac{1-p'}{p'}\frac{1-F(\tilde{s}|0)}{1-F(\tilde{s}|1)}\right)^{-1}\right)\right\} \\ + (1-\alpha)[p''(1 - F(\tilde{s}|1)) + (1-p'')(1 - F(\tilde{s}|0))] \\ \times \max\left\{ -L + \frac{1+L}{1+\frac{1-p''}{p''}\frac{1-F(\tilde{s}|0)}{1-F(\tilde{s}|1)}}, V_{t-1}\left(\left(1 + \frac{1-p''}{p''}\frac{1-F(\tilde{s}|0)}{1-F(\tilde{s}|1)}\right)^{-1}\right)\right\} \end{array} \right] \\ &+ \left[\begin{array}{c} \alpha [p'F(\tilde{s}|1) + (1-p')F(\tilde{s}|0)] V_{t-1}\left(\left(1 + \frac{1-p'}{p'}\frac{F(\tilde{s}|0)}{F(\tilde{s}|1)}\right)^{-1}\right) \\ + (1-\alpha)[p''F(\tilde{s}|1) + (1-p'')F(\tilde{s}|0)] V_{t-1}\left(\left(1 + \frac{1-p''}{p''}\frac{F(\tilde{s}|0)}{F(\tilde{s}|1)}\right)^{-1}\right) \end{array} \right] \\ \geq &\ -c + [\hat{p}(1 - F(\tilde{s}|1)) + (1-\hat{p})(1 - F(\tilde{s}|0))] \\ &\times \max\left\{ -L + (1+L)\frac{\hat{p}(1-F(\tilde{s}|1))}{\hat{p}F(\tilde{s}|1) + (1-\hat{p})F(\tilde{s}|0)}, V_{t-1}\left(\frac{\hat{p}(1-F(\tilde{s}|1))}{\hat{p}F(\tilde{s}|1) + (1-\hat{p})F(\tilde{s}|0)}\right)\right\} \\ &+ [\hat{p}F(\tilde{s}|1) + (1-\hat{p})F(\tilde{s}|0)] V_{t-1}\left(\frac{\hat{p}F(\tilde{s}|1)}{\hat{p}F(\tilde{s}|1) + (1-\hat{p})F(\tilde{s}|0)}\right) \\ = &\ W_t(\hat{p}), \end{aligned}$$

where the inequality makes use of the definition of the max operator[9], as well as the weak convexity of

[9] Specifically, the property that

$$\begin{aligned} &\alpha \max\{A, B\} + (1-\alpha)\max\{C, D\} \\ \geq &\ \max\{\alpha A + (1-\alpha) C, \alpha B + (1-\alpha) D\}, \end{aligned}$$

for all $\alpha \in [0,1]$ and $A, B, C, D \in \mathbb{R}$.

V_0 and V_{t-1}.[10] Therefore, W_t is weakly convex for all $t \geq 1$. Consequently, for each $t \geq 0$, V_t is also weakly convex, as it is equal to the maximum of two weakly convex functions.

Next observe that we have $V_t(p) \geq V_{t-1}(p)$, as the agent always have the option of not visiting the final expert.

Now, the family of functions $\{V_k(p)\}_{k\in\mathbb{N}}$ are nondecreasing and uniformly bounded by p, so its pointwise limit is unique. Denote it by $V(p)$. Moreover, the family of functions are weakly convex. By Theorem 3.1.4 of Hiriart-Urruty and Lemaréchal (2012), the convergence is uniform and $V(p)$ is also a nondecreasing and weakly convex function. Finally, the uniqueness of the value function follows from Lemma 3. Q.E.D.

Proof of Proposition 4: Take a candidate cutoff strategy $\hat{s}$ of the experts. Denote by $\Psi(\hat{s})$ the set of agent's pure strategy that is a best response to experts' cutoff $\hat{s}$. To ease exposition, define $J : [\underline{s}, \bar{s}] \times [\underline{s}, \bar{s}] \times \triangle\Lambda \to \mathbb{R}$ by

$$J(s, \hat{s}, \alpha) \equiv \int \left(1 + \frac{\sum_{h\in H} \beta_1(h, Y)\, q_0(h; \hat{s}, \beta)}{\sum_{h\in H} \beta_1(h, Y)\, q_1(h; \hat{s}, \beta)} \frac{1}{s} \frac{1-\pi}{\pi}\right)^{-1} d\alpha(\beta).$$

Note that $J(\cdot, \hat{s}, \alpha)$ is strictly increasing. Define $x : [\underline{s}, \bar{s}] \times \triangle\Lambda \to \mathbb{R}_+$ as the unique solution to the equation $J(\cdot, \hat{s}, \alpha) = \frac{l}{1+l}$. Define a correspondence $Z : [\underline{s}, \bar{s}] \rightrightarrows [\underline{s}, \bar{s}]$ as follows:

$$Z(\hat{s}) \equiv \{\max\{\underline{s}, \min\{\bar{s}, x(\hat{s}, \alpha)\}\} : \alpha \in \triangle(\Psi(\hat{s}))\}.$$

The correspondence Z can be interpreted as the set of best responses of an individual expert, given all other experts' cutoff strategy $\hat{s}$, and that the agent playing a best response to it. Specifically, if $x(\hat{s}, \alpha) \in [\underline{s}, \bar{s}]$, then an individual expert finds it optimal to adopt cutoff $x(\hat{s}, \alpha)$, given all other experts' adopting cutoff $\hat{s}$ and the agent playing $\alpha \in \triangle(\Psi(\hat{s}))$. If $x(\hat{s}, b) > \bar{s}$, then given others'

[10] Note that

$$\begin{aligned}
\frac{\hat{p}F(\tilde{s}|1)}{\hat{p}F(\tilde{s}|1) + (1-\hat{p})F(\tilde{s}|0)} &= \frac{\alpha\left[p'F(\tilde{s}|1) + (1-p')F(\tilde{s}|0)\right]}{\hat{p}F(\tilde{s}|1) + (1-\hat{p})F(\tilde{s}|0)} \frac{1}{1 + \frac{1-p'}{p'}\frac{F(\tilde{s}|0)}{F(\tilde{s}|1)}} \\
&+ \frac{(1-\alpha)\left[p''F(\tilde{s}|1) + (1-p'')F(\tilde{s}|0)\right]}{\hat{p}F(\tilde{s}|1) + (1-\hat{p})F(\tilde{s}|0)} \frac{1}{1 + \frac{1-p''}{p''}\frac{F(\tilde{s}|0)}{F(\tilde{s}|1)}}; \text{ and} \\
\frac{\hat{p}(1-F(\tilde{s}|1))}{\hat{p}F(\tilde{s}|1) + (1-\hat{p})F(\tilde{s}|0)} &= \frac{\alpha\left[p'(1-F(\tilde{s}|1)) + (1-p')(1-F(\tilde{s}|0))\right]}{\hat{p}(1-F(\tilde{s}|1)) + (1-\hat{p})(1-F(\tilde{s}|0))} \frac{1}{1 + \frac{1-p'}{p'}\frac{1-F(\tilde{s}|0)}{1-F(\tilde{s}|1)}} \\
&+ \frac{(1-\alpha)\left[p''(1-F(\tilde{s}|1)) + (1-p'')(1-F(\tilde{s}|0))\right]}{\hat{p}(1-F(\tilde{s}|1)) + (1-\hat{p})(1-F(\tilde{s}|0))} \frac{1}{1 + \frac{1-p''}{p''}\frac{1-F(\tilde{s}|0)}{1-F(\tilde{s}|1)}}.
\end{aligned}$$

strategy profile, an individual expert finds it optimal to always recommend N, i.e., adopting cutoff $\bar{s}$. Likewise, if $x(\hat{s}, b) < \bar{s}$, then an individual expert's best response is to always recommend Y, i.e., adopting cutoff $\underline{s}$.

To show equilibrium existence, it suffices to show that the correspondence Z has a fixed point. To this end, we invoke the Kakutani's fixed point theorem. In the remainder of the proof, we show that (i) Z is a non-empty-valued self-map, (ii) Z is convex-valued, and (iii) Z is upper semi-continuous.

(i) We have established the existence of agent's best response in the second benchmark case. The fact that Z is a non-empty valued self-map follows from the definition above.

(ii) $Z(\hat{s})$ is clearly convex-valued if $Z(\hat{s})$ is single-valued. Suppose $z, z^{\#} \in Z(\hat{s})$, and $z^{\&} \in (z, z^{\#})$. Denote by α and $\alpha^{\#}$ the corresponding optimal strategy respectively. Define $T : [0,1] \to \mathbb{R}$ by

$$T(\gamma) \equiv \gamma J\left(z^{\&}, \hat{s}, \alpha\right) + (1-\gamma) J\left(z^{\&}, \hat{s}, \alpha^{\#}\right) - \frac{l}{1+l}.$$

We have $T(1) > 0$ and $T(0) < 0$. As T is continuous and increasing, by the intermediate value theorem, there exists a unique $\gamma^* \in (0,1)$ such that $T(\gamma^*) = 0$. As every pure strategy on the support of α and $\alpha^{\#}$ is a best response to $\hat{s}$, it is clear that $\gamma^*\alpha + (1-\gamma^*)\alpha^{\#}$ is also a best response to $\hat{s}$. Therefore, $z^{\&} \in Z(\hat{s})$.

(iii) Recall the set of agent's best response $\Psi(\hat{s})$ is characterized by two cutoffs: one that calls for stopping the search altogether, and one that calls for taking the operation immediately. Specifically, the corresponding beginning-of-period value function, denoted by $V(p; \hat{s})$, is obtained by substituting $\tilde{s} = \hat{s}$ into (5). Similar to the definition (12) in the proof of Proposition 3, we denote by $W(p; \hat{s})$ the beginning-of-period value function if the agent decides to search this period:

$$\begin{aligned} W(p; \hat{s}) \equiv\; & -c + \left[p\left(1 - F(\hat{s}|1)\right) + (1-p)\left(1 - F(\hat{s}|0)\right)\right] \\ & \times \max\left\{ V\left(\frac{1}{1 + \frac{1-p}{p}\frac{1-F(\hat{s}|0)}{1-F(\hat{s}|1)}}; \hat{s}\right), -L + \frac{1}{1 + \frac{1-p}{p}\frac{1-F(\hat{s}|0)}{1-F(\hat{s}|1)}}(1+L) \right\} \\ & + \left[pF(\hat{s}|1) + (1-p)F(\hat{s}|0)\right] V\left(\frac{1}{1 + \frac{1-p}{p}\frac{F(\hat{s}|0)}{F(\hat{s}|1)}}; \hat{s}\right). \end{aligned} \tag{13}$$

The cutoffs $p_0(\hat{s})$ and $p_1(\hat{s})$ in the best response of the agent are given respectively by $W(p_0(\hat{s}); \hat{s}) = 0$ and $V(p_1(\hat{s}); \hat{s}) = -L + (1+L)p_1(\hat{s})$.[11] We first show the following lemma.

Lemma 7 *$p_0(\hat{s})$ and $p_1(\hat{s})$ are continuous in $\hat{s}$.*

[11] The existence and uniqueness of $p_0(\hat{s})$ and $p_1(\hat{s})$ follow by noting that the proof of Lemma 3 holds regardless of the cutoff strategy of the experts.

Proof. First we show that the family of functions $\{V(p;\hat{s})\}_{\hat{s}\in[\underline{s},\bar{s}]}$ and $\{W(p;\hat{s})\}_{\hat{s}\in[\underline{s},\bar{s}]}$ are Lipschitz continuous with a common modulus $1+L$. Recall from Proposition 3 that each $V(\cdot;\hat{s})$ and $W(\cdot;\hat{s})$ is weakly convex on $[0,1]$. Thus, each $V(\cdot;\hat{s})$ and $W(\cdot;\hat{s})$ is differentiable for almost all $p\in[0,1]$, and the derivative is nondecreasing in p. Suppose c is small enough so that $V(1;\hat{s})>0$. Otherwise, $V(p;\hat{s})=\frac{\partial V(p;\hat{s})}{\partial p}=0$ for all $p\in[0,1]$, and we are done. Next, if p is sufficiently large, according to (5), $V(p;\hat{s})$ is given by

$$\begin{aligned}V(p;\hat{s}) &= W(p;\hat{s}) \\ &= -c+\{-L(1-p)[1-F(\hat{s}|0)]+p[1-F(\hat{s}|1)]\}+[pF(\hat{s}|1)+(1-p)F(\hat{s}|0)]V\left(\frac{1}{1+\frac{1-p}{p}\frac{F(\hat{s}|0)}{F(\hat{s}|1)}};\hat{s}\right)\end{aligned}$$

To obtain an upper bound on the derivative $\frac{\partial V(p;\hat{s})}{\partial p}$, suppose $V(p;\hat{s})$ is differentiable at $p=1$. Differentiate both sides of the equation above with respect to p,

$$\begin{aligned}\frac{\partial V(p;\hat{s})}{\partial p} &= L[1-F(\hat{s}|0)]+[1-F(\hat{s}|1)]+[F(\hat{s}|1)-F(\hat{s}|0)]V\left(\frac{1}{1+\frac{1-p}{p}\frac{F(\hat{s}|0)}{F(\hat{s}|1)}};\hat{s}\right) \\ &\quad +\frac{F(\hat{s}|1)F(\hat{s}|0)}{(pF(\hat{s}|1)+(1-p)F(\hat{s}|0))}\frac{\partial}{\partial p}V\left(\frac{1}{1+\frac{1-p}{p}\frac{F(\hat{s}|0)}{F(\hat{s}|1)}};\hat{s}\right).\end{aligned}$$

Evaluating it at $p=1$ gives

$$\begin{aligned}\left.\frac{\partial V(p;\hat{s})}{\partial p}\right|_{p=1} &\leq L[1-F(\hat{s}|0)]+[1-F(\hat{s}|1)]+[F(\hat{s}|1)-F(\hat{s}|0)]+F(\hat{s}|0)\left.\frac{\partial V(p;\hat{s})}{\partial p}\right|_{p=1} \\ &\Leftrightarrow \left.\frac{\partial V(p;\hat{s})}{\partial p}\right|_{p=1}\leq 1+L.\end{aligned}$$

Thus, the families of functions $\{V(p;\hat{s})\}_{\hat{s}\in[\underline{s},\bar{s}]}$ and $\{W(p;\hat{s})\}_{\hat{s}\in[\underline{s},\bar{s}]}$ are Lipschitz continuous with a common modulus $1+L$.

Take an arbitrary sequence $\{\hat{s}_n\}$ that converges to some $\hat{s}\in[\underline{s},\bar{s}]$. This gives two sequences of cutoffs $\{p_0(\hat{s}_n)\}_n$ and $\{p_1(\hat{s}_n)\}_n$ satisfying

$$\begin{aligned}W(p_0(\hat{s}_n);\hat{s}_n) &= 0\text{, and} \\ V(p_1(\hat{s}_n);\hat{s}_n) &= -L+(1+L)p_1(\hat{s}_n).\end{aligned}\tag{14}$$

Suppose $\{p_0(\hat{s}_n)\}$ and $\{p_1(\hat{s}_n)\}$ are convergent (otherwise take subsequences). As both $\{V(p;\hat{s})\}_{\hat{s}\in[\underline{s},\bar{s}]}$ and $\{W(p;\hat{s})\}_{\hat{s}\in[\underline{s},\bar{s}]}$ are equicontinuous and uniformly bounded (by 1), by the Arzela-Ascoli Theorem, there exists a subsequence $\{\hat{s}_{n_k}\}$ such that $V(p;\hat{s}_{n_k})$ and $W(p;\hat{s}_{n_k})$ converge uniformly. It is clear that the limiting functions are respectively $V(p;\hat{s})$ and $W(p;\hat{s})$, as by Proposition 3, there exists a unique function that satisfies (5) with $\tilde{s} = \hat{s}$. Now pass (14) to the limit,

$$\begin{aligned} W\left(\lim_{k\to\infty} p_0\left(\hat{s}_{n_k}\right); \hat{s}\right) &= 0, \text{ and} \\ V\left(\lim_{k\to\infty} p_1\left(\hat{s}_{n_k}\right); \hat{s}\right) &= -L + (1+L)\lim_{k\to\infty} p_1\left(\hat{s}_{n_k}\right). \end{aligned}$$

By Lemma 3, the only subsequential limits of $\{p_0(\hat{s}_n)\}$ and $\{p_1(\hat{s}_n)\}$ are $p_0(\hat{s})$ and $p_1(\hat{s})$ respectively. Therefore, $\lim_{n\to\infty} p_0(\hat{s}_n) = p_0(\hat{s})$ and $\lim_{n\to\infty} p_1(\hat{s}_n) = p_1(\hat{s})$. ■

Now take a pair of sequences $\{\hat{s}_m\}, \{z_m\}$ such that $\hat{s}_m \to \hat{s}$, $z_m \in Z(\hat{s}_m)$, and $z_m \to z$. To prove the upper semi-continuity of Z, we need to show that $z \in Z(\hat{s})$. Suppose first that $z \in (\underline{s}, \bar{s})$. Then it is without loss to assume $z_m \in (\underline{s}, \bar{s})$ for all $m \in \mathbb{N}$ (otherwise, take a subsequence). Consequently, for all $m \in \mathbb{N}$, $J(z_m, \hat{s}_m, \alpha_m) = \frac{l}{1+l}$ for some mixed strategy α_m with support on $\Psi(\hat{s}_m)$. As this is a game of perfect recall, by the Kuhn's Theorem (Kuhn (1953)), there exists a sequence of behavioral strategies $\{b_m\}$ such that each b_m is equivalent to α_m in the game in which the experts adopt cutoff $\hat{s}_m$. Here, a behavioral strategy $b = (b^0, b^1)$ consists of two components: $b^0 : H \to [0,1]$ specifies the probability of not searching at the beginning of a period following history $h \in H$; $b^1 : H \to [0,1]$ specifies the probability of taking the operation after receiving a recommendation Y at the end of a period following history $h \in H$. With a slight abuse of notation, denote by $q(h; \hat{s}, b)$ the ex-ante probability that history $h \in H$ arises, given a strategy profile $(\hat{s}, b)$, and denote by $J(s, \hat{s}, b)$ the probability that $\omega = 1$ conditional on the expert being consulted, observing s, the agent takes the operation following a Y recommendation, as well as the equilibrium strategy profile being $(\hat{s}, b)$, i.e.,

$$J(s, \hat{s}, b) \equiv J(s, \hat{s}, \alpha(b)),$$

where $\alpha(b)$ is the mixed strategy equivalent to behavioral strategy b. Following computation similar to (7), we have

$$J(s, \hat{s}, b) = \left(1 + \frac{\sum_{h\in H} b^1(h, Y)\, q_0(h; \hat{s}, b)}{\sum_{h\in H} b^1(h, Y)\, q_1(h; \hat{s}, b)} \frac{1}{s} \frac{1-\pi}{\pi}\right)^{-1}. \tag{15}$$

As the set of histories H is countable, following a standard diagonalization argument, one can construct a subsequence $\{b_{m_k}\}$ that converges pointwise to some $b^{\#} : H \to [0,1]^2$. The following claim shows that $b^{\#}$ is a best response to $\hat{s}$ because of the continuity of $p_0(\cdot)$ and $p_1(\cdot)$ established in Lemma 7.

Claim 2 *$b^{\#}$ is a best response to $\hat{s}$.*

Proof. Suppose not. Then there exists a $h \in H$ such that either (i) $p(h;\hat{s}) \in (p_0(\hat{s}), p_1(\hat{s}))$ but either $b^{0\#}(h) > 0$, or $b^{1\#}(h) > 0$, or (ii) $p(h;\hat{s}) < p_0(\hat{s})$ but $b^{0\#}(h) = 0$; or (iii) $p(h;\hat{s}) > p_1(\hat{s})$ but $b^{1\#}(h) = 0$. Suppose case (i) arises. For either $i = 0, 1$, for all m sufficiently large, we have $b^i_m(h) > 0$ and $p(h;\hat{s}_m) \notin (p_0(\hat{s}_m), p_1(\hat{s}_m))$. As $p_0(\cdot)$ and $p_1(\cdot)$ are continuous, taking limit with respect to m gives $p(h;\hat{s}) \notin (p_0(\hat{s}), p_1(\hat{s}))$, a contradiction. Suppose case (ii) arises. As $p(h;s_m) \to p(h;\hat{s})$ and $p_0(s_m) \to p_0(\hat{s})$, we have that for all m sufficiently large, $p(h;\hat{s}_m) < p_0(\hat{s}_m)$, so $b^{0\#}_m(h) = 1$. Thus, $b^{0\#}(h) = 1$, a contradiction. Case (iii) is symmetric to case (ii). ■

It remains to show that for $\omega \in \{0, 1\}$,

$$\sum_{h\in H} b^1_{m_k}(h, Y)\, q_\omega(h;\hat{s}_{m_k}, b_{m_k}) \to \sum_{h\in H} b^{1\#}(h, Y)\, q_\omega\left(h;\hat{s}, b^{\#}\right) \text{ as } k \to \infty. \tag{16}$$

Observe first that the probability $q_\omega(h;\hat{s}_{m_k}, b_{m_k})$ can be decomposed as follows. Suppose $h = \left(r_1, r_2, ..., r_{|h|}\right)$.[12]

$$\begin{aligned} q_\omega(h;\hat{s}_{m_k}, b_{m_k}) &= \Pr(r_1|\hat{s}_{m_k}, \omega)\left(1 - b^1_{m_k}(r_1)\right)\left(1 - b^0_{m_k}(r_1)\right) \\ &\quad \times \Pr(r_2|\hat{s}_{m_k}, \omega)\left(1 - b^1_{m_k}(r_1, r_2)\right)\left(1 - b^0_{m_k}(r_1, r_2)\right) \\ &\quad \times .... \Pr\left(r_{|h|}|\hat{s}_{m_k}, \omega\right)\left(1 - b^1_{m_k}(h)\right)\left(1 - b^0_{m_k}(h)\right). \end{aligned}$$

Note that as $\Pr(r|\hat{s}_{m_k}, \omega)$ is either $F(\hat{s}_{m_k}|\omega)$ or $1 - F(\hat{s}_{m_k}|\omega)$, and $F(\cdot|\omega)$ is continuous, it is clear that $\Pr(r|\hat{s}_{m_k}, \omega) \to \Pr(r|\hat{s}, \omega)$. Together with the fact that $b^0_{m_k}(\cdot)$ and $b^1_{m_k}(\cdot)$ converges pointwise to $b^{0\#}(\cdot)$ and $b^{1\#}(\cdot)$ respectively, $q_\omega(h;\hat{s}_{m_k}, b_{m_k})$ converges to $q_\omega\left(h;\hat{s}, b^{\#}\right)$ for each $h \in H$. We now establish (16) using Lebesgue's dominated convergence theorem. For each $k \in \mathbb{N}$, define $Q_k : \mathbb{N} \to \mathbb{R}$ and $Q : \mathbb{N} \to \mathbb{R}$ as follows.

$$\begin{aligned} Q_k(n, \omega) &\equiv \sum_{h\in H:|h|=n} b^1_{m_k}(h, Y)\, q_\omega(h;\hat{s}_{m_k}, b_{m_k}) \text{ ;and} \\ Q(n, \omega) &\equiv \sum_{h\in H:|h|=n} b^{1\#}(h, Y)\, q_\omega\left(h;\hat{s}, b^{\#}\right). \end{aligned}$$

With these definitions, we can write

$$\sum_{h\in H} b_{m_k}(h, Y)\, q_\omega(h;\hat{s}_{m_k}, b_{m_k}) = \sum_{n=0}^{\infty} Q_k(n, \omega).$$

Convergence (16) can be restated as $\sum_{n=0}^{\infty} Q_k(n, \omega) \to \sum_{n=0}^{\infty} Q(n, \omega)$. We show below that there exists $N, K \in \mathbb{N}$ such that for all $n > N$ and $k > K$, we have $Q_k(n, \omega) \leq B_0 \exp(-B_1 n)$, where $B_0, B_1 > 0$ are constants. First, choose a $\varepsilon < \min\{p_0(\hat{s}), 1 - p_1(\hat{s})\}$, and a K such that for all $k > K$, we have

[12] Recall $|h|$ stands for the length of the history h.

$p_0(\hat{s}_{m_k}) > p_0(\hat{s}) - \varepsilon \equiv P_0$ and $p_1(\hat{s}_{m_k}) < p_1(\hat{s}) + \varepsilon \equiv P_1$. Take any $k > K$. Using the fact that period n is reached only if the agent's posterior belief at the end of period $n-1$ is within $[P_0, P_1]$,

$$\begin{aligned} Q_k(n,\omega) &\leq \sum_{h\in H:|h|=n-1} \Pr\left(p\left(h,\hat{s}_{m_k}\right) \in [P_0,P_1] |\omega\right) \\ &\leq \begin{cases} \sum_{h\in H:|h|=n-1} \Pr\left(p\left(h,\hat{s}_{m_k}\right) \leq P_1|\omega=1\right) & \text{if } \omega=1 \\ \sum_{h\in H:|h|=n-1} \Pr\left(p\left(h,\hat{s}_{m_k}\right) \geq P_0|\omega=0\right) & \text{if } \omega=0 \end{cases}. \end{aligned} \tag{17}$$

An upper bound on $\sum_{h\in H:|h|=n-1} \Pr\left(p\left(h,\hat{s}_{m_k}\right) \leq P_1|\omega=1\right)$ can be obtained by noting that for each $\hat{s}_{m_k}$, the expert's recommendation $r \in \{Y,N\}$ is a Bernoulli random variable, with $\Pr(r=Y) = 1 - F(\hat{s}_{m_k}|\omega)$. The agent's posterior after receiving $n-1$ recommendations is weakly less than P_1 if and only if the number y of recommendation Y is sufficiently small:

$$y \leq \frac{(n-1)\ln\left(\frac{F\left(\hat{s}_{m_k}|1\right)}{F\left(\hat{s}_{m_k}|0\right)}\right) + \ln\left(\frac{1}{P_1}-1\right) + \ln\frac{\pi}{1-\pi}}{\ln\left(\frac{1-F\left(\hat{s}_{m_k}|0\right)}{F\left(\hat{s}_{m_k}|0\right)}\right) - \ln\left(\frac{1-F\left(\hat{s}_{m_k}|1\right)}{F\left(\hat{s}_{m_k}|1\right)}\right)}.$$

Applying Hoeffding's inequality,

$$\begin{aligned} &\Pr\left(y \leq \frac{\ln\left(\frac{F\left(\hat{s}_{m_k}|1\right)}{F\left(\hat{s}_{m_k}|0\right)}\right) + \frac{1}{n-1}\left(\ln\left(\frac{1}{P_1}-1\right) + \ln\frac{\pi}{1-\pi}\right)}{\ln\left(\frac{1-F\left(\hat{s}_{m_k}|0\right)}{F\left(\hat{s}_{m_k}|0\right)}\right) - \ln\left(\frac{1-F\left(\hat{s}_{m_k}|1\right)}{F\left(\hat{s}_{m_k}|1\right)}\right)}(n-1)|\omega=1\right) \\ \leq\ &\exp\left(-2\left[\left(1-F\left(\hat{s}_{m_k}|1\right)\right) - \frac{\ln\left(\frac{F\left(\hat{s}_{m_k}|1\right)}{F\left(\hat{s}_{m_k}|0\right)}\right) + \frac{1}{n-1}\left(\ln\left(\frac{1}{P_1}-1\right) + \ln\frac{\pi}{1-\pi}\right)}{\ln\left(\frac{1-F\left(\hat{s}_{m_k}|0\right)}{F\left(\hat{s}_{m_k}|0\right)}\right) - \ln\left(\frac{1-F\left(\hat{s}_{m_k}|1\right)}{F\left(\hat{s}_{m_k}|1\right)}\right)}\right]^2 (n-1)|\omega=1\right). \end{aligned}$$

As $n,k\to\infty$, the term in the bracket in the last line approaches

$$(1-F(\hat{s}|1)) - \frac{\ln F(\hat{s}|0) - \ln F(\hat{s}|1)}{\ln\left(\frac{1-F(\hat{s}|1)}{F(\hat{s}|1)}\right) - \ln\left(\frac{1-F(\hat{s}|0)}{F(\hat{s}|0)}\right)} \equiv L_{\hat{s}} > 0.$$

Therefore, there exists a pair of sufficiently large integers $N_1, K_1 > K$ such that for all $n > N_1$ and $k > K_1$, we have

$$Q_k(n,1) \leq \sum_{h\in H:|h|=n-1} \Pr\left(p\left(h,\hat{s}_{m_k}\right) \leq P_1|\omega=1\right) \leq \exp\left(-\frac{L_{\hat{s}}^2}{2}(n-1)\right).$$

Define a dominating function $\Psi_1 : \mathbb{N} \to [0,1]$ by

$$\Psi_1(n) \equiv \begin{cases} 1 & \text{if } n < N_1 \\ \exp\left(-\frac{L_{\hat{s}}^2}{2}(n-1)\right) & \text{if } n \geq N_1 \end{cases}.$$

It is clear that

$$\sum_{n=1}^{\infty} \Psi_1(n) = N_1 + \frac{\exp\left(-\frac{L_{\hat{s}}^2}{2}(N_1 - 1)\right)}{1 - \exp\left(-\frac{L_{\hat{s}}^2}{2}\right)} < \infty.$$

Therefore by Lebesgue's dominated convergence theorem, we have $\sum_{n=0}^{\infty} Q_k(n,1) \to \sum_{n=0}^{\infty} Q(n,1)$ as $k \to \infty$.

Convergence $\sum_{n=0}^{\infty} Q_k(n,0) \to \sum_{n=0}^{\infty} Q(n,0)$ can be established in a similar way. Noting that the agent's posterior after receiving $n-1$ recommendations is weakly larger than P_0 if and only if the number y of recommendation Y is sufficiently large:

$$y \geq \frac{(n-1)\ln\left(\frac{F(\hat{s}_{m_k}|1)}{F(\hat{s}_{m_k}|0)}\right) + \ln\left(\frac{1}{P_0} - 1\right) + \ln\frac{\pi}{1-\pi}}{\ln\left(\frac{1-F(\hat{s}_{m_k}|0)}{F(\hat{s}_{m_k}|0)}\right) - \ln\left(\frac{1-F(\hat{s}_{m_k}|1)}{F(\hat{s}_{m_k}|1)}\right)}.$$

Applying Hoeffding's inequality,

$$\begin{aligned}
&\Pr\left(y \geq \frac{\ln\left(\frac{F(\hat{s}_{m_k}|1)}{F(\hat{s}_{m_k}|0)}\right) + \frac{1}{n-1}\left(\ln\left(\frac{1}{P_0} - 1\right) + \ln\frac{\pi}{1-\pi}\right)}{\ln\left(\frac{1-F(\hat{s}_{m_k}|0)}{F(\hat{s}_{m_k}|0)}\right) - \ln\left(\frac{1-F(\hat{s}_{m_k}|1)}{F(\hat{s}_{m_k}|1)}\right)}(n-1)\,|\omega = 0\right) \\
\leq\ &\exp\left(-2\left[\frac{\ln\left(\frac{F(\hat{s}_{m_k}|1)}{F(\hat{s}_{m_k}|0)}\right) + \frac{1}{n-1}\left(\ln\left(\frac{1}{P_0} - 1\right) + \ln\frac{\pi}{1-\pi}\right)}{\ln\left(\frac{1-F(\hat{s}_{m_k}|0)}{F(\hat{s}_{m_k}|0)}\right) - \ln\left(\frac{1-F(\hat{s}_{m_k}|1)}{F(\hat{s}_{m_k}|1)}\right)} - (1 - F(\hat{s}_{m_k}|0))\right]^2 (n-1)\right).
\end{aligned}$$

As $n, k \to \infty$, the term in the bracket in the last line approaches

$$\frac{\ln F(\hat{s}|0) - \ln F(\hat{s}|1)}{\ln\left(\frac{1-F(\hat{s}|1)}{F(\hat{s}|1)}\right) - \ln\left(\frac{1-F(\hat{s}|0)}{F(\hat{s}|0)}\right)} - (1 - F(\hat{s}|0)) \equiv L'_{\hat{s}} > 0.$$

Therefore, there exists a pair of sufficiently large integers N_0, K_0 such that for all $n > N_0$ and $k > K_0 > K$, we have

$$Q_k(n,0) \leq \sum_{h \in H:|h|=n-1} \Pr(p(h, \hat{s}_{m_k}) \geq P_0|\omega = 0) \leq \exp\left(-\frac{L_{\hat{s}}'^2}{2}(n-1)\right).$$

Define a dominating function $\Psi_0 : \mathbb{N} \to [0,1]$ by

$$\Psi_0(n) \equiv \begin{cases} 1 & \text{if } n < N_0 \\ \exp\left(-\frac{L_{\hat{s}}'^2}{2}(n-1)\right) & \text{if } n \geq N_0 \end{cases}.$$

It is clear that

$$\sum_{n=1}^{\infty} \Psi_0(n) = N_0 + \frac{\exp\left(-\frac{L_{\hat{s}}'^2}{2}(N_0 - 1)\right)}{1 - \exp\left(-\frac{L_{\hat{s}}^2}{2}\right)} < \infty.$$

Therefore by Lebesgue's dominated convergence theorem, we have $\sum_{n=0}^{\infty} Q_k(n,0) \to \sum_{n=0}^{\infty} Q(n,0)$.

Next, suppose $z = \bar{s}$. Suppose there exists a subsequence $\{z_{m_k}\} \subset \{z_m\}$ such that for all $k \in \mathbb{N}$, $J(z_{m_k}, \hat{s}_{m_k}, \alpha_{m_k}) = \frac{l}{1+l}$ for some $\alpha_{m_k} \in \triangle(\Psi(\hat{s}_{m_k}))$. Then the argument above applies and we have $z \in Z(\hat{s})$. Suppose the subsequence above does not exist. Then there exists a subsequence $\{z_{m_k}\} \subset \{z_m\}$ such that (i) for all $k \in \mathbb{N}$, $J(z_{m_k}, \hat{s}_{m_k}, \alpha_{m_k}) < \frac{l}{1+l}$ for some $\alpha_{m_k} \in \triangle(\Psi(\hat{s}_{m_k}))$; and (ii) α_{m_k} converges to some $\alpha \in \triangle(\Psi(\hat{s}))$. Then the argument above implies that $J(z_{m_k}, \hat{s}_{m_k}, \alpha_{m_k}) \to J(z, \hat{s}, \alpha) \leq \frac{l}{1+l}$. Thus, $z = \bar{s} \in Z(\hat{s})$.

Finally, suppose $z = \underline{s}$. Again, suppose there exists a subsequence $\{z_{m_k}\} \subset \{z_m\}$ such that (i) for all $k \in \mathbb{N}$, $J(z_{m_k}, \hat{s}_{m_k}, \alpha_{m_k}) > \frac{l}{1+l}$ for some $\alpha_{m_k} \in \triangle(\Psi(\hat{s}_{m_k}))$; and (ii) α_{m_k} converges to some $\alpha \in \triangle(\Psi(\hat{s}))$, as otherwise, we are done. Taking limit as $k \to \infty$, $J(z, \hat{s}, \alpha) \geq \frac{l}{1+l}$. Thus, $z = \underline{s} \in Z(\hat{s})$. Q.E.D.

Proposition 7 *Denote by G^M the game identical to the main model except that there are only M available experts.*

(i) Each game G^M has an equilibrium.

(ii) Suppose $(\hat{s}_M, b_M)$ is an equilibrium of game G^M, where $\hat{s}_M$ is the experts' cutoff, and b_M is the agent's behavioral strategy. Suppose further that $\{(\hat{s}_M, b_M)\}_{M \in \mathbb{N}}$ converges pointwise to (s^*, b^*). Then (s^*, b^*) is an equilibrium of the game with infinitely many experts, as defined in Definition 1.

Proof. (i) The existence of equilibrium in game G^M can be proved by making appropriate modifications to the proof of Proposition 4. First, note that the definition of a strategy in game G^M is different from that in the main model. Specifically, in game G^M, the set of recommendation histories is $H^M = \{Y, N\}^M$. A behavioral strategy b of the agent is a pair of mapping $b^0, b^1 : H^M \to [0,1]$ that respectively specify the probability of exit at the beginning and the end of each period (conditional on being recommended Y). To make use of the proof of Proposition 4, it suffices to note that one can extend the definition of b to $\{Y, N\}^{\infty}$ in the trivial manner: set $b^0(h) = b^1(h) = 1$ for all $h \in H$ such that $|h| > M$.

Second, in the proof of Proposition 4, we have established the following fact. Suppose $\{\hat{s}_m\}$ is a sequence of experts' cutoffs such that $\hat{s}_m \to \hat{s} \in (\underline{s}, \bar{s})$, each behavioral strategy b_m is a best response to $\hat{s}_m$, and $\{b_m\}$ converges pointwise (i.e., history-by-history) to some $b^{\#}$. Then $b^{\#}$ is also a best response to $\hat{s}$. Below we show how to derive the analogous result in game G^M. In game G^M, the set of best response to experts' cutoff $\hat{s}_m$ is time-dependent; for each period $\tau \leq M$, it is optimal to continue the search only if the current posterior belief lies in an interval denoted by $\left[p_0^{M,\tau}(\hat{s}_m), p_1^{M,\tau}(\hat{s}_m)\right]$. These cutoffs can be derived using value functions similar to those defined in the proof of Proposition 3. Specifically, denote by $V^M_{M-\tau}(p;\hat{s}_m)$ the agent's beginning-of-period-τ value function at posterior p and experts' cutoff $\hat{s}_m$; and $W^M_{M-\tau}(p;\hat{s}_m)$ the value function assuming that he is forced to search in period τ. These functions are defined recursively as in (11) and (12) with $\tilde{s}=\hat{s}_m$. Then $p^{M,\tau}_0(\hat{s}_m)$ and $p^{M,\tau}_1(\hat{s}_m)$ satisfy

$$W^M_{M-\tau}\left(p^{M,\tau}_0(\hat{s}_m);\hat{s}_m\right)=0;\text{ and }V^M_{M-\tau}\left(p^{M,\tau}_1(\hat{s}_m);\hat{s}_m\right)=-L+(1+L)\,p^{M,\tau}_1(\hat{s}_m). \tag{18}$$

To obtain the aforementioned result, it suffices to show that for all $\tau\leq M$, the functions $p^{M,\tau}_0(\cdot)$ and $p^{M,\tau}_1(\cdot)$ are continuous. Take a sequence $\{\hat{s}_m\}$ of experts' cutoff such that $\{\hat{s}_m\}_m$, $\left\{p^{M,\tau}_0(\hat{s}_m)\right\}_m$ and $\left\{p^{M,\tau}_1(\hat{s}_m)\right\}_m$ are all convergent for all $\tau\leq M$. Following an argument similar to that in the proof of Lemma 7, the families of functions $\left\{W^M_{M-\tau}(\cdot;\hat{s}_m)\right\}_{m,\tau}$ and $\left\{V^M_{M-\tau}(\cdot;\hat{s}_m)\right\}_{m,\tau}$ are Lipschitz continuous with a common modulus $1+L$. These families are clearly uniformly bounded by 1. Therefore, by Arzela-Ascoli theorem, there exists a subsequence $\{s_{m_k}\}_k$ such that for all $\tau\leq M$, the sequences of functions $\left\{W^M_{M-\tau}(\cdot;\hat{s}_{m_k})\right\}_k$ and $\left\{V^M_{M-\tau}(\cdot;\hat{s}_{m_k})\right\}_k$ converge uniformly. As $W^M_{M-\tau}(\cdot;\hat{s}_{m_k})$ and $V^M_{M-\tau}(\cdot;\hat{s}_{m_k})$ satisfy their respective recursive definitions for all $\tau\leq M$ and $k\in\mathbb{N}$, taking the limit as $k\to\infty$, the limiting functions must satisfy the respective recursive definitions for $\hat{s}$. It is clear from the construction of $V^M_{M-\tau}(\cdot;\hat{s})$ and $W^M_{M-\tau}(\cdot;\hat{s})$ in the proof of Proposition 3 that these functions are unique. Therefore, we have that $W^M_{M-\tau}(\cdot;\hat{s}_{m_k})$ converges uniformly to $W^M_{M-\tau}(\cdot;\hat{s})$, and that $V^M_{M-\tau}(\cdot;\hat{s}_{m_k})$ converges uniformly to $V^M_{M-\tau}(\cdot;\hat{s})$. Now replacing $\hat{s}_m$ with $\hat{s}_{m_k}$ in (18) and taking limit, we get

$$W^M_{M-\tau}\left(\lim_{k\to\infty}p^{M,\tau}_0(\hat{s}_{m_k});\hat{s}\right)=0;\text{ and }V^M_{M-\tau}\left(\lim_{k\to\infty}p^{M,\tau}_1(\hat{s}_{m_k});\hat{s}\right)=-L+(1+L)\lim_{k\to\infty}p^{M,\tau}_1(\hat{s}_{m_k}).$$

As $p^{M,\tau}_0(\hat{s})$ and $p^{M,\tau}_1(\hat{s})$ are unique, we have that $\lim_{k\to\infty}p^{M,\tau}_0(\hat{s}_{m_k})=p^{M,\tau}_0(\hat{s})$, and $\lim_{k\to\infty}p^{M,\tau}_1(\hat{s}_{m_k})=p^{M,\tau}_1(\hat{s})$.

Finally, as explained in the text, the experts' conditional belief is given exactly by (7) in a game with finitely many experts G^M. With observations outlined above, one can follow the proof of Proposition 4 to establish equilibrium existence for game G^M.

(ii) We first establish that the agent's behavioral strategy b^* is best response to experts' cutoff s^*. Denote by $p^{M,\tau}_0(\hat{s})$ and $p^{M,\tau}_1(\hat{s})$ the optimal cutoff posteriors of the agent in game G^M in period τ, given that all experts adopting cutoff strategy $\hat{s}$. Fix a $\tau\leq M$. Following an argument similar to that in Lemma 7, cutoffs $p^{M,\tau}_0(\hat{s})$ and $p^{M,\tau}_1(\hat{s})$ are determined by

$$W^M_{M-\tau}\left(p^{M,\tau}_0(\hat{s}_M);\hat{s}_M\right)=0;\text{ and }V^M_{M-\tau}\left(p^{M,\tau}_1(\hat{s}_M);\hat{s}_M\right)=-L+(1+L)\,p^{M,\tau}_1(\hat{s}_M). \tag{19}$$

As $\left\{V^M_{M-\tau}(\cdot;\hat{s}_M)\right\}_{M\in\mathbb{N},M\geq\tau}$ and $\left\{W^M_{M-\tau}(\cdot;\hat{s}_M)\right\}_{M\in\mathbb{N},M\geq\tau}$ are both equicontinuous and uniformly bounded by 1, there exists a subsequence $\{M^\tau_k\}\subset\mathbb{N}$ such that both sequences converge uniformly.

Carrying out this procedure inductively and applying diagonalization, one can construct a sequence $\{M_k\} \subset \mathbb{N}$ such that for all $\tau \in \mathbb{N}$, $V^{M_k}_{M_k-\tau}(\cdot;\hat{s}_{m_k})$ converges uniformly to $V(\cdot;\hat{s})$. Substitute $M = M_k$ in (19) and taking limit, we get

$$W\left(\lim_{k\to\infty} p_0^{M_k,\tau}(\hat{s}_{M_k});s^*\right) = 0; \text{ and } V\left(\lim_{k\to\infty} p_1^{M_k,\tau}(\hat{s}_{M_k});s^*\right) = -L + (1+L)\left[\lim_{k\to\infty} p_1^{M_k,\tau}(\hat{s}_{M_k})\right].$$

Thus, $\lim_{k\to\infty} p_0^{M_k,\tau}(\hat{s}_{M_k}) = p_0(s^*)$ and $\lim_{k\to\infty} p_1^{M_k,\tau}(\hat{s}_{M_k}) = p_1(s^*)$ for all $\tau \in \mathbb{N}$. Consequently, b is a best response to s^*, following from an argument similar to Claim 2.

The conditional probability formula (15) holds exactly for finitely many experts. Thus,

$$\hat{s}_{M_k} = \max\left\{\underline{s}, \min\left\{\frac{\sum_{h\in H} b^1_{M_k}(h,Y)\, q_0(h;\hat{s}_{M_k},b_{M_k})}{\sum_{h\in H} b^1_{M_k}(h,Y)\, q_1(h;\hat{s}_{M_k},b_{M_k})}\frac{1-\pi}{\pi}l, \bar{s}\right\}\right\}.$$

In the proof of Proposition 4, we have established that for each $\omega \in \{0,1\}$, $\sum_{h\in H} b^1_{M_k}(h,Y)\, q_\omega(h;\hat{s}_{M_k},b_{M_k}) \to \sum_{h\in H} b^{*1}(h,Y)\, q_\omega(h;s^*,b^*)$. Therefore, taking limit on both sides of the equation above,

$$s^* = \max\left\{\underline{s}, \min\left\{\frac{\sum_{h\in H} b^{*1}(h,Y)\, q_0(h;s^*,b^*)}{\sum_{h\in H} b^{*1}(h,Y)\, q_1(h;s^*,b^*)}\frac{1-\pi}{\pi}l, \bar{s}\right\}\right\},$$

which is equivalent to condition (i) in the definition of an equilibrium (Definition 1). ■

Proof of Lemma 4: To ease notation, denote $s^*_n \equiv s^*(c_n,\underline{s})$. Suppose $\{s^*_n\}$ does not converge to $\underline{s}$. Then there exists a $\varepsilon > 0$ and a convergent subsequence $\{s^*_{n_k}\}$ such that $s^*_{n_k} \to \underline{s} + \varepsilon$. Consider the agent's search problem. Denote by $\{p_{1,n_k}\}$ the corresponding sequence of equilibrium (upper) cutoffs. Using (5), his equilibrium value function $V_k(\cdot)$ at search cost c_{n_k} is given by

$$V_k(p) = \max\left\{\begin{array}{c} 0, -c_{n_k} + \left[p\left(1-F\left(s^*_{n_k}|1\right)\right) + (1-p)\left(1-F\left(s^*_{n_k}|0\right)\right)\right] \\ \times \max\left\{V_k\left(\frac{1}{1+\frac{1-p}{p}\frac{1-F\left(s^*_{n_k}|0\right)}{1-F\left(s^*_{n_k}|1\right)}}\right), -L + \frac{1}{1+\frac{1-p}{p}\frac{1-F\left(s^*_{n_k}|0\right)}{1-F\left(s^*_{n_k}|1\right)}}(1+L)\right\} \\ + \left[pF\left(s^*_{n_k}|1\right) + (1-p)F\left(s^*_{n_k}|0\right)\right] V_k\left(\frac{1}{1+\frac{1-p}{p}\frac{F\left(s^*_{n_k}|0\right)}{F\left(s^*_{n_k}|1\right)}}\right) \end{array}\right\}.$$

Taking limit of $k \to \infty$, we get

$$V(p) = \max \left\{ \begin{array}{c} 0, [p(1 - F(\underline{s}+\varepsilon|1)) + (1-p)(1 - F(\underline{s}+\varepsilon|0))] \\ \times \max \left\{ V\left(\frac{1}{1+\frac{1-p}{p}\frac{1-F(\underline{s}+\varepsilon|0)}{1-F(\underline{s}+\varepsilon|1)}} \right), -L + \frac{1}{1+\frac{1-p}{p}\frac{1-F(\underline{s}+\varepsilon|0)}{1-F(\underline{s}+\varepsilon|1)}}(1+L) \right\} \\ + [pF(\underline{s}+\varepsilon|1) + (1-p)F(\underline{s}+\varepsilon|0)] V\left(\frac{1}{1+\frac{1-p}{p}\frac{F(\underline{s}+\varepsilon|0)}{F(\underline{s}+\varepsilon|1)}} \right) \end{array} \right\}.$$

It is straightforward to verify that $V(p) = p$ is a solution. Moreover, it is unique by Proposition 3. Therefore, $p_{1,n_k} \to 1$. However, an upper bound on $s^*_{n_k}$ is given by

$$\frac{1}{1 + \frac{1-p'_1\left(s^*_{n_k}\right)}{p'_1\left(s^*_{n_k}\right)}\frac{1}{s^*_{n_k}}} \leq \frac{l}{1+l} \Leftrightarrow s^*_{n_k} \leq \frac{1-p'_1\left(s^*_{n_k}\right)}{p'_1\left(s^*_{n_k}\right)} l,$$

where $p'_1\left(s^*_{n_k}\right)$ is defined in (10). Substitute $p'_1\left(s^*_{n_k}\right) = \left(1 + \left(\frac{1}{p_{1,n_k}} - 1\right)\frac{1-F\left(s^*_{n_k}|1\right)}{1-F\left(s^*_{n_k}|0\right)}\right)^{-1}$ into the upper bound above,

$$\frac{1-p'_1\left(s^*_{n_k}\right)}{p'_1\left(s^*_{n_k}\right)} l = \left(\left(\frac{1}{p_{1,n_k}} - 1\right)\frac{1-F\left(s^*_{n_k}|1\right)}{1-F\left(s^*_{n_k}|0\right)}\right) l,$$

which converges to 0 as $k \to \infty$, because $p_{1,n_k} \to 1$. Thus, $s^*_{n_k} \to 0$, a contradiction. Q.E.D.

Proof of Lemma 5: (i) To ease notation, denote $s^*_n \equiv s^*(c_n, \underline{s})$. By the previous lemma, $s^*_n \to \underline{s}$. As $s^*_n > \underline{s}$, (8) holds with equality, for all $n \in \mathbb{N}$. Moreover, conditional on the agent taking the operation upon recommendation Y, the most pessimistic and the most optimistic beliefs for the expert are respectively $p'_1(s^*_n)$ and $p_1(s^*_n)$. We therefore have

$$\frac{1}{1 + \frac{1-p'_1(s^*_n)}{p'_1(s^*_n)}\frac{1}{s^*_n}} \leq \frac{l}{1+l} \leq \frac{1}{1 + \frac{1-p_1(s^*_n)}{p_1(s^*_n)}\frac{1}{s^*_n}} \Leftrightarrow p'_1(s^*_n) \leq \frac{l}{s^*_n + l} \leq p_1(s^*_n). \tag{20}$$

The difference between $p_1(s^*_n)$ and $p'_1(s^*_n)$ vanishes as $s^*_n \to \underline{s}$ because

$$p_1\left(s^*_{n_k}\right) - p'_1(s^*_n) = p_1(s^*_n) - \frac{1}{\left(\frac{1}{p_1(s^*_n)} - 1\right)\frac{1-F(s^*_n|1)}{1-F(s^*_n|0)} + 1} = p_1(s^*_n)\frac{\frac{1-F(s^*_n|1)}{1-F(s^*_n|0)} - 1}{\frac{1-F(s^*_n|1)}{1-F(s^*_n|0)} + 1} \to 0 \text{ as } n \to \infty.$$

Therefore, $\lim_{k\to\infty} p_1(s^*_n) = \lim_{k\to\infty} p'_1(s^*_n) = \frac{l}{\underline{s}+l}$.

(ii) For each $n \in \mathbb{N}$, denote by $V_n(\cdot)$ the agent's value function in an equilibrium of the game in which the search cost is c_n and the experts' cutoff is $s^*_n = s^*(c_n, 0)$. Recall the agent's cutoff $p_1(s^*_n)$ is defined as the solution to $V_n(p) = -L + p(1+L)$:

$$\begin{aligned} & -L + p_1(s^*_n)(1+L) \\ = \; & -c_n + [p_1(s^*_n)(1 - F(s^*_n|1)) + (1 - p_1(s^*_n))(1 - F(s^*_n|0))]\left(-L + \frac{1}{1 + \frac{1-p_1(s^*_n)}{p_1(s^*_n)}\frac{1-F(s^*_n|0)}{1-F(s^*_n|1)}}(1+L)\right) \\ & + [p_1(s^*_n)F(s^*_n|1) + (1-p_1(s^*_n))F(s^*_n|0)] V_n\left(\frac{1}{1 + \frac{1-p_1(s^*_n)}{p_1(s^*_n)}\frac{F(s^*_n|0)}{F(s^*_n|1)}}\right). \end{aligned}$$

Upon re-arranging,

$$c_n = -p_1(s_n^*) F(s_n^*|1)(1+L) + [p_1(s_n^*) F(s_n^*|1) + (1-p_1(s_n^*)) F(s_n^*|0)] \left[V_n \left(\frac{1}{1+\frac{1-p_1(s_n^*)}{p_1(s_n^*)} \frac{F(s_n^*|0)}{F(s_n^*|1)}} \right) + L \right]. \tag{21}$$

As $V_n(p) \leq 1$ for all $n \in \mathbb{N}$ and $p \in [0,1]$, and that $p_1(s_n^*) \geq \frac{l}{s_n^*+l}$ (see (20)), equation (21) implies that

$$c_n \leq (1+L) \frac{s_n^* F(s_n^*|0)}{s_n^*+l}. \tag{22}$$

Now for each $n \in \mathbb{N}$, consider the following (necessarily suboptimal) search strategy of the agent: sample a fixed number M_n of experts and take the operation in the end if and only if all of them recommend Y. Here, M_n is chosen such that the posterior reaches $p_1(s_n^*)$ if all M_n experts recommends Y. The agent's equilibrium payoff is bounded from below by the expected payoff of this strategy, which is given by

$$\pi (1-F(s_n^*|1))^{M_n} + (1-\pi)(1-F(s_n^*|0))^{M_n}(-L) - c_n M_n.$$

We consider the limit of (a) $c_n M_n$, and (b) $\pi(1-F(s_n^*|1))^{M_n} + (1-\pi)(1-F(s_n^*|0))^{M_n}(-L)$ respectively. First consider (a). The number M_n must satisfy

$$\frac{1}{1+\frac{1-\pi}{\pi}\left(\frac{1-F(s_n^*|0)}{1-F(s_n^*|1)}\right)^{M_n-1}} \leq p_1(s_n^*) \leq \frac{1}{1+\frac{1-\pi}{\pi}\left(\frac{1-F(s_n^*|0)}{1-F(s_n^*|1)}\right)^{M_n}}$$

$$\Leftrightarrow \quad M_n \in \left[\frac{\ln\left(\frac{\pi}{1-\pi}\left(\frac{1}{p_1(s_n^*)}-1\right)\right)}{\ln\left(\frac{1-F(s_n^*|0)}{1-F(s_n^*|1)}\right)}, \frac{\ln\left(\frac{\pi}{1-\pi}\left(\frac{1}{p_1(s_n^*)}-1\right)\right)}{\ln\left(\frac{1-F(s_n^*|0)}{1-F(s_n^*|1)}\right)} + 1 \right]. \tag{23}$$

Next, as $p_1'(s_n^*) \leq \frac{l}{s_n^*+l}$ (see (20)), using the definition of $p_1'(s_n^*)$ (see (10)), an upper bound on $p_1(s_n^*)$ is given by:

$$p_1(s_n^*) \leq \frac{1}{\frac{s_n^*}{l}\frac{1-F(s_n^*|0)}{1-F(s_n^*|1)}+1}. \tag{24}$$

Consequently, an upper bound on M_n is given by

$$M_n \leq \frac{\ln\left(\frac{\pi}{1-\pi}\left(\frac{1}{p_1(s_n^*)}-1\right)\right)}{\ln\left(\frac{1-F(s_n^*|0)}{1-F(s_n^*|1)}\right)} + 1 \leq \frac{\ln\left(\frac{\pi}{1-\pi}\frac{s_n^*}{l}\frac{1-F(s_n^*|0)}{1-F(s_n^*|1)}\right)}{\ln\left(\frac{1-F(s_n^*|0)}{1-F(s_n^*|1)}\right)} + 1. \tag{25}$$

Using both (22) and (25), we have

$$c_n M_n \leq \left(\frac{1+L}{s_n^*+l}\right)\left[s_n^* \ln\left(\frac{\pi}{1-\pi}\frac{s_n^*}{l}\frac{1-F(s_n^*|0)}{1-F(s_n^*|1)}\right)\right] \frac{F(s_n^*|0)}{\ln\left(\frac{1-F(s_n^*|0)}{1-F(s_n^*|1)}\right)} + (1+L)\frac{s_n^*}{s_n^*+l}F(s_n^*|0).$$

Now we take limit on both sides of the inequality above. By Lemma 4, $s_n^* \to 0$. For the first term on the right hand side, $\lim_{n\to\infty}\left(\frac{1+L}{s_n^*+l}\right) = \frac{1+L}{l}$. For the second (bracketed) term,

$$\begin{aligned}
& \lim_{n\to\infty} s_n^* \ln\left(\frac{\pi}{1-\pi}\frac{s_n^*}{l}\frac{1-F(s_n^*|0)}{1-F(s_n^*|1)}\right) \\
= & \lim_{s^*\to 0}\frac{\ln\left(\frac{\pi}{1-\pi}\frac{s^*}{l}\frac{1-F(s^*|0)}{1-F(s^*|1)}\right)}{s^{*-1}} \\
= & \lim_{s^*\to 0}\frac{\left(\frac{\pi}{1-\pi}\frac{s^*}{l}\frac{1-F(s^*|0)}{1-F(s^*|1)}\right)^{-1}\frac{\pi}{1-\pi}\frac{1}{l}\left(\frac{1-F(s^*|0)}{1-F(s^*|1)}+s^*\frac{f(s^*|1)(1-F(s^*|0))-(1-F(s^*|1))f(s^*|0)}{(1-F(s^*|1))^2}\right)}{-s^{*-2}} \\
= & -\lim_{s^*\to 0} s^*\left(1+s^*\frac{f(s^*|1)(1-F(s^*|0))-f(s^*|0)(1-F(s^*|1))}{(1-F(s^*|0))(1-F(s^*|1))}\right) \\
= & \ 0.
\end{aligned}$$

In the calculation above, the second equality uses L'Hospital rule. For the third term,

$$\lim_{n\to\infty}\frac{F(s_n^*|0)}{\ln\left(\frac{1-F(s_n^*|0)}{1-F(s_n^*|1)}\right)} = \lim_{s^*\to 0}\frac{F(s^*|0)}{\ln\left(\frac{1-F(s^*|0)}{1-F(s^*|1)}\right)} = \lim_{s^*\to 0}\frac{f(s^*|0)}{\frac{f(s^*|1)(1-F(s^*|0))-(1-F(s^*|1))f(s^*|0)}{(1-F(s^*|0))(1-F(s^*|1))}} = \frac{f(0|0)}{f(0|1)-f(0|0)} = -1,$$

where second equality uses L'Hospital rule. For the last term, $\lim_{n\to\infty}(1+L)\frac{s_n^*}{s_n^*+l}F(s_n^*|0) = 0$. Therefore, we have

$$\lim_{n\to\infty} c_n M_n = 0. \tag{26}$$

Next we consider the limit of (b): $\pi(1-F(s_n^*|1))^{M_n} + (1-\pi)(1-F(s_n^*|0))^{M_n}(-L)$. Note that for n sufficiently large,

$$\begin{aligned}
& \pi(1-F(s_n^*|1))^{M_n} + (1-\pi)(1-F(s_n^*|0))^{M_n}(-L) \\
\geq & \ (1-\pi)\left[\frac{p_1(s_n^*)}{1-p_1(s_n^*)}\frac{1-F(s_n^*|0)}{1-F(s_n^*|1)} - L\right](1-F(s_n^*|0))^{M_n} \\
\geq & \ (1-\pi)\left[\frac{l}{s_n^*}\frac{1-F(s_n^*|0)}{1-F(s_n^*|1)} - L\right](1-F(s_n^*|0))^{M_n} \\
\geq & \ (1-\pi)\left[\frac{l}{s_n^*}\frac{1-F(s_n^*|0)}{1-F(s_n^*|1)} - L\right](1-F(s_n^*|0))^{\frac{\ln\left(\frac{\pi}{1-\pi}\frac{s_n^*}{l}\frac{1-F(s_n^*|0)}{1-F(s_n^*|1)}\right)}{\ln\left(\frac{1-F(s_n^*|0)}{1-F(s_n^*|1)}\right)}+1} \equiv \Gamma_n.
\end{aligned}$$

where the first inequality uses (23), the second inequality uses (20), and the last inequality uses the fact that $s_n^* \to 0$, and (25). Below we show that $\lim_{n\to\infty}\Gamma_n = \pi$. To this end, first consider the limit of Λ_n^1 defined by

$$\Gamma_n^1 \equiv \frac{\ln\left(\frac{\pi}{1-\pi}\frac{s_n^*}{l}\frac{1-F(s_n^*|0)}{1-F(s_n^*|1)}\right)}{\ln\left(\frac{1-F(s_n^*|0)}{1-F(s_n^*|1)}\right)}\ln(1-F(s_n^*|0)).$$

We show below that $\lim_{n\to\infty}\Gamma_n^1 = -\infty$. Applying L'Hospital rule, and using the normalization $f(s^*|1) = s^* f(s^*|0)$, we have

$$\begin{aligned}
&\lim_{n\to\infty}\Gamma_n^1 \\
= &\lim_{s^*\to 0}\frac{\ln(1-F(s^*|0))}{\frac{\ln\left(\frac{1-F(s^*|0)}{1-F(s^*|1)}\right)}{\ln\left(\frac{\pi}{1-\pi}\frac{s^*}{l}\frac{1-F(s^*|0)}{1-F(s^*|1)}\right)}} \\
= &\lim_{s^*\to 0}\frac{\left(\ln\left(\frac{\pi}{l(1-\pi)}\right)+\ln s^*+\ln(1-F(s^*|0))-\ln(1-F(s^*|1))\right)^2\frac{f(s^*|0)}{1-F(s^*|0)}}{\left(\ln(1-F(s^*|0))-\ln(1-F(s^*|1))\right)\left(\frac{1}{s^*}\right)-f(s^*|0)\left(\frac{-1}{1-F(s^*|0)}+\frac{s^*}{1-F(s^*|1)}\right)\left(\ln\left(\frac{\pi}{l(1-\pi)}\right)+\ln s^*\right)}.
\end{aligned}$$

Observe that the denominator in the expression above is negative. Moreover,

$$\begin{aligned}
\lim_{s^*\to 0}s^*\ln s^* &= 0;\text{ and} \\
\lim_{s^*\to 0}\frac{\ln(1-F(s^*|0))-\ln(1-F(s^*|1))}{s^*} &= \lim_{s^*\to 0}\frac{f(s^*|0)\left(\frac{-1}{1-F(s^*|0)}+\frac{s^*}{1-F(s^*|1)}\right)}{1} = f(0|0).
\end{aligned}$$

Therefore, $\lim_{n\to\infty}\Gamma_n^1 = -\infty$. Next, consider the limit of the sequence Γ_n^2 defined by

$$\Gamma_n^2 \equiv -\ln s_n^* + \frac{\ln\left(\frac{\pi}{1-\pi}\frac{s_n^*}{l}\frac{1-F(s_n^*|0)}{1-F(s_n^*|1)}\right)}{\ln\left(\frac{1-F(s_n^*|0)}{1-F(s_n^*|1)}\right)}\ln(1-F(s_n^*|0)).$$

We compute $\lim_{n\to\infty}\Gamma_n^2$ below. The limit can be written as:

$$\lim_{n\to\infty}\Gamma_n^2 = \lim_{s^*\to 0}\frac{\ln s^*\ln(1-F(s^*|1))}{\ln(1-F(s^*|0))-\ln(1-F(s^*|1))}+\lim_{s^*\to 0}\frac{\left[\ln\left(\frac{\pi}{l(1-\pi)}\right)+\ln\left(\frac{1-F(s^*|0)}{1-F(s^*|1)}\right)\right]\ln(1-F(s^*|0))}{\ln(1-F(s^*|0))-\ln(1-F(s^*|1))}.$$

Consider the first term. Applying L'Hospital rule,

$$\begin{aligned}
&\lim_{s^*\to 0}\frac{\ln s^*\ln(1-F(s^*|1))}{\ln(1-F(s^*|0))-\ln(1-F(s^*|1))} \\
= &\lim_{s^*\to 0}\frac{\ln s^*}{\frac{\ln(1-F(s^*|0))}{\ln(1-F(s^*|1))}-1} \\
= &\lim_{s^*\to 0}\frac{(\ln(1-F(s^*|1)))}{-(1-F(s^*|1))\ln(1-F(s^*|1))+s^*(1-F(s^*|0))\ln(1-F(s^*|0))} \\
&\times\left[\frac{\ln(1-F(s^*|1))(1-F(s^*|0))(1-F(s^*|1))}{s^*f(s^*|0)}\right]
\end{aligned}$$

The limit of the second (bracketed term) is 0. The limit of the first term above can be computed by L'Hospital rule again:

$$\begin{aligned}
&\lim_{s^*\to 0}\frac{\ln(1-F(s^*|1))}{-(1-F(s^*|1))\ln(1-F(s^*|1))+s^*(1-F(s^*|0))\ln(1-F(s^*|0))} \\
= &\lim_{s^*\to 0}\frac{1}{[\ln(1-F(s^*|1))-\ln(1-F(s^*|0))]+\left(\frac{1-F(s^*|0)}{f(s^*|0)}\right)\frac{\ln(1-F(s^*|0))}{s^*}}\frac{-1}{1-F(s^*|1)} \\
= &\ 1.
\end{aligned}$$

Now consider the second term in the expression for $\lim_{n\to\infty}\Gamma_n^2$ above. Applying L'Hospital rule, and using the normalization $f(s^*|1)=s^*f(s^*|0)$,

$$\begin{aligned}
&\lim_{s^*\to 0}\frac{\left[\ln\left(\frac{\pi}{l(1-\pi)}\right)+\ln\left(\frac{1-F(s^*|0)}{1-F(s^*|1)}\right)\right]\ln\left(1-F\left(s^*|0\right)\right)}{\ln\left(1-F\left(s^*|0\right)\right)-\ln\left(1-F\left(s^*|1\right)\right)}\\
=&\lim_{s^*\to 0}\frac{1}{\frac{-1}{1-F(s^*|0)}-\frac{-s^*}{1-F(s^*|1)}}\left\{\begin{array}{c}\left[\ln\left(\frac{\pi}{l(1-\pi)}\right)+\ln\left(1-F\left(s^*|0\right)\right)-\ln\left(1-F\left(s^*|1\right)\right)\right]\left(\frac{-1}{1-F(s^*|0)}\right)\\+\left[\frac{-1}{1-F(s^*|0)}-\frac{-s^*}{1-F(s^*|1)}\right]\ln\left(1-F\left(s^*|0\right)\right)\end{array}\right\}\\
=&\ln\left(\frac{\pi}{l\left(1-\pi\right)}\right).
\end{aligned}$$

Therefore, we have $\lim_{n\to\infty}\Gamma_n^2=\ln\left(\frac{\pi}{l(1-\pi)}\right)$.

Now observe that by definitions, Γ_n can be expressed as

$$\Gamma_n=(1-\pi)\left(1-F\left(s_n^*|0\right)\right)\left[l\frac{1-F\left(s_n^*|0\right)}{1-F\left(s_n^*|1\right)}\exp\left(\Lambda_n^2\right)-L\exp\left(\Lambda_n^1\right)\right].$$

Taking limit and using results above, we have $\lim_{n\to\infty}\Gamma_n=\pi$.

As a result, together with (26), a lower bound on the limit expected payoff of the search strategy under consideration is π. However, the upper bound of the agent's equilibrium payoff is also π. Therefore, the agent's equilibrium payoff in the limit as $c_n\to 0$ is exactly π.

Moreover, note that as $V_n(\cdot)$, the agent's equilibrium payoff at search cost c_n, is weakly convex, $V_n(\pi)$ is bounded from above by $\frac{\pi-p_0(s_n^*)}{1-p_0(s_n^*)}$. Taking limit as $n\to\infty$,

$$\pi=\lim_{n\to\infty}V_n(\pi)\le\lim_{n\to\infty}\frac{\pi-p_0\left(s_n^*\right)}{1-p_0\left(s_n^*\right)}.$$

Finally it is clear that $\lim_{n\to\infty}\frac{\pi-p_0(s_n^*)}{1-p_0(s_n^*)}\le\pi$ as $p_0\left(s_n^*\right)\in[0,1]$. We therefore have

$$\pi=\lim_{n\to\infty}\frac{\pi-p_0\left(s_n^*\right)}{1-p_0\left(s_n^*\right)}\Leftrightarrow p_0\left(s_n^*\right)=0.$$

(iii) Fix a $\underline{s}>0$ and assume $l<L$. We show that if the sequence $\{p_0\left(s_n^*\right)\}_n$ converges, its limit strictly exceeds 0. Suppose not. Fix an arbitrarily $q\in\left(0,\frac{1}{2}\frac{l}{l+\underline{s}}\right)$, and an integer N such that $p_0\left(s_n^*\right)<q<p_1\left(s_n^*\right)$ for all $n>N$. For $n>N$, the agent's continuation value function evaluated at q, $V_n(q)$, is strictly positive. We first derive an upper bound for $V_n(q)$. In the best conceivable scenario, the agent with $\omega=0$ learns the state immediately and gets a payoff of 0; whereas the agent with $\omega=1$ gets a consecutive sequence of Y recommendations, leading to a posterior $p_1\left(s_n^*\right)$ and take the operation. The number of consecutive Y recommendations needed, denoted by M_n, satisfies

$$\frac{1}{1+\frac{1-q}{q}\left(\frac{1-F(s_n^*|0)}{1-F(s_n^*|1)}\right)^{M_n-1}} \leq p_1\left(s_n^*\right) \leq \frac{1}{1+\frac{1-q}{q}\left(\frac{1-F(s_n^*|0)}{1-F(s_n^*|1)}\right)^{M_n}}$$

$$\Leftrightarrow \quad M_n \in \left[\frac{\ln\left(\frac{q}{1-q}\left(\frac{1}{p_1(s_n^*)}-1\right)\right)}{\ln\left(\frac{1-F(s_n^*|0)}{1-F(s_n^*|1)}\right)}, \frac{\ln\left(\frac{q}{1-q}\left(\frac{1}{p_1(s_n^*)}-1\right)\right)}{\ln\left(\frac{1-F(s_n^*|0)}{1-F(s_n^*|1)}\right)}+1\right].$$

Therefore,

$$V_n\left(q\right) \leq q\left(1-F\left(s_n^*|1\right)\right)^{M_n} - c_n M_n. \tag{27}$$

Below we derive a contradiction by showing that this upper bound cannot be positive for all $q \in (0,1)$. Using (20),

$$M_n \geq \frac{\ln\left(\frac{q}{1-q}\left(\frac{1}{\frac{l}{s_n^*+l}}-1\right)\right)}{\ln\left(\frac{1-F(s_n^*|0)}{1-F(s_n^*|1)}\right)} = \frac{\ln\left(\frac{q}{1-q}\frac{s_n^*}{l}\right)}{\ln\left(\frac{1-F(s_n^*|0)}{1-F(s_n^*|1)}\right)}.$$

Using (21), and (24),

$$\begin{aligned}
c_n \quad &\geq \quad -\frac{1}{\frac{1-F(s_n^*|0)}{1-F(s_n^*|1)}\frac{s_n^*}{l}+1}F\left(s_n^*|1\right)\left(1+L\right) \\
&\quad + \left[\frac{1}{\frac{1-F(s_n^*|0)}{1-F(s_n^*|1)}\frac{s_n^*}{l}+1}F\left(s_n^*|1\right) + \left(1-\frac{1}{\frac{1-F(s_n^*|0)}{1-F(s_n^*|1)}\frac{s_n^*}{l}+1}\right)F\left(s_n^*|0\right)\right]\left[V_n\left(\frac{1}{1+\frac{1-\frac{1}{\frac{s_n^*}{l}+1}}{\frac{1}{\frac{s_n^*}{l}+1}}\frac{F(s_n^*|0)}{F(s_n^*|1)}}\right)+L\right] \\
&\geq \quad \frac{1}{\frac{1-F(s_n^*|0)}{1-F(s_n^*|1)}\frac{s_n^*}{l}+1}\left(-F\left(s_n^*|1\right)\left(1+L\right)+\left[F\left(s_n^*|1\right)+\frac{1-F\left(s_n^*|0\right)}{1-F\left(s_n^*|1\right)}\frac{s_n^*}{l}F\left(s_n^*|0\right)\right]L\right).
\end{aligned}$$

Consequently,

$$c_n M_n \geq \frac{\ln\left(\frac{1-q}{q}\frac{l}{s_n^*}\right)}{\frac{1-F(s_n^*|0)}{1-F(s_n^*|1)}\frac{s_n^*}{l}+1}\frac{-F\left(s_n^*|1\right)+\frac{1-F(s_n^*|0)}{1-F(s_n^*|1)}\frac{s_n^*}{l}F\left(s_n^*|0\right)L}{\ln\left(\frac{1-F(s_n^*|1)}{1-F(s_n^*|0)}\right)}.$$

Taking $\limsup_{n\to\infty}$ on both sides of the inequality above, and using L'Hospital rule,

$$\begin{aligned}
\lim_{n\to\infty}\sup c_n M_n \quad &\geq \quad \frac{\ln\left(\frac{1-q}{q}\frac{l}{\underline{s}}\right)}{\frac{\underline{s}}{l}+1}\lim_{n\to\infty}\frac{-F\left(s_n^*|1\right)+\frac{1-F(s_n^*|0)}{1-F(s_n^*|1)}\frac{s_n^*}{l}F\left(s_n^*|0\right)L}{\ln\left(\frac{1-F(s_n^*|1)}{1-F(s_n^*|0)}\right)} \\
&= \quad \frac{\ln\left(\frac{1-q}{q}\frac{l}{\underline{s}}\right)}{\frac{\underline{s}}{l}+1}\lim_{n\to\infty}\frac{-f\left(s_n^*|1\right)+\left(\begin{array}{c}\frac{(-f(s_n^*|0))(1-F(s_n^*|1))+(1-F(s_n^*|0))f(s_n^*|1)}{(1-F(s_n^*|1))^2}s_n^*F\left(s_n^*|0\right) \\ +\frac{1-F(s_n^*|0)}{1-F(s_n^*|1)}\left(s_n^*f\left(s_n^*|0\right)+F\left(s_n^*|0\right)\right)\end{array}\right)\frac{L}{l}}{\frac{f(s_n^*|0)(1-F(s_n^*|1))-(1-F(s_n^*|0))f(s_n^*|1)}{(1-F(s_n^*|0))(1-F(s_n^*|1))}} \\
&= \quad \frac{1}{\underline{s}+l}\frac{L-l}{1-\underline{s}}\underline{s}\ln\left(\frac{1-q}{q}\frac{l}{\underline{s}}\right).
\end{aligned}$$

Taking $\limsup_{n\to\infty}$ on both sides of (27), and using the fact that $(1 - F(s_n^*|1))^{M_n} \leq 1$,

$$\limsup_{n\to\infty} V_n(q) \leq q - \frac{1}{\underline{s}+l}\frac{L-l}{1-\underline{s}}\underline{s}\ln\left(\frac{1-q}{q}\frac{l}{\underline{s}}\right).$$

Recall $V_n(q) > 0$, so

$$0 \leq q - \ln\left(\frac{1-q}{q}\frac{l}{\underline{s}}\right)\frac{1}{\underline{s}+l}\frac{L-l}{1-\underline{s}}\underline{s}.$$

However, this is a contradiction as q can be an arbitrarily small positive number. Q.E.D.

Proof of Proposition 6: (i) It follows immediately from Lemma 5.

(ii) For each $n \in \mathbb{N}$, denote by $V_n(\cdot)$ the agent's value function in an equilibrium of the game in which the search cost is c_n and the experts' cutoff is $s_n^* \equiv s^*(c_n, \underline{s})$. The agent's equilibrium payoff at c_n, denoted by $U(c_n, \underline{s})$, is given by $V_n(\pi)$. By definition,

$$V_n(p_1(s_n^*)) = -L + (1+L)\,p_1(s_n^*).$$

Moreover, as each $V_n(\cdot)$ is weakly convex, and $V_n(0) = 0$, an upper bound on $V_n(\pi)$ is given by

$$V_n(\pi) \leq \max\left\{\frac{-L + (1+L)\,p_1(s_n^*)}{p_1(s_n^*)}\pi, 0\right\}.$$

By Lemma 5, the right-hand side of the inequality above converges, so we have

$$\limsup_{n\to\infty} U(c_n, \underline{s}) = \limsup_{n\to\infty} V_n(\pi) \leq \max\left\{\left(1 - \frac{L}{l}\underline{s}\right)\pi, 0\right\}.$$

If $1 - \frac{L}{l}\underline{s} \leq 0$, then $\limsup_{n\to\infty} U(c_n, \underline{s}) = 0$ and there is no search in the limit equilibrium. Thus, $\limsup_{n\to\infty} T(c_n, \underline{s}) = 0$.

Suppose $1 - \frac{L}{l}\underline{s} > 0$ so that there is search for n sufficiently large. Denote by E_n the expected payoff of the expert who carries out the operation, in an equilibrium of the game in which the search cost is c_n and the experts' cutoff is s_n^*. The payoff E_n is bounded from as follows:

$$\begin{aligned} E_n &\leq -l + (1+l)\Pr\left(\omega = 1 | p = p_1(s_n^*), s \geq s_n^*\right) \\ &= -l + (1+l)\frac{1}{1 + \frac{1-p_1(s_n^*)}{p_1(s_n^*)}\frac{1-F(s_n^*|0)}{1-F(s_n^*|1)}}, \end{aligned}$$

where $p = p_1(s_n^*)$ stands for the event that the agent enters into the period with a belief $p_1(s_n^*)$. Moreover, using Bayes' rule, starting with a prior belief π, the probability that the agent's posterior reaches $p_1'(s_n^*)$ is at most $\frac{\pi}{p_1'(s_n^*)}$. The total experts' payoff is therefore bounded from above by:

$$\frac{\pi}{p_1'(s_n^*)}E_n \leq \frac{\pi}{p_1'(s_n^*)}\left(-l + (1+l)\frac{1}{1 + \frac{1-p_1(s_n^*)}{p_1(s_n^*)}\frac{1-F(s_n^*|0)}{1-F(s_n^*|1)}}\right).$$

By Lemma 5, the right-hand side of the inequality above converges, so we have

$$\limsup_{n\to\infty} \frac{\pi}{p_1'(s_n^*)} E_n \leq \pi (1 - \underline{s}).$$

Q.E.D.

References

[1] Alger, I. and Salanie, F., 2006. A theory of fraud and overtreatment in experts markets. Journal of Economics and Management Strategy, 15(4), pp.853-881.

[2] Bulow, J. and Klemperer, P., 2002. Prices and the Winner's Curse. RAND journal of Economics, pp.1-21.

[3] Dulleck, U. and Kerschbamer, R., 2006. On doctors, mechanics, and computer specialists: The economics of credence goods. Journal of Economic literature, pp.5-42.

[4] Hiriart-Urruty, J.B. and Lemaréchal, C., 2012. Fundamentals of convex analysis. Springer Science & Business Media.

[5] Hong, H. and Shum, M., 2002. Increasing competition and the winner's curse: Evidence from procurement. The Review of Economic Studies, 69(4), pp.871-898.

[6] Kamenica, E. and Gentzkow, M., 2011. Bayesian persuasion. American Economic Review, 101(6), pp.2590-2615.

[7] Kuhn, H.W., 1953. Extensive games and the problem of information. Contributions to the Theory of Games, 2(28), pp.193-216.

[8] Milgrom, P.R., 1979. A convergence theorem for competitive bidding with differential information. Econometrica, pp.679-688.

[9] Pesendorfer, W. and Swinkels, J.M., 1997. The loser's curse and information aggregation in common value auctions. Econometrica, pp.1247-1281.

[10] Pesendorfer, W. and Wolinsky, A., 2003. Second opinions and price competition: Inefficiency in the market for expert advice. Review of Economic Studies, 70(2), pp.417-437.

[11] Wolinsky, A., 1993. Competition in a market for informed experts' services. RAND Journal of Economics, pp.380-398.

[12] Wolinsky, A., 2005. Procurement via sequential search. Journal of Political Economy, 113(4), pp.785-810.