Should I Stay or Should I Go? Search without Priors

Dirk Bergemann[1] Karl H. Schlag[2]

February 18, 2011

[1]Department of Economics, Yale University, New Haven, CT 06511, USA, dirk.bergemann@yale.edu.

[2]Department of Economics, University of Vienna, 1010 Vienna, Austria, karl.schlag@univie.ac.at.

Abstract

Sequential search without recall is typically accompanied by substantial uncertainty. Classic models reduce this uncertainty to risk by considering a prior over the underlying distributions. We show how to search among a finite number of alternatives without specifying priors. Our objective is to minimize the maximal loss in payoffs as compared to the payoffs attained when the true underlying distributions are known. We find loss can be made small when there are two alternatives if the respective offers are drawn from the same underlying distribution. One needs to randomize appropriately to hedge against uncertainty and not use a reservation price.

Keywords: search, no recall, minimax loss, bounded rationality.

JEL codes: D83, C61, D81

1 Introduction

Uncertainty is a fundamental component of search. There are two classic models of search: (i) costly search with recall (Kohn and Shavell, 1974), such as when searching for a good with the lowest price, and (ii) search without recall (Gal et al., 1981), such as typically assumed in models of job search. No recall means that offers once rejected cannot be accepted later. What is the role of uncertainty? When there is recall and limited number of alternatives one can escape uncertainty by inspecting all alternatives first. When search costs are small as in many applications this behavior is not very costly. In contrast, when there is no recall, as considered in this paper, then one cannot inspect all alternatives to then choose among them and uncertainty weighs stronger.

Classic models do not tell us how to confront uncertainty. The recipe is to reduce uncertainty to risk by considering priors, however there is no recipe on how to practically find this prior. Even when there is prior information, one needs to specify a prior.

Does it really matter which prior one chooses? We will investigate this question. We compare performance of an agent using a subjective prior to the performance she would have had if she would have known the true underlying distributions. Thus we will consider the loss of using an "incorrect" prior and derive the maximal loss. In our analysis we will derive a rule that minimizes this maximal loss. This rule is a natural alternative to subjective expected utility maximization when the decision maker is worried about using an "incorrect" prior. Such a rule will be universal in the sense that it requires only minimal information about the environment. We will compare its behavioral properties to those of the rational decision maker. One property will be that the rule that minimizes the maximum loss is stochastic in an aim to arbitrage against uncertainty. As benchmarks we also derive the deterministic rule that minimizes the maximum loss among all deterministic rules.

For our investigation we consider search without recall as it is here where uncertainty plays the largest role. In a first step we consider two alternatives only. So there will be two offers drawn from the same underlying unknown distribution. One first inspects the first offer and has to decide whether or not to accept this offer before one observes the second offer. Once the first offer is rejected, then due to no recall, one has to choose between the second offer or no offer. Next we analyze the model where the existence of a second offer is unknown at the time when one has to decide on whether or not to accept the first offer. We consider the case where there are three offers. We consider both the case where all offers are drawn from the same distribution as well as the case where the are drawn from separate distributions. In the latter case we contrast the results when offers arrive in a given order to the case in which the agent may choose the order in which offers are inspected.

In terms of methodology, the value of minimax loss can be found by solving a fictitious zero sum game between the decision maker and nature. The decision maker chooses a rule that specifies for each value in the first offer the probability of accepting this offer. Once the first offer is rejected the decision maker accepts the second offer if this lies above the value of the outside option which has a known value. Nature chooses a distribution over offer distributions. We also consider the case where strategy or rule of the decision maker is confined to be deterministic. Here the above saddle point approach is no longer applicable.

Despite the ever growing literature on ambiguity and a concern for robust choices there are only a few papers that analyze how agents choices in markets when they do not have priors. On the supply side, Bergemann and Schlag (2008) and Handel (2009) consider pricing of a monopolist under demand uncertainty. Renou and Schlag (2010) analyze pricing when there is uncertainty about the rationality of the competing firms. On the demand side, Schlag (2004) consider firms repeatedly facing consumers who can only visit at most one firm per round and who are uncertain about market prices. In a more general framework, Riedel (2009) considers search under multiple priors when preferences are defined as in Gilboa and Schmeidler (1989). This approach does not yield sensible results in the setting of this paper (see Conclusion). There is larger related but different literature in behavioral economics in which agents do not actively make a choice, but instead where their behavior is postulated (such as

Gabaix and Laibson, 2004, Radner and Sundararajan, 2005, Spiegler, 2006).

We proceed as follows. In Section 2 we introduce the model using the case where there are only two alternatives. In Section 3 we proceed by considering the case where all values are drawn from the underlying distribution. In Section 4 we turn to the case where each alternative is associated to a separate distribution of values. In Section 5 we conclude.

2 The Model

An agent faces n alternatives. Choice of alternative i yields utility v_i that belongs to $[a, b]$ with $a < b$ and is drawn from the distribution G_i. To make no choice yields utility a. Let V_i be associated random variable. v_i will also be called an offer. $EV_i = \int v_i dG_i(v_i)$ denotes the expected value of V_i. Let $G = (G_1, .., G_n)$ be the joint distribution. There are no search costs.

The agent knows a and b but nothing more about the distributions G_i, $i = 1, .., n$.

Timing is as follows. There are n rounds, where round i is associated to alternative i. Before round i, only after rejecting all previous offers if $i > 1$, the agent observes the offer v_i and has to decide (i) whether to choose alternative i and to receive utility v_i and to end the game, or instead (ii) to reject it and to move on to the next offer contained in alternative $i + 1$. When rejecting offer n the agent receives utility a and the game ends. In particular, there is no recall. Once rejecting alternative i and observing v_{i+1} the agent cannot go back and choose alternative i.

Given the above there is no value in rejecting the nth offer. Hence we will assume from now on that offer n is always accepted. In our further analysis there is no loss to generality from choosing $a = 0$ and $b = 1$.

The rule of the agent is given by $s : \cup_{i=1}^{n} [0, 1]^i \to [0, 1]$ where $s(v_1, .., v_i)$ is the probability of choosing alternative i if alternative i has yields value v_i after rejecting all previous offers $v_1, .., v_{i-1}$. Sometimes we will confine our analysis to deterministic rules $s_d : \cup_{i=1}^{n} [0, 1]^i \to \{0, 1\}$.

Let $\pi(s, G)$ be the expected payoff generated by rule s when facing G. So a rational agent who knows G will choose $s" \in \arg\max_s \pi(s, G)$. When $n = 2$ then this means to accept offer 1 if $v_1 > EV_2$ and to reject alternative 1 and then accept offer 2 if $v_1 < EV_2$. Let $\mathcal{G}$ be a distribution over joint distributions G. Then a rational agent with prior $\mathcal{G}$ will choose $s" \in \arg\max_s \pi(s, \mathcal{G})$ where $\pi(s, \mathcal{G}) = \int \pi(s, G)\, d\mathcal{G}(G)$.

2.1 Concern for Loss

In this paper we consider agents who do not know G and similarly do not have a prior. They only know that the support of any G_i belongs to $[0, 1]$.

Ideally the agent would like to perform best regardless of the specific G. This will not be possible. The only thing we can mention is that it is best when facing any G to to accept an offer equal to 1 and to reject any offer of 0.

We will measure the value or performance of a rule in terms of loss as compared to the case where G is known. Formally,

$$L(s, G) = \max_{s"} \pi(s", G) - \pi(s, G).$$

When $n = 2$ then we obtain

$$\begin{aligned} L(s, G) &= \int_0^1 \max\{v_1, EV_2\}\, dG_1(v_1) \\ &\quad - \int_0^1 \int_0^1 [(s(v_1) v_1 + (1 - s(v_1)) v_2)]\, dG_1(v_1)\, dG_2(v_2). \end{aligned} \tag{1}$$

The first term describes the maximal value of expected utility attainable when G_1 and G_2 are known. The second term describes the expected utility attained when following the rule s. Consequently L measures the foregone utility of not knowing the distributions G_1 and G_2 when choosing s, in particular $L \geq 0$. Equivalently,

$$L(s, G) = \int_{EV_2}^1 ((1 - s(v_1))(v_1 - EV_2))\, dG_1(v_1) + \int_0^{EV_2} s(v_1)(EV_2 - v_1)\, dG_1(v_1).$$

The first term describes the loss in utility due to excess switching while the second the loss in not switching enough.

Our aim is to search for rules that minimize the maximum loss over all possible joint distributions G. So we aim to find $s^* \in \arg\min_s \max_G L(s, G)$. The value of the maximum loss will be defined by $\bar{L}$, so $\bar{L} = \min_s \max_G L(s, G)$. Elements of $\arg\max_G L(s^*, G)$ will be called least favorable distributions. Subscripts d are added when we limit attention to deterministic rules.

We briefly connect the above to subjective expected utility maximization under priors. Let $L(s, \mathcal{G}) = \int L(s, G)\, d\mathcal{G}(G)$. So $L(s, \mathcal{G})$ is the expected loss when compared to the benchmark where the agent learns the joint distribution G realized under $\mathcal{G}$ before choosing s. Following the minimax theorem, we obtain that

$$\bar{L} = \min_s \max_G L(s, G) = \min_s \max_{\mathcal{G}} L(s, \mathcal{G}) = \max_{\mathcal{G}} \min_s L(s, \mathcal{G}) \tag{2}$$

where $\min_s L(s, \mathcal{G})$ is the loss attained by an agent who has prior $\mathcal{G}$. Consequently, $\bar{L}$ is an upper bound on how much the payoff of a decision maker with prior $\mathcal{G}$ can differ from payoff realized when the decision maker knows the true underlying distributions $(G_i)_i$, where these are drawn according to $\mathcal{G}$.

One may interpret s^* as a compromise in the sense that its payoff is never too far from the best payoff under some prior $\mathcal{G}$. We illustrate. As

$$\begin{aligned} \max_{s"} \pi(s", \mathcal{G}) - \pi(s, \mathcal{G}) &\leq \int \max_{s"} \pi(s", G)\, d\mathcal{G}(G) - \pi(s, \mathcal{G}) = \int \left(\max_{s"} \pi(s", G) - \pi(s, \mathcal{G}) \right) d\mathcal{G}(G) \\ &= \int L(s, G)\, d\mathcal{G}(G) \end{aligned}$$

we find that

$$\max_{\mathcal{G}} \left(\max_{s"} \pi(s", \mathcal{G}) - \pi(s, \mathcal{G}) \right) = \max_G \left(\max_{s"} \pi(s", G) - \pi(s, G) \right)$$

and hence

$$s^* \in \arg\min_s \max_{\mathcal{G}} \left(\max_{s"} \pi(s", \mathcal{G}) - \pi(s, \mathcal{G}) \right) \text{ and } \bar{L} = \min_s \max_{\mathcal{G}} \left(\max_{s"} \pi(s", \mathcal{G}) - \pi(s, \mathcal{G}) \right).$$

So any rational agent with prior $\mathcal{G}$ believes that she can obtain at most $\bar{L}$ more than she can when using s^*. On the other hand, for any alternative rule $s' \notin \arg\min_s \max_G L$ there is a prior $\mathcal{G}'$ such that a rational agent with prior $\mathcal{G}'$ believes that she can obtain strictly more than $\bar{L}$ beyond what she believes she will receive when using s'.

We can also connect to the concept of regret as defined Savage (1951) and axiomatized by Milnor (1954). Let

$$\begin{aligned} l(s,G) &= \int_0^1 \int_0^1 \max\{v_1, v_2\}\, dG_1(v_1)\, dG_2(v_2) \\ &\quad - \int_0^1 \int_0^1 \left[(s(v_1)\, v_1 + (1 - s(v_1))\, v_2)\right] dG_1(v_1)\, dG_2(v_2). \end{aligned}$$

So $l(s,G)$ measures how much the payoff achieved by the agent differs from the benchmark in which she knows the two offers before making her decision. Or in other words, the benchmark is the situation in which the agent has perfect recall. Note that $L(s,G) \leq l(s,G)$.

3 Identically Distributed Values

In the following we consider the case where the values of each alternative are drawn from the same distribution. So we assume that $G_1 \equiv ... \equiv G_n$. We will drop indices where appropriate, for instance set $G = G_1$.

3.1 Two Alternatives

Consider now two alternatives. As both values are drawn from the same distribution, the value v_1 of the first alternative can be used to learn about the underlying distribution G. As v_1 is in fact a good estimate of EV, one may wish to always accept alternative 1. In particular, this seems to makes sense in view of a very small cost of accessing the second alternative. This leads to loss

$$L = \int_{EV}^1 (v_1 - EV)\, dG(v_1)$$

which is easily seen to be maximally equal to $\frac{1}{2}$, where the maximum is attained when $G = \frac{1}{2} 1_{v \geq 0} + \frac{1}{2} 1_{v \geq 1}$. If instead one chooses to always reject the first alternative, then this leads to the same loss, in particular maximal loss is again equal to $\frac{1}{2}$. So it does not seem like a good idea to base choice on estimation. An alternative approach to copy the form of the optimal rule, namely to specify a cutoff κ and to choose alternative 1 if and only if $v_1 \geq \kappa$. A salient choice of κ is to set $\kappa = 1/2$. Our result below shows that this is the best deterministic rule in the sense that it minimizes maximum loss among all deterministic rules.

An alternative approach would be to react to the indifference between accepting and rejecting by randomizing. However, simply choosing a constant probability of accepting the first alternative will lead to the same loss as if the first alternative is always accepted. One has to use the information contained in the observed value of alternative 1 by conditioning on v_1. It is natural to let the acceptance probability s be strictly increasing in v_1, always rejecting the lowest offer 0 while always accepting the highest offer 1. The arguable simplest choice of such a strictly increasing function is to set $s(v) = v$ for all $v \in [0,1]$. This turns out to be the best rule for minimizing maximum loss.

Proposition 1 (Two Identical Options) *The rule s^* that is given by $s^*(v_1) = v_1$ minimizes maximum loss, it guarantees loss to be below $\bar{L} = 0.0625$. Restricting attention to deterministic rules only we find that the rule s_d^* given by $s_d^*(v_1) = 1$ if and only if $v_1 \geq 0.5$ minimizes maximum loss, it guarantees loss to be below $\bar{L}_d = 0.125$.*

Note that maximum loss is doubled if one confines attention to deterministic rules. The proof below reveals the following least favorable distributions. For the deterministic rule this is given by $G_0^* = \frac{1}{2}1_{v\geq 1/2} + \frac{1}{2}1_{v\geq 1}$, where minimax loss can also be approximated by considering $G = \frac{1}{2}1_{v\geq 0} + \frac{1}{2}1_{v\geq 1/2-\varepsilon}$ for $\varepsilon > 0$ and taking the limit as $\varepsilon \to 0$. For the stochastic rule this is given by $G_{1/2}^* = \frac{1}{2}1_{v\geq 0} + \frac{1}{2}1_{v\geq 1/2}$ and G_0^*. The proof also reveals that it is hardest for a rational decision maker to learn the correct action when he has a prior $\frac{1}{2}G_0^* + \frac{1}{2}G_{1/2}^*$, in the sense that $\min_s L\left(s, \frac{1}{2}G_0^* + \frac{1}{2}G_{1/2}^*, \frac{1}{2}G_0^* + \frac{1}{2}G_{1/2}^*\right) \geq \min_s L(s, \Gamma, \Gamma)$ for all priors Γ over distributions G.

Proof. Let $\mu = EV$. Then loss is given by

$$
\begin{aligned}
L &= \int_\mu^1 ((1 - s(v_1))(v_1 - \mu))\, dG(v_1) + \int_0^\mu s(v_1)(\mu - v_1)\, dG(v_1) \\
&= \int_0^1 \left[1_{v_1\geq\mu}((1 - s(v_1))(v_1 - \mu)) + 1_{v_1<\mu} s(v_1)(\mu - v_1)\right] dG(v_1).
\end{aligned}
$$

Let $G_{v_1,\mu}$ be the distribution that has mean μ and support on either $\{0, v_1\}$ or $\{v_1, 1\}$. Assume that $s(1) = 1$ and $s(0) = 0$. Then

$$\begin{aligned} L &= \int_0^1 \int_0^1 \left[1_{v\geq\mu}\left(\left(1-s(v)\right)(v-\mu)\right) + 1_{v<\mu} s(v)(\mu - v)\right] dG_{v_1,\mu}(v)\, dG(v_1) \\ &= \int_0^1 L\left(s, G_{v_1,\mu}\right) dG(v_1). \end{aligned}$$

Let $G_{0\beta}$ be such that $C(G_{0\beta}) = \{0, \beta\}$ and $G_{1\beta}$ st. $C(G_{1\beta}) = \{\beta, 1\}$ where $C(G)$ denotes the support of G. Let $\Gamma_2 = \{G : G \in \{G_{0\beta}, G_{1\beta}\} \text{ for some } \beta\}$. Then $\max_G L(s, G) = \max_{G\in\Gamma_2} L(s, G)$.

Consider first deterministic rules with $s(0) = 0$ and $s(1) = 1$. Let $\beta > 0$ be such that $s(\beta) = 0$. Consider $G = \lambda 1_{v\geq 0} + (1-\lambda) 1_{v\geq\beta}$. Then

$$L = \left(1 - \lambda^2\right)\beta - (1-\lambda)\beta = \beta\lambda(1-\lambda) \leq \frac{1}{4}\beta.$$

Now consider instead some $\alpha < 1$ be such that $s(\alpha) = 1$. Given $G = (1-\lambda) 1_{v\geq\alpha} + \lambda 1_{v\geq 1}$ we obtain

$$L = [\lambda + (1-\lambda)(\lambda + (1-\lambda)\alpha)] - [\lambda + (1-\lambda)\alpha] = \lambda(1-\lambda)(1-\alpha) \leq \frac{1}{4}(1-\alpha).$$

As for $v = 1/2$ we have either $s(1/2) = 0$ or $s(1/2) = 1$ it follows from the above that $\bar{L} \geq 1/8$.

Evaluating L when $s = s_d^*$ where $s_d^*(v) = 1$ if and only if $v \geq 1/2$ we obtain that $\bar{L} \leq \frac{1}{8}$.

Thus we find that $\min_{s \text{ determ}} \max_{G\in\Gamma_2} L(s, G) = \max_{G\in\Gamma_2} L(s_d^*, G) = \frac{1}{8}$. Given our previous result we obtain that s_d^* minimizes the maximum loss among all deterministic rules.

Consider now stochastic rules.

To solve $\min_s \max_{G\in\Gamma_2} L(s, G)$ we will fix β, find $G_{0\beta}^*$ and $G_{1\beta}^*$ such that $\left(s_\beta^*, \mu G_{0\beta}^* + (1-\mu) G_{1\beta}^*\right)$ is a saddle point. Let λ_i be such that $1 - \lambda_i = P_{G_{i\beta}^*}(V = \beta)$.

We will set μ such that the decision maker does not update its prior when observing β. This means that $\mu(1-\lambda_0) = (1-\mu)(1-\lambda_1)$.

In order for $s_\beta^*(\beta) \in (0, 1)$ the decision maker has to be indifferent between accepting and rejecting the offer β. Hence

$$\beta = E(\mu G_0 + (1-\mu) G_1) = \mu(1-\lambda_0)\beta + (1-\mu)\left((1-\lambda_1)\beta + \lambda_1\right).$$

Together this means that $\beta = 2(1-\mu)(1-\lambda_1)\beta + (1-\mu)\lambda_1$.

Assume that $s_\beta^*(\beta) = \gamma$. Then

$$\begin{aligned} L\left(s_\beta^*, G_0, G_0\right) &= (1-\lambda_0)(1-\gamma)(\beta - (1-\lambda_0)\beta) = (1-\lambda_0)(1-\gamma)\lambda_0\beta \\ L\left(s_\beta^*, G_1, G_1\right) &= (1-\lambda_1)\gamma((1-\lambda_1)\beta + \lambda_1 - \beta) = (1-\lambda_1)\gamma\lambda_1(1-\beta) \end{aligned}$$

$L\left(s_\beta^*, G_i, G_i\right)$ being maximal implies $\lambda_0 = \lambda_1 = 1/2$ which means that $\mu = \beta = 1/2$. Indifference between $L\left(s_\beta^*, G_0, G_0\right)$ and $L\left(s_\beta^*, G_1, G_1\right)$ implies that $\frac{1}{4}(1-\gamma)\beta = \frac{1}{4}\gamma(1-\beta)$ and hence $\gamma = \beta$. It is easily verified this in fact constitutes a saddle point.

■

3.2 Uncertainty About Whether There Will Be a Second Alternative

Consider now the situation in which the agent does not know whether he will have a second alternative at the point of time when he has to decide whether or not to accept the first alternative. There is a probability δ whose value is unknown to the agent that determines the probability of receiving a second alternative after rejecting the first alternative. Expected utility of rejecting the first offer is now δEV_2. As in the previous section we will assume that both values are drawn from the same distribution, so $G_1 \equiv G_2$. Note that the set of possible rules for the agent remains unchanged.

Loss, which also now depends on δ, is given by

$$L(s, G, \delta) = \int_{\delta EV_2}^{1} ((1 - s(v_1))(v_1 - \delta EV_2))\, dG_1(v_1) + \int_0^{\delta EV_2} s(v_1)(\delta EV_2 - v_1)\, dG_1(v_1).$$

Proposition 2 (Uncertainty about Second Identical Option) *The rule s^{**} that is given by $s^{**}(v_1) = 4\frac{v_1}{1+3v_1}$ minimizes maximum loss, it guarantees loss to be below $\bar{L} = 1/9 = 0.\bar{1}$. The rule s_d^{**} that is given by $s_d^{**}(v_1) = 1$ if and only if $v_1 \geq 1/5$ minimizes maximum loss among all deterministic rules, it guarantees loss to be below $\bar{L}_d = 1/5$.*

The least favorable distribution for the deterministic case is given by $G = \frac{1}{2}1_{v \geq 1/5} + \frac{1}{2}1_{v \geq 1}$ and $\delta = 1$. Maximum loss $\bar{L}_d$ is also approximated by $G = 1_{v \geq 1/5 - \varepsilon}$ and

$\delta = 0$ in the limit as $\varepsilon \to 0$. Least favorable distributions for stochastic rules are $G = \frac{1}{2}1_{v \geq 1/3} + \frac{1}{2}1_{v \geq 1}$ with $\delta = 1$ and $G = 1_{v \geq 1/3}$ with $\delta = 0$.

Note how "desperate" the agent gets if uncertain about whether there will be a second alternative. In the deterministic case the agent accepts any offer above 1/5 when uncertain about whether there will be future alternatives while only accepting offers above 1/2 if there is a second alternative. In the stochastic case, the acceptance probability rises from v_1 to $4\frac{v_1}{1+3v_1}$, this increase largest when $v_1 = 1/3$ where the increase is equal to 1/3.

Proof. Following the arguments made in the proof of Proposition 1 loss will be maximized for some $G \in \Gamma_2$.

Consider first deterministic rules. Let $\alpha < 1$ be such that $s(\alpha) = 1$. Given $G_{1\alpha} = (1-\lambda) 1_{v \geq \alpha} + \lambda 1_{v \geq 1}$ and $\delta = 1$ we obtain

$$L(s, G_{1\alpha}, G_{1\alpha}, 1) = \left(\lambda(1+1-\lambda) + (1-\lambda)^2 \alpha\right) - (\lambda + (1-\lambda)\alpha) = \lambda(1-\lambda)(1-\alpha) \leq \frac{1}{4}(1-\alpha).$$

Let $\beta > 0$ be such that $s(\beta) = 0$. Then $L(s, 1_{v \geq \beta}, 1_{v \geq \beta}, 0) = \beta$. Note that $v = \frac{1}{4}(1-v)$ holds when $v = 1/5$. As $v = 1/5$ must be contained in one of the two above cases, we obtain $\bar{L} \geq \frac{1}{5}$. Evaluating the different cases for L when $s = s_d^{**}$ shows that $\bar{L} \leq 1/5$ which completes the proof for deterministic rules.

Consider now stochastic rules. Fix $\beta \in (0,1)$. Let $s(\beta) = \gamma$. We obtain

$$\begin{aligned} L(s, G_{1\beta}, G_{1\beta}, 1) &= \left(\lambda + (1-\lambda)\lambda + (1-\lambda)^2 \beta\right) \\ &\quad - (\lambda + (1-\lambda)\gamma\beta + (1-\lambda)(1-\gamma)(\lambda + (1-\lambda)\beta)) \\ &= \lambda\gamma(1-\lambda)(1-\beta) \leq \frac{1}{4}\gamma(1-\beta) \end{aligned}$$

and

$$L(s, 1_{v \geq \beta}, 1_{v \geq \beta}, 0) = \beta - \gamma\beta = \beta(1-\gamma)$$

where $\frac{1}{4}\gamma(1-\beta) = \beta(1-\gamma)$ holds when $\gamma = 4\frac{\beta}{1+3\beta}$. Hence, $\bar{L} \geq \frac{1}{4}4\frac{\beta}{1+3\beta}(1-\beta) = \beta\frac{1-\beta}{1+3\beta}$. Moreover, it is easy to verify that

$$L(s, G, G, \delta) \leq \max\{L(s, G_{1\beta}, G_{1\beta}, 1), L(s, 1_{v \geq \beta}, 1_{v \geq \beta}, 0)\} \; \forall G \in \Gamma_2.$$

So if $\gamma = 4\frac{\beta}{1+3\beta}$ we find that $\max_{G \in \Gamma_2, \delta \in [0,1]} L(s, G, G, \delta) = \beta\frac{1-\beta}{1+3\beta}$. As $\max_{\beta \in [0,1]} \beta\frac{1-\beta}{1+3\beta} = 1/9$ the proof is complete. ■

3.3 Costly Search

In the following we assume that there is cost $c > 0$ for receiving a second offer. Then loss is given by

$$L = \int_{\mu-c}^{1} \left(\left(1 - s\left(v_1\right)\right)\left(v_1 - \left(\mu - c\right)\right)\right) dG\left(v_1\right) + \int_{0}^{\mu-c} s\left(v_1\right)\left(\left(\mu - c\right) - v_1\right) dG\left(v_1\right).$$

Proposition 3 *The deterministic rule that minimizes maximum loss among all deterministic rules for $c < 1/2$ is the same one that does this for $c = 0$. Maximal loss $\bar{L}_d(c)$ is equal to $\frac{1}{2}\left(\frac{1}{2} - c\right)^2$ with $\frac{d}{dc}\bar{L}(0) = -1/2$. The stochastic rule that minimizes maximum loss for $c < 0.11$ satisfies $s_c(v) = (1 - v)\frac{(c+v)^2}{v - v^2 + c^2}$ for $c \le v \le 1 - c$, $s(v) = 0$ for $v < c$ and $s(v) = 1$ for $v \ge 1 - c$ where $\frac{d}{dc}s_c(v) > 0$ with $\frac{d}{dc}s_c(v)\,|_{c=0} = 2$. Maximum loss $\bar{L}(c)$ is increasing in c with value 0.06653725 when $c = 0.11$ and $\frac{d}{dc}\bar{L}(0) = 0$.*

Notice the different reaction to cost. The deterministic rule remains unchanged and maximal loss decreases. The stochastic rule increases its probability of staying and maximum loss is increased slightly.

3.4 Three Alternatives

[Material on deterministic rules for three alternatives to be added here.]

3.5 Many Alternatives

[??looks like lower bound for $n = 3$ is 0.075, for $n = 4$ it is 0.065 625, not clear for larger n.]

4 Heterogeneous Offers

In the following we consider the setting where offers are heterogeneous in the sense that each offer is drawn from its own distribution.

4.1 Fixed Order

There are n offers. Assume that the offers are received in a fixed order. Here the analysis is particularly simple. Assume that the first offer has value v_1. When accepting this offer, loss is maximally $1 - v_1$, attained when $G_2 = 1_{v \geq 1}$. When rejecting the first offer, loss is maximally v_1, attained when $G_2 \equiv ... \equiv G_n = 1_{v \geq 0}$. When constrained to deterministic rules one hence accepts the first offer if $1 - v_1 < v_1$, so if $v_1 > 1/2$, rejects if $v_1 < 1/2$. Loss is maximally $1/2$, attained when $v_1 = 1/2$. With stochastic rules one can lower loss by randomizing and hence in some sense hedging against uncertainty. One equalizes maximal loss when facing these two distributions, here by accepting the first offer with probability equal to v_1. In both cases maximal expected loss is then equal to $v_1 (1 - v_1)$ which is maximally equal to $\frac{1}{4}$ when $v_1 = 1/2$. This leads to the following result.

Proposition 4 (Non-Identical Options in Fixed Order) *The rule s^{*n} that is given by $s^{*n}(v_1) = v_1$ minimizes maximum loss, it guarantees loss to be below $\bar{L} = \frac{1}{4}$. The rule s_d^{*n} that is given by $s_d^{*n}(v_1) = 1$ if and only if $v_1 \geq 1/2$ minimizes maximum loss among all deterministic rules, it guarantees loss to be below $\bar{L}_d = 1/2$.*

Regardless of the number of alternatives maximum loss is non negligible. The reason is that nature as adversary has too much power. There are two least favorable distributions G^0 and G^1 where $G_1^0 \equiv G_1^1 = 1_{v \geq 1/2}$, $G_2^0 = 1_{v \geq 0}$ and $G_2^1 = 1_{v \geq 1}$. In the worst case, all uncertainty is contained in the second alternative and observing the value of the first alternative has no added value. The situation changes when the agent is allowed to choose which alternative to sample first.

4.2 Endogenous Order

In the following we continue to assume that each value drawn from its own distribution, but now allow the agent to choose the order in which the samples are chosen. Strategies are now very complicated. They determine which alternative to choose first and as a function of the values observed in the past which alternative to choose next. As we are only concerned with maximal loss and the problem is symmetric from an ex-ante standpoint it will be best (in terms of minimizing maximum loss) to consider rules of the following form. At the beginning, the alternatives are placed in a random order, each order being equally likely. Choices of whether or not to accept an offer do not depend on this order. In the following we will only consider such rules. Hence rules are defined as before. We also provide an alternative interpretation. It is as if offers arrive in a given order, before each offer arrives the distribution that determines its value is chosen without replacement from a set of distributions. The benchmark for determining loss is an agent who knows which distribution has been assigned to which offer. So for instance, when there are two alternative only, then loss is given by $\frac{1}{2}L(s,G)+\frac{1}{2}L(s,G_2,G_1)$ where L is defined as in (1).

In the following when we refer to a deterministic rule then we mean a rule that either accepts or rejects an offer.

Note that there is no value to choosing the order of visits when values are drawn from the same distribution. In fact, one may think that each order of visiting alternatives equally likely reduces the problem to one where all values are chosen from the same distribution. This is however not the case. When all are drawn from the same distribution then the same or similar values will be realized with positive probability by different alternatives.

4.2.1 Two Alternatives

Assume that there are two alternatives.

Assume that one is facing either G^0 or G^1 given above except that one no longer knows which alternative is the deterministic one. Remember that we assume that each alternative is chosen to be first equally likely. The best deterministic rule for this situation would be to accept values ≥ 0.5, hence to accept offers $1/2$ and 1 and to reject 0. This leads to expected loss 0.25. This is in fact our solution. For non deterministic rule will accept offer 0.5 with probability $1/2$ to guarantee loss below 0.125.

Proposition 5 (Two Non-Identical Options in Endogenous Order) *The rule*

s^{*n} that is given by $s^{*n}(v_1) = v_1$ minimizes maximum loss, it guarantees loss to be below $\bar{L} = 0.125$. The rule s_d^{*n} that is given by $s_d^{*n}(v_1) = 1$ if and only if $v_1 \geq 0.5$ minimizes maximum loss among all deterministic rules, it guarantees loss to be below $\bar{L}_d = 0.25$.

While loss can be guaranteed to be below 0.0625 when values are drawn from the same distribution (see Proposition 1), loss can only be guaranteed to be below 0.125 when offers are heterogenous in the sense that their values are drawn from different distributions.

Proof. We first show that loss is maximized when G_1 and G_2 are degenerate. This is because $L = \frac{1}{2}L(s, G) + \frac{1}{2}L(s, G_2, G_1)$ where

$$\begin{aligned} L(s, G) &\leq \int_0^1 \int_0^1 \max\{v_1, v_2\} \, dG_1(v_1) \, dG_2(v_2) \\ &\quad - \int_0^1 \int_0^1 [(s(v_1) v_1 + (1 - s(v_1)) v_2)] \, dG_1(v_1) \, dG_2(v_2) \\ &= l(s, G) \end{aligned}$$

with equality holding whenever $G_1 = 1_{v \geq v_1}$ and $G_2 = 1_{v \geq v_2}$ for some v_1 and v_2.

Assume that $G_1 = 1_{v \geq v_1}$ and $G_2 = 1_{v \geq v_2}$ for some v_1 and v_2 with $v_1 \leq v_2$. Then

$$L = v_2 - \left(\frac{1}{2} (s(v_1) v_1 + (1 - s(v_1)) v_2) + \frac{1}{2} (s(v_2) v_2 + (1 - s(v_2)) v_1) \right).$$

Consider stochastic rules. Note that we may assume that $s(0) = 0$ and $s(1) = 1$. So if $v_1 = v$ and $v_2 = 1$ then

$$L = 1 - \left(\frac{1}{2} (s(v) v + 1 - s(v)) + \frac{1}{2} \right) = \frac{1}{2} s(v) (1 - v).$$

If $v_1 = 0$ and $v_2 = v$ then

$$L = v - \left(\frac{1}{2} v + \frac{1}{2} s(v) v \right) = \frac{1}{2} v (1 - s(v)).$$

It follows that $L \geq \frac{1}{2} v (1 - v)$ with equality attained when $s(v) = v$. Hence $\bar{L} \geq \frac{1}{8}$.

Consider now $s(v) = v$. Given $v_2 \geq v_1$ we obtain

$$L = v_2 - \left(\frac{1}{2} \left(v_1^2 + (1 - v_1) v_2 \right) + \frac{1}{2} \left(v_2^2 + (1 - v_2) v_1 \right) \right) = \frac{1}{2} (v_2 - v_1) (1 - (v_2 - v_1)) \leq \frac{1}{8}.$$

Now consider a deterministic rule with $s(v) = 1$. The above shows that $L \geq \frac{1}{2}(1 - v)$. If instead, $s(v) = 0$ then $L \geq \frac{1}{2} v$. Hence $\bar{L}_d \geq \frac{1}{4}$. It is now easily verified that s_d^{*n} ensures that $L \leq 1/4$. ■

4.2.2 Three Alternatives

A rule s is now a mapping from $[0,1] \cup [0,1]^2 \to [0,1]$ where $s(v_1)$ is the probability of accepting the first offer v_1 and $s(v_1, v_2)$ is the probability of accepting the second offer v_2 after rejecting a first offer equal to v_1.

Proposition 6 *For $n = 3$, within the set of deterministic rules, the rule s^{*w} defined by $s^{*w}(v_1) = 1$ if and only if $v_1 \geq 2/3$ and $s^{*w}(\cdot, v_2) = 1$ if and only if $v_2 \geq 1/3$ minimizes maximum loss, it guarantees loss to be below $2/9 = 0.\bar{2}$.*

Note that maximum loss can be minimized without conditioning in round 2 on what value was observed in round 1.

Proof. Consider any deterministic rule. Let κ_1 be such that κ_1 is not accepted in the first round. Then loss is at least $\frac{1}{3}\kappa_1$, induced when one value is κ_1 and the other two are equal to 0. Let κ_2 be such that κ_2 is not accepted in round. Then loss is at least $\frac{2}{3}(1-\kappa_2)$, induced when values are κ_2, κ_2 and 1.

We can assume that $\kappa_1 \geq \kappa_2$. Hence maximum loss is at least $\max\left\{\frac{1}{3}\kappa_2, \frac{2}{3}(1-\kappa_2)\right\}$ which is bounded below by $2/9$.

It is easily verified that maximum of the above rule is $2/9$. Worst cases are where one value is 1 and the other two are both $\frac{2}{3}$ or both $\frac{1}{3}$. Worst case is approximated for $\varepsilon \to 0$ when two values are 0 and the third is either $\frac{1}{3} - \varepsilon$ or $\frac{2}{3} - \varepsilon$. This completes the proof. ■

4.2.3 Time Consistency

In the following we consider time consistency of choices. The benchmark for comparison used above is the situation in which the agent knew the distributions at the beginning. One might consider the following alternative setting in which at each choice the benchmark is the case where the agent knows the distribution. For instance, consider degenerate distributions where each alternative is associated to a unique value. Assume that values arrive in order v_1, v_2, v_3 with $v_1 > v_2 > v_3$. An agent knowing these values will accept the first offer, if by mistake rejecting the first offer she will accept the second offer. The benchmark for comparison in round 1 is v_1, the benchmark for comparison in round 2 is v_2. Previously we have assumed that the benchmark for round 2 was v_1. The loss function for round 1, when $v_1 = \max\{v_2, v_3\}$, is given by

$$
\begin{aligned}
L &= \frac{1}{3}(1 - s(v_1))\frac{1}{2}(s(v_1, v_2)(v_1 - v_2) + (1 - s(v_1, v_2))(v_1 - v_3)) \\
&+ \frac{1}{3}(1 - s(v_1))\frac{1}{2}(s(v_1, v_3)(v_1 - v_3) + (1 - s(v_1, v_3))(v_1 - v_2)) \\
&+ \frac{1}{3}\left(s(v_2)(v_1 - v_2) + (1 - s(v_2))\left(\frac{1}{2}s(v_2, v_3)(v_1 - v_3) + \frac{1}{2}(1 - s(v_2, v_1))(v_1 - v_3)\right)\right) \\
&+ \frac{1}{3}\left(s(v_3)(v_1 - v_3) + (1 - s(v_3))\left(\frac{1}{2}s(v_3, v_2)(v_1 - v_2) + \frac{1}{2}(1 - s(v_3, v_1))(v_1 - v_2)\right)\right).
\end{aligned}
$$

This is the same loss function we have used above.

The loss function for round 2, using the best outcome for an agent knowing the distribution in round 2 as benchmark, anticipating behavior in round 1, for is $v_1 \geq v_2 \geq v_3$ is then given by

$$
\begin{aligned}
L_2 &= \frac{1}{3}(1 - s(v_1))\frac{1}{2}(1 - s(v_1, v_2))(v_2 - v_3) \\
&+ \frac{1}{3}(1 - s(v_1))\frac{1}{2}s(v_1, v_3)(v_2 - v_3) \\
&+ \frac{1}{3}(1 - s(v_2))\left(\frac{1}{2}s(v_2, v_3)(v_1 - v_3) + \frac{1}{2}(1 - s(v_2, v_1))(v_1 - v_3)\right) \\
&+ \frac{1}{3}(1 - s(v_3))\left(\frac{1}{2}s(v_3, v_2)(v_1 - v_2) + \frac{1}{2}(1 - s(v_3, v_1))(v_1 - v_2)\right).
\end{aligned}
$$

Consider rules in round 2 that do not depend on the value in round 1. Then we find that the best choice is to set $s(\cdot, v_2) = 1$ if and only if $v_2 \geq 1/2$, the same choice as if there were only a single offer. For round 1 we find the best choice to be $s(v_1) = 1$ if and only if $v_1 \geq \kappa$ for some $\kappa \geq 1/2$. Regret is maximally equal to $1/3$. Hence we find that once the agent changes the benchmark in round 2, regret from the perspective of round 1 increases from $2/9$ (see Proposition 6) to $1/3$. The worst case is when $v_1 = 1/2 - \varepsilon$, $v_2 = v_3 = 0$.

4.2.4 Many Alternatives

Learning is very difficult when alternatives are chosen from different distributions. The reason is alternatives passed by need not reappear. In the worst case, each alternative has a deterministic value with a unique highest offer where due to the design this will be only offered a single time. Yet this best alternative is known when deriving the benchmark. Consequently, we do not expect that maximal loss converges zero as the number of alternatives increases. To show formally how difficult learning is we will restrict attention to subclasses of deterministic rules.

Proposition 7 *Consider a deterministic rule where there is strictly decreasing sequence of cutoffs* $(\kappa_i)_{i=1}^n$ *such that offer* v_i *in round* i *is accepted if and only if* $v_i > \kappa_i$. *Then maximal loss is bounded below by* $\frac{1}{4}\frac{n-2}{n-1}$ *when* n *is even and by* $\frac{n-1}{4n}$ *when* n *is odd.*

Proof. Assume that n is even. Assume that all values are 0 except one that is equal to $\kappa_{n/2}$. Then expected loss is equal to $\frac{n/2-1}{n}\kappa_{n/2}$. Now assume that all values are equal to $\kappa_{n/2}$ except one that is equal to 1. Then expected loss is equal to $\frac{1}{2}\left(1-\kappa_{n/2}\right)$. Equating these two expression we obtain the bound $\frac{1}{4}\frac{n-2}{n-1}$.

Assume that n is odd and consider $\kappa_{(n+1)/2}$. Then we obtain analogously $\frac{(n-1)/2}{n}\kappa_{(n+1)/2}$ and $\frac{(n-1)/2}{n}\left(1-\kappa_{n/2}\right)$ and hence maximum loss is bounded below by $\frac{n-1}{4n}$. ■

In the following we derive a lower bound on the maximal loss of any rule. To do this we use the minimax theorem. We present a prior over distributions under which the loss of a Bayesian is not small. By (2) this is a lower bound on the loss of any rule. More explicity, we choose a particular prior $\mathcal{G}$, derive $\min_s L(s, \mathcal{G})$ and then use the fact that $\bar{L} \geq \min_s L(s, \mathcal{G})$ for all $\mathcal{G}$. Thie prior $\mathcal{G}$ puts weights on two distributions G_1 and G_2. G_1 yields 0 in $n-1$ offers and once 1/2. When facing G_1 it is best to accept offer 1/2 the first time it is experienced. G_2 yields 0 in $n-2$ offers and 1/2 and 1 in the remaining two offers. When facing G_2 it is best to wait for offer 1. When endowed with a prior over these two distributions under which G_1 and G_2 are equally likely, it is best to accept the first offer that is not 0. An agent who knows whether she is facing G_1 or G_2 can reject 1/2 when facing G_2 and wait for offer 1. Loss is hence equal to $\frac{1}{2}1+\frac{1}{2}\frac{1}{2}-\left(\frac{1}{2}\frac{1}{2}+\frac{1}{2}\frac{3}{4}\right)=\frac{1}{8}$. The above proves the following.

Proposition 8 *Maximal loss of any rule for any* $n \geq 3$ *is bounded below by* 0.125.

5 Conclusion

In this preliminary research paper we have shown that it is feasible to derive insights and conclusions about search without adding priors. We find our measure of loss to be particular appealing. For a Bayesian it is a measure of how difficult learning is. For a decision maker it shows how to be closest to any Bayesian it may be interacting with. While we are open in principle to alternative measures of loss we have not come across any that has similar properties. For instance, our loss function has the property that one focusses on environments in which mistakes of not correctly anticipating the environment are costly. A related but different measure is that of regret, due to Savage (1951) and axiomatized by Milnor (1954). Most of our present analysis also holds for this measure. The benchmark is not the outcome of an agent who knows the distribution but the outcome of an agent who knows the values of each alternative. We find that this benchmark is only very hypothetical in the framework of search as search is intricatly linked to randomness. So we chose to define loss that essentially questions how to model this randomness, but does not ignore the randomness. An alternative measure of loss often suggested in the literature is that of negative utility, which leads to the decision making criterion of maximin utility, as suggested by Wald (1950) and also axiomatized by Milnor (1954). These are also the preferences considered by Riedel (2009) when there are no restrictions on the set of priors. Reminiscent of many other applications (e.g. see Bergemann and Schlag, 2008), its results are trivial and hence lack insights. The worst outcome is that all alternatives yield value 0 in which case it does not matter what the decision maker does. In other words, everything is optimal under the maximin utility criterion.

References

[1] D. Bergemann, K. H Schlag (2008), Pricing without priors, J. Eur. Econ. Ass. 6, 560-569.

[2] X. Gabaix and D. Laibson (2004), Competition and Consumer Confusion, Mimeo.

[3] S. Gal, M. Landsberger, and B. Levykson (1981), A Compound Strategy for Search in the Labor Market, Int. Econ. Rev., Vol. 22, No. 3, pp. 597-608.

[4] I. Gilboa, D. Schmeidler (1989), Maxmin expected utility with non-unique priors, J. Math. Econ. 18, 133-173.

[5] Handel, B. (2009), Market Power and Market Ignorance, Mimeo, Northwestern University.

[6] M. Kohn and S. Shavell (1974). The theory of search. Journal of Economic Theory, 9:93–123.

[7] K. Misra, B.R. Handel, and J.R. Roberts (2009), Robust Firm Pricing with Panel Data, ERID Working Paper Number 81.

[8] J. Milnor (1954), Games against Nature, in: R. Thrall, C. Coombs and R. Davis (Eds), Decision Processes. John Wiley, London, pp. 49-60.

[9] R. Radner and A. Sundararajan (2005), Dynamic Pricing of Network Goods with Boundedly Rational Consumers, Working Paper CeDER-05-08, Center for Digital Economy Research.

[10] L. Renou and K.H. Schlag (2010), Minimax Regret and Strategic Uncertainty, J. Econ. Theory 145(1), 264-286.

[11] F. Riedel (2009), Optimal Stopping with Multiple Priors, Econometrica, Vol. 77, No. 3, 857–908

[12] L. J. Savage (1951), The theory of statistical decision, J. Amer. Statist. Ass. 46, 55-67.

[13] Schlag, K.H. (2004), Competing for Boundedly Rational Consumers, Mimeo.

[14] R. Spiegler (2006) Competition over Agents with Boundedly Rational Expectations, Theoretical Economics, 1(2): 207-231.

[15] Wald, A. (1950), *Statistical decision functions*, New York: John Wiley & Sons.