

Robust sequential search

KARL H. SCHLAG

Department of Economics, University of Vienna

ANDRIY ZAPECHELNYUK

School of Economics and Finance, University of St Andrews

We study sequential search without priors. Our interest lies in decision rules that are close to being optimal under each prior and after each history. We call these rules robust. The search literature employs optimal rules based on cutoff strategies, and these rules are not robust. We derive robust rules and show that their performance exceeds $1/2$ of the optimum against binary independent and identically distributed (i.i.d.) environments and $1/4$ of the optimum against all i.i.d. environments. This performance improves substantially with the outside option value; for instance, it exceeds $2/3$ of the optimum if the outside option exceeds $1/6$ of the highest possible alternative.

KEYWORDS. Sequential search, search without priors, robustness, dynamic consistency, competitive ratio.

JEL CLASSIFICATION. C44, D81, D83.

1. INTRODUCTION

Suppose that you check stores one by one in search of the cheapest place to buy some good. Your decision of when to stop searching depends on the distribution of prices you expect to encounter in unvisited stores. The methodology of Bayesian decision making proposes to turn this into an optimization problem. The input is your prior belief about possible environments, mathematically formulated as a distribution over distributions. This is a complex and usually intractable intertemporal decision problem. Special cases can be solvable, but solutions are fragile as they depend on your beliefs about what you do not know (see [Gastwirth \(1976\)](#)).

We are interested in a robust approach to this problem that does not depend on specific prior beliefs of a decision maker. So instead of focusing on the optimal performance for a specific prior, we aim to perform relatively close to the optimum in each environment, and, hence, under each prior over environments. Furthermore, we are interested in maintaining this property over time, not just at the outset. We formalize a performance criterion that fulfills these desiderata. Decision rules that are optimal under this criterion are called *robust*. We present robust rules and show how well they perform.

Karl H. Schlag: karl.schlag@univie.ac.at

Andriy Zapechelnuk: az48@st-andrews.ac.uk

The authors would like to thank Dirk Bergemann, Jeffrey Ely, Olivier Gossner, Johannes Hörner, Bernhard Kasberger, Alexei Parakhonyak, and anonymous referees for helpful comments and suggestions.

The practical relevance of robust decision making is apparent. How can a shopper know the distribution of prices offered in the next store? How does she form a prior about such distributions? Even if a prior is formed, will the shopper be able to overcome the complexity of Bayesian optimization? Will the decision rule still be good if the prior puts little or no weight on the environment that is realized? How will the shopper argue about the optimality of a particular decision rule in front of her peers if they do not have the same prior as she does? These questions can be addressed by a decision rule that performs relatively well for any prior. Such a rule can be proposed as a compromise among Bayesian decision makers who have different priors. It is a shortcut to avoid the cumbersome calculations involved when computing the Bayesian optimal rule. Moreover, as a single rule that does not depend on individual (unobservable) beliefs, it is a useful benchmark for empirical studies.

The setting we consider in this paper is as follows. Alternatives arrive according to some independent and identically distributed (i.i.d.) process. We refer to this process as an environment. An individual knows what alternatives can arrive, but does not know the environment she faces. She has to decide after each draw whether to stop the search or to draw another alternative. There is free recall: when the individual stops, she chooses the best alternative found so far. Values are discounted over time; thus, waiting for better alternatives is costly.

We measure the performance of a given decision rule as follows. For each environment and each history, we compute the ratio of the rule's payoff to the maximal possible payoff. We then find the smallest ratio among all environments and all histories. We call this the *performance ratio* of the rule. The performance ratio describes what fraction of the maximal payoff can be guaranteed, regardless of the environment and regardless of which alternatives have realized over time. A rule that achieves the largest possible performance ratio is called *robust*. It is as if we are looking for an epsilon optimal rule with the smallest possible epsilon. To choose a robust decision rule is our recommendation to an individual who does not know the environment and wishes to avoid forming a prior and optimizing against this prior.

In this paper, we first consider binary environments and then general environments. An environment is called *binary* if it can generate at most one alternative whose value is above the outside option. This alternative is called the high alternative. Consider an individual who knows she is facing a binary environment. So she stops searching once she sees the high alternative. The only question is when to stop if the high alternative has not yet arrived. We find a robust decision rule for such environments. The corresponding performance ratio is larger than $1/2$, so the individual can always guarantee at least half of the maximal payoff. Moreover, if there is an upper bound on the possible values of the high alternative, then the robust performance ratio is strictly increasing in the outside option, attaining $2/3$ and $3/4$ when the outside option is $1/6$ and $1/3$ of that upper bound, respectively.

Next, we consider general environments. Here we allow for any i.i.d. distribution over a given set of alternatives. We show that the robust performance ratio is always at least $1/4$. Surprisingly, this ratio is the same in general environments as it is in binary environments, provided that there is an upper bound on possible alternatives and the

outside option is not too small. The decision rule that supports these findings prescribes to stop after any given history with a probability that is increasing in the value of the best realized alternative.

An important feature of robust rules is that they randomize whenever it is worth waiting for a higher alternative. This stands in contrast to Bayesian rules that optimize against a given prior. These rule are generically deterministic. We show that no deterministic rule can perform better than the rule that does not search at all.

Related literature

A popular criterion for decision making under multiple priors is maximin utility (Wald (1950), Gilboa and Schmeidler (1989)). Unlike our approach, under this criterion there is no concern for being close to the optimum irrespective of the prior. Maximin utility aims to do best for a specific prior where payoffs are lowest. In our search setting, the maximin utility rule prescribes not to search at all.

Our method of evaluating and comparing decision rules is closely related to the minimax regret criterion. In this literature, the degree of suboptimality (referred to as regret) is measured either in terms of differences (Savage (1951)) or, as popular in the computer science literature, in terms of ratios (Sleator and Tarjan (1985); see also the axiomatization of Terlizzese (2008)), which can also be found in the robust contract literature (e.g., Chassang (2013)). We prefer ratios to obtain a scale-free measure and, thus, to be able to compare the performance across different specifications of the environment.

However, our evaluation method differs conceptually from that used in the minimax regret literature. We evaluate the performance not only *ex ante*, but also after each additional piece of information has been gathered. This is also done in a follow-up paper by Schlag and Sobolev (2020) that studies finite-horizon search in a more specific setting. This method stands in contrast to the traditional approaches. One of these approaches evaluates strategies retrospectively, after all uncertainty is resolved. This tradition goes back to Savage (1951). Search models that follow this tradition appear in Bergemann and Schlag (2011b) and Parakhonyak and Sobolev (2015). An alternative approach is to evaluate strategies *ex ante*, by the present value of their expected payoffs, where the searcher is able to commit to her strategy. This approach is adopted in the secretary problem (Fox and Marnie (1960)) that studies sequential search within a nonrandom set of exchangeable alternatives (for a review, see Ferguson (1989)). An analysis of robust search with *ex ante* commitment in the setting of this paper is difficult and remains unsolved. Bergemann and Schlag (2011b) and Parakhonyak and Sobolev (2015) study a special case with two periods, and Babaioff et al. (2009) study asymptotic performance of approximately optimal algorithms in a related problem with no recall, so these results are not comparable to our paper.

The term *robustness* goes back to Huber (1964, 1965). It is defined as a statistical procedure whose “performance is insensitive to small deviations of the actual situation from the idealized theoretical model” (Huber (1965)). Prasad (2003) and Bergemann and Schlag (2011a) formalize this notion for decision making. They measure insensitivity under small deviations as performance being close to that of the optimal policy. The

same approach has been applied to large deviations, where the performance is evaluated under a large class of distributions, as in statistical treatment choice (Manski (2004), Schlag (2006), and Stoye (2009)), auctions (Kasberger and Schlag (2017)), and search in markets (Bergemann and Schlag (2011b) and Parakhonyak and Sobolev (2015)). The term “robustness” has been used in the same spirit, to achieve an objective independently of modeling details, in robust mechanism design (Bergemann and Morris (2005)), and in the field of control theory (Zhou et al. (1995)).

The term “robustness” has been used in a different spirit to describe optimal decisions under maximin utility, as in Hansen et al. (2001), Ben-Tal et al. (2009), Chas-sang (2013), Carroll (2015), and Carrasco et al. (2018). It also appears in Kajii and Morris (1997), where the concept of robustness is related to closeness in the strategy space rather than in the payoff space.

We proceed as follows. In Section 2, we introduce our model and focus on stationary decision rules. In Section 3, we consider binary environments, while in Section 4, we consider general environments. In Section 5, we study general decision rules. Section 6 concludes. The proofs are provided in the Appendix.

2. MODEL

2.1 Setting

An individual chooses among alternatives that arrive sequentially. Each alternative is identified with its value to the individual. The individual starts with an outside option x_0 that is given and is strictly positive, so $x_0 > 0$. Alternatives x_1, x_2, \dots are realizations of an infinite sequence of i.i.d. random variables. In each round $t = 0, 1, 2, \dots$, after having observed x_t , the individual decides whether to stop the search or to wait for another alternative. There is *free recall*: when the individual decides to stop, she chooses the highest alternative she has seen so far. The highest alternative up to t is referred to as *best-so-far alternative* and is denoted by y_t , so

$$y_t = \max\{x_0, x_1, \dots, x_t\}.$$

Payoffs are discounted over time with a discount factor $\delta \in (0, 1)$. From the perspective of round 0, the payoff of stopping after t rounds is $\delta^t y_t$. The discount factor incorporates various multiplicative costs of search, such as the individual’s impatience and a decay of values that have not been accepted.

Alternatives belong to a given set X with $X \subset \mathbb{R}_+$, $0 \in X$, and $\bar{x} = \sup X > x_0$. For instance, this set can be \mathbb{R}_+ , \mathbb{N}_0 , $[0, \bar{x}]$, or $\{0, \bar{x}\}$. We refer to X as the set of feasible alternatives. Inclusion of 0 in X is for notational convenience. Nothing changes if we replace 0 by some \underline{x} as long as the outside option satisfies $x_0 \geq \underline{x}$. Inclusion of 0 is natural in applications where search may not generate a new alternative in each round. Here, the absence of a new alternative is modeled as the zero-valued alternative.

The ratio x_0/\bar{x} plays an important role in our analysis; $1 - x_0/\bar{x}$ can be considered as a measure of potential relative gains from search. If X is unbounded, so $\bar{x} = \infty$, then it

is understood that $x_0/\bar{x} = 0$. We assume for clarity of exposition that

$$\frac{x_0}{\bar{x}} \leq \frac{\delta^2}{2 - \delta}. \quad (\text{A}_1)$$

This means that x_0/\bar{x} is not too large or the discount factor is not too small. For example, if $x_0/\bar{x} = 1/2$ or $x_0/\bar{x} = 1/6$, then δ should exceed approximately 0.8 or 0.5, respectively. Clearly, assumption (A₁) is vacuous if $\bar{x} = \infty$. Though we focus on the case when (A₁) holds, we also provide insights for the case when (A₁) does not hold.

Alternatives are independently drawn from X according to a probability distribution F with finite support. We refer to F as an *environment*. Let \mathcal{F}_X be the set of all such environments. The assumption of finite support is made to simplify the definition of histories that can occur with positive probability. The main results extend to arbitrary distributions with finite mean.

The decision making of the individual is formally captured by a *decision rule* that specifies the probability of stopping in every round $t = 0, 1, 2, \dots$ and after every possible history of alternatives in that round. A decision rule is called *stationary* if the stopping probability depends only on the best-so-far alternative, but not on the history that has generated this best-so-far alternative. So a stationary decision rule is a mapping p that specifies the stopping probability $p(y)$ for each best-so-far alternative $y \in X \cup \{x_0\}$ such that $y \geq x_0$.

To simplify exposition, in the following discussion we restrict attention to stationary decision rules. Later, in Section 5, we show that our results continue to hold for general decision rules.

2.2 Performance criterion

We consider an individual who knows all of the above except for the distribution F according to which alternatives are drawn. What she knows about F is that it is contained in a given set of feasible environments \mathcal{F} , where $\mathcal{F} \subset \mathcal{F}_X$. In this paper, we pay special attention to two types of feasible environments. In Section 3, we consider so-called binary environments that can have at most one value above x_0 , and in Section 4, we allow for all environments in \mathcal{F}_X .

Our individual can rule out environments that do not belong to \mathcal{F} , but she does not assess likelihoods of environments that belong to \mathcal{F} . Instead this individual searches for a decision rule that performs well regardless of which environment in \mathcal{F} she faces. We introduce a performance criterion according to which she chooses her decision rule.

The basic idea is as follows. The individual evaluates payoffs when facing a given environment F as in the standard expected utility model. However, unlike the standard model, she has no prior over the different environments belonging \mathcal{F} . Instead, she evaluates a decision rule in a given environment according to how far it is from the best rule for this environment. She then chooses the rule that minimizes the maximum “distance” across all environments in \mathcal{F} . We now introduce the criterion formally.

Let us connect an environment to the best-so-far alternatives it can realize. We say that a best-so-far alternative y is *consistent* with environment F if it can be obtained

under F with a positive probability in some round. Let $Y(F)$ be the set of best-so-far alternatives consistent with F . Note that $y \in Y(F)$ if either $y = x_0$ or $y > x_0$ that can occur under F with positive probability, so $Y(F) = \{x_0\} \cup (\text{supp}(F) \cap (x_0, \infty))$.

Next we introduce payoffs. For a given environment $F \in \mathcal{F}$ and a given best-so-far alternative $y \in Y(F)$, let $U_p(F, y)$ be the expected payoff of a decision rule p under F when the best-so-far alternative is y . According to rule p , the individual stops and gets y with probability $p(y)$, and draws a new alternative with probability $1 - p(y)$. In the latter case, the new best-so-far alternative becomes $\max\{y, x\}$, where x is the value of the new alternative, so

$$U_p(F, y) = p(y)y + (1 - p(y))\delta \int_X U_p(F, \max\{y, x\}) dF(x). \quad (1)$$

Let $V(F, y)$ be the highest possible expected payoff that can be achieved under F when the best-so-far alternative is y , so

$$V(F, y) = \sup_p U_p(F, y).$$

We also refer to $V(F, y)$ as the *optimal payoff*. Note that the optimal payoff is always strictly positive, as $V(F, y) \geq y \geq x_0 > 0$. By Weitzman (1979), the rule that attains $V(F, y)$ under F is a cutoff rule. It prescribes to stop whenever the best-so-far alternative y exceeds a *reservation value* c_F implicitly given as the unique solution of the equation

$$c_F = \delta \left(\int_0^{c_F} c_F dF(x) + \int_{c_F}^{\infty} x dF(x) \right). \quad (2)$$

It follows that

$$V(F, y) = \max\{y, c_F\}. \quad (3)$$

We measure the performance of a decision rule p by the smallest fraction of the optimal payoff attained by p across all environments $F \in \mathcal{F}$ and all best-so-far alternatives $y \in Y(F)$. We call this fraction the *performance ratio* and denote it by $R_p(\mathcal{F})$, so

$$R_p(\mathcal{F}) = \inf_{F \in \mathcal{F}} \inf_{y \in Y(F)} \frac{U_p(F, y)}{V(F, y)}.$$

Note that this performance ratio is guaranteed in each round of search and after each history of alternatives that can be generated with positive probability. Note also that the value of the performance ratio would be the same if we included not only all environments in \mathcal{F} , but also all distributions (priors) over environments in \mathcal{F} .

The highest possible performance ratio is called *robust* and is denoted by $R^*(\mathcal{F})$, so

$$R^*(\mathcal{F}) = \sup_p R_p(\mathcal{F}).$$

A decision rule p^* is called robust if it attains the robust performance ratio, so

$$R_{p^*}(\mathcal{F}) = R^*(\mathcal{F}).$$

Note that $R^*(\mathcal{F})$ depends only on the information available from the start—the set of feasible environments \mathcal{F} —and, implicitly, on the set of feasible alternatives X and the discount factor δ .

2.3 Randomization

We point out the importance of randomization for the design of robust rules. Intuitively it makes sense to randomize whenever feasible environments are sufficiently diverse, as this is how the individual can mitigate the trade-off between stopping the search when it is optimal to continue and continuing the search when it is optimal to stop.

Specifically, we now bound the performance ratio of deterministic rules and conclude that no deterministic rule can outperform the rule that does not search. Consider a deterministic rule p , so $p(x_0) \in \{0, 1\}$. If $p(x_0) = 1$, then p takes the outside option in round 0 in all environments, so it does not search. In this case, the performance ratio is $R_p(\mathcal{F}) = x_0 / \sup_{F \in \mathcal{F}} V(F, x_0)$. Alternatively, if $p(x_0) = 0$, then p keeps searching indefinitely and yields the payoff ratio equal to 0 when facing an environment that never generates alternatives better than x_0 . In this case, $R_p(\mathcal{F}) = 0$.

So by using deterministic rules, one cannot guarantee more than $x_0 / \sup_{F \in \mathcal{F}} V(F, x_0)$. This is the performance ratio of the rule p_1 that does not search, so $p_1(y) = 1$ for all y . This ratio can be arbitrarily small if the outside option x_0 is small or if there are feasible environments that can generate very high alternatives.

3. BINARY ENVIRONMENTS

Suppose that the individual faces an environment that is known to generate at most one alternative above the outside option. The individual knows what she is looking for, she just does not know whether she will find it and, if so, how valuable it will be. We call such environments binary.

Note that any alternative that lies below the outside option can be treated as if it had value 0 as such alternatives would never be chosen. Hence, we can act as if a binary environment generates only two different alternatives, 0 and some value z above x_0 . An environment is called *binary*, denoted by $F_{(z, \sigma)}$, if it is a lottery over two values, 0 and z , with probabilities $1 - \sigma$ and σ , respectively, where $z \in X$ such that $z > x_0$, and $\sigma \in [0, 1]$. Let \mathcal{B}_X be the set of all binary environments over X , so

$$\mathcal{B}_X = \{F_{(z, \sigma)} : z \in X \text{ s.t. } z > x_0, \sigma \in [0, 1]\}.$$

In this section, we assume that the set of feasible environments \mathcal{F} is equal to \mathcal{B}_X . So the individual knows that alternatives are drawn from X and that she faces a binary environment in \mathcal{B}_X . Note that X may contain more than one nonzero alternative. So the individual who faces some environment in \mathcal{B}_X may not know the value of the high alternative z , although she knows that there is at most one such alternative. Only in the special case when $X = \{0, z\}$ does the individual know the value of the high alternative.

How should the individual behave when facing an unknown binary environment? Suppose that she sees an alternative z that lies above x_0 . Then her best-so-far alternative

is z and she knows that this is the highest possible alternative. Thus, she stops searching and sets $p(z) = 1$. So in the following discussion we can assume that

$$p(z) = 1 \quad \text{for all } z > x_0. \quad (4)$$

We investigate only how optimally to choose $p(x_0)$, which is the probability of stopping in each round when the high alternative has not arrived yet.

We present a robust rule for binary environments. We use the notation

$$\eta(x) = \frac{1}{2} + \frac{1}{8}(x + \sqrt{x(x+8)}). \quad (5)$$

THEOREM 1. *Let $0 < x_0 < \bar{x} \leq \infty$ and let (A_1) hold. Then the robust performance ratio is*

$$R^*(\mathcal{B}_X) = \eta(x_0/\bar{x}).$$

It is attained by the robust decision rule p_b^ given by*

$$p_b^*(y) = \begin{cases} \frac{1-\delta}{2-\delta + \frac{1}{2}\left(\frac{y}{\bar{x}} - \sqrt{\frac{y}{\bar{x}}\left(\frac{y}{\bar{x}} + 8\right)}\right)} & \text{if } y = x_0, \\ 1 & \text{if } y > x_0. \end{cases}$$

Most proofs are provided in the [Appendix](#).

Intuitively, the robust decision rule is derived as follows. Consider bounded environments, so $\bar{x} < \infty$. There are two worst-case environments. One never generates alternatives above x_0 , and, hence, it is optimal to stop. The other randomizes between 0 and the highest feasible alternative \bar{x} in such a way that it is optimal to continue. The stopping probability $p_b^*(x_0)$ equalizes the payoff ratios in these environments.

Note that the robust performance ratio as shown in Theorem 1 depends only on x_0/\bar{x} . It does not depend on how many feasible alternatives there are above x_0 . In particular, this ratio remains unchanged if there is only one feasible alternative above x_0 , so $X = \{0, \bar{x}\}$. Moreover, the robust performance ratio does not depend on the discount factor δ . This comes from the fact that both the payoff U_p of the rule and the optimal payoff V are evaluated using the same discount factor. When δ is larger, the impact from the additional search of the robust rule cancels out with that of the optimal rule in the worst case.

Observe that the robust performance ratio is at least $1/2$, so one can guarantee a half of the optimal payoff without having any information about the value of the high alternative. This performance bound is tight when the value of the high alternative is unbounded, so $\bar{x} = \infty$.

Consider the case where the set of feasible alternatives is unbounded, so $\bar{x} = \infty$. Then (A_1) holds for all δ and the performance ratio is $\eta(0) = 1/2$. The robust rule p_b^* prescribes to stop with probability $(1-\delta)/(2-\delta)$ as long as $y = x_0$, independently of the value of the outside option x_0 .

TABLE 1. Illustrative robust performance ratio.

x_0/\bar{x}	1/50	1/20	1/10	1/6	1/5	1/4	1/3	1/2
$\eta(x_0/\bar{x})$	0.552	0.585	0.625	0.666	0.685	0.71	0.75	0.82

Now consider the case where the set of feasible alternatives is bounded, so $\bar{x} < \infty$. Assume that (A₁) holds. The robust performance ratio is increasing in the ratio of the outside option x_0 to the highest feasible alternative \bar{x} . For example, one can guarantee at least 2/3 and 3/4 of the optimal payoff if x_0/\bar{x} exceeds 1/6 and 1/3, respectively. Table 1 shows the performance ratio $\eta(x_0/\bar{x})$ for a few values of x_0/\bar{x} , provided the discount factor is not small, so $x_0/\bar{x} \leq \delta^2/(2 - \delta)$.

Finally, consider the case where (A₁) does not hold, so $x_0/\bar{x} > \delta^2/(2 - \delta)$. Here, decision rule p_b^* has a performance ratio larger than $\eta(x_0/\bar{x})$. This is because the rule p_b^* attains the performance ratio $\eta(x_0/\bar{x})$ when treating the probability σ of the high alternative as a real-valued parameter. When (A₁) is true, then we verify that $\sigma \in [0, 1]$ holds in the worst-case environment. However, when (A₁) is false, then $\sigma > 1$, which is not feasible. In this case, the worst-case payoff ratio has to be computed by including the constraint $\sigma \leq 1$. Including an additional constraint means that the performance ratio of p_b^* can only get larger.

4. GENERAL ENVIRONMENTS

4.1 Setting

Suppose now that the individual faces an environment that is known to generate alternatives that belong to X . The individual knows what alternatives can or cannot appear, but she does not know the likelihood of any of the alternatives. We refer to such environments as general.

In the following discussion, we assume that the set of feasible environments \mathcal{F} is not restricted to binary environments, but can contain any environments in \mathcal{F}_X . We assume that X contains at least two elements that are strictly greater than x_0 ; in particular, $\bar{x} > x_0$.

How should the individual behave when she sees an alternative z that lies above the outside option x_0 ? If $z \geq \delta\bar{x}$, then it is best to stop, as no feasible alternative is worth waiting for. However, if $z < \delta\bar{x}$, then the individual faces a trade-off between stopping when a better alternative may still come and continuing to search when it is optimal to stop. This stands in contrast to the behavior in binary environments where one should always stop after seeing any alternative above x_0 .

4.2 Simple lower bound

We start by presenting a lower bound on the robust performance ratio in any environments. An interesting property of this result is the simplicity of the rule that attains this bound. It has the stopping probability that is independent of the best-so-far alternative.

THEOREM 2. *Let $\mathcal{F} \subset \mathcal{F}_X$. The robust performance ratio satisfies*

$$R^*(\mathcal{F}) \geq \frac{1}{4}.$$

The lower bound 1/4 is attained by the decision rule p_g given by

$$p_g(y) = \frac{1 - \delta}{2 - \delta} \quad \text{for all } y \geq x_0.$$

REMARK 1. The lower bound 1/4 on the robust performance ratio is tight. The robust performance ratio is equal to 1/4 (thus, rule p_g is robust) when the set X is unbounded, so $\bar{x} = \infty$.

The robust performance ratio is clearly the lowest when \bar{x} is the highest. We sketch the argument regarding why the robust performance ratio of 1/4 is attained when $\bar{x} = \infty$. In this setting, regardless of how large the best-so-far alternative already is, an infinitely larger alternative can still appear. Hence, the value of the best-so-far alternative plays no role when designing a robust decision rule. We can thus limit attention to rules that have a constant probability of stopping q for some $q \in [0, 1]$. The worst-case environment generates a high alternative z that occurs with extremely small probability σ , but that is sufficiently large so that it is worth waiting for. A greater q means a greater probability of stopping before z realizes, but also a shorter delay before stopping and obtaining z after it has realized. The performance ratio is attained as $z \rightarrow \infty$ and $\sigma \rightarrow 0$ such that the optimal payoff of waiting for the first realization of z goes to infinity. Because $y/z \rightarrow 0$, the individual essentially cares only about getting z . The performance ratio takes the form

$$\frac{q}{1 - \delta(1 - q)} \left(1 - \frac{q}{1 - \delta(1 - q)} \right). \quad (6)$$

To understand (6), it is useful to think of $1 - \delta$ as the exogenous probability that the search stops in the current round and yields zero payoff. Then the expression

$$\frac{q}{1 - \delta(1 - q)} = q + \delta(1 - q)q + \delta^2(1 - q)^2q + \dots$$

can be seen as the expected probability of getting z after it has been realized. To interpret the second factor in (6), consider the following expressions. First,

$$\delta\sigma(1 + \delta(1 - \sigma) + \delta^2(1 - \sigma)^2 + \dots) = \frac{\delta\sigma}{1 - \delta(1 - \sigma)}$$

is the expected probability of not stopping before z realizes under the optimal rule, where the probability of stopping in each round is $1 - \delta$. Second,

$$(1 - q)\delta\sigma(1 + (1 - q)\delta(1 - \sigma) + (1 - q)^2\delta^2(1 - \sigma)^2 + \dots) = \frac{(1 - q)\delta\sigma}{1 - (1 - q)\delta(1 - \sigma)}$$

is the expected probability of not stopping before z realizes under the rule q , where the probability of stopping in each round is $(1 - \delta)q$. As σ tends to 0, the ratio of the latter to the former is

$$\frac{(1 - q)\delta\sigma}{1 - (1 - q)\delta(1 - \sigma)} \cdot \frac{1 - \delta(1 - \sigma)}{\delta\sigma} \xrightarrow{\sigma \rightarrow 0} \frac{(1 - q)(1 - \delta)}{1 - \delta(1 - q)} = 1 - \frac{q}{1 - \delta(1 - q)}.$$

We can now interpret (6). The first factor in (6) is the expected probability of obtaining z after it has realized. The second factor in (6) is the expected probability of not stopping before z realizes for the first time. Setting $q/(1 - \delta(1 - q))$ equal to $1/2$, so $q = (1 - \delta)(2 - \delta)$, maximizes (6), leading to the performance ratio $1/4$.

4.3 Bounded environments

In the following discussion, we present the robust performance ratio for the case where the outside option x_0 is not too small in general environments in which alternatives are bounded, so $\bar{x} < \infty$. We call these environments bounded. Recall the definition of η given by (5) in Section 3 and define

$$f(x, t) = \begin{cases} \frac{(1 - \delta)(1 - t)}{(1 - \delta)(1 - t) + (\sqrt{t} - \sqrt{x})^2} & \text{if } x < t, \\ 1 & \text{if } x \geq t. \end{cases}$$

THEOREM 3. *Let $\bar{x} < \infty$ and let (A_1) hold. There exists a constant $\lambda \in (1/90, 7/100)$ such that if*

$$\frac{x_0}{\bar{x}} \geq \lambda,$$

then the robust performance ratio is

$$R^*(\mathcal{F}_X) = \eta\left(\frac{x_0}{\bar{x}}\right).$$

It is attained by the robust decision rule p_g^ given by*

$$p_g^*(y) = f\left(\frac{y}{\bar{x}}, \eta\left(\frac{x_0}{\bar{x}}\right)\right) \quad \text{for } y \in [x_0, \bar{x}].$$

Theorem 3 shows that if the outside option is not too small relative to the highest possible alternative, in the sense that $x_0/\bar{x} \geq \lambda$, then the robust performance ratio in general environments is the same as it is in binary environments. Remarkably, the constant λ is very small. We prove that $\lambda < 7/100$. Moreover, we numerically (up to precision 10^{-8}) find that

$$\lambda \approx 0.01120000,$$

where $1/90 < 0.0112 < 1/89$ (see Remark 3 in Appendix A.5). Thus, as x_0/\bar{x} increases from 0 to a mere $1/89$, the robust performance ratio climbs from at least $1/4$ (by Theorem 2) to at least $1/2$ (by Theorem 3). In particular, one can guarantee at least $2/3$ and




FIGURE 1. The black line shows the robust performance ratio in binary and general environments with $x_0/\bar{x} \geq \lambda$ (see Theorem 3). The grey line shows the robust performance ratio in binary environments with $x_0/\bar{x} < \lambda$ (see Theorem 1). The dotted line shows the hypothetical robust performance ratio in general environments with $x_0/\bar{x} < \lambda$ (as conjectured from Theorem 2 and Remark 1).

3/4 of the optimum if the outside option exceeds 1/6 and 1/3 of the highest feasible alternative, respectively. Figure 1 illustrates the robust performance ratio in different settings.

Notice that the rule p_b^* that attains the robust performance ratio in binary environments no longer has this property in general environments. This is because it stops immediately when an alternative above the outside option arrives. In contrast, the robust rule p_g^* randomizes whenever it is worth waiting for a higher alternative. Of course, as binary environments belong to the set of general environments, p_g^* is also robust in binary environments as long as $x_0/\bar{x} \geq \lambda$.

Why is the robust performance ratio in general environments the same as it is in binary environments when $x_0/\bar{x} \geq \lambda$? When using rule p_g^* , the worst-case general environments are binary environments that randomize between 0 and \bar{x} . These are the same worst-case environments we found in the binary environment setting. We verify that the performance ratio of p_g^* is equal to $\eta(x_0/\bar{x})$, which is the robust performance ratio in binary environments. This means that p_g^* is robust in general environments because the performance ratio can only become worse if one adds more environments. However, when $x_0/\bar{x} < \lambda$, then the performance ratio of p_g^* is strictly smaller in general environments than it is in binary environments. This is because the worst-case environment now randomizes between 0 and an alternative strictly below \bar{x} .

On a final note, we comment on re-optimization. When the individual draws an alternative z that lies above the outside option x_0 , then the future looks like the original problem except it is as if now the outside option is larger, namely, equal to z . One might be tempted to switch to a decision rule that is designed for this outside option. However, such re-optimization is not allowed in our model. This is investigated in a follow-up paper by Schlag and Sobolev (2020).

5. NONSTATIONARY DECISION RULES

5.1 *Setting*

So far we have restricted attention to stationary decision rules. The behavior prescribed by these rules does not depend on time or history. It depends only on the highest alternative drawn so far. We now demonstrate that our results remain unchanged if we allow for nonstationary decision rules.

We update our definitions introduced in Section 2 to incorporate the dependence of decision rules on histories.

Let $h_t = (x_0, x_1, \dots, x_t)$ be a history of alternatives up to round t for $t = 0, 1, 2, \dots$. A decision rule p prescribes for each history h_t a probability $p(h_t)$ of stopping after that history. A decision rule p is *stationary* if $p(x_0, x_1, \dots, x_t) = p(x_0, x'_1, \dots, x'_t)$ whenever $\max\{x_0, x_1, \dots, x_t\} = \max\{x_0, x'_1, \dots, x'_t\}$. Let \mathcal{P} be the set of all decision rules and let \mathcal{P}_{st} be the set of stationary decision rules.

We say that a history h_t is *consistent* with an environment F if it has a strictly positive probability under F . Let $\mathcal{F} \subset \mathcal{F}_X$ be a set of feasible environments. For an environment $F \in \mathcal{F}$, let $H(F)$ be the set of all finite histories consistent with F .

Consider an environment $F \in \mathcal{F}$ and a history $h_t \in H(F)$. With abuse of notation let $h_{t+1} = (h_t, x_{t+1})$. Let $U_p(F, h_t)$ be the expected payoff of a decision rule p under F when the history is h_t , so

$$U_p(F, h_t) = p(h_t) \max\{x_0, x_1, \dots, x_t\} + (1 - p(h_t)) \delta \int_X U_p(F, (h_t, x_{t+1})) dF(x_{t+1}).$$

Let $V(F, h_t)$ be the optimal payoff under F when the history is h_t , so

$$V(F, h_t) = \sup_{p \in \mathcal{P}} U_p(F, h_t).$$

Note that $V(F, h_t) \geq \max\{x_0, x_1, \dots, x_t\} \geq x_0 > 0$.

The *performance ratio* $R_p(\mathcal{F})$ of a decision rule p is defined as the lowest payoff ratio over all feasible environments and all histories that are consistent with each of those environments:

$$R_p(\mathcal{F}) = \inf_{F \in \mathcal{F}} \inf_{h_t \in H(F)} \frac{U_p(F, h_t)}{V(F, h_t)}.$$

So the performance ratio is the largest fraction of the optimal payoff that a decision rule guarantees no matter what environment in \mathcal{F} the individual faces, in each round of search and after each history of alternatives that can be generated with a positive probability. As in Section 2, the highest possible performance ratio is called *robust*, and is given by

$$R^*(\mathcal{F}) = \sup_{p \in \mathcal{P}} R_p(\mathcal{F}).$$

A decision rule p^* is called *robust* if it attains the robust performance ratio, so $R_{p^*}(\mathcal{F}) = R^*(\mathcal{F})$.

5.2 Randomization

Before describing the results, let us show that randomization is essential for robust search. In other words, deterministic rules are not robust, provided there is an alternative worth waiting for. This insight appeared informally in Section 2.3. Here we prove it formally.

A decision rule p is called *deterministic* if for each history, the individual either stops searching or continues to search with certainty. Formally, this means that $p(h_t) \in \{0, 1\}$ for every possible history h_t . Let p_1 be the deterministic rule that stops searching after each history, so $p_1(h_t) = 1$ for all h_t .

Let $\mathcal{F} \subset \mathcal{F}_X$ be the set of feasible environments. Let F_0 and $F_{\bar{x}}$ be the environments that generate alternative 0 and alternative \bar{x} with certainty, respectively. We show that deterministic decision rules are not robust when these two environments are feasible and \bar{x} is worth waiting for, in the sense that $\delta\bar{x} > x_0$.

PROPOSITION 1. *Let $F_0, F_{\bar{x}} \in \mathcal{F}$ and let $\delta\bar{x} > x_0$. Then*

$$R_p(\mathcal{F}) \leq R_{p_1}(\mathcal{F}) = \frac{x_0}{\delta\bar{x}} < R^*(\mathcal{F}).$$

We explain the above result. The rule p_1 that does not search has a performance ratio equal to $x_0/(\delta\bar{x})$. This is because the worst-case environment for this rule is $F_{\bar{x}}$. This ratio can be arbitrarily small if the outside option x_0 is close to 0 or if the highest alternative \bar{x} is very large. The first inequality in Proposition 1 shows that no deterministic rule can outperform the rule p_1 that does not search. The second inequality shows that deterministic rules are not robust.

REMARK 2. Proposition 1 sheds light on the performance of Bayesian rules. These are the rules used by Bayesian decision makers who maximize their expected payoffs for some prior. Any such rule prescribes to stop the search if the best-so-far alternative is better than the expected continuation payoff under the given prior and to continue the search otherwise. Indifference between stopping and continuing under a given prior is nongeneric in the sense that it does not hold under an open set of priors in the neighborhood of that prior. Hence, Bayesian rules are generically deterministic. So by Proposition 1, Bayesian rules are generically not robust.

5.3 Binary environments

Here we show that there is no loss of generality to restrict attention to stationary rules when investigating robust performance in binary environments. In other words, any rule can be outperformed by an appropriately chosen stationary rule.

PROPOSITION 2. *For each $p \in \mathcal{P}$ there exists $\tilde{p} \in \mathcal{P}_{st}$ such that $R_p(\mathcal{B}_X) \leq R_{\tilde{p}}(\mathcal{B}_X)$.*

By Proposition 2, it is immediate that Theorem 1 extends to nonstationary decision rules.

COROLLARY 1. *Theorem 1 holds when the set of decision rules is \mathcal{P} .*

5.4 General environments

Let us now consider the performance of nonstationary rules in general environments as defined in Section 4.

Theorem 2 clearly continues to hold as it identifies a lower bound that remains valid if the set of decision rules becomes richer.

COROLLARY 2. *Theorem 2 holds when the set of decision rules is \mathcal{P} .*

We apply Corollary 1 to prove the next result.

COROLLARY 3. *Theorem 3 holds when the set of decision rules is \mathcal{P} .*

PROOF. By Corollary 1, the robust performance ratio in binary environments is $\eta(x_0/\bar{x})$. So the robust performance ratio in general environments can only be smaller than $\eta(x_0/\bar{x})$. However, by Theorem 3, the stationary decision rule p_g^* attains $\eta(x_0/\bar{x})$ in general environments. Hence, the rule p_g^* is robust in general environments, with or without the restriction to stationary rules. \square

Finally, we hasten to point out that we do not extend Remark 1 to nonstationary rules in this paper. Whether the performance ratio of 1/4 is robust when the set X of alternatives is unbounded remains an open question.

6. CONCLUSION

It is difficult to search under the classic objectives of expected utility maximization when the distribution of alternatives is not known. In fact, the literature has not produced satisfactory insights into how to search in this setting. In this paper, we identify that this difficulty is due to the desire to achieve the very highest payoff for the given beliefs. Namely, we find that it is easier to search if one reduces the target and replaces “very highest” by “relatively high.” The ease refers to the ability to derive a solution for a very general setting, the simplicity of our algorithm, and the minimality of assumptions one needs to impose on the environment.

Many interesting topics remain that have not been addressed in this paper. Are there good rules that allow the individual to reoptimize after each new alternative arrives? How do we search if costs are additive? What is the robust performance ratio when there is no free recall? What if alternatives do not arrive according to an i.i.d. process?

The methodology developed in this paper is applicable to a spectrum of dynamic decision making problems and should spark future research.

APPENDIX: PROOFS

A.1 Auxiliary definitions and results

Let $r_p(y)$ be given by

$$r_p(y) = \inf_{F \in \mathcal{B}_X} \frac{U_p(F, y)}{V(F, y)} = \inf_{z \in X, z > x_0, \sigma \in [0, 1]} \frac{U_p(F(z, \sigma), y)}{V(F(z, \sigma), y)} \quad \text{for } y \geq x_0. \tag{7}$$

So $r_p(y)$ is the smallest payoff ratio of a stationary rule p under binary environments when the best-so-far alternative is y . Observe that in the model with binary environments, the performance ratio of a decision rule p is given by

$$R_p(\mathcal{B}_X) = r_p(x_0). \tag{8}$$

This is because whenever a best-so-far alternative is $y > x_0$, the individual knows that no better alternative will ever arrive and, thus, stops immediately. The payoff ratio is 1 in this case. So the performance ratio $R_p(\mathcal{B}_X)$ is determined by the smallest payoff ratio when the best-so-far alternative is x_0 .

The ratio $r_p(x_0)$ is our main instrument for finding the performance ratio of p not only in binary environments \mathcal{B}_X , but also in general environments \mathcal{F}_X .

We now find $r_p(y)$ for a given decision rule p . To simplify the exposition of the proofs, we introduce some notation. For $x_0 \leq y < z$ and $s \in [0, 1)$, let κ_p and m_p be given by

$$\kappa_p(y) = \frac{p(y)}{1 - \delta + \delta p(y)} \tag{9}$$

and

$$m_p(y, z, s) = \frac{(1 - s)\kappa_p(y)y + (1 - \kappa_p(y))\kappa_p(z)sz}{1 - s\kappa_p(y)}. \tag{10}$$

Note that $1 - s\kappa_p(y) > 0$, because $\kappa_p(y) \in [0, 1]$ and $s < 1$. For $\sigma \in [0, 1]$, let

$$s_\sigma = \frac{\delta\sigma}{1 - \delta + \delta\sigma}. \tag{11}$$

The next lemma shows that $\kappa_p(y)y$ and $m_p(y, z, s_\sigma)$ are the payoffs of rule p when the environment is $F_{(z, \sigma)}$ for the cases of $z \leq y$ and $z > y$, respectively.

LEMMA 1. *We have*

$$U_p(F_{(z, \sigma)}, y) = \begin{cases} \kappa_p(y)y & \text{if } z \leq y, \\ m_p(y, z, s_\sigma) & \text{if } z > y. \end{cases} \tag{12}$$

PROOF. Let $z \leq y$. By (1), the payoff of decision rule p in a binary environment $F_{(z, \sigma)}$ is given by $U_p(F_{(z, \sigma)}, y) = p(y)y + (1 - p(y))\delta U_p(F_{(z, \sigma)}, y)$. Solving this equation for $U_p(F_{(z, \sigma)}, y)$ and using (9) yields

$$U_p(F_{(z, \sigma)}, y) = \frac{p(y)y}{1 - \delta + \delta p(y)} = \kappa_p(y)y \quad \text{if } z \leq y. \tag{13}$$

Next, let $z > y$. By (1),

$$U_p(F_{(z, \sigma)}, y) = p(y)y + (1 - p(y))\delta(\sigma U_p(F_{(z, \sigma)}, z) + (1 - \sigma)U_p(F_{(z, \sigma)}, y)). \tag{14}$$

Inserting $y = z$ into (13) yields $U_p(F_{(z,\sigma)}, z) = \kappa_p(z)z$. Inserting this into (14) and solving for $U_p(F_{(z,\sigma)}, y)$ yields

$$U_p(F_{(z,\sigma)}, y) = \frac{p(y)y + (1 - p(y))\delta\sigma\kappa_p(z)z}{1 - \delta(1 - \sigma)(1 - p(y))}.$$

Finally, we use (9) to replace $p(y)$ by $(1 - \delta)\kappa_p(y)/(1 - \delta\kappa_p(y))$ and we use (11) to replace σ by $\frac{(1-\delta)s_\sigma}{\delta(1-s_\sigma)}$. Note that $s_\sigma \leq \delta < 1$, so $\delta(1 - s_\sigma) > 0$ and $\kappa_p(y) \geq 0$, so $1 - \delta\kappa_p(y) > 0$. After simplification, we obtain

$$U_p(F_{(z,\sigma)}, y) = \frac{(1 - s_\sigma)\kappa_p(y)y + (1 - \kappa_p(y))\kappa_p(z)s_\sigma z}{1 - s_\sigma\kappa_p(y)} = m_p(y, z, s_\sigma) \quad \text{if } z > y. \quad \square$$

We now characterize $r_p(y)$. Below we repeatedly use the infimum operator. Whenever the infimum is taken over an empty set, we follow the convention that $\inf(\emptyset) = +\infty$.

PROPOSITION 3. *Let $y \in X \cup \{x_0\}$ and $z \in X$ such that $x_0 \leq y < z$. For each rule p ,*

$$r_p(y) \leq \inf_{s \in (\frac{y}{z}, \delta]} \frac{m_p(y, z, s)}{sz}. \tag{15}$$

Moreover, if $p(y)$ is weakly increasing in y , then

$$r_p(y) = \min \left\{ \kappa_p(y), \inf_{z \in X, z > x_0, s \in (\frac{y}{z}, \delta]} \frac{m_p(y, z, s)}{sz} \right\}. \tag{16}$$

PROOF. Let p be a decision rule, let $y \in [x_0, \bar{x}]$, and let $F_{(z,\sigma)} \in \mathcal{B}_X$. First, we find $V(F_{(z,\sigma)}, y)$. Solving (2) for $c_{F_{(z,\sigma)}}$ and using (11), we obtain

$$c_{F_{(z,\sigma)}} = \frac{\delta\sigma z}{1 - \delta + \delta\sigma} = s_\sigma z.$$

By (3),

$$V(F_{(z,\sigma)}, y) = \max\{y, c_{F_{(z,\sigma)}}\} = \max\{y, s_\sigma z\}. \tag{17}$$

Next,

$$\begin{aligned} r_p(y) &= \inf_{z \in X, z > x_0, \sigma \in [0, 1]} \frac{U_p(F_{(z,\sigma)}, y)}{V(F_{(z,\sigma)}, y)} \\ &= \min \left\{ \inf_{\substack{z \in X, \sigma \in [0, 1] \\ \text{s.t. } x_0 < z \leq y}} \frac{U_p(F_{(z,\sigma)}, y)}{V(F_{(z,\sigma)}, y)}, \inf_{\substack{z \in X, \sigma \in [0, 1] \\ \text{s.t. } c_{F_{(z,\sigma)}} \leq y < z}} \frac{U_p(F_{(z,\sigma)}, y)}{V(F_{(z,\sigma)}, y)} \right\} \\ &\quad \left\{ \inf_{\substack{z \in X, \sigma \in [0, 1] \\ \text{s.t. } y < c_{F_{(z,\sigma)}} < z}} \frac{U_p(F_{(z,\sigma)}, y)}{V(F_{(z,\sigma)}, y)} \right\} \\ &= \min \left\{ \inf_{\substack{z \in X, s \in [0, \delta] \\ \text{s.t. } x_0 < z \leq y}} \frac{\kappa_p(y)y}{y}, \inf_{\substack{z \in X, s \in [0, \delta] \\ \text{s.t. } sz \leq y < z}} \frac{m_p(y, z, s)}{y} \right\}, \end{aligned}$$

$$\inf_{\substack{z \in X, s \in [0, \delta] \\ \text{s.t. } y < sz}} \left. \frac{m_p(y, z, s)}{sz} \right\}. \tag{18}$$

The first equality is by (7), the second equality is by partitioning the set $\{(z, \sigma) : z \in X, z > x_0, \sigma \in [0, 1]\}$ into three subsets, and the third equality is by (17), Lemma 1, and the property that for each $s \in [0, \delta]$, there is a unique $\sigma \in [0, 1]$ such that $s = s_\sigma$. The inequality (15) follows from (18).

To prove (16), consider $sz \leq y < z$. By (10), we have

$$m_p(y, z, s) \geq \frac{(1-s)\kappa_p(y)y + (1-\kappa_p(y))\kappa_p(y)sy}{1-s\kappa_p(y)} = \kappa_p(y)y,$$

where the inequality is because $p(z)$ is increasing by assumption and, hence, $\kappa_p(z)z \geq \kappa_p(y)y$, and the equality is by the simplification of the expression. Moreover, by (10), we have $m_p(y, z, 0) = \kappa_p(y)y$. Therefore,

$$\inf_{\substack{z \in X, s \in [0, \delta] \\ \text{s.t. } sz \leq y < z}} \frac{m_p(y, z, s)}{y} = \frac{\kappa_p(y)y}{y} = \kappa_p(y). \tag{19}$$

So (16) follows from (18) and (19). □

We now aim to compute $r_p(x_0)$ using Proposition 3. It is apparent from Proposition 3 that the difficulty in finding $r_p(x_0)$ is the minimization of $m_p(x_0, z, s)/(sz)$ with respect to (z, s) . For our proofs, we minimize $m_p(x_0, z, s)/(sz)$ with respect to s for $s \in (0, 1)$. We then show that the solution is feasible, so it attains $r_p(y)$, provided (A₁) holds.

LEMMA 2. *Let $z \in X$ with $z > x_0$. Suppose that $x_0 \leq y < z \leq \bar{x}$. Then*

$$\inf_{s \in (0, 1)} \frac{m_p(y, z, s)}{sz} = \begin{cases} \frac{1}{z} (\kappa_p(y)\sqrt{y} + \sqrt{(1-\kappa_p(y))(\kappa_p(z)z - \kappa_p(y)y)})^2 & \text{if } y < \kappa_p(z)z, \\ \kappa_p(z) & \text{if } y \geq \kappa_p(z)z. \end{cases}$$

Moreover, the value of s that minimizes $m_p(y, z, s)/(sz)$ is given by

$$s^*(y, z) = \begin{cases} \left(\kappa_p(y) + \sqrt{(1-\kappa_p(y))\left(\frac{\kappa_p(z)z}{y} - \kappa_p(y)\right)} \right)^{-1} & \text{if } y < \kappa_p(z)z, \\ 1 & \text{if } y \geq \kappa_p(z)z. \end{cases} \tag{20}$$

PROOF. Assume $x_0 \leq y < z \leq \bar{x}$. Let us find a solution $s^*(y, z)$ to the problem

$$\inf_{s \in (0, 1)} \frac{m_p(y, z, s)}{sz} = \inf_{s \in (0, 1)} \frac{(1-s)\kappa_p(y)y + (1-\kappa_p(y))\kappa_p(z)sz}{(1-s\kappa_p(y))sz}.$$

We have

$$\frac{\partial}{\partial s} \left(\frac{m_p(y, z, s)}{sz} \right) = \frac{\kappa_p(y)}{(1-s\kappa_p(y))^2 s^2 z} G(s), \tag{21}$$

where

$$G(s) = (\kappa_p(z)z(1 - \kappa_p(y)) - \kappa_p(y)y)s^2 + 2\kappa_p(y)ys - y.$$

The expression $\kappa_p(y)/((1 - s\kappa_p(y))^2s^2z)$ is positive when $s \in (0, 1]$ and $z > y \geq x_0 > 0$. So the sign of (21) is determined by $G(s)$. Observe that $G(s)$ is increasing for $s \in (0, 1]$, because

$$\frac{dG(s)}{ds} = 2\kappa_p(z)z(1 - \kappa_p(y))s + 2\kappa_p(y)y(1 - s) \geq 0.$$

Also observe that $G(0) = -y$ and $G(1) = (1 - \kappa_p(y))(\kappa_p(z)z - y)$.

Suppose that $\kappa_p(z)z - y \leq 0$, so $G(1) \leq 0$. It follows that $G(s) \leq 0$ for all $s \in [0, 1]$. Thus, $s^*(y, z) = 1$ is a solution.

Alternatively, suppose that $\kappa_p(z)z - y > 0$, so $G(1) \geq 0$. Then there exists a solution of $G(s) = 0$ on $(0, 1]$. The quadratic equation $G(s) = 0$ has two solutions,

$$s_1 = \frac{1}{\kappa_p(y) + \sqrt{D}} \quad \text{and} \quad s_2 = \frac{1}{\kappa_p(y) - \sqrt{D}},$$

where

$$D = (1 - \kappa_p(y)) \left(\frac{\kappa_p(z)z}{y} - \kappa_p(y) \right).$$

Because $1 \geq \kappa_p(z)z > y$, we have $\sqrt{D} > 1 - \kappa_p(y) > 0$, so $s_1 \in (0, 1)$ and $s_2 \notin [0, 1]$. It follows that $s^*(y, z) = s_1$ is a solution.

Thus, we have shown (20). It remains to substitute $s = s^*(y, z)$ into $m_p(y, z, s)/(sz)$. For the case of $\kappa_p(z)z - y \leq 0$, we obtain

$$\inf_{s \in (0, 1]} \frac{m_p(y, z, s)}{sz} = \frac{m_p(y, z, 1)}{z} = \kappa_p(z).$$

For the case of $\kappa_p(z)z - y > 0$, we obtain, after simplification,

$$\inf_{s \in (0, 1]} \frac{m_p(y, z, s)}{sz} = \frac{m_p(y, z, s_1)}{z} = \frac{1}{z} (\kappa_p(y)\sqrt{y} + \sqrt{(1 - \kappa_p(y))(\kappa_p(z)z - \kappa_p(y)y)})^2. \quad \square$$

A.2 Proof of Theorem 1

Let $q \in [0, 1]$, and let p_q be a decision rule given by $p_q(x_0) = q$ and $p_q(y) = 1$ for all $y > x_0$. To prove Theorem 1, we use the notation and results from Appendix A.1. We show that under assumption (A1), the rule p_{q^*} with

$$q^* = \frac{1 - \delta}{2 - \delta + \frac{1}{2} \left(\frac{x_0}{\bar{x}} - \sqrt{\frac{x_0}{\bar{x}} \left(\frac{x_0}{\bar{x}} + 8 \right)} \right)} \tag{22}$$

maximizes $R_{p_q}(\mathcal{B}_X)$ among all $q \in [0, 1]$. So this rule is robust. We then verify that $R_{p_{q^*}}(\mathcal{B}_X) = \eta(x_0/\bar{x})$.

Recall the notation r_p , κ_p , and m_p for a rule p introduced in Appendix A.1 by (7), (9), and (10), respectively. By (8) and Proposition 3,

$$R_{p_q}(\mathcal{B}_X) = r_{p_q}(x_0) = \min \left\{ \kappa_{p_q}(x_0), \inf_{z \in X, z > x_0, s \in (\frac{x_0}{z}, \delta]} \frac{m_{p_q}(x_0, z, s)}{sz} \right\}.$$

Let $z \in X$ with $z > x_0$ and let $s \in (0, 1]$. Note that $\kappa_{p_q}(z) = 1$. Let

$$k = \kappa_{p_q}(x_0) = \frac{q}{1 - \delta + \delta q}. \tag{23}$$

By (10), we have

$$m_{p_q}(x_0, z, s) = \frac{(1 - s)\kappa_{p_q}(x_0)x_0 + (1 - \kappa_{p_q}(x_0))\kappa_{p_q}(z)sz}{1 - s\kappa_{p_q}(x_0)} = \frac{(1 - s)kx_0 + (1 - k)sz}{1 - sk}.$$

Observe that $m_{p_q}(x_0, z, s)/(sz)$ is decreasing in z , because

$$\frac{d}{dz} \left(\frac{m_{p_q}(x_0, z, s)}{sz} \right) = \frac{d}{dz} \left(\frac{(1 - s)kx_0 + (1 - k)sz}{(1 - sk)sz} \right) = -\frac{(1 - s)kx_0}{(1 - sk)sz^2} \leq 0,$$

so $m_{p_q}(x_0, z, s)/(sz)$ attains its infimum as $z \rightarrow \bar{x}$. We thus obtain

$$R_{p_q}(\mathcal{B}_X) = \min \left\{ k, \inf_{s \in (\frac{x_0}{\bar{x}}, \delta]} \frac{(1 - s)k\frac{x_0}{\bar{x}} + (1 - k)s}{(1 - sk)s} \right\}.$$

Note that each $k \in [0, 1]$ corresponds to a unique $q \in [0, 1]$, so

$$R^*(\mathcal{B}_X) = \sup_{q \in [0, 1]} R_{p_q}(\mathcal{B}_X) = \sup_{k \in [0, 1]} \min \left\{ k, \inf_{s \in (\frac{x_0}{\bar{x}}, \delta]} \frac{(1 - s)k\frac{x_0}{\bar{x}} + (1 - k)s}{(1 - sk)s} \right\}. \tag{24}$$

Consider the case of $\bar{x} = \infty$, so $x_0/\bar{x} = 0$. By (24), we obtain

$$R^*(\mathcal{B}_X) = \sup_{k \in [0, 1]} \min \left\{ k, \inf_{s \in (0, \delta]} \frac{(1 - k)s}{(1 - sk)s} \right\} = \sup_{k \in [0, 1]} \min\{k, 1 - k\} = \frac{1}{2}.$$

The maximum is attained at $k^* = 1/2$. By (23), $k^* = 1/2$ corresponds to $q^* = (1 - \delta)/(2 - \delta)$. We have thus shown (22) for the case of $x_0/\bar{x} = 0$.

Next, to address the case of $\bar{x} < \infty$, we first find

$$\rho^*(x_0, \bar{x}) = \sup_{k \in [0, 1]} \min \left\{ k, \inf_{s \in (0, 1]} \frac{m_{p_q}(x_0, \bar{x}, s)}{s\bar{x}} \right\}.$$

We later show that $R^*(\mathcal{B}_X) = \rho^*(x_0, \bar{x})$. By Lemma 2 with $y = x_0$ and $z = \bar{x}$, we have

$$\inf_{s \in (0, 1]} \frac{m_{p_q}(x_0, \bar{x}, s)}{s\bar{x}} = (k\sqrt{\bar{x}} + \sqrt{(1 - k)(1 - kx)})^2,$$

where we denote

$$x = \frac{x_0}{\bar{x}}.$$

Thus, we obtain

$$\rho^*(x_0, \bar{x}) = \sup_{k \in [0,1]} \min \{k, (k\sqrt{x} + \sqrt{(1-k)(1-kx)})^2\}. \tag{25}$$

Observe that

$$\frac{d}{dk}(k\sqrt{x} + \sqrt{(1-k)(1-kx)}) = -\frac{(\sqrt{(1-k)x} - \sqrt{1-kx})^2}{2\sqrt{(1-k)(1-kx)}} \leq 0. \tag{26}$$

So the term $(k\sqrt{x} + \sqrt{(1-k)(1-kx)})^2$ is decreasing in k . We thus conclude that the solution of (25) is the solution to the equation

$$k = (k\sqrt{x} + \sqrt{(1-k)(1-kx)})^2. \tag{27}$$

Because the left-hand side of (27) is strictly increasing and the right-hand side of (27) is decreasing in k for $k \in [0, 1]$, the solution of (27) is unique. By substituting

$$k^* = \frac{1}{2} + \frac{1}{8}(x + \sqrt{x(x+8)}) \tag{28}$$

into (27), we verify that k^* is the solution of (27). By (23), k^* corresponds to q^* as given by (22). Finally, by (5) and (25),

$$\rho^*(x_0, \bar{x}) = k^* = \eta(x).$$

It remains to show that that $R^*(\mathcal{B}_X) = \rho^*(x_0, \bar{x})$. Observe that by (24), we have $R^*(\mathcal{B}_X) = \rho^*(x_0, \bar{x})$ if $s^*(k^*)$ is feasible in (24); specifically, if $s^*(k^*) \in (x_0/\bar{x}, \delta]$. By Lemma 2, for $k \in [0, 1]$,

$$s^*(k) = \left(k + \sqrt{(1-k)\left(\frac{\bar{x}}{x_0} - k\right)}\right)^{-1} = \frac{\sqrt{x}}{k\sqrt{x} + \sqrt{(1-k)(1-kx)}}. \tag{29}$$

To verify that $s^*(k^*) > x_0/\bar{x} = x$, observe that $s^*(k)$ is increasing in k , so

$$s^*(k^*) \geq s^*(0) = \sqrt{x} > x.$$

Finally, we verify that $s^*(k^*) \leq \delta$, provided (A_1) holds, so $x \leq \delta^2/(2 - \delta)$. It is easy to see from (29) that $s^*(k)$ is increasing in both k and x . Moreover, by (28), k^* is increasing in x . Consequently, $s^*(k^*)$ is increasing in x . So to verify that $s^*(k^*) \leq \delta$ for all $x \leq \delta^2/(2 - \delta)$, it suffices to verify it for $x = \delta^2/(2 - \delta)$. Substituting $x = \delta^2/(2 - \delta)$ into (28) yields $k^* = 1/(2 - \delta)$. Next, substituting both $x = \delta^2/(2 - \delta)$ and $k^* = 1/(2 - \delta)$ into (29) yields $s^*(k^*) = \delta$. We, thus, have verified that $s^*(k^*)$ is in $(x_0/\bar{x}, \delta]$ whenever (A_1) holds.

A.3 Connecting general and binary environments

As a preliminary step before proving Theorems 2 and 3, we show that the performance ratio of a *regular* decision rule p is equal to $r_p(x_0)$, where r_p is defined in Appendix A.1.

DEFINITION 1. Rule p is *regular* if it satisfies two conditions:

- (i) $p(y)$ is weakly increasing in y
- (ii) $r_p(y)$ is weakly increasing in y .

Regularity is an intuitive condition. Condition (i) means that the individual is more likely to accept a greater best-so-far alternative. Condition (ii) is a “free-disposal” property. If this property does not hold, so a better best-so-far alternative leads to a lower guaranteed payoff ratio, then the individual could be better off by destroying some part of the value of the best-so-far alternative.

The next proposition shows that when determining the performance ratio of any regular decision rule in general environments, we can restrict attention to binary environments and to a single best-so-far alternative $y = x_0$. This is a dramatic simplification of the problem, because binary environments are easy to deal with and we do not need to worry about a multitude of possible values of best-so-far alternatives.

PROPOSITION 4. *Suppose that p is regular. Then $R_p(\mathcal{F}_X) = r_p(x_0)$.*

We hasten to point out that by Proposition 4, the restriction to binary environments is without loss of generality only when dealing with regular rules. The rule p_b^* identified in Theorem 1 is not regular, so Proposition 4 does not apply to this rule. In fact, it is easy to verify that p_b^* is not robust in general environments.

PROOF OF PROPOSITION 4. The proof relies on two lemmas. We first present the statements of these two lemmas and show how they prove Proposition 4. Then we prove the two lemmas.

Let Y_0 be the set of all possible best-so-far alternatives:

$$Y_0 = \bigcup_{F \in \mathcal{F}_X} Y(F) = \{y \in X \cup \{x_0\} : y \geq x_0\}.$$

Lemma 3 shows that the smallest payoff ratio does not change if we expand the set $Y(F)$ for each $F \in \mathcal{F}_X$ to Y_0 .

LEMMA 3. *We have $\inf_{F \in \mathcal{F}_X, y \in Y(F)} \frac{U_p(F, y)}{V(F, y)} = \inf_{F \in \mathcal{F}_X, y \in Y_0} \frac{U_p(F, y)}{V(F, y)}$.*

The proof is deferred to the end of this subsection.

Lemma 4 proves that when determining the worst-case payoff ratio of a regular decision rule in general environments, we can restrict attention to binary environments.

LEMMA 4. *Suppose that p is regular. Then, for each $y \in Y_0$,*

$$\inf_{F \in \mathcal{F}_X} \frac{U_p(F, y)}{V(F, y)} = r_p(y). \tag{30}$$

The proof is deferred to the end of this subsection.

From Lemmas 3 and 4, we obtain

$$R_p(\mathcal{F}_X) = \inf_{F \in \mathcal{F}_X} \left(\inf_{y \in Y(F)} \frac{U_p(F, y)}{V(F, y)} \right) = \inf_{y \in Y_0} \left(\inf_{F \in \mathcal{F}_X} \frac{U_p(F, y)}{V(F, y)} \right) = \inf_{y \in Y_0} r_p(y) = r_p(x_0),$$

where the first equality is by the definition of $R_p(\mathcal{F}_X)$, the second equality is by Lemma 3, the third equality is by Lemma 4, and the fourth equality is because r_p is regular. This completes the proof of Proposition 4. \square

We now prove Lemmas 3 and 4.

PROOF OF LEMMA 3. For each $y \in Y_0$, let $\hat{\mathcal{F}}_X(y)$ be the set of all environments in which y is a feasible best-so-far alternative:

$$\hat{\mathcal{F}}_X(y) = \{F \in \mathcal{F}_X : y \in Y(F)\}.$$

Using this notation, we can equivalently express $R_p(\mathcal{F}_X)$ as

$$R_p(\mathcal{F}_X) = \inf_{y \in Y_0} \left(\inf_{F \in \hat{\mathcal{F}}_X(y)} \frac{U_p(F, y)}{V(F, y)} \right) = \inf_{y \in Y_0} \left(\inf_{F \in \text{Closure}(\hat{\mathcal{F}}_X(y))} \frac{U_p(F, y)}{V(F, y)} \right). \tag{31}$$

It remains to show that $\text{Closure}(\hat{\mathcal{F}}_X(y)) = \mathcal{F}_X$ for each $y \in Y_0$. Intuitively, if there is a distribution $F \notin \hat{\mathcal{F}}_X(y)$, so y is not in the support of F , then there is another distribution that is arbitrarily close to F and yet places a positive, albeit arbitrarily small, probability on y .

Since $x_0 \in Y(F)$ for all $F \in \mathcal{F}_X$, we have $\hat{\mathcal{F}}_X(x_0) = \mathcal{F}_X$. Now consider $y \in Y_0 \setminus \{x_0\}$. Let $F \in \mathcal{F}_X$. Let D_y be the Dirac distribution that assigns probability 1 on y . Let $(F_k)_{k=1}^\infty$ be a sequence of distributions given by

$$F_k = \frac{1}{k} D_y + \left(1 - \frac{1}{k}\right) F, \quad k \in \mathbb{N},$$

so $\lim_{k \rightarrow \infty} F_k = F$. Notice that $F_k \in \mathcal{F}_X$ for all $k \in \mathbb{N}$. By construction, we have $\{y\} \subset Y(D_y) \subset Y(F_k)$ and, thus, $F_k \in \hat{\mathcal{F}}_X(y)$ for all $k \in \mathbb{N}$, so $F \in \text{Closure}(\hat{\mathcal{F}}_X(y))$. Since the above statement is true for all $F \in \mathcal{F}_X$, it follows that $\text{Closure}(\hat{\mathcal{F}}_X(y)) = \mathcal{F}_X$. The statement of the lemma is then immediate by (31). \square

PROOF OF LEMMA 4. Let p be a regular rule. Fix any best-so-far alternative $y \in Y_0$ and any environment $F \in \mathcal{F}_X$. We show that there exist a probability $\sigma \in [0, 1]$ and alternatives $w, z \in Y(F)$ with $y \leq w \leq z \leq \bar{x}$ such that

$$\frac{U_p(F, y)}{V(F, y)} \geq \frac{U_p(F, w)}{V(F, w)} \geq \frac{U_p(F(z, \sigma), w)}{V(F(z, \sigma), w)}. \tag{32}$$

Then, by (32), the definition of r_p , and the regularity of rule p , we obtain

$$\frac{U_p(F, y)}{V(F, y)} \geq \frac{U_p(F(z, \sigma), w)}{V(F(z, \sigma), w)} \geq r_p(w) \geq r_p(y).$$

This implies (30), because $\mathcal{B}_X \subset \mathcal{F}_X$ and F is an arbitrary environment in \mathcal{F}_X . To complete the proof, it remains to show (32).

Let w be an alternative that satisfies

$$U_p(F, w) = \min_{x \in Y(F), x \geq y} U_p(F, x). \tag{33}$$

By assumption, F has a finite support, so $Y(F)$ is finite. Therefore, w is well defined. Because $U_p(F, w) \leq U_p(F, y)$ by (33), and $V(F, w) \geq V(F, y)$ by (3) and the constraint $w \geq y$, we obtain the first inequality in (32).

Let us show the second inequality in (32). Recall that by (1),

$$U_p(F, w) = p(w)w + (1 - p(w))\delta \int_0^{\bar{x}} U_p(F, \max\{w, x\}) dF(x). \tag{34}$$

Let c_F be the reservation value of F as given by (2). Subsequently, we omit the subscript, so $c = c_F$. As $\delta < 1$, it is easy to verify that

$$0 \leq c/\delta < \bar{x}. \tag{35}$$

Suppose that $w \geq c$. Then, by (3), $V(F, w) = w$. Moreover, $V(F_{(w,1)}, w) = w$. Also, by (33) and (34),

$$U_p(F, w) \geq p(w)w + (1 - p(w))\delta U_p(F, w).$$

Solving this inequality yields

$$U_p(F, w) \geq \frac{p(w)w}{1 - \delta(1 - p(w))} = \kappa_p(w)w = U_p(F_{(w,1)}, w),$$

where the first equality is by (9) and the second equality is by Lemma 1 (see Appendix A.1). Thus, we have obtained the second inequality in (32) with $F_{(z,\sigma)} = F_{(w,1)}$.

Next suppose that $w < c$. To show the second inequality in (32), we first bound the expression $\int_0^{\bar{x}} U_p(F, \max\{w, x\}) dF(x)$ from below. For each $x \in Y(F)$, let

$$u(x) = \begin{cases} U_p(F, w), & x \leq c, \\ U_p(F, x), & x > c. \end{cases} \tag{36}$$

Since $w < c$, by (33) and (36), we have

$$U_p(F, \max\{w, x\}) \geq u(x) \quad \text{for all } x \in [0, \bar{x}]. \tag{37}$$




FIGURE 2. Two cases that arise when solving (38) and (39). The dotted lines show the values of $u(x)$ for x in the discrete support of F . The solid lines show $u_{\text{conv}}(x)$.

Let us find alternatives w' and z in $Y(F)$ with $w \leq w' \leq c/\delta \leq z \leq \bar{x}$ and a probability $\sigma \in [0, 1]$ such that

$$\int_0^{\bar{x}} u(x) dF(x) \geq (1 - \sigma)u(w') + \sigma u(z), \tag{38}$$

$$(1 - \sigma)w' + \sigma z = \frac{c}{\delta}. \tag{39}$$

We apply the convexification method as in [Kamenica and Gentzkow \(2011\)](#). Consider the greatest convex function $u_{\text{conv}} : [0, \bar{x}] \rightarrow \mathbb{R}$ such that $u_{\text{conv}}(x) \leq u(x)$ for all $x \in Y(F)$, as shown in Figure 2. This is known as the convex closure of u . Clearly, to satisfy (38) and (39), the alternatives w' and z must be the closest values in $Y(F)$ to the left and to the right of c/δ , respectively, such that $u(w') = u_{\text{conv}}(w')$ and $u(z) = u_{\text{conv}}(z)$. Then σ is derived from (39). In particular, if $u(c/\delta) = u_{\text{conv}}(c/\delta)$, then $w' = z = c/\delta$ and σ is arbitrary.

Let (w', z, σ) satisfy (38) and (39). Observe that the straight line through points $(w', u(w'))$ and $(z, u(z))$ is weakly below the graph of u . Moreover, this straight line has a nonnegative slope, because by (33), we have $u(w) \leq u(x)$ for all $x \geq w$. So we can conclude that

$$\begin{aligned} u(w') &\leq u(x) \quad \text{for all } x \geq w', \\ u(z) &\leq u(x) \quad \text{for all } x \geq z. \end{aligned} \tag{40}$$

We consider two cases illustrated by Figure 2(a) and (b), with $w' > c$ and $w' = c$. It is clear that $w' < c$ cannot occur, because $u(x)$ is constant for $x \leq c$ (see Figure 2).

Case 1. Suppose that $w' > c$. We have

$$u(z) = p(z)z + (1 - p(z))\delta \int_0^{\bar{x}} u(\max\{z, x\}) dF(x) \geq p(z)z + (1 - p(z))\delta u(z),$$

where the equality is by (34) and (36), and the inequality is by (40). Solving the above inequality for $u(z)$ yields

$$u(z) \geq \frac{p(z)z}{1 - \delta(1 - p(z))} = \kappa_p(z)z = U_p(F_{(z,1)}, z), \quad (41)$$

where the first equality is by (9) and the second equality is by Lemma 1 (see Appendix A.1). We analogously prove that

$$u(w) \geq U_p(F_{(w,1)}, w). \quad (42)$$

We thus obtain

$$\begin{aligned} U_p(F, w) &= p(w)w + (1 - p(w))\delta \int_0^{\bar{x}} U_p(F, \max\{x, x'\}) dF(x') \\ &\geq p(w)w + (1 - p(w))\delta \int_0^{\bar{x}} u(x) dF(x) \\ &\geq p(w)w + (1 - p(w))\delta((1 - \sigma)u(w') + \sigma u(z)) \\ &\geq p(w)w + (1 - p(w))\delta((1 - \sigma)U_p(F_{(w',1)}, w') + \sigma U_p(F_{(z,1)}, z)) \\ &= (1 - \sigma)U_p(F_{(w',1)}, w) + \sigma U_p(F_{(z,1)}, w), \end{aligned} \quad (43)$$

where the first line is by (34), the second line is by (37), the third line is by (38), the fourth line is by (41) and (42), and the last line is again by (34). Alternatively, by (3), the assumption that $w < c < w' \leq z$, and (39),

$$V(F, w) = c = (1 - \sigma)\delta w' + \sigma\delta z = (1 - \sigma)V(F_{(w',1)}, w) + \sigma V(F_{(z,1)}, w). \quad (44)$$

It follows from (43) and (44) that

$$\begin{aligned} \frac{U_p(F, w)}{V(F, w)} &\geq \frac{(1 - \sigma)U_p(F_{(w',1)}, w) + \sigma U_p(F_{(z,1)}, w)}{(1 - \sigma)V(F_{(w',1)}, w) + \sigma V(F_{(z,1)}, w)} \\ &\geq \min \left\{ \frac{U_p(F_{(w',1)}, w)}{V(F_{(w',1)}, w)}, \frac{U_p(F_{(z,1)}, w)}{V(F_{(z,1)}, w)} \right\}. \end{aligned}$$

Thus, we have obtained (32) with $F_{(z,\sigma)}$ equal to either $F_{(w',1)}$ or $F_{(z,1)}$.

Case 2. Suppose that $w' = c$. By (39), we have $z > c/\delta > c$. As in Case 1, we obtain the inequality (41), so

$$u(z) \geq \kappa_p(z)z.$$

Also, by (36), we have $u(w') = U_p(F, w)$. So by (38),

$$\int_0^{\bar{x}} u(x) dF(x) \geq (1 - \sigma)u(w') + \sigma u(z) \geq (1 - \sigma)U_p(F, w) + \sigma \kappa_p(z)z. \quad (45)$$

Therefore, by (34), (36), and (45),

$$U_p(F, w) \geq p(w)w + (1 - p(w))\delta((1 - \sigma)U_p(F, w) + \sigma \kappa_p(z)z).$$

Solving the inequality for $U_p(F, w)$ yields

$$\begin{aligned} U_p(F, w) &\geq \frac{p(w)w + (1 - p(w))\delta\sigma\kappa_p(z)z}{1 - \delta(1 - \sigma)(1 - p(w))} \\ &= \frac{(1 - s)\kappa_p(w)w + (1 - \kappa_p(w))\kappa_p(z)sz}{1 - s\kappa_p(w)} \\ &= m_p(w, z, s) = U_p(F_{(z,\sigma)}, w), \end{aligned} \tag{46}$$

where the first equality is by substituting $\sigma = \frac{(1-\delta)s}{\delta(1-s)}$ and using the definition of κ_p in (9), the second equality is by the definition of m_p in (10), and the third equality is by Lemma 1 (see Appendix A.1). Alternatively, by (3), the assumption that $w \leq w' = c < z$, and (39),

$$V(F, w) = c = V(F_{(z,\sigma)}, w). \tag{47}$$

Hence, the second inequality in (32) follows from (46) and (47) with the specified $F_{(z,\sigma)}$. This completes the proof. \square

A.4 Proving Theorem 2 and Remark 1

PROOF OF THEOREM 2. We need to show that the decision rule p_g given by the constant stopping probability

$$p_g(y) = \frac{1 - \delta}{2 - \delta} \quad \text{for all } y \geq x_0$$

always yields a performance ratio of at least 1/4.

Using (9), we obtain $\kappa_{p_g}(x) = 1/2$ for all x . Inserting this into (10), we obtain for all $z > y \geq x$,

$$\frac{m_{p_g}(y, z, s)}{sz} = \frac{(1 - s)\frac{1}{2}y + \frac{1}{4}sz}{\left(1 - \frac{1}{2}s\right)sz} = \frac{2(1 - s)y + sz}{2(2 - s)sz} \geq \frac{1}{2(2 - s)} \geq \frac{1}{4}. \tag{48}$$

Then, by Proposition 3, $\kappa_{p_g}(y) = 1/2$, and (48) we obtain

$$r_{p_g}(y) \geq \min\left\{\frac{1}{2}, \frac{1}{4}\right\} = \frac{1}{4}. \tag{49}$$

Observe from (48) that the ratio $m_{p_g}(y, z, s)/(sz)$ is increasing in y for each (z, σ) such that $z \in X$, $z > x_0$, and $s \in (\frac{y}{z}, \delta]$. It follows that $r_{p_g}(y)$ is increasing, so p_g is regular. So $R_{p_g}(\mathcal{F}_X) = r_{p_g}(x_0)$ by Proposition 4 and $r_{p_g}(x_0) \geq 1/4$ by (49). \square

PROOF OF REMARK 1. Let $\bar{x} = \sup X = \infty$. As shown above, the rule p_g yields at least 1/4. We now show that no rule can achieve more than 1/4, thus proving that p_g is robust and yields $R_{p_g}(\mathcal{F}_X) = 1/4$.

Consider an arbitrary rule p and let κ_p be given by (9). Consider an increasing sequence $(z_n)_{n \in \mathbb{N}}$ such that $z_1 > x_0$ and $z_n \in X$ for all $n \in \mathbb{N}$, and $\lim_{n \rightarrow \infty} z_n = \infty$. Let $K_p = \limsup_{n \rightarrow \infty} \kappa_p(z_n)$. First, by Lemma 3 and the definitions of R_p and r_p , we have $R_p(\mathcal{F}_X) \leq r_p(z_n)$ for all $n \in \mathbb{N}$, so

$$R_p(\mathcal{F}_X) \leq \limsup_{n \rightarrow \infty} r_p(z_n).$$

Next, for each $y \geq x_0$ we have

$$\begin{aligned} r_p(y) &\leq \inf_{n \in \mathbb{N}} \left(\inf_{s \in (\frac{y}{z_n}, \delta]} \frac{m_{p_g}(y, z_n, s)}{sz_n} \right) \\ &= \inf_{n \in \mathbb{N}} \left(\inf_{s \in (\frac{y}{z_n}, \delta]} \frac{(1-s)\kappa_p(y)\frac{y}{z_n} + (1-\kappa_p(y))\kappa_p(z_n)s}{(1-s\kappa_p(y))s} \right) \\ &\leq \inf_{s \in (0, \delta]} \frac{(1-\kappa_p(y))K_p}{1-s\kappa_p(y)} \leq (1-\kappa_p(y))K_p, \end{aligned}$$

where the first inequality is by Proposition 3, the equality is by (10), the second inequality is by taking $n \rightarrow \infty$, and the third inequality is by taking $s \rightarrow 0$.

Finally, using the fact that $\kappa_p(y) \in [0, 1]$ for all y and, thus, $K_p \in [0, 1]$, we obtain

$$\lim_{n \rightarrow \infty} r_p(z_n) \leq \lim_{n \rightarrow \infty} (1-\kappa_p(z_n))K_p = (1-K_p)K_p \leq \frac{1}{4}.$$

We thus obtain that $R_p(\mathcal{F}_X) \leq 1/4$. □

A.5 Proof of Theorem 3

Let T_1 be a constant given by

$$T_1 = \left(\frac{(773 + 9\sqrt{1290})^{1/3}}{18} + \frac{79}{18(773 + 9\sqrt{1290})^{1/3}} - \frac{2}{9} \right)^2 \approx 0.6026. \tag{50}$$

Let λ_1 be a constant in $[0, 1]$ that satisfies the equation $T_1 = \eta(\lambda_1)$, so

$$\lambda_1 = \frac{(2T_1 - 1)^2}{T_1} \approx 0.0699 < 7/100. \tag{51}$$

Let $\bar{x} = \sup X < \infty$. By rescaling the values, we assume without loss of generality that $\bar{x} = 1$. Fix x_0 that satisfies assumption (A₁) and, in addition, $x_0 \geq \lambda_1$. So

$$0 < \lambda_1 \leq x_0 \leq \frac{\delta^2}{2-\delta} < 1. \tag{52}$$

Let

$$t = \eta(x_0) = \frac{1}{2} + \frac{1}{8}(x_0 + \sqrt{x_0(x_0 + 8)}). \tag{53}$$

We interpret t as the target performance ratio for outside option x_0 . Consider the decision rule p_g^* given for each $y \geq x_0$ by

$$p_g^*(y) = f(y, t) = \begin{cases} \frac{(1 - \delta)(1 - t)}{(1 - \delta)(1 - t) + (\sqrt{t} - \sqrt{y})^2} & \text{if } y < t, \\ 1 & \text{if } y \geq t. \end{cases}$$

Because $R^*(\mathcal{B}_X) = \eta(x_0)$ by Theorem 1 and $\mathcal{B}_X \subset \mathcal{F}_X$, we have

$$R_{p_g^*}(\mathcal{F}_X) \leq R^*(\mathcal{F}_X) \leq R^*(\mathcal{B}_X) = \eta(x_0).$$

So to prove Theorem 3, it suffices to show that $R_{p_g^*}(\mathcal{F}_X) \geq \eta(x_0)$ whenever $x_0 \geq \lambda_1$. For this purpose, we show that p_g^* is regular, and that $r_{p_g^*}(x_0) \geq \eta(x_0)$. Then we apply Proposition 4 to obtain that $R_{p_g^*}(\mathcal{F}_X) = r_{p_g^*}(x_0) \geq \eta(x_0)$.

Because the rule p_g^* is fixed, we omit the subscript referring to p_g^* in some notation. In particular, we write $r(y)$ for $r_{p_g^*}(y)$ and $\kappa(y)$ for $\kappa_{p_g^*}(y)$. By (9) we have

$$\kappa(y) = \frac{p_g^*(y)}{1 - \delta + \delta p_g^*(y)} = \begin{cases} \frac{1 - t}{1 - t + (\sqrt{t} - \sqrt{y})^2} & \text{if } y < t, \\ 1 & \text{if } y \geq t. \end{cases} \tag{54}$$

Let us introduce the notation

$$M(y, z) = \inf_{s \in (0,1]} \frac{m_{p_g^*}(y, z, s)}{sz} = \inf_{s \in (0,1]} \frac{(1 - s)\kappa(y)y + s(1 - \kappa(y))\kappa(z)z}{(1 - s\kappa(y))sz}.$$

By Lemma 2,

$$M(y, z) = \begin{cases} \frac{1}{z}(\kappa(y)\sqrt{y} + \sqrt{(1 - \kappa(y))(\kappa(z)z - \kappa(y)y)})^2 & \text{if } y < \kappa(z)z, \\ \kappa(z) & \text{if } y \geq \kappa(z)z. \end{cases} \tag{55}$$

The rest of the proof is divided into four steps:

Step (i) We show that $r(y) = \inf_{z \in (y,1]} M(y, z)$.

Step (ii) The term $M(y, z)$ is increasing in y for each $z \in (y, 1]$.

Step (iii) If $x_0 \geq \lambda_1$, then $M(x_0, z) \geq \eta(x_0)$ for each $z \in (x_0, 1]$.

Step (iv) If $x_0 = 1/90$, then $M(x_0, z) < \eta(x_0)$ for some $z \in (x_0, 1]$.

Steps (i) and (ii) imply that $r(y)$ is increasing in y , so rule p_g^* is regular. Steps (i) and (iii) imply that $r(x_0) \geq \eta(x_0)$ whenever $x_0 \geq \lambda_1$. Step (iv) implies that the tight lower bound λ for the result $r(x_0) \geq \eta(x_0)$ whenever $x_0 \geq \lambda$ must satisfy $1/90 < \lambda \leq \lambda_1$. We now prove Steps (i)–(iv).

Proof of Step (i). By Proposition 3,

$$r(y) = \min \left\{ \kappa(y), \inf_{z \in X, s \in (\frac{y}{z}, 1]} \frac{m_{p_g^*}(y, z, s)}{sz} \right\} \geq \min \left\{ \kappa(y), \inf_{z \in (y, 1]} M(y, z) \right\}.$$

However, by (55), we have

$$\lim_{z \rightarrow y} M(y, z) = \kappa(y).$$

We thus obtain $r(y) = \inf_{z \in (y, 1]} M(y, z)$.

Proof of Step (ii). We now show that $M(y, z)$ is weakly increasing in y for $x_0 \leq y \leq z$. If $t \leq y \leq z$, then by (54), we have $\kappa(y) = \kappa(z) = 1$, so $M(y, z) = y/z$ is increasing in y . If $\kappa(z)z \leq y \leq z$, then $M(y, z) = \kappa(z)$ is constant in y . It remains to show that $M(y, z)$ is increasing in y for $x_0 \leq y < \min\{t, \kappa(z)z\}$. As follows from (55), it suffices to show that

$$\Phi(y) = \kappa(y)\sqrt{y} + \sqrt{(1 - \kappa(y))(\kappa(z)z - \kappa(y)y)} \tag{56}$$

is increasing in y . Because $y < t$, by (54), we have

$$\kappa(y) = \frac{1 - t}{1 - t + (\sqrt{t} - \sqrt{y})^2} = \frac{1 - t}{1 + y - 2\sqrt{yt}}. \tag{57}$$

Substituting $\kappa(y)$ into (56) and simplifying the expression yields

$$\Phi(y) = \frac{(\sqrt{y} - \sqrt{t})\psi(y, L) - (1 - t)\sqrt{y}}{1 + y - 2\sqrt{yt}},$$

where

$$L = \kappa(z)z \quad \text{and} \quad \psi(y, L) = \sqrt{(1 + y - 2\sqrt{yt})L - (1 - t)y}.$$

Let us take the derivative of $\Phi(y)$. After simplification we obtain

$$\Phi'(y) = \frac{(1 - t)((1 - y)\psi(y, L) - (1 + y - 2\sqrt{yt})L + 2y - (1 + y)\sqrt{yt})}{2\psi(y, L)(1 + y - 2\sqrt{yt})^2\sqrt{y}}. \tag{58}$$

To show that $\Phi'(y) \geq 0$, we fix y and evaluate $\Phi'(y)$ for two extreme values of L . By assumption, $L = \kappa(z)z$ and $y < \kappa(z)z \leq 1$. When $L = 1$, we have $\psi(y, 1) = 1 - \sqrt{yt}$, and it can be easily verified that $\Phi'(y) = 0$ in this case. When $L = y$, we have $\psi(y, y) = \sqrt{yt} - y$, and it can be easily verified that $\Phi'(y) = 0$ in this case as well. Moreover, the denominator in (58) is strictly positive, as $0 < x_0 \leq y < t < 1$, so

$$\psi(y, L) \geq \psi(y, y) = \sqrt{yt} - y > 0 \quad \text{and} \quad 1 + y - 2\sqrt{yt} > 1 + y - 2y = 1 - y > 0.$$

Finally, because $\psi(y, L)$ is concave in L , the numerator in (58) is concave in L . It follows that $\Phi'(y) \geq 0$ for each $L \in [y, 1]$. We thus conclude that $\Phi(y)$ is weakly increasing in y when $x_0 \leq y < \min\{t, \kappa(z)z\}$.

Proof of Step (iii). We now show that $M(x_0, z) \geq t$ for all $z \in (x_0, 1]$. First, by (53), t is increasing in x_0 and satisfies $\eta(x_0) \geq 1/2$, so, solving (53) for x_0 yields $x_0 = (2t - 1)^2/t$. We thus obtain

$$\kappa(x_0) = \kappa\left(\frac{(2t - 1)^2}{t}\right) = \frac{1 - t}{1 - t + \left(\sqrt{t} - \frac{2t - 1}{\sqrt{t}}\right)^2} = \frac{1 - t}{1 - t + \frac{(1 - t)^2}{t}} = t. \tag{59}$$

Substituting $y = x_0 = (2t - 1)^2/t$ and $\kappa(x_0) = t$ into (55), we obtain for each $z \in (x_0, 1]$

$$M(x_0, z) = \begin{cases} \frac{1}{z}((2t - 1)\sqrt{t} + \sqrt{(1 - t)(\kappa(z)z - (2t - 1)^2)})^2 & \text{if } \kappa(z)z > x_0, \\ \kappa(z) & \text{if } \kappa(z)z \leq x_0. \end{cases} \quad (60)$$

Consider first the interval $\{z \in (x_0, 1] : \kappa(z)z \leq x_0\}$. When z belongs to this interval, we have $M(x_0, z) = \kappa(z) \geq \kappa(x_0) = t$.

Consider now the interval $\{z \in (x_0, 1] : \kappa(z)z > x_0\}$. When z belongs to this interval, we have, by (60),

$$\begin{aligned} M(x_0, z) \geq t &\iff ((2t - 1)\sqrt{t} + \sqrt{(1 - t)(\kappa(z)z - (2t - 1)^2)})^2 \geq tz \\ &\iff (2t - 1)\sqrt{t} + \sqrt{(1 - t)(\kappa(z)z - (2t - 1)^2)} \geq \sqrt{tz} \\ &\iff \sqrt{(1 - t)(\kappa(z)z - (2t - 1)^2)} \geq \sqrt{t}(\sqrt{z} - (2t - 1)) \\ &\iff (1 - t)(\kappa(z)z - (2t - 1)^2) \geq t(\sqrt{z} - (2t - 1))^2. \end{aligned}$$

Let us change variable $z = g^2$ and let

$$W(g) = (1 - t)(\kappa(g^2)g^2 - (2t - 1)^2) - t(g - (2t - 1))^2. \quad (61)$$

So $M(x_0, z) \geq t$ if and only if $W(g) \geq 0$.

Suppose that $\kappa(g^2)g^2 > x_0$ and $g^2 \geq t$, so $\kappa(g^2) = 1$. In this case,

$$\begin{aligned} W(g) &= (1 - t)(g^2 - (2t - 1)^2) - t(g - (2t - 1))^2 \\ &= (g - (2t - 1))((1 - t)(g + (2t - 1)) - t(g - (2t - 1))) \\ &= (g - (2t - 1))(2t - 1)(1 - g) \geq 0, \end{aligned}$$

because $g - (2t - 1) \geq g - (2g^2 - 1) = (1 + 2g)(1 - g) \geq 0$ and $t \geq 1/2$.

Last, suppose that $\kappa(g^2)g^2 > x_0$ and $g^2 < t$. Substituting $\kappa(g^2)$ given by (57) into (61) and collecting the coefficients with respect to variable g , we obtain

$$W(g) = \frac{\tilde{W}(g)\kappa(g^2)}{1 - t},$$

where

$$\begin{aligned} \tilde{W}(g) &= -tg^4 + (2t^{\frac{3}{2}} + 4t^2 - 2t)g^3 + (-8t^{\frac{5}{2}} + 4t^{\frac{3}{2}} - 3t^2 + t)g^2 \\ &\quad + (8t^{\frac{5}{2}} - 8t^{\frac{3}{2}} + 2t^{\frac{1}{2}} + 4t^2 - 2t)g + (2t - 1)^2. \end{aligned}$$

Note that $1 - t > 0$, because $t = \eta(x_0)$ and by (52), we have $x_0 < 1$. Also, $\kappa(g^2) \geq 1/2 > 0$. So $W(g) \geq 0$ if and only if $\tilde{W}(g) \geq 0$. We have

$$\frac{d^2\tilde{W}(g)}{dg^2} = -12tg^2 + (12t^{\frac{3}{2}} + 24t^2 - 12t)g - 16t^{\frac{5}{2}} + 8t^{\frac{3}{2}} - 6t^2 + 2t.$$

This expression is quadratic and concave in g . So it is globally nonpositive if the discriminant of the quadratic equation is nonpositive:

$$-48t^2(15t + 4t^{\frac{3}{2}} - 2t^{\frac{1}{2}} - 12t^2 - 5) \leq 0. \tag{62}$$

Inequality (62) has three roots 0, 1, and T_1 , where constant T_1 is given by (50), $T_1 \approx 0.6026$. Moreover, on the interval of $t \in [1/2, 1)$, inequality (62) holds for $t \in [T_1, 1)$ and does not hold for $t \in [1/2, T_1)$. Recall that by (50) and (53), $t \geq T_1$ if and only if $x_0 \geq \lambda_1$. We thus conclude that if $x_0 \geq \lambda_1$, then $\tilde{W}(g)$ is concave on the interval of g that satisfies $\kappa(g^2)g^2 > x_0$ and $g^2 < t$.

To complete the proof of Step (iii), observe that the concavity of $\tilde{W}(g)$ on the interval of g that satisfies $\kappa(g^2)g^2 > x_0$ and $g^2 < t$ implies quasiconcavity of $M(x_0, z)$ on the interval of z that satisfies $\kappa(z)z > x_0$ and $z < t$. Also, it is easy to verify that $M(x_0, z)$ is continuous in z . Finally, we have already obtained that $M(x_0, z) \geq t$ at the boundaries, when $z = t$ and when z satisfies $\kappa(z)z = x_0$. Consequently, we obtain that $M(x_0, z) \geq t$ on the interval of z that satisfies $\kappa(z)z > x_0$ and $z < t$.

Proof of Step (iv). Let $x_0 = 1/90$ and $z = 1/10$. Inserting these values into (9) and (53), we obtain $\kappa(z)$ and $t = \eta(x_0)$. Then inserting these into (60), we obtain that $M(x_0, z) < \eta(x_0)$.

This completes the proof of Theorem A.5.

REMARK 3. We numerically find the tight lower bound λ for the result $r(x_0/\bar{x}) \geq \eta(x_0/\bar{x})$ whenever $x_0/\bar{x} \geq \lambda$. This bound is $\lambda \approx 0.01120000$, so $1/90 < \lambda < 1/89$. To obtain λ , we make the following adjustment in Step (iii). We numerically evaluate $\tilde{W}(g)$ with precision 10^{-8} to show that $\tilde{W}(g) \geq 0$ on the interval of g that satisfies $\kappa(g^2)g^2 > x_0$ and $g^2 < t$ if and only if $t \geq T \approx 0.53884276$. Because t satisfies $t = \eta(x_0/\bar{x})$, the constraint $t \geq T$ is equivalent to $x_0/\bar{x} \geq \lambda \approx 0.01120000$, where λ is obtained from the equation $T = \eta(\lambda)$.

A.6 Proof of Proposition 2

Let $p \in \mathcal{P}$. Recall that $F_{(z,\sigma)} \in \mathcal{B}_X$ is a lottery over 0 and z with probabilities $1 - \sigma$ and σ , respectively, where $z > x_0$ and $\sigma \in [0, 1]$. Because there can be at most one value above x_0 , we have $R_p(\mathcal{B}_X) \leq R_{\bar{p}}(\mathcal{B}_X)$, where $\bar{p} \in \mathcal{P}$ is given by

$$\bar{p}(h_t) = \begin{cases} p(h_t) & \text{if } x_1 = \dots = x_t = 0, \\ 1 & \text{otherwise.} \end{cases} \tag{63}$$

In what follows, we consider arbitrary sequences $q = (q_0, q_1, \dots)$, where $q_t = p(h_t)$ for $h_t = (x_0, 0, \dots, 0)$.

Let $F_{(z,\sigma)} \in \mathcal{B}_X$. Let $U_q(F_{(z,\sigma)}, t)$ be the individual's expected payoff in round t after history $h_t = (x_0, 0, \dots, 0)$, when playing sequence q and facing $F_{(z,\sigma)}$. It is given by

$$U_q(F_{(z,\sigma)}, t) = q_t x_0 + (1 - q_t) \delta(\sigma z + (1 - \sigma) U_q(F_{(z,\sigma)}, t + 1)). \tag{64}$$

Let $W_q(F_{(z,\sigma)})$ be the worst expected payoff across all rounds when facing $F_{(z,\sigma)}$, so

$$W_q(F_{(z,\sigma)}) = \inf_{t=0,1,\dots} U_q(F_{(z,\sigma)}, t). \tag{65}$$

Let q' be an arbitrary sequence of probabilities. This q' is called a *benchmark* and is fixed for the rest of the proof. Let \bar{q}^∞ be a sequence with a constant stopping probability $\bar{q} \in [0, 1]$, so $\bar{q}^\infty = (\bar{q}, \bar{q}, \dots)$. We now show that there exists $\bar{q} \in [0, 1]$ such that the constant sequence \bar{q}^∞ has a weakly higher performance ratio in binary environments than q' . Note that we need to compare only the worst expected payoffs $W_{q'}$ and $W_{\bar{q}^\infty}$, as the performance ratio is determined by the worst payoff ratio across all histories, and the optimal payoff V does not depend on the decision rule or history.

Let F_0 be the environment that generates 0 with certainty, so $F_0 = F_{(z,0)}$ for any z . Let \bar{q} be a solution of the equation

$$U_{q'}(F_0, 0) = \bar{q}x_0 + (1 - \bar{q})\delta U_{q'}(F_0, 0), \tag{66}$$

so

$$\bar{q} = \frac{(1 - \delta)U_{q'}(F_0, 0)}{x_0 - \delta U_{q'}(F_0, 0)}.$$

By (64), $U_{q'}(F_0, 0) \in [0, x_0]$, so $\bar{q} \in [0, 1]$. Let us show that for all $F_{(z,\sigma)} \in \mathcal{B}_X$, we have

$$W_{\bar{q}^\infty}(F_{(z,\sigma)}) \geq W_{q'}(F_{(z,\sigma)}). \tag{67}$$

Let $F_{(z,\sigma)} \in \mathcal{B}_X$. By (64), observe that for any sequence $q = (q_0, q_1, \dots)$ and any t ,

$$\begin{aligned} &U_q(F_{(z,\sigma)}, t) - x_0 \\ &= (1 - q_t)(\delta\sigma z + \delta(1 - \sigma)U_q(F_{(z,\sigma)}, t + 1) - x_0) \\ &= (1 - q_t)(\delta\sigma z - (1 - \delta(1 - \sigma))x_0 + \delta(1 - \sigma)(U_q(F_{(z,\sigma)}, t + 1) - x_0)). \end{aligned}$$

Iterating the above expression for $t + 1, t + 2, \dots$, we obtain

$$U_q(F_{(z,\sigma)}, t) - x_0 = (\delta\sigma z - (1 - \delta(1 - \sigma))x_0) \sum_{k=0}^{\infty} \left(\delta^k (1 - \sigma)^k \prod_{s=t}^{t+k} (1 - q_s) \right). \tag{68}$$

Suppose that $\delta\sigma z - (1 - \delta(1 - \sigma))x_0 = 0$. Then $U_q(F_{(z,\sigma)}, t) - x_0 = 0$ for every q and every t . In particular, $W_{\bar{q}^\infty}(F_{(z,\sigma)}) = W_{q'}(F_{(z,\sigma)}) = x_0$. So (67) holds in this case.

Next suppose that $\delta\sigma z - (1 - \delta(1 - \sigma))x_0 \neq 0$. Define

$$\psi_t^\sigma(q) = \frac{U_q(F_{(z,\sigma)}, t) - x_0}{\delta\sigma z - (1 - \delta(1 - \sigma))x_0}. \tag{69}$$

So, by (68),

$$\psi_t^\sigma(q) = \sum_{k=0}^{\infty} \left(\delta^k (1 - \sigma)^k \prod_{s=t}^{t+k} (1 - q_s) \right). \tag{70}$$

If q is a constant sequence so that $q = (q_0, q_0, \dots)$, then $\psi_t^\sigma(q)$ is constant in t , so we can omit the subscript t . It is given by

$$\psi^\sigma(q) = \sum_{k=0}^{\infty} \delta^k (1 - \sigma)^k (1 - q_0)^{k+1} = \frac{1 - q_0}{1 - \delta(1 - \sigma)(1 - q_0)}. \tag{71}$$

By (65) and (69), when $\delta\sigma z - (1 - \delta(1 - \sigma))x_0 > 0$, we have

$$W_{\bar{q}^\infty}(F_{(z,\sigma)}) \geq W_{q'}(F_{(z,\sigma)}) \iff \psi^\sigma(\bar{q}^\infty) \geq \inf_t \psi_t^\sigma(q').$$

Similarly, when $\delta\sigma z - (1 - \delta(1 - \sigma))x_0 < 0$, we have

$$W_{\bar{q}^\infty}(F_{(z,\sigma)}) \geq W_{q'}(F_{(z,\sigma)}) \iff -\psi^\sigma(\bar{q}^\infty) \geq \inf_t (-\psi_t^\sigma(q')).$$

Therefore, to prove (67) for all $F_{(z,\sigma)} \in \mathcal{B}_X$, it remains to show that for each $\sigma \in [0, 1]$,

$$\inf_t \psi_t^\sigma(q') \leq \psi^\sigma(\bar{q}^\infty) \leq \sup_t \psi_t^\sigma(q'). \tag{72}$$

Fix $\sigma \in [0, 1]$. To prove (72), we first find the upper and lower bounds on the values of $\psi_0^0(q)$ achievable by choosing a sequence q subject to the constraint

$$\inf_t \psi_t^\sigma(q') \leq \psi_s^\sigma(q) \leq \sup_t \psi_t^\sigma(q') \quad \text{for all } s = 0, 1, 2, \dots \tag{73}$$

To do this, we solve

$$\min_q \psi_0^0(q) \quad \text{subject to (73),} \tag{74}$$

$$\max_q \psi_0^0(q) \quad \text{subject to (73).} \tag{75}$$

LEMMA 5. *There exist a solution q_{\min}^σ of (74) and a solution q_{\max}^σ of (75) that are constant sequences.*

We postpone the proof of this lemma to the end of this section and first complete the proof of Proposition 2.

Let q_{\min}^σ and q_{\max}^σ be the constant sequences given by Lemma 5. This means that $\psi^0(q_{\min}^\sigma) \leq \psi_0^0(q) \leq \psi^0(q_{\max}^\sigma)$ for any q that satisfies (73). It is easy to see from (66) and (69) that \bar{q}^∞ satisfies (73). We thus obtain

$$\psi^0(q_{\min}^\sigma) \leq \psi^0(\bar{q}^\infty) \leq \psi^0(q_{\max}^\sigma). \tag{76}$$

Next, by (71), for any constant sequence $q = (q_0, q_0, \dots)$, both $\psi^0(q)$ and $\psi^\sigma(q)$ are strictly decreasing in q_0 . This implies $q_{\min}^\sigma \geq \bar{q}^\infty \geq q_{\max}^\sigma$, which in turn implies $\psi^\sigma(q_{\min}^\sigma) \leq \psi^\sigma(\bar{q}^\infty) \leq \psi^\sigma(q_{\max}^\sigma)$ for $\sigma > 0$. Moreover, because q_{\min}^σ and q_{\max}^σ satisfy the constraint (73), we obtain

$$\inf_t \psi_t^\sigma(q') \leq \psi^\sigma(q_{\min}^\sigma) \leq \psi^\sigma(\bar{q}^\infty) \leq \psi^\sigma(q_{\max}^\sigma) \leq \sup_t \psi_t^\sigma(q').$$

So (72) holds. This completes the proof of Proposition 2.

PROOF OF LEMMA 5. We prove that a solution q_{\max}^σ of the maximization problem (75) is a constant sequence. The proof that q_{\min}^σ is a constant sequence is analogous and, thus, is omitted.

Fix $\sigma \in [0, 1]$. We use the notation

$$\underline{\psi}^\sigma(q') = \inf_t \psi_t^\sigma(q') \quad \text{and} \quad \bar{\psi}^\sigma(q') = \sup_t \psi_t^\sigma(q').$$

Let \tilde{q} be the solution of the equation

$$\bar{\psi}^\sigma(q') = (1 - \tilde{q})(1 + \delta(1 - \sigma)\bar{\psi}^\sigma(q')). \tag{77}$$

We now show that the constant sequence $\tilde{q}^\infty = (\tilde{q}, \tilde{q}, \dots)$ is a solution of the maximization problem (75). To prove this, we solve a finite-horizon problem described below. We assume that the individual makes decisions in rounds $t = 0, 1, \dots, T$, after which the individual's behavior is fixed by $q_t = \tilde{q}$ for all $t > T$. Because the maximal value of $\psi_0^0(q)$ in the problem (75) can differ from that in the problem with horizon T by at most δ^T , we find the solution to the infinite-horizon problem (75) as the limit of the solutions to the finite-horizon problem as $T \rightarrow \infty$.

Fix $T \in \mathbb{N}$ and consider the problem

$$\begin{aligned} & \max_q \psi_0^0(q) \quad \text{subject to} \\ & \underline{\psi}^\sigma(q') \leq \psi_t^\sigma(q) \leq \bar{\psi}^\sigma(q') \quad \text{for all } t = 0, 1, 2, \dots, \\ & q_t = \tilde{q} \quad \text{for all } t = T + 1, T + 2, \dots \end{aligned} \tag{78}$$

We now show that \tilde{q}^∞ is a solution of (78). We proceed by backward induction, starting from round $k = T$, and then continuing to rounds $k = T - 1$, etc.

Let $k \in \{0, 1, \dots, T\}$ and suppose $q_t = \tilde{q}$ for each $t > k$. Using (70), rewrite $\psi_k^\sigma(q)$ as

$$\psi_k^\sigma(q) = (1 - q_k)(1 + \delta(1 - \sigma)\psi_{k+1}^\sigma(q)). \tag{79}$$

By (71), for all $t > k$,

$$\psi_t^0(q) = \frac{1 - \tilde{q}}{1 - \delta(1 - \tilde{q})} \quad \text{and} \quad \psi_t^\sigma(q) = \frac{1 - \tilde{q}}{1 - \delta(1 - \sigma)(1 - \tilde{q})} = \bar{\psi}^\sigma(q'), \tag{80}$$

where the last equality is by the definition of \tilde{q} in (77). So, by (79), (80), and the constraint in (78), we obtain that

$$\psi_k^\sigma(q) = (1 - q_k)(1 + \delta(1 - \sigma)\bar{\psi}^\sigma(q')) \leq \bar{\psi}^\sigma(q'). \tag{81}$$

This implies by (77) that

$$q_k \geq \tilde{q}. \tag{82}$$

So if $q_k \neq \tilde{q}$, it must be the case that $q_k > \tilde{q}$. We now prove that if $q_k > \tilde{q}$, then $\psi_0^0(q)$ can be weakly increased by reducing q_k .

Suppose that $q_k > \tilde{q}$. First, we deal with the case of $k > 0$. To show that $\psi_0^0(q)$ can be increased by reducing q_k , we keep q_t fixed for all t different from $k - 1$ and k , and vary q_{k-1} and q_k such that $\psi_{k-1}^\sigma(q)$ remains constant; that is,

$$d\psi_{k-1}^\sigma(q) = -(1 + \delta(1 - \sigma)\psi_k^\sigma(q))dq_{k-1} + (1 - q_{k-1})\delta(1 - \sigma)\frac{\partial\psi_k^\sigma(q)}{\partial q_k}dq_k = 0. \quad (83)$$

By (80) and (81), we have

$$\begin{aligned} \psi_k^\sigma(q) &= (1 - q_k)(1 + \delta(1 - \sigma)\bar{\psi}^\sigma(q')) \\ &= (1 - q_k)\left(1 + \delta(1 - \sigma)\frac{1 - \tilde{q}}{1 - \delta(1 - \sigma)(1 - \tilde{q})}\right) = \frac{1 - q_k}{1 - \delta(1 - \sigma)(1 - \tilde{q})}. \end{aligned}$$

So

$$\frac{\partial\psi_k^\sigma(q)}{\partial q_k} = -\frac{1}{1 - \delta(1 - \sigma)(1 - \tilde{q})}.$$

Thus, from (83), we obtain

$$\frac{dq_{k-1}}{dq_k} = -\frac{\delta(1 - \sigma)(1 - q_{k-1})}{1 - \delta(1 - \sigma)(q_k - \tilde{q})}.$$

Inserting $\sigma = 0$ into (70) and (80), by the induction assumption that $q_{k+1} = \tilde{q}$,

$$\psi_k^0(q) = (1 - q_k)(1 + \delta\psi_{k+1}^0(q)) = (1 - q_k)\left(1 + \frac{\delta(1 - \tilde{q})}{1 - \delta(1 - \tilde{q})}\right) = \frac{1 - q_k}{1 - \delta(1 - \tilde{q})}$$

and

$$\psi_{k-1}^0(q) = (1 - q_{k-1})(1 + \delta\psi_k^0(q)) = (1 - q_{k-1})\frac{1 - \delta(q_k - \tilde{q})}{1 - \delta(1 - \tilde{q})}.$$

Thus, by (70) with $\sigma = 0$,

$$\frac{\partial\psi_0^0(q)}{\partial q_k} = \delta^k \left(\prod_{s=0}^{k-1} (1 - q_s) \right) \frac{\partial\psi_k^0(q)}{\partial q_k} = -\delta^k \left(\prod_{s=0}^{k-1} (1 - q_s) \right) \frac{1}{1 - \delta(1 - \tilde{q})}$$

and

$$\begin{aligned} \frac{\partial\psi_0^0(q)}{\partial q_{k-1}} &= \delta^{k-1} \left(\prod_{s=0}^{k-2} (1 - q_s) \right) \frac{\partial\psi_{k-1}^0(q)}{\partial q_{k-1}} = -\delta^{k-1} \left(\prod_{s=0}^{k-2} (1 - q_s) \right) \frac{1 - \delta(q_k - \tilde{q})}{1 - \delta(1 - \tilde{q})} \\ &= -\delta^k \left(\prod_{s=0}^{k-1} (1 - q_s) \right) \frac{1 - \delta(q_k - \tilde{q})}{(1 - \delta(1 - \tilde{q}))\delta(1 - q_{k-1})}. \end{aligned}$$

Therefore, if $q_{k-1} < 1$, then

$$\frac{d\psi_0^0(q)}{dq_k} = \frac{\partial\psi_0^0(q)}{\partial q_k} + \frac{\partial\psi_0^0(q)}{\partial q_{k-1}} \frac{dq_{k-1}}{dq_k}$$

$$\begin{aligned}
 &= -\delta^k \left(\prod_{s=0}^{k-1} (1 - q_s) \right) \left(\frac{1}{1 - \delta(1 - \tilde{q})} + \frac{1 - \delta(q_k - \tilde{q})}{(1 - \delta(1 - \tilde{q}))\delta(1 - q_{k-1})} \frac{dq_{k-1}}{dq_k} \right) \\
 &= -\frac{\delta^k}{1 - \delta(1 - \tilde{q})} \left(\prod_{s=0}^{k-1} (1 - q_s) \right) \left(1 - \frac{(1 - \delta(q_k - \tilde{q}))(1 - \sigma)}{1 - \delta(1 - \sigma)(q_k - \tilde{q})} \right) \\
 &= -\frac{\delta^k}{1 - \delta(1 - \tilde{q})} \left(\prod_{s=0}^{k-1} (1 - q_s) \right) \frac{\sigma}{1 - \delta(1 - \sigma)(q_k - \tilde{q})} \leq 0.
 \end{aligned}$$

Alternatively, if $q_{k-1} = 1$, then $\psi_0^0(q)$ is independent of q_k , so $d\psi_0^0(q)/dq_k = 0$. Thus, if $q_k > \tilde{q}$, then decreasing q_k increases $\psi_0^0(q)$ without violating the constraint in (78), as long as $q_k \geq \tilde{q}$.

Next, we deal with the case of $k = 0$. By (70) and (80), we have

$$\frac{d\psi_0^0(q)}{dq_0} = -1 - \delta\psi_1^0(q) < 0.$$

So again, if $q_0 > \tilde{q}$, then decreasing q_0 increases $\psi_0^0(q)$ without violating the constraint in (78) as long as $q_0 \geq \tilde{q}$.

We thus proved that if q is a solution of (78) with $q_k > \tilde{q}$ and $q_t = \tilde{q}$ for all $t > k$, then there exists a solution with $q_t = \tilde{q}$ for all $t \geq k$. As this is true for each $k = T, T - 1, \dots, 1, 0$ by induction, we obtain that \tilde{q}^∞ is a solution of (78), so $q_{\max}^\sigma = \tilde{q}^\infty$. \square

A.7 Proof of Proposition 1

The equality $R_{p_1}(\mathcal{F}) = x_0/(\delta\bar{x})$ follows from the obvious fact that the worst-case environment for the rule p_1 is $F_{\bar{x}}$, with $U_{p_1}(F_{\bar{x}}, x_0) = x_0$ and $V(F_{\bar{x}}, x_0) = \delta\bar{x}$.

Let p be an arbitrary deterministic rule. Let us show that $R_p(\mathcal{F}) \leq R_{p_1}(\mathcal{F})$. Let h_t^0 be the history of x_0 followed by t zeros, so $h_t^0 = (x_0, 0, 0, \dots, 0)$. Suppose that there exists $t \in \{0, 1, 2, \dots\}$ such that p stops searching after t zero-valued alternatives, so $p(h_t^0) = 1$. In the environment $F_{\bar{x}}$, the payoff of rule p in round t is $U_p(F_{\bar{x}}, h_t^0) = x_0$ and the optimal payoff in round t is $V(F_{\bar{x}}, h_t^0) = \delta\bar{x}$. Moreover, by (3), $V(F, h_0) = V(F, h_t^0)$. Consequently,

$$R_p(\mathcal{F}) \leq \inf_{F \in \mathcal{F}} \frac{U_p(F_{\bar{x}}, h_t^0)}{V(F_{\bar{x}}, h_t^0)} = \frac{x_0}{\delta\bar{x}} = R_{p_1}(\mathcal{F}).$$

Alternatively, suppose that p never stops searching as long as only 0s occurred in the past, so $p(h_t^0) = 0$ for all t . Then, in the environment F_0 , the payoff of rule p in round 0 is $U_p(F_0, h_0) = 0$ and the optimal payoff in round 0 is $V(F_0, h_0) = x_0$. Consequently,

$$R_p(\mathcal{F}) \leq \frac{U_p(F_0, h_0)}{V(F_0, h_0)} = \frac{0}{x_0} = 0 < R_{p_1}(\mathcal{F}).$$

We thus conclude that $R_p(\mathcal{F}) \leq R_{p_1}(\mathcal{F})$ for each deterministic rule p .

Finally, we show that

$$R^*(\mathcal{F}) > \frac{x_0}{\delta \bar{x}}. \quad (84)$$

Suppose that $\bar{x} = \infty$, so $x_0/(\delta \bar{x}) = 0$. Then (84) trivially holds, because $R^*(\mathcal{F}) \geq 1/4$ by Theorem 2.

Next, suppose that $\bar{x} < \infty$. To simplify notation, let $x = x_0/\bar{x}$. Let $\delta_0(x)$ be the smallest $\delta \in (0, 1)$ that satisfies (A₁) for a given $x \in (0, 1)$. Hence $\delta_0(x)$ is the solution of the equation $x = \frac{\delta^2}{2-\delta}$, so

$$\delta_0(x) = \frac{\sqrt{x(x+8)} - x}{2} = \frac{(\sqrt{x(x+8)} - x)(\sqrt{x(x+8)} + x)}{2(\sqrt{x(x+8)} + x)} = \frac{4}{x + \sqrt{x(x+8)}}.$$

We thus obtain

$$\frac{x}{\delta} \leq \frac{x}{\delta_0(x)} = \frac{x(x + \sqrt{x(x+8)})}{4}. \quad (85)$$

Suppose that $x < 7/100$. Then substituting $x = 7/100$ into the right-hand side of (85) yields

$$\frac{x}{\delta} \leq \frac{x(x + \sqrt{x(x+8)})}{4} < \frac{14}{\sqrt{5649} - 7} < \frac{1}{4}.$$

So (84) holds, because $R^*(\mathcal{F}) \geq 1/4$ by Theorem 2.

Suppose now that $x \geq 7/100$. Recall that $x < 1$ by (A₁). By (85), we have

$$\frac{x}{\delta} \leq \frac{x(x + \sqrt{x(x+8)})}{4} < \frac{x + \sqrt{x(x+8)}}{4} \leq \frac{1}{2} \left(1 + \frac{x + \sqrt{x(x+8)}}{4} \right) = \eta(x).$$

So (84) holds, because $R^*(\mathcal{F}) \geq \eta(x)$ when $x \geq 7/100$ by Theorem 3.

REFERENCES

- Babaioff, Moshe, Michael Dinitz, Anupam Gupta, Nicole Immorlica, and Kunal Talwar (2009), "Secretary problems: Weights and discounts." In *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA'09*, 1245–1254, Philadelphia, PA, USA. [1433]
- Ben-Tal, Aharon, Laurent El Ghaoui, and Arkadi Nemirovski (2009), *Robust Optimization*. Princeton University Press. [1434]
- Bergemann, Dirk and Stephen Morris (2005), "Robust mechanism design." *Econometrica*, 73, 1771–1813. [1434]
- Bergemann, Dirk and Karl H. Schlag (2011)a, "Robust monopoly pricing." *Journal of Economic Theory*, 146, 2527–2543. [1433]

- Bergemann, Dirk and Karl H. Schlag (2011)b, “Should I stay or should I go? Search without priors.” Report. [1433, 1434]
- Carrasco, Vinicius, Vitor Farinha Luz, Nenad Kos, Matthias Messner, Paulo Monteiro, and Humberto Moreira (2018), “Optimal selling mechanisms under moment conditions.” *Journal of Economic Theory*, 177, 245–279. [1434]
- Carroll, Gabriel (2015), “Robustness and linear contracts.” *American Economic Review*, 105, 536–563. [1434]
- Chassang, Sylvain (2013), “Calibrated incentive contracts.” *Econometrica*, 81, 1935–1971. [1433, 1434]
- Ferguson, Thomas S. (1989), “Who solved the secretary problem?” *Statistical Science*, 4, 282–289. [1433]
- Fox, J. H. and L. G. Marnie (1960), “In Martin Gardner’s column: Mathematical games.” *Scientific American*, 202, 150–153. [1433]
- Gastwirth, Joseph L. (1976), “On probabilistic models of consumer search for information.” *The Quarterly Journal of Economics*, 90, 38–50. [1431]
- Gilboa, Itzhak and David Schmeidler (1989), “Maxmin expected utility with a non-unique prior.” *Journal of Mathematical Economics*, 18, 141–153. [1433]
- Hansen, Lars Peter, and Thomas J. Sargent (2001), “Robust control and model uncertainty.” *American Economic Review*, 91, 60–66. [1434]
- Huber, Peter J. (1964), “Robust estimation of a location parameter.” *The Annals of Mathematical Statistics*, 35, 73–101. [1433]
- Huber, Peter J. (1965), “A robust version of the probability ratio test.” *The Annals of Mathematical Statistics*, 36, 1753–1758. [1433]
- Kajii, Atsushi and Stephen Morris (1997), “The robustness of equilibria to incomplete information.” *Econometrica*, 65, 1283–1309. [1434]
- Kamenica, Emir and Matthew Gentzkow (2011), “Bayesian persuasion.” *American Economic Review*, 101, 2590–2615. [1455]
- Kasberger, Bernhard and Karl H. Schlag (2017), “Robust bidding in first-price auctions: How to bid without knowing what others are doing.” Report. [1434]
- Manski, Charles F. (2004), “Statistical treatment rules for heterogeneous populations.” *Econometrica*, 72, 1221–1246. [1434]
- Parakhonyak, Alexei and Anton Sobolev (2015), “Non-reservation price equilibrium and search without priors.” *Economic Journal*, 125, 887–909. [1433, 1434]
- Prasad, Kislaya (2003), “Non-robustness of some economic models.” *The B.E. Journal of Theoretical Economics*, 3. [1433]
- Savage, Leonard J. (1951), “The theory of statistical decision.” *Journal of the American Statistical Association*, 46, 55–67. [1433]

- Schlag, Karl H. (2006), “Eleven—tests needed for a recommendation.” Report. [1434]
- Schlag, Karl H. and Anton Sobolev (2020), “Robust search: How to learn from the past without priors?” Working paper. [1433, 1442]
- Sleator, Daniel D. and Robert E. Tarjan (1985), “Amortized efficiency of list update and paging rules.” *Communications of the ACM*, 28, 202–208. [1433]
- Stoye, Jörg (2009), “Minimax regret treatment choice with finite samples.” *Journal of Econometrics*, 151, 70–81. [1434]
- Terlizzese, Daniele (2080), “Relative minimax.” EIEF Working Papers Series 0804. [1433]
- Wald, Abraham (1950), *Statistical Decision Functions*. Wiley, New York. [1433]
- Weitzman, Martin L. (1979), “Optimal search for the best alternative.” *Econometrica*, 47, 641–654. [1436]
- Zhou, Kemin, John C. Doyle, and Keith Glover (1995), *Robust and Optimal Control*. Prentice Hall. [1434]

Co-editor Ran Spiegler handled this manuscript.

Manuscript received 22 October, 2019; final version accepted 1 December, 2020; available online 6 December, 2020.