

When is it Important to Know You've Been Rejected? A Search Problem with Probabilistic Appearance of Offers

Sanmay Das

Center for Biological and Computational Learning, Massachusetts Institute of Technology, Cambridge MA 02139
sanmay@csail.mit.edu

John N. Tsitsiklis

Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge MA 02139 jnt@mit.edu

A problem that often arises in the process of searching for a job or for a candidate to fill a position is that applicants do not know if they will receive an offer from any given firm with which they interview, and, conversely, firms do not know whether applicants will definitely take positions they are offered. In this paper, we model the search process as an optimal stopping problem with probabilistic appearance of offers, from the perspective of a single decision-maker who wants to maximize the realized value of the offer she accepts. Our main results quantify the *value of information* in the following sense: how much better off is the decision-maker if she knows each time whether an offer appeared or not, compared to the case where she is only informed when offers actually appear? We show that for some common distributions of offer values, she can expect to receive very close to her optimal value even in the lower information case, as long as she knows the probability that any given offer will appear. However, her expected value in the low information case (as compared to the high information case) can fall dramatically when she does not know the appearance probability *ex ante* but must infer it from data. This suggests that hiring and job-search mechanisms may not suffer from serious losses in efficiency or stability from participants hiding information about their decisions, unless agents are uncertain of their own attractiveness as employees or employers.

Key words: search problems; secretary problems; value of information

History:

1. Introduction.

Many job markets are structured in a manner where potential employees submit their applications to a number of employing firms simultaneously, and then wait to hear back from these firms. Firms themselves often make exploding offers that employees have to decide on in a short time-frame. Sometimes the firms will tell potential employees as soon as they are no longer under consideration, and in other cases they wait until the end of the search process to provide this information to applicants. The central question that we address in this paper is this: *How much better off is an*

applicant if she is told every time she has been rejected by a firm, as opposed to only knowing when she receives offers?

In order to study this problem, we construct a stylized model in which the decision problem faced by agents is a version of the problem variously referred to in the literature as the Cayley-Moser problem, the (job) search problem, the house hunting problem and the problem of selling an asset (Ferguson (1989)). In the original problem, a job applicant knows that there will be exactly n job opportunities, which will be presented to her sequentially. At the time each job is presented, she observes the utility she would receive from taking that job offer (one can think of it purely in terms of wages), and must decide immediately whether to accept the job offer or not. If she declines the offer, she may not go back to it. If she accepts it, she may not pick any of the subsequent offers. What is the strategy that maximizes her expected utility? This problem has been addressed for various distributions of offer values, and much of that work is summarized by Gilbert and Mosteller (1966).

The problem we consider is a variant of the above problem in which the total number of possible offers is known, but each offer appears only with a certain probability. This problem is motivated in part by models of two-sided matching markets like labor markets or dating markets. In particular, a problem considered by Das and Kamenica (2005) is one in which men are asked out on dates by women, and must respond immediately, but, while they have priors on the values of going out with particular women, they do not know the order in which women are going to appear, so they are not aware of whether or not a better option might come along in the future. This is because a better woman than the one currently asking a man out might either have already appeared in the ordering and not asked him out, or might appear later and not ask him out, or might appear later and ask him out. A similar problem can arise in faculty hiring processes for universities and colleges. Universities may not know whether applicants will definitely take positions that are offered, and, conversely, applicants do not know if they will receive an offer from any given university with which they interview. This paper only looks at one side of this process without considering the dynamics involved when multiple agents interact, potentially strategically. Another motivation comes from thinking of the offers as investment opportunities (Gilbert and Mosteller (1966)). In particular, the continuous-time variant we discuss can be interpreted in terms of investment opportunities that arrive as a Poisson process where the decision-maker wants to choose the best one. To simplify the analysis, we assume that the probability that a particular offer appears, p , is the same across all offers and is independent of the actual value of the offer. The value of p may or may not be known

to the applicant, and can be thought of as a measure of the “attractiveness” of the applicant or decision-maker.

Most of the previous research on search models focuses on solving an agent’s infinite horizon optimal stopping problem when there is either a cost to generating the next offer, or a discount factor associated with future utility (the book by DeGroot (1970) provides an account of much of this line of research). The problem we study here is a finite-horizon search problem with no cost to seeing more offers and no search frictions. The basic questions we pose and attempt to answer relate to how much the expected utility of the decision-maker changes between different information sets and different mechanisms. The question with regard to information sets can be thought of as follows. Suppose you interview with n firms that might want to hire you. Then the companies get ordered randomly and come along in that order and decide whether or not to make you an offer. How much would you pay to go from a situation in which you only saw which companies made you an offer (the *low information* variant) to a situation in which you saw, for each company, whether or not they chose to make you an offer (the *high information* variant)? Generalizing the two informational cases to continuous time provides good approximations for large n and insight into the value of information in these cases. It also allows us to make an interesting connection to a closely related problem called the *secretary problem*. We will also discuss the difference in expected utility between two different mechanisms. The exploding offer mechanism can lead to a substantial decline in the expected utility of a job-seeker compared to a mechanism in which she sees all the offers she will receive simultaneously and can choose from among them. What if you could pay to see the entire set of offers you would get simultaneously so that you could pick among them? How much should you be willing to pay? We will explicitly compare the expected loss in value in going from this *simultaneous choice* mechanism to the *sequential choice* mechanism that generates the stopping problem.

1.1. Related Work.

In the classical secretary problem (CSP), a decision-maker has to hire one applicant out of a pool of n applicants who will appear sequentially. Again, the decision-maker must decide immediately upon seeing an applicant whether to hire her or not. The key difference between secretary problems and search problems, as Ferguson (1989) notes, is that in secretary problems “the payoff depends on the observations only through their relative ranks and not otherwise on their actual values.” The most studied types of secretary problems are games with 0-1 payoffs, with the payoff of 1 being received if and only if the decision-maker hires the best applicant. The decision-maker’s optimal policy is thus one that maximizes the probability of selecting the best applicant.

A historical review of the early literature on secretary problems, including important references, can be found in the paper by Gilbert and Mosteller (1966), as can solutions to many extensions of the basic problem, including the search problem (with finite and known n and no search costs) for various different distributions over the values of applicants. Many interesting variants of the original problem, mostly focusing on maximizing the probability of hiring the best applicant, have appeared in intervening decades. For instance, Cowan and Zabczyk (1978) introduce a continuous-time version of the problem with applicants arriving according to a Poisson process, which is closely related to the continuous-time problem we describe in Section 4. Their work has been extended by Bruss (1987) and by Kurushima and Ano (2003). Stewart (1981) studies a secretary problem with an unknown number of applicants which is also related to the problem we consider, but differs in the sense that he assumes n to be a random variable and the arrival times of offers to be i.i.d. exponential random variables, so that the decision-maker must maintain a belief distribution on n in order to optimize.

There has been considerable interest in explicitly modeling two-sided search and matching problems in the economics community. In particular, Burdett and Wright (1998) study two-sided search with nontransferable utility, which is relevant to our model because we assume exogenous offer values, implying that an employer cannot make her offer more attractive by, for example, offering a higher salary. The book by Roth and Sotomayor (1990) and the chapter by Mortensen and Pissarides (1999) both provide excellent background on this line of literature in economics.

1.2. Contributions.

This paper introduces a model of search processes where offers appear probabilistically and sequentially without explicit costs to sampling more offers, but with a limited number of possibilities that cannot be recalled. This is a good model for various job search and hiring processes where offers are “exploding” and search takes place during a fixed hiring season. Our main contributions can be summarized as follows:

a) We introduce two possible search processes, a “high information” process in which agents find out whether an offer appears or does not appear (this can also be thought of as agents being accepted or rejected) at each point in time, and a “low information” process in which agents only receive signals when an offer appears, so they do not know how many times they might have been rejected already.

b) We solve for the expected values of the low and high information processes for uniform and exponentially distributed offer values when agents know the underlying probability of offer appearance. We show that the expected utility in the low information process comes very close to

the expected utility in the high information process, and that the gap is widest in a critical range of expected number of offers between four and six.

c) We show that when agents do not know the true probability of offer appearance the expected utility in the low information process declines substantially relative to the high information process. This shows that the most important informational value of rejections lies in helping decision makers estimate their own “attractiveness,” when this attractiveness is measured in terms of the probability of offer appearance.

d) We introduce continuous time versions of the search processes, characterized by Poisson appearance of offers, and obtain closed form solutions for expected values of the high information processes. The solutions have a surprisingly simple form, which helps us gain insight into the dependence of the expected value on the offer arrival rate.

e) We evaluate the “competitive ratio” (in the sense used in computer science (Borodin and El-Yaniv 1998, e.g.)), which quantifies the relative reduction in the expected value, compared to the case where all offers are received simultaneously. We compare the competitive ratios of expected values in the stopping problem and the “simultaneous choice” problem to the ratios of expected values in the high and low information cases.

2. The Model.

We consider a search process in which a decision-maker (job-seeker) has to choose among n potential total offers, which appear sequentially. At each point in time, an offer either appears (with probability p), in which case its value w is revealed to the applicant, or does not appear (with probability $1 - p$). If an offer does not appear, the applicant may or may not be told this fact. For the purposes of this paper, we assume that all offers have an identical probability of appearance p , and that the values w are independently and identically distributed. We will consider two cases for the distribution of w , namely uniform and exponential. The job-seeker must decide immediately upon seeing an offer whether to accept it or not. If she accepts the offer, she receives utility w , and if she rejects it she may not recall that offer in the future.

We consider a number of variants of this process for the two distributions mentioned above. The two axes along which we parameterize the process are (a) whether or not the decision-maker knows the probability p of getting an offer; and (b) whether or not the decision-maker receives a signal when an offer does *not* appear. In the first case, the question is whether or not the decision-maker has to learn p . The second case essentially embodies two informational variants of the decision problem. In the high information variant, the decision-maker is told at each of the n stages whether

an offer appeared or not. Therefore, she always knows the exact total number of possible offers that may yet appear. In the low information variant, the decision-maker is only informed when an offer appears — if the offer does not appear the decision-maker is not informed of this event. Thus, the decision-maker does not know how many offers are potentially left out of the n total offers. We will begin by showing results about the informational variants assuming that the decision-maker knows p . In each case we will consider two distributions over the offers w_i , one a uniform $[0, 1]$ and the other an exponential distribution with rate parameter α . For calibration, when we report numerical results, we assume $\alpha = 2$ so that the expected values of draws from both distributions are the same (0.5).

2.1. An Example Where $n=2$.

As a motivating example, let us consider the case where $n = 2$, offer values are uniformly distributed in $[0, 1]$, and offers arrive with probability p . Later we will derive the expected values for general n . We can compute the expected value for an agent participating in the search process in the high and low information cases. In general, we will denote the expected value of the high information search process with n possible offers as H_n and the value of the low information process with n possible offers as L_n .

First, in the high information case, the agent knows that there are two time periods t in total, and she knows which time period she is in. At $t = 1$ the reservation value of an agent is her expected value if she declines the offer, which is just her expected value in the one period process. In the one period process, the agent should always accept any offer she receives, so the expected value is just the product of the probability that an offer appears and the expected value of that offer, or $0.5p$. Therefore, at $t = 1$, the agent should accept an offer only if it is greater than $0.5p$. Since offer values are distributed uniformly in $[0, 1]$, the probability that this is the case is $1 - 0.5p$. The expected value of the offer given that she does accept it is $(1 + 0.5p)/2$. The expected continuation value of the process if she rejects the offer is $0.5p$. Given that an offer arrives at $t = 1$ with probability p , the expected value of the search process is:

$$\begin{aligned} H_2 &= p \left((1 - 0.5p) \frac{1 + 0.5p}{2} + 0.5p(0.5p) \right) + (1 - p)(0.5p) \\ &= \frac{1}{8}p^3 - \frac{1}{2}p^2 + p \end{aligned}$$

The low information case is somewhat more complicated. The major difference from the high information case is that the decision-maker's threshold for stopping at the first offer to appear




Figure 1 Expected value of the difference between the high and low information cases as a function of p for $n = 2$ and values independently drawn from a uniform $[0, 1]$ distribution.

changes. When the decision-maker sees the first offer (assuming she ever sees an offer and has to make a decision), she does not know if the offer is first in the ordering or if the offer is second in the ordering and the first offer did not appear. The probability that she will see another offer is then the probability that a second offer will appear given that one has appeared. Suppose we denote realized appearance/non-appearance outcomes by vectors of zeros and ones where the zeros indicate non-appearance and the ones indicate appearance. The total space of outcomes is $\{[0\ 0], [0\ 1], [1\ 0], [1\ 1]\}$. The appearance of one offer reduces the possible space of outcomes to $\{[0\ 1], [1\ 0], [1\ 1]\}$. The probability that a second offer appears given that a first has appeared is then $p^2 / ((1-p)p + p(1-p) + p^2) = p / (2-p)$. Therefore the threshold for the decision-maker to stop at the first offer to appear is $p / (4-2p)$.

We compute the expected value of the process by analyzing each of the four possible realizations (see Appendix A) and conclude that:

$$L_2 = \frac{-5p^4 + 26p^3 - 48p^2 + 32p}{8(2-p)^2}$$

2.1.1. The Value of Information Simplifying the difference in expected values between the high and low information processes for $n = 2$, $D_2 = H_2 - L_2$, we find that:

$$D_2 = \frac{(p-1)p^3}{8(p-2)}$$

By setting the derivative to 0, we find that the difference is largest for $p = 0.7847$. Figure 1 shows the values of D_2 for p between 0 and 1.

3. The Search Process for General n .

This section provides the recursive solutions for the expected values of participating in the high and low information search processes. Solving for the expected value of the high information case is trivial, but it will serve as a point of comparison for the low information cases and will allow us to generalize to an interesting continuous time variant.

3.1. The High Information Case

When offers are drawn from a uniform $[0, 1]$ distribution, the recursive solution to the expected value in the high information process is given by:

$$H_n = p \left(\frac{1 + H_{n-1}^2}{2} \right) + (1-p)H_{n-1} \quad (1)$$

with the base case $H_1 = 0.5p$.

When offers are drawn from an exponential distribution with rate parameter α , the expected value is given by:

$$H_n = H_{n-1} + p \frac{1}{\alpha} e^{-\alpha H_{n-1}} \quad (2)$$

with the base case $H_1 = p \frac{1}{\alpha}$.

Complete derivations of these equations are given in Appendix B.

3.2. The Low Information Case

In the low information process with n total possible offers, any state at which the decision-maker has to take a decision can be completely characterized by n and by the number of offers that have appeared thus far, denoted by k . The expected value of not stopping at offer k (state (n, k)) is given by the product of the probability that state $[n, k+1]$ will be reached (if not, the decision-maker sees no more offers and gets utility 0) and the expected value $L(n, k+1)$.

The probability that state $(n, k+1)$ is reached given that (n, k) was reached is:

$$q_k = \frac{\sum_{i=k+1}^n \binom{n}{i} p^i (1-p)^{n-i}}{\sum_{i=k}^n \binom{n}{i} p^i (1-p)^{n-i}}$$

The continuation value of the process (the expected value of not stopping) is $q_k L(n, k+1)$. We know that $L(n, n) = 0.5$ for offers distributed uniformly in $[0, 1]$ and $L(n, n) = 1/\alpha$ for offers distributed exponentially with rate parameter α , so we can compute the expected value recursively. Let $z_k = q_k L(n, k+1)$ and w be the value of the k th offer to appear. Then, for the case where offers are distributed uniformly on $[0, 1]$:

$$L(n, k) = \Pr(w > z_k) E[w | w > z_k] + \Pr(w < z_k) z_k$$

$$\begin{aligned}
&= (1 - z_k) \frac{1 + z_k}{2} + z_k^2 \\
&= \frac{1}{2}(1 + z_k^2)
\end{aligned}$$

Similarly, for the case where offers are distributed exponentially with rate parameter α :

$$L(n, k) = \Pr(w > z_k)E[w|w > z_k] + \Pr(w < z_k)z_k = \frac{1}{\alpha}e^{-\alpha z_k} + z_k$$

The expected value of the n offer low information process is then $L_n = L(n, 0)$.

3.3. The Value of Information

Figure 2 shows the value of information for various different n and for the two distributions we consider. We can see that the critical region where the value of information is highest is reached at lower p for higher n – this happens when the expected value of the process is in an intermediate range. A rule of thumb is that the value of information is highest when the expected number of offers, np , is in the range of 4 to 6. We formalize this in a continuous time setting in the next section. The most important observation is that the information does not appear to be critical to making a good decision. Even in the worst of all the cases in Figure 2, the loss from participating in the low information process is only about 3%. Therefore, it seems clear that participants do not suffer great declines in expected utility from not being told when they are rejected, as long as they know the true probability p of offers appearing. In Section 5 we consider the case where p is unknown and show that the loss can be significantly higher.

4. Continuous Time Variants.

The natural continuous time limits of the process introduced in Section 2 involves Poisson arrivals of offers over a limited time horizon. We assume that offers arrive according to a Poisson process with arrival rate λ in the time interval $[0, 1]$. Again, the offer payoffs are sampled from either a uniform $[0, 1]$ distribution or an exponential distribution with rate parameter α , and the decision-maker has to decide upon seeing each offer whether to stop and accept that offer or continue searching. These continuous time variants allow us to abstract away from the particular number of possible offers and think in terms of the expected number of offers. We show that the high information processes have closed form solutions for the expected value at any point in time that allow us to gain insight into the dependence of the expected value on the expected number of offers. In this section we study and solve for the expected values of a decision-maker in the high and low information continuous time search processes, and discuss the relation between these processes and the discrete variants discussed above.




Figure 2 The ratio of the expected values of the low and high information processes for different values of n and p , for offer values drawn from the uniform $[0, 1]$ distribution (left) and the exponential distribution with rate parameter 2 (right).

4.1. The High Information Variant

In the high information variant, each time an offer appears, the decision-maker gets to see both the value of the offer, say w , and the precise time of appearance, t . The decision-maker should stop if w is greater than the continuation value $v(t)$. At any time t , to derive the continuation value we need to consider when the next offer will be received. At time t , the probability density function of the time of the next offer arrival (if any) is $\lambda e^{-\lambda(x-t)}$ for $x \leq 1$ (any density after 1 effectively “gets lost”). The value of receiving an offer at time x can be derived as in Section 3.

4.1.1. Uniform Distribution Let w be the random value of an offer received at time x . The value of receiving such an offer is:

$$\begin{aligned}
 & \Pr(w > v(x))E[w|w > v(x)] + \Pr(w < v(x))v(x) \\
 &= (1 - v(x))\left(v(x) + \frac{1 - v(x)}{2}\right) + v^2(x) && \text{(because } w \sim U[0, 1]) \\
 &= \frac{1}{2}(1 - v^2(x)) + v^2(x) \\
 &= \frac{1}{2}(1 + v^2(x))
 \end{aligned}$$

The continuation value at time t must satisfy:

$$v(t) = \int_t^1 \lambda e^{-\lambda(x-t)} \frac{1}{2}(1 + v^2(x)) dx$$

Therefore,

$$e^{-\lambda t}v(t) = \frac{1}{2}\lambda \int_t^1 e^{-\lambda x}(1+v^2(x)) dx$$

Differentiating with respect to t ,

$$(-\lambda v(t) + v'(t))e^{-\lambda t} = -\frac{1}{2}\lambda e^{-\lambda t}(1+v^2(t))$$

Since $v(1) = 0$ and $v \in [0, 1]$,

$$v'(t) = -\frac{1}{2}\lambda(v(t) - 1)^2$$

Or

$$\left(-\frac{1}{v(t) - 1}\right)' = -\frac{1}{2}\lambda$$

Integrating from t to 1,

$$\frac{1}{v(1) - 1} - \frac{1}{v(t) - 1} = \frac{1}{2}\lambda(1 - t)$$

Which gives us the solution:

$$v(t) = \frac{1 - t}{\frac{2}{\lambda} + 1 - t} \tag{3}$$

Therefore the value of a process with arrival rate λ is $v(0) = \lambda/(\lambda + 2)$.

4.1.2. Exponential Distribution The logic is exactly the same as above, except that with an exponential distribution with rate parameter α the continuation value at time t must satisfy

$$v(t) = \int_t^1 \lambda e^{-\lambda(x-t)} \left(\frac{1}{\alpha} e^{-\alpha v(x)} + v(x)\right) dx$$

Differentiating with respect to t , we get:

$$\Rightarrow v'(t) = -\frac{\lambda}{\alpha} e^{-\alpha v(t)}$$

or

$$v(t) = \frac{1}{\alpha} \log(-\lambda t + c)$$

where c is a constant of integration. Using the boundary condition $v(1) = 0$

$$v(t) = \frac{1}{\alpha} \log(-\lambda t + \lambda + 1) \tag{4}$$

Therefore, in this case the value of a process with arrival rate λ is $v(0) = \log(1 + \lambda)/\alpha$.

4.2. The Low Information Variant

In the low information variant of the continuous time process, the decision-maker knows only the number of offers she has received, not the precise time t at which any of the offers were received. Therefore, any time that a decision has to be made, the state is completely characterized by the number of offers received so far. Let the value of a process in which k offers have been received so far (but the decision-maker has not yet seen the value of the k th offer) be denoted by $v[k]$. Let w be the (unknown) value of the current offer. The continuation value of the process can then be computed in a manner exactly analogous to the discrete time case. Let

$$q_k = \Pr(\text{At least one more offer will be received} \mid k \text{ offers were received})$$

$$z_k = q_k v[k + 1]$$

Then

$$v[k] = \Pr(w > z_k) E[w \mid w > z_k] + \Pr(w < z_k) z_k$$

For offers distributed uniformly in $[0, 1]$, we have

$$v[k] = \frac{1}{2}(1 + z_k^2) \tag{5}$$

For offers distributed exponentially with rate parameter α , we have

$$v[k] = \frac{1}{\alpha} e^{-\alpha z_k} + z_k \tag{6}$$

There are two differences from the discrete case. First, q_k must be computed differently, because we now have Poisson arrivals. Let $f(k)$ be the Poisson probability mass function (the probability of getting exactly k offers) and $F(k)$ be the cumulative distribution function, for a particular value of λ . Then

$$\begin{aligned} q_k &= \frac{1 - F(k)}{1 - F(k - 1)} \\ &= 1 - \frac{f(k)}{1 - F(k - 1)} \end{aligned}$$

These are easily computed since we know that $f(k) = e^{-\lambda} \frac{\lambda^k}{k!}$ and $F(k) = \sum_{i=0}^k e^{-\lambda} \frac{\lambda^i}{i!}$.

The second difference from the binomial case is that we do not have an obvious base case, such as the case where n offers out of n are received, from which we can start a backwards recursion. However, we can show that $\lim_{k \rightarrow \infty} q_k = 0$.

$$\lim_{k \rightarrow \infty} q_k = \lim_{k \rightarrow \infty} \frac{1 - F(k)}{1 - F(k - 1)}$$

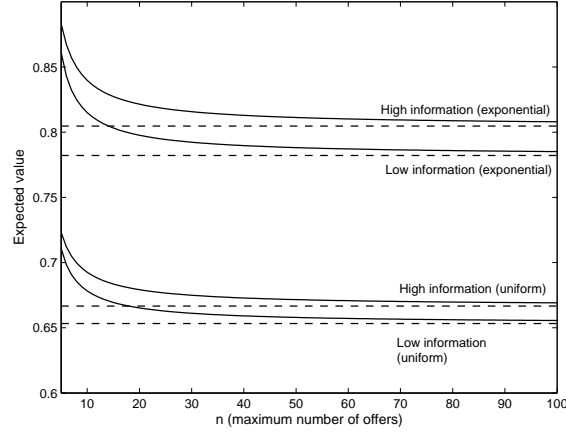


Figure 3 Expected values of the high and low information processes in continuous and discrete time holding $\lambda = pn$ constant (at $\lambda = 4$). Dashed lines represent the values of the continuous time processes and solid lines the values of the discrete time processes

$$\begin{aligned}
 &= \lim_{k \rightarrow \infty} \frac{-f'(k)}{-f'(k-1)} && \text{(Applying L'Hospital's Rule)} \\
 &= \lim_{k \rightarrow \infty} \frac{\lambda}{k} = 0
 \end{aligned}$$

Therefore, it is reasonable to approximate the actual value by assuming some threshold K such that $q_K = 0$ (the threshold K may depend on the particular value of λ). To convey a sense of the practical value of the threshold K we should note that a threshold such as $K = 200$ enables us to compute the expected values to a high degree of precision for λ as high as 100, since the probability of getting more than 200 offers is completely negligible for $\lambda = 100$. For higher λ values one would need to use higher thresholds.

4.3. Relation to the Discrete Time Process

Figure 3 shows that the expected values of the discrete time processes converge to the expected values of the continuous time variants as $n \rightarrow \infty$, while holding $\lambda = pn$ constant (other values of λ yield similar graphs). We can also show formally that the expected value of the continuous time high information process serves as a lower bound for the expected value of the discrete time high information process when offer values are distributed uniformly in $[0, 1]$.

THEOREM 1. *The value of the high information discrete-time process for given p and n is greater than the value of the high information continuous-time process with $\lambda = pn$, when offer values are drawn from a uniform $[0, 1]$ distribution.*

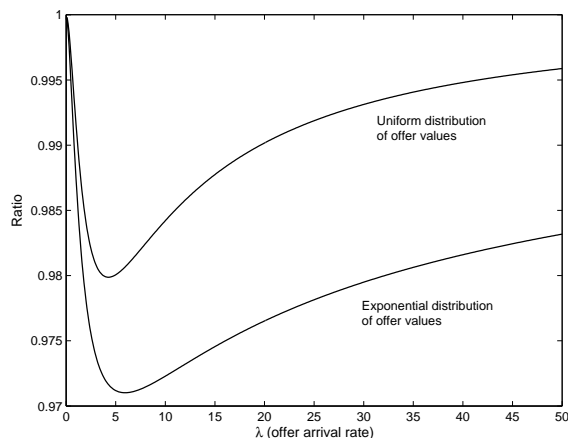


Figure 4 Ratio between expected values of the low and high information cases as a function of λ for the continuous time processes.

See Appendix C for the proof.

We also conjecture that Theorem 1 remains valid for the case of an exponential distribution (see Appendix C for further details) and that the low information expected values for the continuous time variants may also serve as lower bounds for the discrete time cases. The intuition is that the continuous time versions have a higher variance for the number of offers appearing (np as opposed to $np(1-p)$), which is why they yield lower expected values, especially for high values of p (corresponding to lower n since the product is held constant).

Interestingly, a difficult variant of the secretary problem (with the goal of maximizing the probability of selecting the best candidate) has been proposed and solved in continuous time by Cowan and Zabczyk (1978), and generalized by others (Kurushima and Ano (2003), Bruss (1987)). Our problem bears the same relation to this problem as the search problem with non-probabilistic appearance of offers (Gilbert and Mosteller (1966)) (recovered by using $p = 1$ in our case) does to the classical secretary problem.

4.4. The Value of Information

As n increases, the continuous time processes become a better approximation to the discrete time cases, and give us an opportunity to study general behavior without worrying about the specific interactions of n and p . Figure 4 shows the difference in expected value between the high and low information processes in continuous time expressed as a ratio. We can see that information is most important in a critical range of λ (between around $\lambda = 3$ and $\lambda = 10$, peaking between 4 and 6) for both distributions and the importance of information drops off quickly thereafter. Information is

also not particularly important if the expected total number of offers is very small. This confirms our intuitions from the discrete time cases.

5. What if p is Unknown?

In some search problems of the kind we have been discussing, the decision-maker may not have a good estimate of the probability p that any given offer will appear. In this case the decision-maker must update her estimate of p while also making decisions as before, with each decision based on her current estimate. This can greatly change the complexion of the problem, and especially of the value of information, because now knowing when an offer will not appear is not only useful for the decision problem, it is also useful for the problem of learning p to help in future decisions.

We will assume that a decision-making agent starts with a prior on p . In the experiments we report here, this prior always starts as a uniform $[0, 1]$ distribution. First, let us consider the high information case and two possible ways of representing and updating the agent's beliefs about p .

5.1. The High Information Case

5.1.1. Using a Beta Prior One possibility is to use a parameterized distribution. The ideal one for this case is the Beta distribution, because the two possible events at each time are success and failure, and the Beta distribution is its own conjugate and is particularly easy to update for this case. If the prior distribution on p before seeing the outcome of a binary event is a $\beta(i, j)$ distribution, then the posterior becomes $\beta(i + 1, j)$ in the event of a success and $\beta(i, j + 1)$ in the event of a failure. The $\beta(1, 1)$ distribution is uniform $[0, 1]$, and so the agent can start with that as the initial prior. Then, in order to compute the expected value of the game at any time after s successes and f failures have been seen, the agent only needs to additionally know the distribution of offer values and the total possible number of offers. However, the dynamic programming recursions are somewhat different than those in earlier sections. An agent who receives an offer and rejects it has a different expected value than an agent who does not receive an offer, due to the informational difference in her next estimate of p .

The value function is parameterized by n , the maximum number of possible offers remaining, s , the number of successes seen so far, and f , the number of failures seen so far.

For offer values distributed uniformly in $[0, 1]$ the expected value of the game is given by:

$$V(n, s, f) = \int_0^1 \eta(x, s + 1, f + 1) \left(x \frac{1}{2} (1 + V^2(n - 1, s + 1, f)) + (1 - x)V(n - 1, s, f + 1) \right) dx$$

where $\eta(x, s + 1, f + 1)$ represents the density function of the Beta $(s + 1, f + 1)$ distribution at x , that is the posterior after seeing s successes and f failures when starting with a Beta $(1, 1)$ prior.

Similarly, for offer values distributed exponentially with rate parameter α , the expected value is given by:

$$V(n, s, f) = \int_0^1 \eta(x, s+1, f+1) \left(x \left(\frac{1}{\alpha} e^{-\alpha V(n, s+1, f)} + V(n, s+1, f) \right) + (1-x)V(n-1, s, f+1) \right) dx$$

To actually compute these values, we can use a discrete approximation to the integral along the probability axis. V can be computed recursively backwards.

5.1.2. Using a Discrete Non-parametric Prior Another option is to simply use a discrete prior to begin with, and use the appropriate belief vector for subcomputations. The key to making this computation efficient is to note that an agent's beliefs will always be the same when s successes and f failures have been observed, regardless of the path. Therefore, the posterior at this time can be computed as:

$$\Pr(p = x | s, f) = \frac{\Pr(s \text{ successes out of } s+f | p=x) \Pr(p=x)}{\Pr(s \text{ successes out of } s+f)}$$

Here $\Pr(p=x)$ is the original prior.

5.2. The Low Information Case

In the low information case, the only information available to update the decision-maker's beliefs about p is the number of offers made so far. In this case, she must update as follows:

$$\Pr(p = x | s \text{ offers}) = \frac{\Pr(\text{at least } s \text{ offers} | p=x) \Pr(p=x)}{\Pr(\text{at least } s \text{ offers})}$$

The probability of getting at least s offers given that $p=x$ can be computed using the cumulative distribution function of the binomial distribution. Also note that the agent's beliefs about p will be the same every time that s successes have been observed.

5.3. Evaluating Performance

In order to estimate the expected utility received, we need to specify the form of learning the agent uses, the information available to the agent, and the true probability p of offer appearance. Then for particular values of p and n we can proceed by evaluating the expected value of a Markov chain in which states are characterized by the number of successes and failures seen so far. In either the high or low information cases, the agent will have a certain reservation value at each state that is completely dependent on the number of successes (in both cases) and failures (in the high information case) observed thus far. Then the expected value of being in that state can




Figure 5 Ratio of expected values of the high information and low information search processes when p is unknown, the agent starts with a uniform prior over $[0, 1]$ on p , and offers are drawn from a uniform $[0, 1]$ distribution (left) or an exponential distribution with rate parameter $\alpha = 2$ (right). Note that the Y axis is significantly different in the two cases.

be computed based just on the agent’s reservation value and the true underlying distribution of offer values and probability of offer appearance. For further details see Appendix D. There is no difference in the expected values for the high information game when using the Beta prior and when using the nonparametric prior, so we report results only from the use of the Beta prior. We first report results when agents start with a uniform prior over $[0, 1]$ for p .

Figure 5 shows results in terms of the value of information (corresponding to those in Figure 2 for the case of known p) for the uniform $[0, 1]$ distribution and the exponential distribution with $\alpha = 2$. There are three cases shown in each graph, corresponding to three true underlying probabilities. The first important thing to note is that there are much larger differences in the expected value between the high and low information cases than there were in the case where agents knew p beforehand. For the case of the uniform distribution, in both cases expected values are increasing and are bounded by 1, so the difference does not become as dramatic as for the exponential distribution. The reason why the ratios of expected values are so different is because in the high information case it is “easy” to learn p by updating your estimate based on seeing both when offers appear and they do not. In the low information case, the only information available does not help the agent nearly as much in updating her estimates.

A second interesting effect we see in the graphs is that the ratio declines precipitously for higher true values of p , especially for the exponential distribution. The reason for this huge decline is the tradeoff that an agent must make in her estimate – if there is a larger n then the agent is of course likely to receive more offers, so her threshold should be higher. However, the higher value




Figure 6 Ratio of expected values of the high information and low information search processes when p is unknown, the agent starts with a uniform prior over $[0.4, 0.6]$ on p , and offers are drawn from an exponential distribution with rate parameter $\alpha = 2$.

of n could also “explain away” the appearance of more offers, so that the agent does not realize that the true underlying p is higher. Consider an agent receiving her fourth offer when $n = 50$. Her threshold for accepting the offer cannot depend on p because she does not know p . This leads to decisions that look relatively “better” for different true underlying values. The same rule makes the agent perform better (relative to the high information case) for $p = 0.5$ than for $p = 0.1$ when $n = 5$, but much worse when $n = 20$. When $p = 0.1$ and $n = 5$, the agent is not sufficiently willing to accept offers, because a large part of the mass of her probability beliefs is on $p > 0.1$. However, when $p = 0.5$ and $n = 20$, the agent becomes too conservative and not risky enough in *rejecting* offers, because the appearance of offers does not necessarily tell her that p is higher, it might just be a function of the fact that there are a large number of total possible offers. She thus becomes more likely to take an offer that is not actually of high enough value.

A question that arises in this context is that of what happens when the agent has a less diffuse prior. In many ways this might correspond to a more realistic situation. Suppose she knows that her true probability of receiving offers is definitely between 0.4 and 0.6 when it is actually 0.5. We studied this question by calculating the ratios of expected values of the low and high information processes when the agent starts with a uniform prior on $[0.4, 0.6]$ (modeled using discrete probability masses, and using the nonparametric technique in the high information case as well as the low information case). The results are shown in Figure 6. We can see that the ratio actually appears to remain constant (and significantly higher than before) as n increases, showing that the expected value goes down much less as we move to the low information case, as we would expect given that the case of known p is the limit of concentrating the prior.

6. Comparison of Mechanisms: Sequential vs. Simultaneous Choice

So far, we have considered the loss from lack of information within a particular mechanism, a sequential choice mechanism which introduces a stopping problem for the decision maker. In this section we ask a different set of questions – namely, what is the loss from using the sequential choice mechanism itself? This has been an important consideration for previous work on secretary problems and on optimal stopping more generally. We will focus on the difference between the high information case with sequential choice and what we call the simultaneous choice case, in which all offers appear simultaneously, and the decision maker can simply choose the best one. In continuous time, the simultaneous choice case is simply one in which all the appearances are realized, and then at time 1, the decision maker gets to choose the best out of all the realized options. It can also be thought of as allowing the decision-maker to backtrack to previous choices.

First let us consider the continuous time case. What is the expected value of participating in a simultaneous choice process with arrival rate λ ? It is the sum over all k of the probability that exactly k offers appear and the expected value given that exactly k offers appear. Appendix E derives these values for the case where offer values are distributed uniformly in $[0, 1]$ and the case where offer values are distributed exponentially with rate parameter α . In the uniform case this expected value is $1 - \frac{1-e^{-\lambda}}{\lambda}$ and for the exponential case it is $\frac{1}{\alpha}[\gamma + \Gamma(0, \lambda) + \log(\lambda)]$ where γ is the Euler constant and Γ represents the (upper) incomplete gamma function.

We already know the expected values of the sequential choice high information processes for both distributions. Figure 7 shows the differences in expected values between the simultaneous and sequential choice cases. Note that the difference can be an order of magnitude higher in this case than it was between the high and low information variants with known p (Figure 4), revealing that the difference in expected value changes much more dramatically when going from one mechanism to another than it does when going from the higher to lower information variant of the sequential choice process. However, the difference can be of the same order of magnitude when going from high to low information in the case where p is unknown. Also note that the shape of the graph is very similar to Figure 4, and the greatest differences are achieved for similar values of λ .

6.1. Some More Search Processes

These results bring up some more questions, which we will pose and answer for the uniform distribution in order to illustrate the differences between the mechanisms we have discussed and some other possible variants. Therefore, results in this section are confined to cases where offer values are generated from a uniform $[0, 1]$ distribution.




Figure 7 Ratio of expected values of the simultaneous choice mechanism and the sequential choice mechanisms with high information as a function of λ for the continuous time processes.

The first question that arises is how the expected values of the processes we are considering compare to the expected values in a comparable non-probabilistic case, in which the total number of appearances is fixed and the decision-maker knows this number? Gilbert and Mosteller discuss the latter case and present a recurrence relation that is also easily derived by setting $p = 1$ in equation 1:

$$H_{n+1} = \frac{1}{2}(1 + H_n^2)$$

Figure 8 shows the ratios of expected values in three different processes. The first is the high information continuous time process with arrival rate λ . In the other two cases, let us postulate the existence of a Gamesmaster, who first stores all the offers generated according to the Poisson process, and then informs the decision-maker of the total number of offers that appeared. The Gamesmaster then presents the offers to the decision-maker, either sequentially or simultaneously. Obviously, the expected value of the simultaneous process is highest, since it is the best decision that the job-seeker can make *retrospectively* (or if she were omniscient with respect to what offers she would receive). The expected value of the sequential process with a known number of offers is also bound to be significantly higher since it eliminates uncertainty about the exact number of offers the decision-maker will receive. Figure 8 shows the ratios of expected values of these three processes. The continuous-time process has a substantially lower expected value than the sequential process with a known number of offers for values of λ below 10, but approaches it much more rapidly than either of the sequential mechanisms approaches the simultaneous mechanism in terms of expected value. The dropoff in expected value between the continuous-time and the sequential




Figure 8 Ratios of expected values in three processes: the high information continuous-time process with Poisson arrival rate λ (denoted “High”), and two processes in which the number of offers are known beforehand after being generated by a Poisson distribution with parameter λ . The decision maker has no recall and must solve a stopping problem in the sequential choice process (denoted “Seq”), but chooses among all realized offers in the simultaneous choice process (denoted “Sim”).

process with known n is particularly dramatic for very small λ , indicating that knowing the exact number of offers you will receive is much more important if you only expect to receive 1-3 offers.

Figure 8 focuses on processes generated from an underlying process with Poisson offers arriving in continuous time, and therefore we (as the experimental designers) possess a fundamental uncertainty about the number of offers arising in each case. In contrast to this, Figure 9 shows the difference in expected values between two sequential processes, one with a fixed and known number of offers pn and the other one with n possible offers that each appear with probability p . While the expected value ratios are substantially smaller when pn is smaller, this is mostly because of the large probability of getting no offers. The tradeoff of possibly getting more offers is clearly not worth it in expectation, but much more so for lower values of pn . An interesting question to ask in this case is, for example, whether it is better to have one offer for sure, or 10 possible offers, each with a 20% chance of appearing (the latter, by a hair: it has expected value 0.5183, as opposed to 0.5 for the former).

7. Conclusions

This paper is intended to highlight the importance of the information structure in search processes, particularly processes that run over a fixed period of time, such as academic job markets. It is common practice in markets of this kind for employers or job candidates to not keep the other side fully informed about the decisions they have made. For example, universities will often not




Figure 9 Ratio of expected values in the high information probabilistic process (denoted “High” with probability p and n total possible offers) and a process in which the number of offers is known beforehand and is equal to pn (denoted “Seq”).

send rejections to candidates until they have completed their search, even if they were no longer seriously considering a candidate much earlier in the process. In order to study the expected loss of participating in such a process compared to a process in which both sides immediately make decisions and have to inform each other about those decisions, we have introduced a stylized model of this process that analyzes it from a one-sided perspective. Our main result is that the loss from participating in the low information process is not significant unless the decision-maker is not well-informed about her own “attractiveness,” measured by the probability of receiving an offer. This suggests that the costs to changing the structure of markets that operate in the “low information” manner may not be worthwhile. If applicants are poorly informed about their own attractiveness to employers, one could imagine mechanisms to improve signaling rather than restructuring the market (of course, this assumes that employers, who participate in these processes repeatedly, can estimate their attractiveness to employees well).

The model we have introduced simplifies the problem along some dimensions. We do not incorporate two-sided strategic considerations, which may become important; for example, less attractive employers may be more inclined to make exploding offers, while more attractive employers are unlikely to do so. Further, the assumption that the probability p of receiving an offer is independent of the value of the offer may be unrealistic for some markets. Future studies should focus on these directions for extending our model.

Acknowledgments

We would like to thank Andrea Caponnetto and Tommy Poggio for useful suggestions. This research was partially supported by the National Science Foundation under contract ECS-0312921 and partly by grants to CBCL from Merrill-Lynch, the National Science Foundation, the Center for e-Business at MIT, the Eastman Kodak Company, Honda R&D Co, and Siemens Corporate Research, Inc.

Appendix A: Low Information Expected Value for $n = 2$

The four possible cases for the low information process when $n = 2$ can be analyzed as follows (where, as in Section 2.1 0 denotes non-appearance of an offer and 1 denotes appearance):

1. [0 0] : Occurs with probability $(1 - p)^2$ and has value 0.
2. [0 1] : Occurs with probability $(1 - p)p$. The offer which appears is accepted with probability $1 - \frac{1}{2} \frac{p}{2-p}$, and if rejected, the utility received is 0. Therefore, the expected value is:

$$\begin{aligned} & \Pr\left(w > \frac{p}{2(2-p)}\right) E[w|w > \frac{p}{2(2-p)}] \\ &= \left(1 - \frac{p}{2(2-p)}\right) \left(\frac{p}{2(2-p)} + \frac{1}{2}\left(1 - \frac{p}{2(2-p)}\right)\right) \\ &= \frac{3p^2 - 16p + 16}{8(2-p)^2} \end{aligned}$$

3. [1 0] : Precisely the same argument as the previous case, with the same probability and expected value.
4. [1 1] : Occurs with probability p^2 . In this case, if the first offer to appear is rejected, the second offer is automatically going to be selected. Therefore the expected value will be the sum of the above expected value and the expected value of the second given that the first is rejected (weighted by the probability of the first being rejected). The additional term is then:

$$\begin{aligned} & (1 - \Pr(w > \frac{1}{2} \frac{p}{2-p}))(1/2) \\ &= \frac{p}{4(2-p)} \end{aligned}$$

Adding this to the expected value for the previous case and simplifying gives:

$$\frac{p^2 - 12p + 16}{8(2-p)^2}$$

Then the total expected value is:

$$\begin{aligned} L_2 &= p(1-p) \frac{3p^2 - 16p + 16}{8(2-p)^2} + p^2 \frac{p^2 - 12p + 16}{8(2-p)^2} \\ &= \frac{-5p^4 + 26p^3 - 48p^2 + 32p}{8(2-p)^2} \end{aligned}$$

Appendix B: Derivation of Dynamic Programming Equations

This section derives the equations for computing the expected value of participating in the high information search process for general n and arbitrary p . In both cases, the base case is the expected value when $n = 1$, which is given by the product of the probability of an offer appearing (p) and the expected value of the offer given that it does appear (0.5 when offers are distributed uniformly in $[0, 1]$ and $1/\alpha$ when offers are

distributed exponentially with rate parameter α . Also, in all cases when there are n possible offers remaining, the threshold for accepting an offer should be the expected value of the search process with $n - 1$ possible offers. Let w denote the value of the offer:

$$H_n = p[\Pr(w > H_{n-1})E(w|w > H_{n-1}) + (1 - \Pr(w > H_{n-1}))H_{n-1}] + (1 - p)H_{n-1}$$

B.1. Uniform $[0, 1]$ Distribution

In this case,

$$\Pr(w > H_{n-1}) = 1 - H_{n-1}$$

$$E(w|w > H_{n-1}) = H_{n-1} + \frac{1 - H_{n-1}}{2} = \frac{1 + H_{n-1}}{2}$$

This gives us:

$$\begin{aligned} H_n &= p((1 - H_{n-1})\frac{1 + H_{n-1}}{2} + H_{n-1}^2) + (1 - p)H_{n-1} \\ &= p\frac{1 + H_{n-1}^2}{2} + (1 - p)H_{n-1} \end{aligned}$$

and we know $H_1 = 0.5p$.

B.2. Exponential Distribution with Rate Parameter α

In this case,

$$\begin{aligned} \Pr(w > H_{n-1}) &= \int_{H_{n-1}}^{\infty} \alpha e^{-\alpha x} dx \\ &= e^{-\alpha H_{n-1}} \end{aligned}$$

$$\begin{aligned} E(w|w > H_{n-1}) &= \int_0^{\infty} \alpha e^{-\alpha x} (x + H_{n-1}) dx && \text{(Using the memorylessness property)} \\ &= \frac{1}{\alpha} + H_{n-1} \end{aligned}$$

Therefore,

$$\begin{aligned} H_n &= p[e^{-\alpha H_{n-1}}(\frac{1}{\alpha} + H_{n-1}) + (1 - e^{-\alpha H_{n-1}})H_{n-1}] + (1 - p)H_{n-1} \\ &= p[\frac{1}{\alpha}e^{-\alpha H_{n-1}} + H_{n-1}] + (1 - p)H_{n-1} \\ &= p\frac{1}{\alpha}e^{-\alpha H_{n-1}} + H_{n-1} \end{aligned}$$

and we know $H_1 = p\frac{1}{\alpha}$.

Appendix C: Proof of Theorem 1 (Lower Bound)

In this section we consider the high information cases in both discrete and continuous time. We show that, in addition to being an approximation of the value of the process for large n , the values of the continuous time processes function as a lower bound for the values of discrete time processes where $pn = \lambda$ for the case where offer values are distributed uniformly in $[0, 1]$. This is Theorem 1, as initially stated in Section 4.

Let us denote the value of the discrete-time process by $H[i]$, where i is the number of offers that have appeared in the past, and the continuation value of the continuous time process at time t by $v(t)$. We want to show that, when $\lambda = pn$, $H[0] > v(0)$. We shall proceed by induction, showing that, for given p and n ,

$$H[i] > v(i/n), \quad \forall i < n$$

We know that

$$v(t) = \frac{1-t}{2/\lambda + 1-t} = \frac{\lambda(1-t)}{2 + \lambda(1-t)}$$

For $i = n-1$, we have $H[n-1] = 0.5p$ because the value is sampled from the uniform $[0, 1]$ distribution, and

$$\begin{aligned} v\left(\frac{n-1}{n}\right) &= \frac{\lambda\left(1 - \frac{n-1}{n}\right)}{2 + \lambda\left(1 - \frac{n-1}{n}\right)} \\ &= \frac{\frac{\lambda}{n}}{2 + \frac{\lambda}{n}} \\ &= \frac{p}{2+p} && \text{(because } \lambda = np\text{)} \\ &< \frac{1}{2}p && \text{(because } p \in [0, 1]\text{)} \\ &= H[n-1] \end{aligned}$$

Now, given that $H[i] > v(i/n)$ we have to show that $H[i-1] > v((i-1)/n)$ for integral $i \geq 1$, which will complete the proof. Let $X = v(i/n)$. Then

$$\begin{aligned} H[i-1] &= p\left(\frac{1}{2}(1 + H[i]^2)\right) + (1-p)H[i] \\ &> \frac{1}{2}p + \frac{1}{2}pX^2 + (1-p)X && \text{(inductive hypothesis)} \\ &= \frac{1}{2}p(1 + X^2 - 2X) + pX + X - pX \\ &= \frac{1}{2}p(1 - X)^2 + X \end{aligned}$$

In order to complete the induction step, it is therefore sufficient to show that

$$\frac{1}{2}p(1 - X)^2 > v((i-1)/n) - X$$

Simplifying the right hand side, we get

$$\begin{aligned}
v((i-1)/n) - X &= \frac{2\lambda n}{(2n + \lambda n - \lambda i + \lambda)(2n + \lambda n - \lambda i)} \\
&= \frac{2\lambda n}{(2n + \lambda n - \lambda i)^2 + \lambda(2n + \lambda n - \lambda i)} \\
&< \frac{2\lambda n}{(2n + \lambda n - \lambda i)^2} \\
&= \frac{1}{2}p(1 - X)^2
\end{aligned}$$

which completes the proof.

CONJECTURE 1. *The value of the high information discrete-time process for specified p and n is greater than the value of the high information continuous-time process with $\lambda = pn$ when offer values are drawn from an exponential distribution.*

This conjecture may not be provable by induction. While the base case is simple enough to prove, the problem is that the difference between two “consecutive” instances of the continuous time process is not always smaller than the differences between the corresponding cases of n and $n - 1$ in the discrete time case.

Appendix D: Expected Values with Unknown p

We can evaluate the expected value of the search process for a given true underlying p and n and a given initial prior by describing the process as a Markov chain whose state consists of the number of past successes and failures (s and f , respectively).

In the high information case, the reservation value of an agent is dependent on s , f , and n , while in the low information case, the reservation value only depends on s and n . Suppressing the dependence on n , denote the reservation value in the high information case by $R_h(s, f)$ and in the low information case by $R_l(s)$. The reservation value at state s is the expected value of the process if the agent does not accept an offer that appears. This is important because the appearance of the offer is itself informative.

Let w be the value of an offer that does appear. Let V_s denote the value of state $(s + 1, f)$ and V_f denote the value of state $(s, f + 1)$. The value of state (s, f) is 0 when $s + f \geq n$.

Then in the high information case, the value of state (s, f) is:

$$p(\Pr(w > R_h(s, f))E[w|w > R_h(s, f)] + \Pr(w < R_h(s, f))V_s) + (1 - p)V_f$$

In the low information case, the value of state (s, f) is (the decision-making agent does not have access to f , but we use it when evaluating the chain):

$$p(\Pr(w > R_l(s))E[w|w > R_l(s)] + \Pr(w < R_l(s))V_s) + (1 - p)V_f$$

The actual reservation values at any given state can be precomputed and stored in a table, since they are completely independent of the value of the state. Then the Markov chain can be evaluated based on this table and the known true probability p .

Appendix E: Expected Values of Simultaneous Choice Processes

For all continuous time models, offers arrive as a Poisson process, and the probability of exactly k offers is given by $\frac{e^{-\lambda}\lambda^k}{k!}$.

For offer values distributed uniformly in $[0, 1]$, if k choices are available, the expected value is $\frac{k}{k+1}$ (from the order statistic of the uniform distribution). Then the expected value of the process is:

$$\begin{aligned} \sum_{i=0}^{\infty} \Pr(i \text{ successes}) \frac{i}{i+1} &= \sum_{i=0}^{\infty} \frac{e^{-\lambda}\lambda^i i}{i!(i+1)} \\ &= \frac{e^{-\lambda}}{\lambda} \sum_{i=0}^{\infty} \left[\frac{(i+1)\lambda^{i+1}}{(i+1)!} - \frac{\lambda^{i+1}}{(i+1)!} \right] \\ &= 1 - \frac{1 - e^{-\lambda}}{\lambda} \end{aligned}$$

The expression for the expected value for offers distributed exponentially with rate parameter α is slightly more complex. First note that the distribution function for the maximum of k such random variables is:

$$f(x) = k[1 - e^{-\alpha x}]^{k-1} \alpha e^{-\alpha x}$$

Therefore the expected value of the maximum is:

$$k\alpha \int_0^{\infty} [e^{-\alpha x}(1 - e^{-\alpha x})^{k-1} x] dx = \frac{\mathcal{H}_k}{\alpha}$$

where \mathcal{H}_i represents the i th harmonic number.

Then the expected value is given by:

$$\begin{aligned} \sum_{i=0}^{\infty} \Pr(i \text{ successes}) \frac{H_i}{\alpha} &= \frac{e^{-\lambda}}{\alpha} \sum_{i=0}^{\infty} \frac{\lambda^i H_i}{i!} \\ &= \frac{1}{\alpha} [\gamma + \Gamma(0, \lambda) + \log(\lambda)] \end{aligned}$$

where γ is the Euler constant and Γ represents the (upper) incomplete gamma function.

References

- Borodin, Allan, Ran El-Yaniv. 1998. *Online Computation and Competitive Analysis*. Cambridge University Press, Cambridge, UK.
- Bruss, F.Thomas. 1987. On an optimal selection problem by Cowan and Zabczyk. *Journal of Applied Probability* **24** 918–928.
- Burdett, K., R. Wright. 1998. Two-sided search with nontransferable utility. *Review of Economic Dynamics* **1** 220–245.
- Cowan, A. J. Zabczyk. 1978. An optimal selection problem associated with the Poisson process. *Theory of Probability and its Applications* **23** 584–592.
- Das, Sanmay, Emir Kamenica. 2005. Two-sided bandits and the dating market. *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence*. Edinburgh, UK, 947–952.

- DeGroot, Morris H. 1970. *Optimal Statistical Decisions*. McGraw-Hill, New York.
- Ferguson, Thomas S. 1989. Who solved the secretary problem? *Statistical Science* **4**(3) 282–289.
- Gilbert, John, Frederick Mosteller. 1966. Recognizing the maximum of a sequence. *Journal of the American Statistical Association* **61** 35–73.
- Kurushima, Aiko, Katsunori Ano. 2003. A note on the full-information poisson arrival selection problem. *Journal of Applied Probability* **40** 1147–1154.
- Mortensen, Dale T., Christopher A. Pissarides. 1999. New developments in models of search in the labor market. *Handbook of labor economics*, vol. 3B. Elsevier Science, North-Holland, Amsterdam, 2567–2627.
- Roth, A. E., M. Sotomayor. 1990. *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. Econometric Society Monograph Series, Cambridge University Press, Cambridge, UK.
- Stewart, T.J. 1981. The secretary problem with an unknown number of options. *Operations Research* **29**(1) 130–145.