Abstract

Thesis title: Sequential Search With Ordinal Ranks and Cardinal Values: An Infinite Discounted Secretary Problem

Asa Benjamin Palley, Master of Science, 2009

Advised by: Professor Peter Cramton
Department of Economics

We consider an extension of the classical secretary problem where a decision maker observes only the relative ranks of a sequence of up to N applicants, whose true values are i.i.d. $U[0, 1]$ random variables. Applicants arrive according to a homogeneous Poisson Process, and the decision maker seeks to maximize the expected time-discounted value of the applicant who she ultimately selects. This provides a straightforward and natural objective while retaining the structure of limited information based on relative ranks. We derive the optimal policy in the sequential search, and show that the solution converges as $N \to \infty$. We compare these results with a closely related full information problem in order to quantify these informational limitations.

AMS classification: Primary 60G40; Secondary 90C39

Keywords: Sequential Search, Secretary Problem, Marriage Problem, Optimal Stopping, Dynamic Programming, Discounting, Relative Ranks

Sequential Search With Ordinal Ranks and Cardinal Values: An Infinite Discounted Secretary Problem

by

Asa Benjamin Palley

Thesis submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Master of Science
2009

Thesis Committee:

Professor Peter Cramton, Chair/Advisor
Professor Tobias von Petersdorff
Professor Subramanian Raghavan

Table of Contents

Introduction	1
The Model	5
The Optimal Policy	6
The Full Information Model	15
The Full Information Optimal Policy	16
Discussion	21
Appendix: Proofs	24
References	27

List of Figures

Figure 1	13
Figure 2	13
Figure 3	14
Figure 4	14
Figure 5	20
Figure 6	20
Figure 7	21

1 Introduction

The classical secretary problem (see Freeman (1983) and Ferguson (1989) for excellent reviews) considers a decision maker (DM) who observes a random sequence of N applicants for a job and wants to hire the very best one, assuming they can be strictly ordered and N is fixed and known ahead of time. After she evaluates each applicant, she must decide whether to end the search and accept the most recent applicant or continue the search and permanently reject him. She must make this decision based solely on the rank of the applicant relative to those already rejected, and seeks to adopt a decision rule to maximize the probability that she chooses the best of the N applicants. Gilbert and Mosteller (1966) show that the optimal policy for large N is to always skip the first N/e ($\sim 37\%$) of the applicants, and accept the next applicant who ranks higher than all previous applicants, a surprisingly concise and heuristic strategy.

This model provides a useful framework to examine the dynamics of sequential search. Many real-world decision makers face a similar situation where they initially have little or no information about their options, and gain useful context and experience as they search. Examples might include a venture capitalist evaluating investment proposals, a young couple looking to buy their first home, or an eBay bidder deciding which auction she should participate in. Lindley (1961) refers to the secretary problem as the marriage problem: A bachelorette goes on dates with potential mates and must decide whether to propose marriage or reject the current date and consider another. We assume that she cannot go back and propose to someone she has already rejected, and that marriage proposals are always accepted. She wants to choose the very best mate, but can only compare the current potential mate to ones she previously dated and rejected. So she may meet someone who is better than anyone she has ever dated before but worry that proposing to him will deny her the opportunity to marry someone even better later on.

Solving for the optimal decision rule also gives us a solid benchmark to analyze decision-making by experimental participants. Seale and Rapoport (1997, 2000) test stopping times under the classical secretary problem and find that participants may tend to stop too soon. They propose that this may be a result of endogenous search costs that affect the decision maker but are not explicitly accounted for in the experiment's payoffs. Stein et al. (2003) analyze classes of heuristic decision rules that real-world decision-makers might use to solve the secretary problem.

However, the classical formulation makes several unrealistic assumptions:

1. The decision maker is only concerned with selecting the very best of all the applicants. Choosing the second-best applicant is no better than choosing the worst applicant; both yield zero value.

2. There is no cost to searching.

3. The number of applicants is fixed and known ahead of time.

A large body of literature tries to address these issues by extending the problem in a number of different ways. To improve the first assumption, Chow et al. (1964) specify that the objective is to minimize the expected overall rank of the selected applicant. Later, Mucci (1973) generalizes this objective to maximizing the expectation of any monotonic payoff function on the overall ranks. Bearden et al. (2006) investigate experimental decision-making under this framework and still find a bias towards stopping too early.

Rasmussen and Robbins (1975) solve a secretary problem when the number of applicants is finite but unknown. Gianini and Samuels (1976) consider an infinite problem, where the DM can observe an unlimited number of applicants and the payoff depends on the selected applicant's rank as well as an increasing loss function through the search. Gianini (1977) shows that the optimal result can be found by solving the finite problem and taking the limit as the number of applicants goes to infinity.

In a closely related problem studied by Moser (1956) that we will refer to as the full information problem, the DM observes a true value for each applicant when he arrives and seeks to maximize the value of the applicant that she selects. We refer to the classical secretary problem as the no-information problem because the DM's objective function and information are based solely on the ordinal ranks of the applicants and need not also have underlying cardinal values. Many extensions of the secretary problem fall somewhere in between, and so may be referred to as partial information problems. Bruss and Ferguson (1993) solve a full information problem with an objective of minimizing the expected rank of the selected applicant. Bearden (2006) proposes a hybrid problem where applicants have random values but the DM only observes an indicator variable for whether the applicant is the best so far. This construction leads to another threshold rule where the DM always rejects the first $\sqrt{N}$ applicants. Mahdian et al. (2008) allow the DM to fully observe each applicant's value when he arrives but specify that the distribution of values is unknown.

We propose a new partial-infomation problem that resolves all three of these issues and also allows for convenient comparisons to the full information problem. We address the first unrealistic assumption by using Bearden's (2006) more gradated objective function that allows for a continuum of values rather than an all-or-nothing payoff. By using cardinal values, we can generate lists which are consistent with our underlying ordinal rankings but express an *intensity* of preference as well. Our formulation maintains the structure of the problem by only providing the DM with ordinal information but allows for a more realistic model of payoffs and time-costs of continuing the search. An appealing feature of the classical secretary problem is that the interviewer only observes relative ranks and must make decisions based on this limited information. This means that the DM is building up a database of information over time as she continues to interview more and more applicants, and this experience allows her to better evaluate later applicants. One of the reasons the secretary problem is so interesting is that it allows us to examine this tradeoff between an information gain from interviewing more applicants and a cost (which appears as an opportunity cost in the classical secretary problem and also as an explicit time-cost in our formulation) of continuing the search. Almost all of the extensions of the secretary problem in the literature still lead to a threshold strategy of the form “always skip the first $f(N)$ applicants, then select the next applicant who satisfies some minimum qualification”. These cutoff rules have attractive heuristic interpretations but are not asymptotically robust. Skipping the first $.37N$ or $\sqrt{N}$ applicants with certainty is not realistic for arbitrarily large N. We provide the DM with slightly better information and show that the optimal policy does not have to follow this kind of threshold strategy.

Later, we eliminate the third unrealistic assumption by relaxing the finite applicant limit N. We find that the DM may still skip a certain number of applicants with certainty, but as $N \to \infty$ this exploration phase depends only on her impatience and search costs, as measured through her discount rate r. For very low values of r, she may be willing to skip a large number of applicants in order to build up information, but this approach is not optimal for high r. The optimal policy also differs from the previous literature in that after the exploration phase is over, the DM follows a more nuanced search strategy. Her selectivity gradually decreases over time as the marginal informational benefit of each additional applicant diminishes. The acceptance decision doesn’t just depend on whether the applicant is the best or not, but takes into account his relative rank and how far she is into the search.

2 The Model

1. There is one job opening available.

2. The number N of applicants is fixed and known ahead of time.

3. Applicants arrive sequentially at random times according to a homogeneous Poisson Process, and are interviewed by the decision maker (DM). The interarrival times T_i between the arrival of the $i-1^{th}$ applicant and the i^{th} applicant are i.i.d. exponential random variables with density $g(t) = \lambda e^{-\lambda t}$. T_1 can be thought of as the time between the announcement of the job opening and the arrival of the first applicant. So the length of time elapsed before the arrival of the m^{th} applicant is $t_m = \sum_{i=1}^{m} T_i$.

4. The true value X_m of the m^{th} applicant is an i.i.d. uniform random variable on [0,1]. However, the DM observes only the relative rank of all of the applicants who have been interviewed so far, and must make an immediate decision about whether to accept or reject the applicant.

5. If an applicant is rejected, he cannot be recalled and accepted later. The search ends when an applicant has been accepted. Let S be the index of the applicant who is selected. Then the stopping time is $t_S = \sum_{i=1}^{S} T_i$.

6. The DM's objective is to adopt a strategy to maximize the expected discounted value of the applicant selected:

$$\max_S \; E[e^{-rt_S} X_S]$$

subject to her limited information about X_m, where r is the DM's constant time-discounting rate. Alternately, we could think of the discount factors e^{-rT_i} as all-inclusive search-and-interview costs, which are random and vary from applicant to applicant.

3 The Optimal Policy

We derive the DM's stopping criteria after interviewing the m^{th} applicant, $m < N$:

Since the DM does not observe the true valuation of X_m, she must make her decision based on the rank of the current applicant relative to the m total applicants who have been interviewed. We follow a mathematically convenient (but conversationally somewhat counterintuitive) convention for ranking these applicants: If we observe m applicants, we say that the worst applicant has a rank of 1 and let $X_{(1)}$ represent their value. Likewise, the best of the m applicants has rank m and value $X_{(m)}$. The k^{th} order statistic $X_{(k)}$ of m i.i.d. $U[0,1]$ random variables (here we follow the convention that $X_{(1)} \leq X_{(2)} \leq \cdots \leq X_{(m)}$, meaning that higher values of k correspond to higher values of $X_{(k)}$) has a beta distribution $B(k, m-k+1)$, with expected value

$$E[X_{(k)}] = \frac{k}{m+1} \tag{1}$$

The DM should stop and accept the current applicant if and only if the expected value of the current applicant is at least as great as the expected discounted value of rejecting him and continuing the search. Note that the length of time t_m elapsed before the arrival of the m^{th} applicant will not affect the DM's choice because the discounting factor e^{-rt_m} will be applied to the eventual payoff regardless.

To assist the DM in this decision, define $V(m)$ as the expected discounted value of rejecting the current applicant and continuing the search after the m^{th} applicant is interviewed. $V(m)$ can be thought of as the residual value of the search process to the DM when she continues to follow an optimal decision rule. $V(m)$ does not include discounting for time that has already elapsed, since these are sunk costs for the DM and should not affect her decision, as noted above.

$$V(m) \equiv E[e^{-r \sum_{i=m+1}^{S} T_i} X_S \mid S \geq m+1] \tag{2}$$

So the DM will stop if and only if

$$E[X_m \mid X_m \text{ is the } k^{th} \text{ order statistic }] \geq V(m) \tag{3}$$

$$\frac{k}{m+1} \geq V(m) \tag{4}$$

$$k \geq (m+1)V(m) \tag{5}$$

If we denote the rank of the m^{th} applicant by $k_m \in \mathbb{N}$, this means that **the DM stops and accepts the current applicant when:**

$$k_m \geq \lceil (m+1)V(m) \rceil \equiv k^*(m) \tag{6}$$

and rejects that current applicant and continues when:

$$k_m \leq \lceil (m+1)V(m) \rceil - 1 \tag{7}$$

Conditioning on whether the DM stops or continues after interviewing the next applicant leads to a first-order recursion:

$$
\begin{aligned}
V(m) = {} & E[e^{-r\sum_{i=m+1}^{S} T_i} X_S \mid S = m+1]P(S = m+1) \\
& +E[e^{-r\sum_{i=m+1}^{S} T_i} X_S \mid S > m+1]P(S > m+1)
\end{aligned} \tag{8}
$$

Since the values X_i are independent and identically distributed, it's easy to show that the rank k_{m+1} of the $m+1^{th}$ applicant has an equal probability of taking any of the values $1, 2, \ldots, m+1$. Our decision rule tells us that $S > m+1 \iff k_{m+1} \leq \lceil (m+2)V(m+1) \rceil - 1$, so

$$
P(S > m+1) = \frac{\lceil (m+2)V(m+1) \rceil - 1}{m+1} \tag{9}
$$

$$
P(S = m+1) = 1 - P(S > m+1) = \frac{m+2-\lceil (m+2)V(m+1) \rceil}{m+1} \tag{10}
$$

Conditioning (8) on the length of time T_{m+1} elapsed between the arrival of the m^{th} and $m+1^{th}$ applicants, we have

$$
\begin{aligned}
V(m) = {} & E[E[e^{-rT_{m+1}} X_{m+1} \mid k_{m+1} \geq \lceil (m+2)V(m+1) \rceil] \mid T_{m+1}]\big(1 - P(S > m+1)\big) \\
& +E[E[e^{-rT_{m+1}}\, e^{-r\sum_{i=m+2}^{S} T_i} X_S \mid S \geq m+2] \mid T_{m+1}]P(S > m+1)
\end{aligned} \tag{11}
$$

$$V(m) = E[e^{-rT_{m+1}}E[X_{m+1} \mid k_{m+1} \geq \lceil (m+2)V(m+1)\rceil]]\big(1 - P(S > m+1)\big)$$
$$+E[e^{-rT_{m+1}}E[e^{-r\sum_{i=m+2}^{S} T_i} X_S \mid S \geq m+2]]P(S > m+1) \tag{12}$$

$$V(m) = E[e^{-rT_{m+1}}]\Big(E[X_{m+1} \mid k_{m+1} \geq \lceil (m+2)V(m+1)\rceil]\big(1 - P(S > m+1)\big)$$
$$+E[e^{-r\sum_{i=m+2}^{S} T_i} X_S \mid S \geq m+2]P(S > m+1)\Big) \tag{13}$$

We calculate $E[X_{m+1} \mid k_{m+1} \geq \lceil (m+2)V(m+1)\rceil]$ by conditioning on the rank of the accepted applicant, and again using (1) and the fact that we are equally likely to observe each of these ranks:

$$\begin{aligned} &E[X_{m+1} \mid k_{m+1} \geq \lceil (m+2)V(m+1)\rceil] \\ &= \sum_{k=\lceil (m+2)V(m+1)\rceil}^{m+1} E[X_{m+1} \mid k_{m+1} = k]P(k_{m+1} = k \mid k_{m+1} \geq \lceil (m+2)V(m+1)\rceil) \\ &= \sum_{k=\lceil (m+2)V(m+1)\rceil}^{m+1} \frac{k}{m+2} \cdot \frac{1}{m+2-\lceil (m+2)V(m+1)\rceil} \end{aligned} \tag{14}$$

We also note that

$$E[e^{-rT_{m+1}}] = \int_0^\infty e^{-rt}\lambda e^{-\lambda t}dt = \frac{\lambda}{\lambda + r} \tag{15}$$

Substituting (9), (10), (14), and (15) into (13) and simplifying, we get

$$V(m) = \frac{\lambda}{\lambda + r}\left(\sum_{k=\lceil (m+2)V(m+1)\rceil}^{m+1} \frac{k}{(m+1)(m+2)} + \frac{\lceil (m+2)V(m+1)\rceil - 1}{m+1}\, V(m+1)\right) \tag{16}$$

The coefficient $\frac{\lambda}{\lambda+r}$ represents the expected discount factor from waiting for the next applicant. The first term represents the expected payoff if that applicant ranks highly enough to be accepted, and the second term represents the expected payoff if the next applicant is rejected and the search continues further, with both terms weighted by their probabilities of occurring.

We can then solve for the continuation values $V(m)$ and critical ranks $k^*(m)$ via backwards induction from a boundary condition: If the DM chooses to continue after observing the $N-1^{th}$ applicant, she must accept the N^{th} and final applicant.
So after interviewing the $N-1^{th}$ applicant:

$$V(N-1) = E[e^{-rT_N} X_N] \tag{17}$$

$$= E[E[e^{-rT_N} X_N \mid T_N]]$$

$$= E[e^{-rT_N} E[X_N]]$$

$$= E[e^{-rT_N}]E[X_N]$$

$$V(N-1) = \frac{\lambda}{2(\lambda + r)} \tag{18}$$

The recursion is too complex to derive an explicit solution for the infinite problem, where we let the applicant limit $N \to \infty$. However, we can show that the sequence of values converges as we relax this boundary condition.

Theorem 1: Let $V_N(m)$ be the expected value of continuing the search after observing the m^{th} applicant when the maximum number of applicants is $N \geq m+1$. Then $V_N(m)$ converges as $N \to \infty$ for all $m \geq 0$. We write this succinctly as $\{\mathbf{V}_N\} \to \{\mathbf{V}_\infty\}$, where $\{\mathbf{V}_N\}$ represents the finite sequence of N continuation values that guide the DM in the finite problem and $\{\mathbf{V}_\infty\}$ represents the infinite sequence of continuation values that guide the DM in the infinite problem.

This means that we can choose a large enough value of N to approximate the sequence of values and cutoffs that guide the DM in the infinite problem. Figure 1 shows how the sequences $\{\mathbf{V}_N\}$ of continuation values increase to $\{\mathbf{V}_\infty\}$ as we relax the applicant limit N. Observe that for any fixed m, $V_N(m)$ is increasing to the limit $V_\infty(m)$. Also note that the continuation values $V_\infty(m)$ increase (at a diminishing rate) as m increases. This reflects the informational improvements from the DM's added experience. She uses this database of rejected applicants to better evaluate future applicants, which leads to better decisions and improvements in the expected outcome. Figure 2 shows the corresponding selectivity of the DM, expressed at each stage of the search by the fraction of applicants who will be rejected if they follow the optimal strategy with the continuation values from Figure 1. The search is largely driven by the continuation values $\{\mathbf{V}_\infty\}$, but we see that the DM starts to lower her selectivity as she nears the applicant limit N. The DM is more willing to compromise when she knows that she only has the opportunity to choose among a handful of remaining applicants. This strategy is driven by the fear of exhausting the applicant list and being forced to accept the very last applicant regardless of their rank. In the infinite problem, this boundary condition never shows up and so we don't see the continuation values and selectivity drop off. However, we do observe a gradual decrease in selectivity as m increases even in the infinite problem. We can attribute this to the declining marginal benefits of additional experience. When m is low, the

DM can afford to be more selective because she knows that rejecting applicants will help her make a better choice later on. When m is high, rejecting more applicants adds little to the DM's ability to evaluate future applicants, so there is less incentive for her to reject the current applicant. Figures 3 and 4 demonstrate the effects of different discount rates. When the discount rate is high (indicating that the DM is impatient), the continuation values are lower and she is less selective. Likewise, a lower discount rate leads to higher continuation values and selectivity.

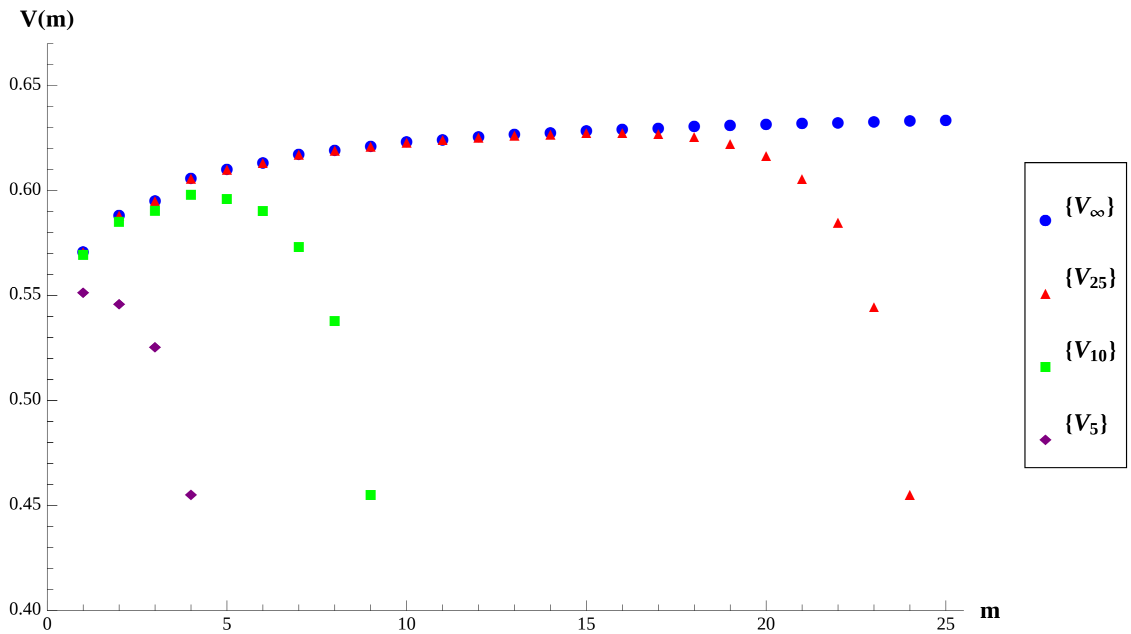

Figure 1: **Continuation values for finite and infinite problems when** $\lambda = 1$ **and** $r = 0.1$

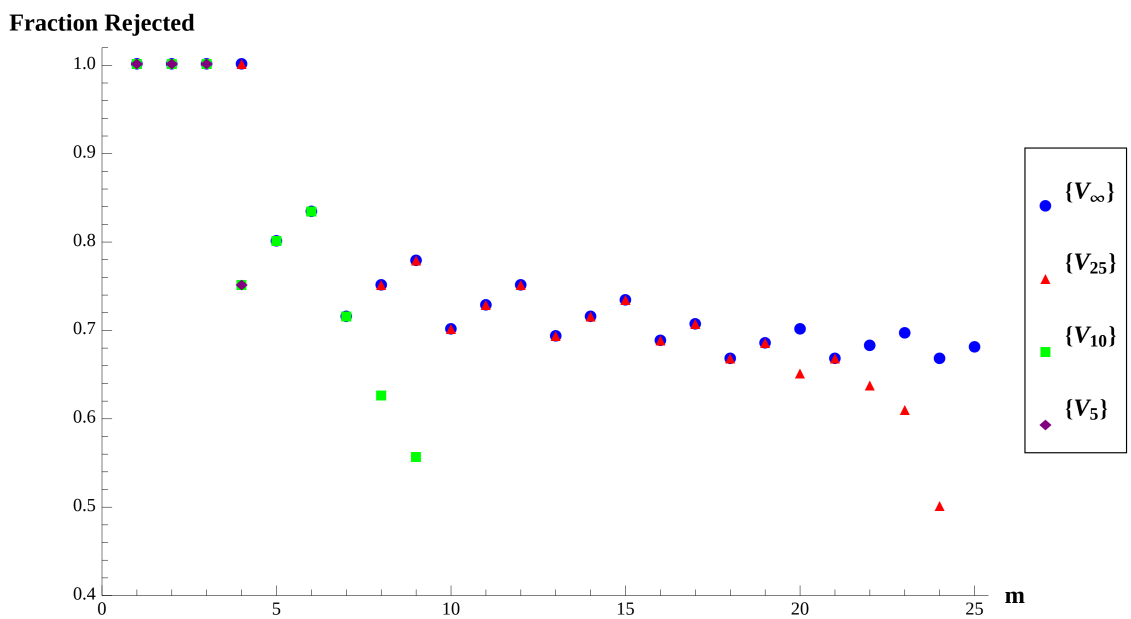

Figure 2: **Selectivity (proportion** $\frac{k^*(m)}{m}$ **of applicants rejected after the** m^{th} **interview) for finite and infinite problems when** $\lambda = 1$ **and** $r = 0.1$

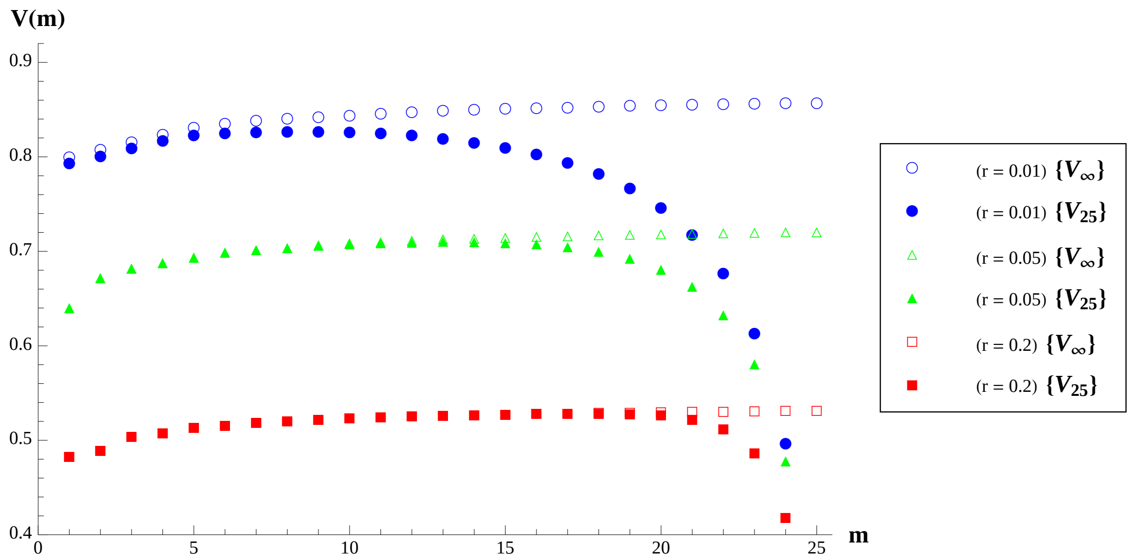

Figure 3: **Continuation values for finite and infinite problems for various r values when $\lambda = 1$**

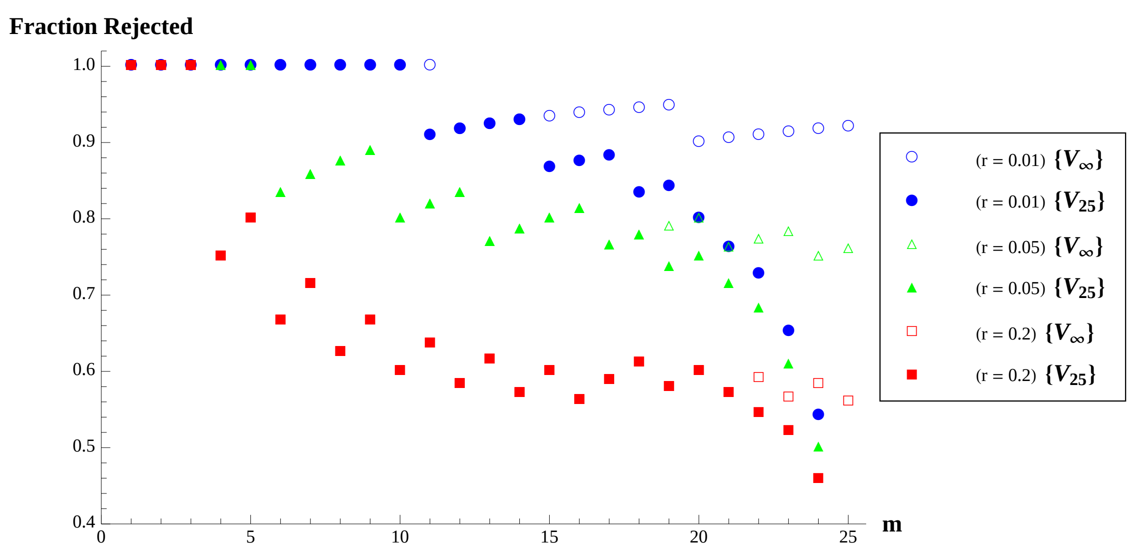

Figure 4: **Selectivity (proportion $\frac{k^*(m)}{m}$ of applicants rejected after the m^{th} interview) for finite and infinite problems for various r values when $\lambda = 1$**

4 The Full Information Model

Now we introduce a full information analogue of the partial information sequential search that we solved above. Here the DM always observes the true value of each applicant as he arrives, so she no longer gains information from her past search experience. As a result, her decision at each stage is driven purely by the time-discounting and the limited number of applicants N. The full information model will provide us with a useful benchmark that we can use to examine the effects of the informational limitations when the DM can only observe ordinal rankings. Our full information search generalizes Moser's (1956) results by introducing interarrival times and discounting and yields a recursion formula that agrees with Moser's for $r = 0$.

1. There is one job opening available.

2. The number of applicants N is fixed and known ahead of time.

3. Applicants arrive sequentially at random times according to a homogeneous Poisson Process, and are interviewed by the DM. The interarrival times T_i between the arrival of the $i-1^{th}$ applicant and the i^{th} applicant are i.i.d. exponential random variables with density $g(t) = \lambda e^{-\lambda t}$. The length of time elapsed before the arrival of the m^{th} applicant is $t_m = \sum_{i=1}^{m} T_i$.

4. The true value X_m of the m^{th} applicant is an i.i.d. uniform random variable on [0,1] *and is fully observable by the DM* .

5. If an applicant is rejected, he cannot be recalled and accepted later. The search ends when an applicant has been accepted. S again denotes the index of the applicant who is selected. $t_S = \sum_{i=1}^{S} T_i$.

6. The DM's objective is to adopt a strategy to maximize the expected discounted value of the applicant selected:

$$\max_{S} \; E[e^{-rt_S} X_S]$$

where r is the DM's constant time-discounting rate.

5 The Full Information Optimal Policy

As before, the DM will assign a sequence of values, which we will write as $V^{FI}(m)$, that capture the expected discounted value of rejecting the m^{th} (current) applicant and continuing the search, excluding the discount factor for the time t_m that has already elapsed and assuming the DM continues to follow an optimal policy. There is no longer any uncertainty about the value of the current applicant, so the decision rule is straightforward:

Stop if and only if

$$X_m \geq V^{FI}(m) \tag{19}$$

Conditioning on whether the DM stops or continues after interviewing the next applicant leads to a simple first-order recursion:

$$V^{FI}(m) = E[e^{-r \sum_{i=m+1}^{S} T_i} X_S \mid S \geq m+1] \tag{20}$$

$$\begin{aligned} V^{FI}(m) = {} & E[e^{-r \sum_{i=m+1}^{S} T_i} X_S \mid S = m+1] \cdot P(S = m+1) \\ & + E[e^{-r \sum_{i=m+1}^{S} T_i} X_S \mid S > m+1] \cdot P(S > m+1) \end{aligned} \tag{21}$$

Conditioning (21) on the length of time T_{m+1} elapsed between the arrival of the m^{th} and $m+1^{th}$ applicants,

$$\begin{aligned} V^{FI}(m) = {} & E[E[e^{-rT_{m+1}} X_{m+1} \mid S = m+1] \mid T_{m+1}] \cdot P(S = m+1) \\ & + E[E[e^{-rT_{m+1}} e^{-r \sum_{i=m+2}^{S} T_i} X_S \mid S > m+1] \mid T_{m+1}] \cdot P(S > m+1) \end{aligned} \tag{22}$$

Our decision rule tells us that $S = m+1 \iff X_{m+1} \geq V^{FI}(m+1)$, so

$$P(S > m+1) = P(X_{m+1} < V^{FI}(m+1)) = V^{FI}(m+1) \tag{23}$$

$$P(S = m+1) = 1 - P(S > m+1) = 1 - V^{FI}(m+1) \tag{24}$$

and

$$\begin{aligned} V^{FI}(m) = {} & E[e^{-rT_{m+1}} E[X_{m+1} \mid X_{m+1} \geq V^{FI}(m+1)]] \cdot (1 - V^{FI}(m+1)) \\ & + E[e^{-rT_{m+1}} E[e^{-r \sum_{i=m+2}^{S} T_i} X_S \mid S \geq m+2]] \cdot V^{FI}(m+1) \end{aligned} \tag{25}$$

$$V^{FI}(m) = E[e^{-rT_{m+1}}] \Big((\frac{1}{2} + \frac{V^{FI}(m+1)}{2})(1 - V^{FI}(m+1) + V^{FI}(m+1) \cdot V^{FI}(m+1) \Big) \tag{26}$$

$$V^{FI}(m) = \frac{\lambda}{2(\lambda + r)} \Big(V^{FI}(m+1)^2 + 1 \Big) \tag{27}$$

We get the same boundary condition as before by considering the DM's choice after interviewing the $N-1^{th}$ applicant. If she continues the search she must accept the N^{th} and final applicant, meaning

$$V^{FI}(N-1) = E[e^{-rT_N} X_N] \tag{28}$$

$$= E[e^{-rT_N}]E[X_N]$$

$$V^{FI}(N-1) = \frac{\lambda}{2(\lambda + r)} \tag{29}$$

Again, we allow $N \to \infty$ in order to analyze the DM's strategy in the infinite full information problem, where there is no limit on the number of applicants.

Theorem 2: Let $V_N^{FI}(m)$ be the expected value of continuing the search after observing the m^{th} applicant when the maximum number of applicants is $N \geq m+1$ in the full information problem. Then $V_N^{FI}(m)$ converges as $N \to \infty$ for all $m \geq 0$.

We write this succinctly as $\{\mathbf{V}_N^{FI}\} \to \{\mathbf{V}_\infty^{FI}\}$, where $\{\mathbf{V}_N^{FI}\}$ represents the finite sequence of continuation values that guide the DM in the finite problem and $\{\mathbf{V}_\infty^{FI}\}$ represents the infinite sequence of continuation values that guide the DM in the infinite problem.

Theorem 3: $V_\infty^{FI}(m) = \frac{\lambda+r}{\lambda} - \frac{\sqrt{2\lambda r + r^2}}{\lambda} \equiv V_\infty^{FI} \;\; \forall m \geq 0$. This is a very strong result, condensing the entire search process in the infinite full information problem down to a single continuation value V_∞^{FI}. After observing any applicant at any time, the DM simply compares his value to this continuation value and accepts if and only if his value is at least as large as V_∞^{FI}.

Figure 5 illustrates how the sequences $\{\mathbf{V}_N^{FI}\}$ of continuation values in the finite full information problems converge to $\{\mathbf{V}_\infty^{FI}\}$. Note that these values also define our stopping rule, so higher continuation values imply higher selectivity. As noted before, we see that $V_\infty^{FI}(m)$ takes the constant value $\frac{\lambda+r}{\lambda} - \frac{\sqrt{2\lambda r + r^2}}{\lambda}$ for all m. Figure 6 compares the infinite full information problem with our infinite partial information problem from before. Since continuation values are fixed in the full information case, it serves as a useful benchmark to look at how the DM gains information through the search when she can only observe the relative ranks of previously observed applicants. Figure 7 provides a table comparing expected values at the beginning of the search, before the DM has interviewed any applicants. We see that these *ex ante* expected search payoffs are a bit higher in the infinite problem than in the finite problem, and this improvement is larger for lower discount rates because the boundary condition is more likely to affect the DM. We also see that the expected payoffs are significantly higher with full information than when the DM only observes relative ranks (partial information). This increase is largest when the discount rate is high, because the DM tends to interview only a few applicants before ending her search. This means that she does not have the time to build up experience to help her make a decision. So having full information from the beginning of the search yields a huge improvement in her ability to make a good decision. Full information still benefits a DM with a low discount rate, but she is patient so she can afford to build up a large database of rejected applicants that help her make a better decision later on. As a result, her expected search payoff doesn't show as large of an increase when she is given full information.

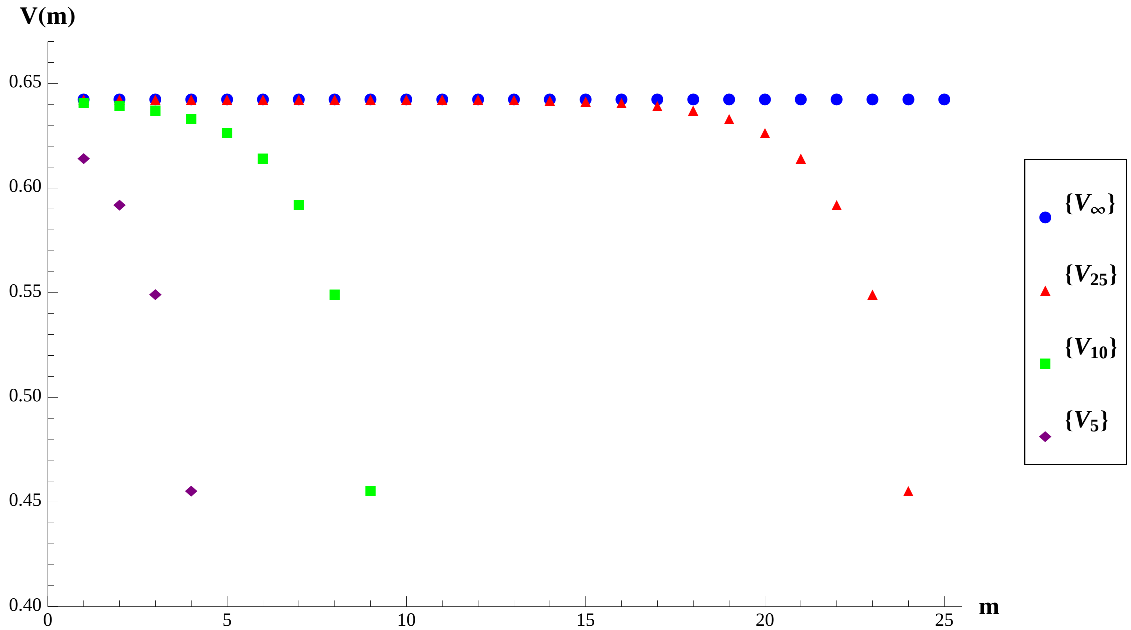

Figure 5: **Continuation values for finite and infinite full information problems when $r = 0.1$ and $\lambda = 1$.**

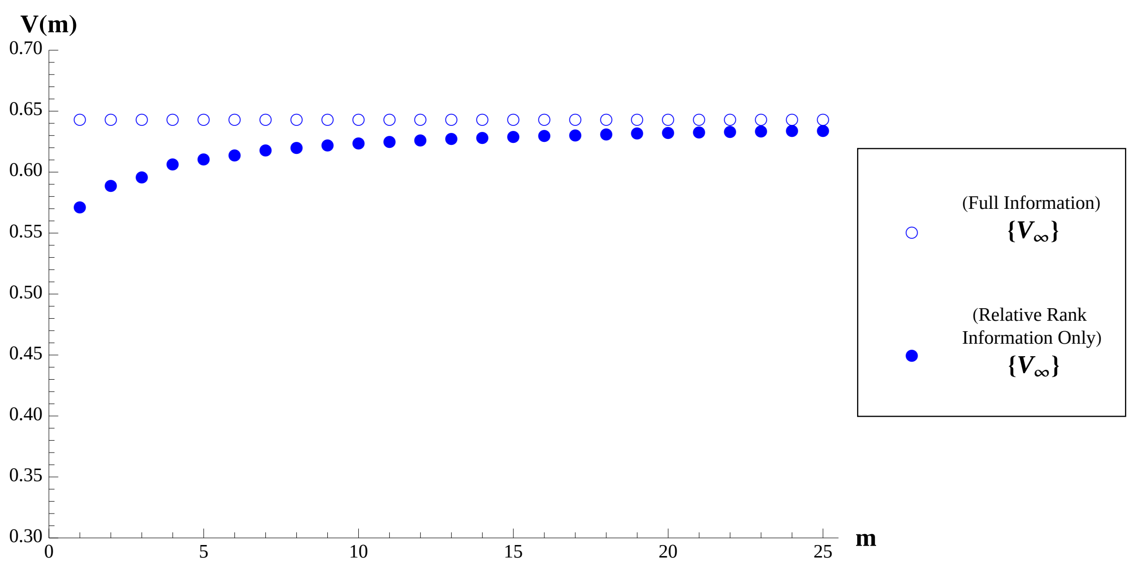

Figure 6: **Comparison of continuation values from the infinite problem with full information and with only relative ranks when $r = 0.1$ and $\lambda = 1$.**

Ex Ante Expected Payoffs		Discount Rate:		
		r = 0.01	r = 0.05	r = 0.2
Partial Information	(Finite Problem, N = 10): $V_{10}(0)$	.7241	.6012	.4167
	(Infinite Problem): $V_\infty(0)$	.7898	.6079	.4167
Full Information	(Finite Problem, N = 10): $V_{10}(0)$	.8263	.7239	.5366
	(Infinite Problem): $V_\infty(0)$	.8682	.7298	.5367

Figure 7: **Table of *ex ante* expected search payoffs $V(0)$ with full information and with only relative ranks for various r values when $\lambda = 1$.**

6 Discussion

Our hybrid secretary problem addresses the three major shortcomings of the classical secretary problem and yields a more realistic optimal policy. The other advantage of our model is that it allows us to isolate and quantify the DM's three primary considerations when constructing a strategy:

1. The DM has to make decisions based on limited information, but accumulates more and more information as she continues to search. Every new applicant who is rejected adds to the database of experience that allows the DM to better place future applicants (this is what motivates the DM to skip the first 37% of the applicants with certainty in the classical problem). So the DM will be more selective earlier in the search because of the informational benefit of rejecting applicants. In the infinite problem, we can quantify the marginal information improvement from rejecting the m^{th} applicant by comparing $V_\infty(m)$ and $V_\infty(m+1)$. We can also look at the cost of the limited information at any point in our search (in either the finite or infinite problems), by comparing

$V(m)$ with the corresponding $V^{FI}(m)$ from the full information problem. In particular, the quantity $V_\infty^{FI}(0) - V_\infty(0)$ gives us the *ex ante* value of having full information compared with the purely ordinal information in our infinite discounted problem.

2. When N is finite, the DM worries about reaching the end of the list of applicants. This forces her to be less selective as she gets closer and closer to this boundary. We can quantify the cost of this limit by comparing $V_N(m)$ with $V_\infty(m)$ from the full information problem (as well as the optimal cutoffs $k_N^*(m)$ and $k_\infty^*(m)$).

3. There is a time cost to the search and the ultimate payoff is discounted according to how long the search took. So the DM has an incentive to be less selective and stop sooner than she otherwise might have in order to reduce this cost. We can quantify the effects and expected reduction in payoff of this discounting by comparing $V(m)$ when $r > 0$ with $V(m)$ when $r = 0$ (as well as the optimal cutoffs $k^*(m)$ in each case).

This framework also outlines a complementary structure for the limited information models that often show up in the economics literature. In auction theory (see Krishna, 2002), valuations are frequently defined through interdependent or common value models, where players are endowed with private signals and must infer information about other players' signals from their behavior. Under the structure of our hybrid secretary problem, the player would also be updating her own information as she makes decisions.

For example, we could explore the connections between matching theory and decision making under uncertainty by studying behavior in a dynamic matching game with interdependent values. Classic two-sided matching theory (see Roth and Sotomayor, 1990) provides a simple but powerful framework to define a stable match and an algorithm that guarantees a solution. Remarkably, these results only require that the individuals provide a strict ordinal ranking of their preferences over all individuals on the other side of the market. This type of model is useful for exploring behavior and designing incentive-compatible mechanisms in a number of situations, including medical labor markets, the college admissions problem, and public school choice. However, the model assumes that individuals have full information about their preferences. This means that their rankings are static and only submitted once at the start of the matching process, with no dynamic interactions beyond computation of the equilibrium. By limiting players' information to ordinal rankings *of the individuals whom they have already observed*, we would be able to consider matches generated over time that might not otherwise meet the stability criteria of classic static matching theory.

Our hybrid secretary problem can also be extended in a number of different directions. We could generalize the underlying value distributions and arrival process and relax the assumption of independence between applicants. This would increase the mathematical complexity of the problem, but would allow us to examine the effects of correlation between applicants. Another interesting study would test the experimental behavior of decision makers in our framework, which would complement previous research on decision making in the classical secretary problem. The more nuanced objective and strategy sets would give us a better look into the decision making process, and the controls λ, N, and r might provide explanations for the biases that previous studies have suggested. The increased size of the strategy sets would also allow for a broader spectrum of heuristics to compare to the mathematically optimal strategy.

Appendix: Proofs

Proof of Theorem 1: Fix any integer $m \geq 0$. $V_N(m)$ can be calculated via backwards induction for any finite $N \geq m+1$ (the lowest N for which $V_N(m)$ is defined). Recall that $V_N(m)$ is the expected value of rejecting the current applicant and continuing the search after observing the m^{th} of N total applicants and consider the sequence $\{V_N(m)\}_{N=m+1}^{\infty}$. $V_{N+1}(m) \geq V_N(m)$ because we are simply adding an additional applicant (whose true value is independent of any previous applicants) on to the end of the list of remaining selections for the DM. So the DM could follow the exact same stopping strategy as before, when there were N total applicants, and simply accept the N^{th} with certainty. This would mimic the search process and expected payoffs exactly. However, the DM is maximizing her search policy over a larger set of possible strategies than before, and she may choose a different policy. But if she wants to do so it must be the case that $V_{N+1}(m)$ is at least at high as $V_N(m)$. This holds for all N, so $V_N(m)$ is monotonically nondecreasing in N. Also note that $V_N(m)$ is trivially bounded below by 0 and above by 1 since $0 \leq e^{rt} \leq 1$ and $0 \leq X_S \leq 1$. So if we let $N \to \infty$, $V_N(m)$ converges to a limit $V_\infty(m)$ by the Monotonic Convergence Theorem. Since this holds for all nonnegative integers m, we can construct a limiting sequence $\{\mathbf{V}_\infty\}$ whose terms $V_\infty(m)$ are the continuation values in the infinite discounted secretary problem. $\square$

Proof of Theorem 2: Fix any integer $m \geq 0$. Then $V_N^{FI}(m)$ is easily calculated via backwards induction for any finite N. Starting at $N = m+1$ (this is the lowest N for which $V_N^{FI}(m)$ is defined), increment N upwards and consider the sequence of values of $V_N^{FI}(m)$. The first term of the sequence is then $V_{m+1}^{FI}(m) = \frac{\lambda}{2(\lambda+r)}$, the boundary condition.

We can write the recursion from our decision rule as

$$V_N^{FI}(m) = f(V_N^{FI}(m+1)), \text{where } f(x) = \frac{\lambda}{2(\lambda+r)}(x^2+1)$$

Now observe that for fixed boundaries N and $N+1$, backwards induction gives us

$$V_N^{FI}(m) = \underbrace{f(f(\dots f(}_{N-m-1}\frac{\lambda}{2(\lambda+r)}\,)\dots))$$

$$V_{N+1}^{FI}(m) = \underbrace{f(f(f(\dots f(}_{N-m}\frac{\lambda}{2(\lambda+r)}\,)\dots)))$$

So we have

$$V_{N+1}^{FI}(m) = f(V_N^{FI}(m))$$

The first term of our sequence satisfies $0 < V_{m+1}^{FI}(m) = \frac{\lambda}{2(\lambda+r)} < 1$ and we show by induction on N that all other terms in the sequence also satisfy $0 < V_N^{FI}(m) < 1$. Assume $0 < V_N^{FI}(m) < 1$. Then $0 < V_N^{FI}(m)^2 < 1$ and $1 < V_N^{FI}(m)^2 + 1 < 2$, so multiplying through by $\frac{\lambda}{2(\lambda+r)}$ gives us $0 < \frac{\lambda}{2(\lambda+r)} < f(V_N^{FI}(m)) = V_{N+1}^{FI}(m) < \frac{\lambda}{\lambda+r} < 1$. Observe that f is continuous and $f'(x) = \frac{\lambda}{\lambda+r}x$. So by the Mean Value Theorem, for all $N \geq m+1$, there exists c between $V_N^{FI}(m)$ and $V_{N+1}^{FI}(m)$ (so $0 < c < 1$) such that

$$f'(c) = \frac{f(V_{N+1}^{FI}(m)) - f(V_N^{FI}(m))}{V_{N+1}^{FI}(m) - V_N^{FI}(m)}$$

$$\frac{\lambda}{\lambda+r}c = \frac{V_{N+2}^{FI}(m) - V_{N+1}^{FI}(m)}{V_{N+1}^{FI}(m) - V_N^{FI}(m)}$$

$0 < \frac{\lambda}{\lambda+r} \leq 1$, so $0 < \frac{\lambda}{\lambda+r}c < 1$. In particular, this satisfies

$$|V_{N+2}^{FI}(m) - V_{N+1}^{FI}(m)| \leq \frac{\lambda}{\lambda+r}c\,|V_{N+1}^{FI}(m) - V_N^{FI}(m)|$$

Letting $N \to \infty$, our sequence is contractive and therefore convergent. $\square$

Proof of Theorem 3: Conditioning on the DM's choice after observing the m^{th} applicant in the infinite full information problem (the recursion follows exactly as it did in the finite full information problem) gives us $V_\infty^{FI}(m) = f(V_\infty^{FI}(m+1))$. Also, observe that

$$V_\infty^{FI}(m) = \lim_{N\to\infty} V_N^{FI}(m) = \lim_{N\to\infty} \underbrace{f(f(\dots f(}_{N-m-1} \frac{\lambda}{2(\lambda+r)})\dots))$$

$$= \lim_{N\to\infty} \underbrace{f(f(\dots f(}_{N-m-2} \frac{\lambda}{2(\lambda+r)})\dots)) = \lim_{N\to\infty} V_N^{FI}(m+1)$$

$$= V_\infty^{FI}(m+1)$$

Substituting into the recursion, we get

$$V_\infty^{FI}(m) = f(V_\infty^{FI}(m)) = \frac{\lambda}{2(\lambda+r)} V_\infty^{FI}(m)^2 + \frac{\lambda}{2(\lambda+r)}$$

$$\frac{\lambda}{2(\lambda+r)} V_\infty^{FI}(m)^2 - V_\infty^{FI}(m) + \frac{\lambda}{2(\lambda+r)} = 0$$

$$= \frac{\lambda+r}{\lambda}\left(1 \pm \sqrt{1 - (\frac{\lambda}{\lambda+r})^2}\right)$$

$$= \frac{\lambda+r}{\lambda} \pm \frac{\sqrt{2\lambda r + r^2}}{\lambda}$$

Note that the first term is ≥ 1 and the second term is ≥ 0 and $0 \leq V(m) \leq 1$. So the only solution is $V_\infty^{FI}(m) = \frac{\lambda+r}{\lambda} - \frac{\sqrt{2\lambda r+r^2}}{\lambda}$. $\square$

References

Babaioff, M., Immorlica, N., Kempe, D., Kleinberg, R., 2008, "Online Auctions and Generalized Secretary Problems," *ACM SIGecom Exchanges* 7(2).

Bearden, J.N., 2006, "A New Secretary Problem with Rank-Based Selection and Cardinal Payoffs," *Journal of Mathematical Psychology* 50, 58-59.

Bearden, J.N., Murphy, R.O., 2007, "On Generalized Secretary Problems," In Abdellaoui, M., Luce, R.D., Machina, M.J., Munier, B. (Eds.) *Uncertainty and Risk 41* pp. 187-205. Springer Berlin Heidelberg.

Bearden, J.N., Rapoport, A., Murphy, R.O., 2006, "Sequential Observation and Selection with Rank-Dependent Payoffs: An Experimental Study," *Management Science* 52 (9), 1437-1449.

Bruss, F.T., Ferguson, T.S., 1993, "Minimizing the Expected Rank with Full Information," *Journal of Applied Probability* 30(3), 616-626.

Chow,Y.S., Moriguti, S., Robbins, H., Samuels, S.M., 1964, "Optimal Selection Based On Relative Ranks," *Israel Journal of Mathematics* 2(2), 81-90.

Cowan, R., Zabczyk, J., 1978, "An Optimal Selection Problem Associated with the Poisson Process," *Theory of Probability and its Applications* 23(3), 584-592.

Ferguson, T.S., 1989, "Who Solved The Secretary Problem?" *Statistical Science* 4(3), 282-296.

Freeman, P.R., 1983, "The Secretary Problem and its Extensions," *International Statistical Review* 51, 189-206.

Gianini, J., Samuels, S.M., 1976, "The Infinite Secretary Problem," *The Annals of Probability* 4(3), 418-432.

Gianini, J., 1977, "The Infinite Secretary Problem as the Limit of the Finite Problem," *The Annals of Probability* 5(4), 636-644.

Gilbert, J.P., Mosteller, F., 1966, "Recognizing the Maximum of a Sequence," *Journal of the American Statistical Association* 61, 35-73.

Krishna,V., 2002, "Auction Theory," *Academic Press*, San Diego.

Lindley, D.V., 1961, "Dynamic Programming and Decision Theory," *Applied Statistics* 10(1), 39-51.

Mahdian, M., McAfee, R.P., Pennock, D., 2008, "The Secretary Problem with a Hazard Rate Condition," In C. Papadimitriou and S. Zhang (Eds.): *WINE 2008, LNCS 5385* pp. 708-715. Springer Berlin / Heidelberg.

Moser, L., 1956, "On a Problem of Cayley," *Scripta Mathematica* 22, 289-292.

Mucci, A.G., 1973, "Differential Equations and Optimal Choice Problems," *The Annals of Statistics* 1(1), 104-113.

Roth, A.E, Sotomayor, M.A.O., 1990, "Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis", *Econometric Society Monographs* 18, Cambridge University Press, Cambridge.

Rusmussen, W.T., Robbins, H., 1975, "The Candidate Problem with Unknown Population Size," *Journal of Applied Probability* 12(4), 692-701.

Samuel-Cahn, E., 2007, "Optimal Stopping for I.I.D. Random Variables Based on the Sequential Information of the Location of Relative Records Only," *Sequential Analysis* 26, 395-401.

Seale, D.A., Rapoport, A., 1997, "Sequential Decision Making with Relative Ranks: An Experimental Investigation of the Secretary Problem," *Organizational Behavior and Human Decision Processes* 69(3), 221-236.

Seale, D.A., Rapoport, A., 2000, "Optimal Stopping Behavior with Relative Ranks: The Secretary Problem with Unknown Population Size," *Journal of Behavioral Decision Making* 13, 391-411.

Stein, W.E., Seale, D.A., Rapoport, A., 2003, "Analysis of Heuristic Solutions to the Best Choice Problem," *European Journal of Operations Research* 151, 140-152.