

REPORT DOCUMENTATION PAGE			Form Approved OMB NO. 0704-0188		
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA, 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) 09-05-2017		2. REPORT TYPE Final Report		3. DATES COVERED (From - To) 1-Apr-2014 - 30-Mar-2017	
4. TITLE AND SUBTITLE Final Report: SOME STOCHASTIC OPTIMIZATION PROBLEMS			5a. CONTRACT NUMBER W911NF-14-1-0166		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER 611102		
6. AUTHORS Sheldon Ross			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAMES AND ADDRESSES University of Southern California Contracts & Grants 3720 S. Flower St. Los Angeles, CA 90089 -0701			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211			10. SPONSOR/MONITOR'S ACRONYM(S) ARO		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S) 65391-MA.12		
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for Public Release; Distribution Unlimited					
13. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.					
14. ABSTRACT We study a variety of stochastic optimization models related to assignment problems, scheduling models, queueing loss models, bandit problems, and others. We have obtained optimal policies in many cases, good heuristic ones in others, and have presented ways to analyze properties of these policies. In most cases, we have considered both static and dynamic policies. We have also analyzed stochastic systems related to queueing models with setup times, win probabilities in random knockout tournaments, and have obtained a new variation of the friendship paradox that "your friends tend to have more friends than you do". In addition, we have made contributions to					
15. SUBJECT TERMS stochastic optimality, queues with setup, socially and individually optimal policies, scheduling with impatient jobs, bounding Poisson approximation errors, friendship paradox, random knockout tournaments, bandit problem					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON Sheldon Ross
a. REPORT UU	b. ABSTRACT UU	c. THIS PAGE UU			19b. TELEPHONE NUMBER +12-137-4048

Report Title

Final Report: SOME STOCHASTIC OPTIMIZATION PROBLEMS

ABSTRACT

We study a variety of stochastic optimization models related to assignment problems, scheduling models, queueing loss models, bandit problems, and others. We have obtained optimal policies in many cases, good heuristic ones in others, and have presented ways to analyze properties of these policies. In most cases, we have considered both static and dynamic policies. We have also analyzed stochastic systems related to queueing models with setup times, win probabilities in random knockout tournaments, and have obtained a new variation of the friendship paradox that "your friends tend to have more friends than you do". In addition, we have made contributions to stochastic model theory by showing how to improve Poisson approximation bounds when applied to the probability that at least one of a set of events occurs. In all cases, we have presented a variety of possible applications of our model.

Enter List of papers submitted or published that acknowledge ARO support from the start of the project to the date of this printing. List the papers, including journal references, in the following categories:

(a) Papers published in peer-reviewed journals (N/A for none)

<u>Received</u>	<u>Paper</u>	
05/05/2017	4 Sheldon Ross. A Sequential Scheduling Problem with Impatient Jobs, Naval Research Logistics, (07 2015): 659. doi:	361,395.00
05/05/2017	11 Babak Haji, Sheldon Ross. A Queueing Loss Model with Heterogeneous Skill Based Servers and Discriminating Arrivals, Naval Research Logistics, (): . doi:	1,040,682.00
05/05/2017	10 Yang Cao. Multiple Service Preemptive Scheduling with Impatience, Prob. in Eng. and Inf. Sciences, (): 226. doi:	1,040,681.00
05/05/2017	8 Ilan Adler, Yang Cao, Richard Karp, Erol Pekoz, Sheldon Ross. Random Knockout Tournaments, Operations Research, (): . doi:	1,040,677.00
05/05/2017	9 Maher Nouiehed, Sheldon Ross. A Feldman Bandit Problem Conjecture, Journal of Applied Probability, (): . doi:	1,040,679.00
05/05/2017	7 Yang Cao, Sheldon Ross. The Friendship Paradox, The Mathematical Scientist, (): 61. doi:	374,483.00
05/05/2017	6 Mark Brown, Sheldon Ross. Optimality Results for Coupon Collecting, Journal of Applied Probability, (): 930. doi:	374,481.00
05/05/2017	5 Sheldon Ross. Improved Chen-Stein Bounds on the Probability of a Union, Journal of Applied Probability, (07 2015): 1265. doi:	361,396.00
07/30/2015	1 . M/G/infinity with Exponentially Distributed Setup Times, Operations Research Letters, (04 2015): 26. doi:	331,881.00
07/30/2015	2 Sheldon Ross. An Individual and Socially Optimal Policy Minimizing Expected Flow Times, Probability in the Engineering and Informational Sciences, (04 2015): 147. doi:	361,389.00
07/30/2015	3 David Wu, Sheldon Ross. A Stochastic Assignment Problem, Naval Research Logistics, (04 2015): 23. doi:	361,391.00
TOTAL:	11	

Number of Papers published in peer-reviewed journals:

(b) Papers published in non-peer-reviewed journals (N/A for none)

Received Paper

TOTAL:

Number of Papers published in non peer-reviewed journals:

(c) Presentations

Number of Presentations: 0.00

Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

Received Paper

TOTAL:

Number of Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

Peer-Reviewed Conference Proceeding publications (other than abstracts):

Received Paper

TOTAL:

Number of Peer-Reviewed Conference Proceeding publications (other than abstracts):

(d) Manuscripts

Received Paper

TOTAL:

Number of Manuscripts:

Books

Received Book

TOTAL:

Received Book Chapter

TOTAL:

Patents Submitted

Patents Awarded

Awards

Sheldon Ross was named the 2016 Outstanding Alumni of Mathematics Department of Purdue University

Graduate Students

<u>NAME</u>	<u>PERCENT SUPPORTED</u>	<u>DISCIPLINE</u>
Yang Cao	25	
FTE Equivalent:	0.25	
Total Number:	1	

Names of Post Doctorates

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
FTE Equivalent:	
Total Number:	

Names of Faculty Supported

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
FTE Equivalent:	
Total Number:	

Names of Under Graduate students supported

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
FTE Equivalent:	
Total Number:	

Student Metrics

This section only applies to graduating undergraduates supported by this agreement in this reporting period

The number of undergraduates funded by this agreement who graduated during this period: 0.00

The number of undergraduates funded by this agreement who graduated during this period with a degree in science, mathematics, engineering, or technology fields:..... 0.00

The number of undergraduates funded by your agreement who graduated during this period and will continue to pursue a graduate or Ph.D. degree in science, mathematics, engineering, or technology fields:..... 0.00

Number of graduating undergraduates who achieved a 3.5 GPA to 4.0 (4.0 max scale):..... 0.00

Number of graduating undergraduates funded by a DoD funded Center of Excellence grant for Education, Research and Engineering:..... 0.00

The number of undergraduates funded by your agreement who graduated during this period and intend to work for the Department of Defense 0.00

The number of undergraduates funded by your agreement who graduated during this period and will receive scholarships or fellowships for further studies in science, mathematics, engineering or technology fields:..... 0.00

Names of Personnel receiving masters degrees

<u>NAME</u>
Total Number:

Names of personnel receiving PhDs

<u>NAME</u> Babak Haji Total Number:	1
---	----------

Names of other research staff

<u>NAME</u> FTE Equivalent: Total Number:	<u>PERCENT SUPPORTED</u>
---	--------------------------

Sub Contractors (DD882)

Inventions (DD882)

Scientific Progress

Technology Transfer

1 Problems Studied

We studied a variety of stochastic optimization models related to assignment problems, scheduling models, queueing loss models, bandit problems, and others. For these models, we have obtained optimal policies in many cases, good heuristic ones in others, and have presented ways to analyze the properties of these policies. In many cases, we have considered both static and dynamic policies. We have also analyzed stochastic systems related to queueing models with setup times, win probabilities in random knockout tournaments, and have obtained a new variation of the friendship paradox that “your friends tend to have more friends than you do”. In addition, we have made general contributions to stochastic model theory by showing how to improve Poisson approximation bounds when applied to the probability that at least one of a set of events occurs.

In [1] we studied a stochastic assignment problem in which an organization has different type jobs to fill, with each job type requiring a given quota of workers. Workers arrive in sequence. After observing which job types an arrival is qualified for, a decision must be made as to their assignment. Assuming that each arrival is independently qualified for job type i with a specified probability $p(i)$, the objective is to minimize the expected number of workers observed until all jobs are filled. The paper [1] allows for both static priority policies as well as dynamic policies. A static priority policy is one determined by a priority ordering of the job types, with the interpretation being that a worker is assigned to the first type job in this ordering that both has not yet met its quota and for which the worker is qualified, whereas a dynamic policy takes remaining quotas into account when decisions are made. The model of [1] can also be applied to channel scheduling in telecommunication systems. Assume in a time division multiple access network that multiple users share one channel, with each user having a given number of packets that it wants to transmit. In each time slot, each user will have its own probability that it will be connected and so able to transmit. However, although multiple users can be connected, only one is allowed to transmit at a time, with a central controller making the assignment decisions. Under the objective of minimizing the mean time until all users have transmitted all their packets, this model is easily seen to be equivalent to our problem.

In [2] we considered a model in which a fixed number of customers are to be served one at a time. Assuming that each customer has its own (a) randomly distributed time that it will wait in queue before departing the system; (b) distribution of service time; and (c) reward that is paid upon service completion, the problem is to dynamically choose the order of customer service so as to maximize the expected total return. It was noted that this model has important applications, two of which are as follows.

(a) One is to a model in which there are multiple machines that need to be upgraded (or maintained) one at a time. The time to upgrade each machine is random with a specified distribution depending on the machine. While waiting to be upgraded a machine may fail. In addition, a cost is incurred if a machine fails before its upgrade is begun. Both the distribution of failure time and the cost incurred depend on the machine. Suppose that a machine cannot fail while it is being upgraded, and a failed machine cannot be upgraded. Each time an upgrade is completed a decision must be made as to which of the yet to be upgraded but still working machines should be upgraded next, with the objective being to minimize the total expected cost.

(b) When all failure distributions are the same, the model of [2] can be considered a model for optimally dispensing randomly arriving organs to patients needing transplants. Without a transplant, each of the patients can only survive for a randomly distributed time. Organs arrive according to a renewal process and each patient has its own societal value if they receive an organ. The objective is to maximize the expected sum of the societal values received.

In [3] we consider a model in which there are n servers with different exponential service distributions. We suppose that all servers are initially busy and that m customers are waiting in an ordered line. We also suppose that upon completing a service, that server is offered to the first person in line; if that person refuses service then the server is offered to the second person in line; and so on. The objective of each person is to minimize their individual expected time until they complete service. Using the theory of one stage lookahead stopping rule policies, we determine the optimal policy for each customer and also show that this policy is socially optimal in the sense that it minimizes the expected sum of the times that each customer spends in the

system.

In the paper [4] we improve the well-know Chen-Stein Poisson approximation bounds when applied to the probability of a union. When the probability is small, the improvement in the distance from the lower to the upper bound is roughly by a factor of 2. Further improvements are determined when the events of the union are negatively dependent. The significance of this result resides in the fact that many important applications of probability involve finding the probability that at least one of a given set of events occur. Let us mention a few.

- In the generalized coupon collecting problem we suppose that each coupon collected is one of k types with specified probabilities, and we are concerned with the number of coupons that need be collected until there are at least $r(i)$ coupons of type i for each i . In this example, the probability that it takes more than n coupons is the probability that for some i there is fewer than $r(i)$ type i coupons among the first n collected.
- The generalized birthday problem is interested in the number of coupons one needs collect until one has $r(i)$ type i coupons for any i . Here the probability that it takes n or fewer coupons is the probability of the union of the events that there are at least $r(i)$ type i coupons in the first n collected.
- In reliability systems with multiple components we often suppose there are specified subsets of components such that the system fails if and only if all components in at least one of these subsets are failed. Hence, the probability that the system fails is the probability of a union.
- In pattern problems we are often concerned with the time until a certain pattern appears. If the pattern is of length k , then the probability it would appear within the first n data values is the probability of the union of the events $A(i)$ for i going from k to n , where $A(i)$ is the event that the data values $i-k+1$ to i constitute the pattern.

The paper [5] is concerned with an infinite server queueing model where an arriving customer results in one of the servers beginning a set up period that it needs to go through before beginning to serve. Whereas the setup distribution is assumed to be exponential, the service time distribution is allowed to be general. If a server completes service when there are no customers waiting then that server becomes idle. If a server completes service when there is at least one waiting customer then with probability $1-p$ the server becomes idle, and with probability p the server begins serving one of the waiting customers and one of the servers in setup is switched back to idle.

The paper [6] clarified and generalized the friendship paradox, which loosely says that "your friends tend to have more friends than you do". Consider any arbitrary friendship graph, in which each of n nodes represents a different person and where there is an undirected edge between two nodes if those two persons are friends. Let $f(i)$ denote the number of friends of person i (i.e., the degree of node i). Suppose that X is a randomly chosen person, equally likely to be any of the nodes. Now, suppose that each of the n people writes, on separate sheets of paper, the names of each of their friends, and let Y be the name on a randomly chosen sheet of paper. The friendship paradox is a paraphrase of the result that $E[f(Y)] \geq E[f(X)]$. In [4] we studied whether there was a stronger stochastic inequality between $f(Y)$ and $f(X)$, as well as the relation between $f(X)$ and $f(Z)$, where Z is a randomly chosen friend of X .

The paper [7] supposed that n customers are waiting to be served by any of m servers. Customer i will only wait in queue for an exponentially distributed time with rate λ_i before departing the system, and will yield a reward of r_i upon being served. Service times at server j are exponential with rate μ_j . The problem of interest is to decide which of the customers to serve at each time point so as to maximize the expected total reward. Here preemption of currently being served customers is allowed.

The paper [8] considered multinomial trials with possible outcomes $1, \dots, s$, with trial outcome probabilities given by the vector $\mathbf{p} = (p_1, \dots, p_s)$, $\sum_{i=1}^s p_i = 1$. For a specified vector of integers $\mathbf{r} = (r_1, \dots, r_s)$, let $N = N_{\mathbf{r}}(\mathbf{p})$ denote the number of trials needed until there have been at least r_i type i outcomes for all $i = 1, \dots, s$. The problem of interest was to determine the probability

vector $\mathbf{p}$ that minimizes $E[N_{\mathbf{r}}(\mathbf{p})]$. One application of this model would be to a parallel system of s components that is subjects to shocks. Supposing that each shock is type i with probability p_i and that component i fails when there are a total of r_i type i shocks, then $E[N_{\mathbf{r}}(\mathbf{p})]$ represents the expected number of shocks until a parallel system composed of these components fails.

The paper [9] considered a tournament among players $1, \dots, n$, in which each match involves two players. The tournament is assumed to be of knock-out type in that the losers of matches are eliminated and do not move on to the next round, and the tournament continues until all but one player is eliminated, with that player being declared the winner of the tournament. The match format is specified by the set of positive integers $r, m_1, \dots, m_r$ with the interpretation that there are a total of r rounds, with round i consisting of m_i matches, $\sum_{i=1}^r m_i = n - 1$. Because $\sum_{j=1}^{i-1} m_j$ players have been eliminated by the end of round $i - 1$, we must have that $m_i \leq (n - \sum_{j=1}^{i-1} m_j)/2$. We suppose that the constitution of the matches in a round is totally random. That is, for instance, the $2m_1$ players that play in round 1 are randomly chosen from all n players and then randomly arranged into m_1 match pairs. The winners of these m_1 matches, along with the $n - 2m_1$ players that did not play a match in round 1 then move to round 2, and so on. We suppose that the players have respective values $v_1, \dots, v_n$, and that a match involving players i and j is won by player i with probability $v_i/(v_i + v_j)$. The objective of the paper is to determine the win probabilities of the different players.

The paper [10] considered the Bernoulli bandit problem where one of the arms has win probability α and the others β , with the identity of the α arm specified by initial probabilities. That is, each time one uses arm 1 it leads to a win with probability p_1 and each time one plays arm 2 it leads to a win with probability x_1 , where the set of values $\{x_1, x_2\}$ is known to equal $\{\alpha, \beta\}$. With $u = \max(\alpha, \beta)$, $v = \min(\alpha, \beta)$, call an arm with win probability u a good arm. Whereas it is known that the strategy of always playing the arm with the largest probability of being a good arm maximizes the expected number of wins in the first n games for all n , the paper was interested in whether it is true that that strategy also maximizes the probability of at least k wins in the first n games played for every k and n .

The paper [11] supposes that customers arrive according to a Poisson process with each customer having a vector indicating which of the possible servers is eligible to server that customer. It is supposed that, independent of all else, a customer is eligible for server i with probability p_i . It is also supposed that server i has service rate μ_i . Any arrival not finding an idle server that is eligible to serve them is lost. The problem of interest is to determine to which eligible server an arrival should be given so as to minimize the rate of lost customers.

2 Important Results

In [1] we show that if all quotas are equal, then the policy that assigns priorities in increasing order of the probabilities $p(i)$ stochastically minimizes, among all priority policies, the time until all quotas have been filled. We also show that if all $p(i)$ are equal, the policy that assigns job priorities in decreasing order of their quotas, stochastically minimizes the time until all quotas are met.

The paper [1] also considers dynamic policies, We show that one should never assign a worker to job type i if they are also qualified for job type j when the remaining quota for job i is less than that for j and $p(j)$ is smaller than $p(i)$. Using this we introduce a heuristic policy and then show how to improve it by using the policy improvement algorithm of dynamic programming.

In [2] we define a list policy as an ordering of the customers, with the instruction that the next customer to be served should be the first one on the list that is still in the system (that is, that customer hasn't either departed or already been served). We give a sufficient condition that results in a specified list policy being optimal, and present some heuristic policies for when this condition does not hold.

In [5] we prove that the number of servers in setup is independent of the number in service, with the latter having a Poisson distribution. Moreover, given the numbers in setup and in service, the remaining service times (as well as the service times already provided) are independent and are distributed according to the equilibrium service distribution.

In [6] we proved that $f(Y)$ is likelihood ratio larger than $f(X)$. We also showed that $f(Z)$ is stochastically, but not likelihood ratio, larger than $f(X)$,

and that there is no general relation between $E[f(Y)]$ and $E[f(Z)]$.

In the impatient server model of [], the Ph.D graduate student Yang Cao, working under the supervision of Professor Ross, showed that there are times when it is optimal to preempt customers from service. Sufficient conditions implying the optimality of the list policy $1, 2, \dots, n$ were also determined.

In [9] we suppose, without loss of generality that the players are numbered so that $v_1 \geq v_2 \geq \dots \geq v_n$. Letting P_i be the probability that player i wins the tournament, $i = 1, \dots, n$, we prove that

$$P_1 \geq \frac{v_1}{\sum_{j=1}^n v_j}, \quad P_n \leq \frac{v_n}{\sum_{j=1}^n v_j}.$$

Let

$$p_i = \frac{1}{n-1} \sum_{j \neq i} \frac{v_i}{v_i + v_j}$$

be the probability that i would win a match against a randomly chosen opponent, we proved that P_i is smaller than it would be if it were the case that i would win each game it plays with probability p_i . That is, we proved that

$$P_i \leq \prod_{s=1}^r \left(\frac{2m_s}{r_s} p_i + 1 - \frac{2m_s}{r_s} \right)$$

where $r_s = n - \sum_{j=1}^{s-1} m_j$ is the number of players that advance to round s . We also prove that $P_1 \geq P_2 \geq \dots \geq P_n$. Among other results whereas we easily show that P_i is an increasing function of v_i , we also obtain via a counterexample, the surprising result that it is not necessarily true that P_i is a decreasing function of $v_j, j \neq i$. That is, there are cases where increased strengths of your opponents can increase your chances of winning the tournament. The paper [9] also considers the special case where $v_1 > v_2 = \dots = v_n$ and shows that when $n = 2^s + k, 0 \leq k < 2^s$, the best format for the strongest player is the so-called *balanced format* that has k matches in the first round and then has all remaining players competing in each subsequent round. We also show that whenever the number of remaining players, say t , is even there is an optimal (from the point of view of the best remaining player) format that calls for $t/2$ matches in the next round. We also show, in this special case, that the worst format for the best player is to have exactly one match each round. Analogous results for the worst player are also shown.

For the parallel system in which component i fails when there have been a total of $r(i)$ type i shocks, coupling arguments were used in [9] to show that the probability vector of shock types that minimizes the expected number of shocks until all components have failed would have $p(j) > p(i)$ whenever $r(j) > r(i)$. We also showed that if $r(i) = va(i)$, where $a(1), \dots, a(s)$ are positive numbers summing to 1, then as v goes to infinity the minimizing probability vector converges to the vector $a(1), \dots, a(s)$. We also studied in detail the case $s = 2$ and were able to approximate the optimal policy. Specifically, when $s = 2$ an expression for $E[N]$ as well as an asymptotic formula for it are derived. In addition, a further specialization of the $s = 2$ case in which $r_1 = 1$, it was shown that the optimal value of p_1 is very close to $\frac{\log(r) - \log \log(r)}{r}$.

Whereas, it is known for the problem of [10], that the strategy of always playing the arm with the largest probability of being a good arm maximizes the expected number of wins in the first n games for all n , the paper [10] conjectures that it also stochastically maximizes the number of wins. That is, it conjectures that this strategy maximizes the probability of at least k wins in the first n games for all k, n . The conjecture is proven when $k = 1$, when $k = n$, and when there are only 2 arms and $k = n - 1$.

A list policy in the queueing loss model of [11] is a permutation $i_1, \dots, i_n$ of $1, \dots, n$ with the interpretation that a customer is assigned to the first idle and eligible server according to the ordering $i_1, \dots, i_n$. In [xx] showed that there are situations where no list policy is optimal. We also showed that if $u(i)$ is decreasing and $p(i)$ is increasing in i , then the list policy $1, 2, \dots, n$ is an optimal policy. We also consider two heuristic policies. The first policy considers for each pair of servers, say i and j , the optimal policy when these are the only servers in the problem. If it is optimal to assign to server i when both servers are both idle and eligible then we say that i wins against j . We then consider the list policy $i_1, \dots, i_n$ of all n servers where i_1 is the server who had the most wins in the $\binom{n}{2}$ problems having only two servers, i_2 is the one with the second largest number of wins, and so on. We also consider a second heuristic, which uses a list policy that lists the servers in decreasing order of the ratio $u(i)/p(i)$. To see how well these heuristics perform, we first present some results that enable us put conditions on possible optimal policies, which enables us to reduce the set of policies that might be optimal to policies in the class C . We then considered the case of $n = 3$ servers.

In this case we are able to numerically determine the rate of lost customers for any specific policy in the class C , and thus determine the optimal policy and its rate of lost customers. We then compared this with the rates of lost customers for our two heuristics as well as for a random rule that always chooses at random among the idle and eligible servers. We show that the first heuristic is usually, but not always, optimal and always performed at least as well as the second, and that both heuristics performed at least as well as the random rule. We also numerically compared the rate of lost customers for the three rules when there are $n = 5$ servers. Although we were unable to explicitly determine the optimal policy in these cases, our numerical work indicated that whereas the first heuristic always performed at least as well as the second, the two heuristics were quite close in their performance, and both always outperformed the random rule.

3 Bibliography

1. D. Wu, and S. Ross, A Stochastic Assignment Problem, *Naval Research Logistics*, Vol. 62, number 1, 23-31, 2015
2. S. Ross, A Sequential Scheduling Problem with Impatient Jobs, *Naval Research Logistics*, Volume 62, Issue 8, pages 659-663, 2015
3. S. Ross, An Individual and Socially Optimal Policy Minimizing Expected Flow Times, *Prob. in Eng. and Inf. Sci.*, Vol. 29, No. 2, 147-151, 2015
4. S. Ross, Improved Chen-Stein Bounds on the Probability of a Union, *Jour. of Appl. Probability*, Volume 53, Issue 4, 1265-1270, 2016
5. B. Haji, S. Ross, $M/G/\infty$ with Exponentially Distributed Setup Times, *Operations Research Letters*, Vol. 43, number 1, 2015
6. Y. Cao, S. Ross, The Friendship Paradox, *The Mathematical Scientist*, Vol. 41 Issue 1, pp. 61-64, June 2016
7. Y. Cao, Multiple Server Preemptive Scheduling with Impatience, *Probability in Engineering and Information Sciences*, April, 2017

8. M. Brown, S. Ross, Optimality Results for Coupon Collection, *Jour. of Appl. Probability*, Volume 53, Issue 3, 930-937, 2016
9. I. Adler, Y. Cao, R. Karp, and E. Pekoz, S. Ross, Random Knockout Tournaments, accepted for publication *Operations Research*, 2017
10. M. Nouiehed, S. Ross, A Feldman Bandit Problem Conjecture, submitted to *Jour. of Appl. Probability*, 2017
11. B. Haji, S. Ross, A Queueing Loss Model with Heterogeneous Skill Based Servers and Discriminating Arrivals, submitted to *Naval Research Logistics*, 2017