Probability in the Engineering and Informational Sciences, **23**, 2009, 637–647. Printed in the U.S.A.
doi:10.1017/S0269964809990040

A MODEL FOR LOCKING IN GAINS WITH AN APPLICATION TO CLINICAL TRIALS

Sheldon M. Ross
Epstein Department of Industrial and Systems Engineering
University of Southern California
Los Angeles, CA
E-mail: smross@usc.edu

Henk Tijms
Department of Operations Research
Vrije University
De Boelelaan 1105
Amsterdam, HV1081, The Netherlands
E-mail: tijms@feweb.vu.nl

Shinyi Wu
Epstein Department of Industrial and Systems Engineering
University of Southern California
Los Angeles, CA
E-mail: shinyiwu@usc.edu

We consider a model in which each round consists of a sequence of games, with each game resulting in either a positive or a zero score. If a zero score occurs, then the current round is ended with no points being accumulated during that round. If a game ends with a positive score, then the player can either end that round or play another game in the round. If she elects to end the round, then the sum of all scores earned in games played during that round are added to her cumulative score and a new round begins.

Under the assumption that successive game scores are independent and identically distributed random variables whose conditional distribution, given that it is positive, is exponential, we consider this problem under such objectives as minimizing the expected number of rounds until a cumulative score exceeds a given goal g and maximizing the probability that a cumulative score of at least g is obtained by the end of round n. We present the model in the hypothetical context of a clinical trial of a treatment for reducing glycated hemoglobin in diabetic patients.

 0269-9648/09 $25.00

1. A MODEL FOR LOCKING IN GAINS

Consider a researcher attempting to document the efficiency of a new procedure for reducing glycated hemoglobin in patients with diabetes, a disease affecting more than 20 million Americans at an annual cost of over 116 billion dollars [1]. Patients are dealt with sequentially. Each time a patient is given the procedure, the result will be either that the patient abandons the study (which occurs with probability α) or has a glycated hemoglobin reduction having distribution $F(x) = 1 - e^{-\lambda x}$ (which occurs with probability $1 - \alpha$). After a patient's treatment reduction is determined, a decision must be made on whether to give that patient another treatment or to go on to the next patient. If the decision is made to move on to a new patient, then the total hemoglobin reduction of the old patient is recorded. However, if a patient abandons the study, then all data concerning that patient is lost.

In Section 2 we consider the problem of determining when to move on to new patients so as to minimize the expected number of patients required for the cumulative reduction of all (nonabandoning) patients to exceed a specified goal g. In Section 3 we consider the same problem, but this time with the objective of minimizing a weighted average of the number of patients and the number of times the procedure is used (i.e., this section supposes not only a cost per procedure but also a recruiting cost for each patient). In Section 4 we suppose that we deal with multiple patients simultaneously in a group-visit setting.

The group visit is a treatment modality increasingly used for service delivery to diabetic patients to improve care efficiency and clinical outcomes (see, for instance, [2–4,7–9]). In Section 5 we again suppose that we deal with patients one at a time and consider the problem of determining when to voluntarily move on to the next patient when our objective is to maximize the probability of obtaining a cumulative reduction of at least g when there is a fixed number of patients.

2. MINIMIZING THE EXPECTED NUMBER OF PATIENTS TO REACH A GOAL

Let us use the terminology that a new round begins when a new patient is given his or her first treatment. The result of each treatment can be regarded either as a failure (patient abandoning the study) or as yielding a random score (equal to the glycated hemoglobin reduction resulting from the treatment). After each nonfailed treatment, the clinician can decide either to give another treatment to that patient or to end that round. When the round is voluntarily ended, the total points (e.g., total of the hemoglobin reduction scores) accumulated during that round is added to the clinic's cumulative reduction total and a new round begins. When a failure occurs, all scores so far accumulated during that round are lost and a new round begins. (Thus, voluntarily ending a round locks in all the scores accumulated during that round.) Assuming that each treatment results in a failure with probability α and that the distribution of scores is $F(x) = 1 - e^{-\lambda x}$, we are interested in determining a good strategy for minimizing the expected number of rounds that it takes to amass a cumulative reduction of at least g.

The dynamic programming state of the system when a decision is to be made is the pair (x, y) with the interpretation that x was the additional amount needed before the current round had begun (so the current round began with a total reduction of $g - x$) and y is the total of the scores so far amassed in the round. The decision is then to either perform another procedure or to end the round. Letting $V_c(x, y)$ and $V_e(x, y)$ be the minimal expected additional numbers of rounds needed when the state is (x, y) and the decision is made to continue (for $V_c(x, y)$) or to end (for $V_e(x, y)$) the round, then $V(x, y)$, the minimal expected additional number of rounds needed from state (x, y), satisfies

$$V(x, y) = \begin{cases} 0 & \text{if} \quad y \geq x \\ \min\{V_c(x, y), V_e(x, y)\} & \text{if} \quad y < x, \end{cases}$$

where

$$V_c(x, y) = \alpha \left(\frac{1}{1 - \alpha} + \int_0^x \lambda e^{-\lambda w} V(x, w)\, dw \right) + (1 - \alpha) \int_0^{x-y} \lambda e^{-\lambda w} V(x, w + y)\, dw$$

$$V_e(x, y) = \frac{1}{1 - \alpha} + \int_0^{x-y} \lambda e^{-\lambda w} V(x - y, w)\, dw.$$

Although it is easily proven that

$$x' \geq x, x' - y' \geq x - y \Rightarrow V(x', y') \geq V(x, y),$$

it is difficult to explicitly solve the optimality equation. As a result, we propose analyzing a reasonable heuristic policy.

Suppose that the decision has been made to only end a round when the total of the scores amassed during that round is at least s. If we think of the result of each procedure as being a game, then, with probability α, each game results in a failure, which ends the round with all previously amassed scores during that round being lost. The reduction score in any successful game (i.e., one that is not a failure) is exponential with rate λ. Let us imagine that the score in a successful game is accumulated at a rate of one per unit time while the game goes on and that the total time of the game is exponential with rate λ. Consequently, if the total amassed score so far during a round is $y, y < s$, then the failure rate at that time is $\lambda\alpha$. However, this constant failure rate implies that once points begin to be accumulated, the amount amassed before a failure occurs is exponential with rate $\lambda\alpha$. Consequently, if we let $T(s)$ denote the total of the usable scores amassed during a round, then

$$T(s) = \begin{cases} 0 & \text{with probability} \quad \alpha + (1 - \alpha)(1 - e^{-\alpha\lambda s}) \\ s + Y & \text{with probability} \quad (1 - \alpha)e^{-\alpha\lambda s}, \end{cases} \tag{1}$$

where Y, which represents the amount by which the total score of a successful round exceeds s, is exponential with rate λ.

https://doi.org/10.1017/S0269964809990040 Published online by Cambridge University Press

Note that

$$\max_s E[T(s)] = \max_s (1-\alpha)\left(s+\frac{1}{\lambda}\right)e^{-\alpha\lambda s} = \frac{1-\alpha}{\alpha\lambda}e^{-(1-\alpha)}$$

with the maximizing value of s being $s^* = (1-\alpha)/\alpha\lambda$.

Remark: For any distribution F, the problem of maximizing the expected total reduction in a round is known as the *burglar's problem* (see [5]). The optimal policy is a one-stage look-ahead stopping rule policy that ends a round whenever the accumulated amount s is such that $s \geq (1-\alpha)(s+\mu_F)$, where μ_F is the mean of the distribution F. Thus, the optimal policy is to end a round whenever $s \geq [(1-\alpha)/\alpha]\mu_F$.

For the criterion of minimizing the expected number of rounds needed to reach a total locked-in score of g, consider the following heuristic policy: If a round is to begin when one has a cumulative score of $g-x$, then the round should be ended when the amount accumulated in the round is at least $\min(s,x)$. To approximate the mean number of rounds needed under the heuristic policy until the player's locked-in score reaches g, note that the number of rounds until the player's locked-in score exceeds $g-s$ has the same distribution as $1+N(g-s)$, where $N(t)$ is the number of renewals by time t of a renewal process whose interarrivals have the distribution given by (1). Now, for a renewal process $\{N(t)\}$ with interarrival times distributed as X,

$$E[X](E[N(t)]+1) = t + E[Y(t)],$$

where $Y(t)$ is the time from t until the next renewal. Approximating $E[Y(t)]$ by its limiting value (see [6]) yields that

$$E[X](E[N(t)]+1) \approx t + \frac{E[X^2]}{2E[X]}$$

or

$$E[N(t)]+1 \approx \frac{t}{E[X]} + \frac{E[X^2]}{2E^2[X]}.$$

Thus, if we let $R_s(g-s)$ be the number of rounds until the player's locked-in score exceeds $g-s$ when the heuristic policy is used, then

$$E[R_s(g-s)] \approx \frac{g-s}{E[T(s)]} + \frac{E[T^2(s)]}{2E^2[T(s)]}.$$

Because

$$E[T(s)] = (1-\alpha)e^{-\alpha\lambda s}(s+1/\lambda)$$

and

$$\begin{aligned} E[T^2(s)] &= (1-\alpha)e^{-\alpha\lambda s}E[(s+Y)^2] \\ &= (1-\alpha)e^{-\alpha\lambda s}(\mathrm{Var}(s+Y) + (E[s+Y])^2) \\ &= (1-\alpha)e^{-\alpha\lambda s}(1/\lambda^2 + (s+1/\lambda)^2), \end{aligned}$$

this yields that

$$E[R_s(g-s)] \approx \frac{(g-s)}{(1-\alpha)(s+1/\lambda)} e^{\alpha\lambda s} + \frac{1/\lambda^2 + (s+1/\lambda)^2}{2(1-\alpha)(s+1/\lambda)^2} e^{\alpha\lambda s}. \tag{2}$$

The player's cumulative locked-in score when it first exceeds $g-s$ will exceed it by an amount that approximately has the equilibrium distribution of $T(s)$. Calling this distribution G_e and using that

$$P(T(s) > y) = (1-\alpha)e^{-\alpha\lambda s} \quad \text{if } y < s,$$

it follows that

$$G_e(x) = \frac{1}{E[T(s)]} \int_0^x P(T(s) > y)\, dy = \frac{\lambda x}{1+\lambda s} \quad \text{if } x < s.$$

Because the number of additional rounds needed under the heuristic policy when the current cumulative score is $g-s+x$, $x < s$, is a geometric random variable with parameter $(1-\alpha)e^{-\alpha\lambda(s-x)}$, it follows from the preceding that the expected additional number of rounds needed when the cumulative locked in score first exceeds $g-s$ is approximately

$$\begin{aligned} \frac{1}{1-\alpha} \int_0^s e^{\alpha\lambda(s-x)}\, dG_e(x) &= \frac{1}{1-\alpha} \frac{\lambda}{1+\lambda s} \int_0^s e^{\alpha\lambda(s-x)}\, dx \\ &= \frac{1}{1-\alpha} \frac{1}{\alpha(1+\lambda s)} (e^{\alpha\lambda s} - 1). \end{aligned} \tag{3}$$

Putting it all together yields that $E[R_s(g)]$, the expected number of rounds that it takes, under the heuristic policy, to obtain a cumulative locked-in score of at least g, is such that

$$\begin{aligned} E[R_s(g)] &\approx E[R_s(g-s)] + \frac{1}{1-\alpha} \frac{1}{\alpha(1+\lambda s)} (e^{\alpha\lambda s} - 1) \\ &\approx \frac{(g-s)}{(1-\alpha)(s+1/\lambda)} e^{\alpha\lambda s} + \frac{1/\lambda^2 + (s+1/\lambda)^2}{2(1-\alpha)(s+1/\lambda)^2} e^{\alpha\lambda s} \\ &\quad + \frac{1}{1-\alpha} \frac{1}{\alpha(1+\lambda s)} (e^{\alpha\lambda s} - 1). \end{aligned} \tag{4}$$

Letting $s = s^* = (1-\alpha)/\alpha\lambda$, which is the heuristic we recommend, yields that

$$\begin{aligned} E[R_{s^*}(g)] &\approx \left(g - \frac{1-\alpha}{\alpha\lambda} \right) \frac{\alpha\lambda}{1-\alpha} e^{1-\alpha} + \frac{1+\alpha^2}{2(1-\alpha)} e^{1-\alpha} + \frac{1}{1-\alpha}(e^{1-\alpha} - 1) \\ &= e^{1-\alpha} \left\{ \frac{g\alpha\lambda}{1-\alpha} + \frac{(1+\alpha)^2}{2(1-\alpha)} \right\} - \frac{1}{1-\alpha}. \end{aligned} \tag{5}$$

Example 1: If $\lambda = 1, \alpha = 0.1$, and $g = 100$, then $s^* = 9, E[T(s^*)] = 9e^{-.9} = 3.659$, and the preceding yields the approximation $E[R_{s^*}(g)] \approx 27.871$, whereas a simulation of 10^5 runs yielded the result $E[R_{s^*}(g)] = 27.925$.

One might wonder at this point if the optimal policy is a control limit policy like our heuristic; that is, whenever a new round begins, does it suffice to specify a value v, depending on the current locked in score, with the instruction that the round should be continued until either the patient abandons or the total score in the round exceeds v? It turns out that such a policy need not be optimal. For a counterexample, consider the parameters of our preceding example and suppose that a round is to begin with a locked-in total of $g - 20$. The probability of success if one tries for the entire additional amount of 20 in a single round is $.9e^{-2}$, indicating that the mean number of additional rounds if that policy is employed is $(1/.9)e^2 = 8.210$. On the other hand, if one continues each round until the total 9 is exceeded, then a simulation indicates that the expected number of rounds needed is 7.092. Thus, if a control limit policy is optimal, then when the current locked in total is $g - 20$, it sets a critical value less than 20. However, whatever value is set, suppose that the reduction from the most recent procedure was such that the total not yet locked in amount during the round is $20 - \epsilon$. If one elects to continue the round, then with probability .9, no additional rounds are needed, whereas with probability .1 a new round would begin with an additional amount 20 still needed. Because in the latter case there is a policy that will require an expected number of 7.092 additional rounds, it follows that the optimal expected number of additional rounds needed if we continue the current round is less than .71, which dominates ending the round. Thus, the optimal policy is not a control limit policy.

3. MINIMIZING ROUNDS AND PROCEDURES

Whereas in the preceding section we only concerned ourselves with finding a good policy for minimizing the expected number of rounds needed to reach a locked-in reduction of at least g, in this section we will suppose that we are also interested in keeping the number of procedures low. To do so, imagine that a cost $C > 0$ is incurred whenever a new round begins and that an additional cost $c > 0$ is incurred each time the procedure is used.

Suppose, as earlier, that the decision has been made to only end a round when the total of the scores amassed during that round is at least the minimum of s and the additional amount needed to obtain a total locked-in reduction of at least g. Let $N(s)$ denote the number of procedures used during a round in which a voluntary stop only occurs when the total reduction is at least s. Let Y_i denote the reduction in the *ith* nonfailed game and let

$$M = \min\left(n : \sum_{i=1}^{n} Y_i > s\right).$$

https://doi.org/10.1017/S0269964809990040 Published online by Cambridge University Press

Because the Y_i are independent exponential random variables with rate λ, it follows that M—which represents the number of procedures needed to obtain a reduction of at least s—is distributed as 1 plus the number of events by time s of a Poisson process with rate λ. Now,

$$P(N(s) > k|M = n) = \begin{cases} (1-\alpha)^k & \text{if} \quad k < n \\ 0 & \text{if} \quad k \geq n. \end{cases}$$

Hence,

$$E[N(s)|M = n] = \sum_{k=0}^{n-1} (1-\alpha)^k = \frac{1-(1-\alpha)^n}{\alpha},$$

yielding that

$$\begin{aligned} E[N(s)] &= \frac{1}{\alpha}[1 - E[(1-\alpha)^M]] \\ &= \frac{1}{\alpha}\left[1 - (1-\alpha)\sum_{i=0}^{\infty}\frac{(1-\alpha)^i e^{-\lambda s}(\lambda s)^i}{i!}\right] \\ &= \frac{1}{\alpha}[1-(1-\alpha)e^{-\alpha\lambda s}] \end{aligned}$$

Now, if we let $N_i(s)$ denote the number of procedures used in round i, then with $R_s(g-s)$ equal to the number of rounds needed until the player's locked-in score exceeds $g-s$, it follows from Wald's equation that the expected number of procedures used in that time is

$$E\left[\sum_{i=1}^{R_s(g-s)} N_i(s)\right] = E[R_s(g-s)]E[N(s)] = \frac{1}{\alpha}[1-(1-\alpha)e^{-\alpha\lambda s}]E[R_s(g-s)],$$

where $E[R_s(g-s)]$ is approximated by (2). Because the expected additional number of procedures needed from that point on is as approximated by (3), we see that $E[P_s(g)]$, the expected number of procedures needed, is such that

$$E[P_s(g)] \approx E[N(s)]E[R_s(g-s)] + \frac{1}{1-\alpha}\frac{1}{\alpha(1+\lambda s)}(e^{\alpha\lambda s} - 1).$$

Letting $E[\mathrm{TC}(s)]$ be the expected total cost incurred under the s-policy then

$$E[\mathrm{TC}(s)] = CE[R_s(g)] + cE[P_s(g)].$$

Because the dominant term in $E[R_s(g)]$ is $(g-s)/E[T(s)]$ whereas the dominant term in $E[P_s(g)]$ is $E[N(s)](g-s/E[T(s)])$, it follows that the dominant term of $E[\mathrm{TC}(s)]$ is

$$\frac{g-s}{E[T(s)]}\{C+cE[N(s)]\}\approx\frac{C+cE[N(s)]}{E[T(s)]}g$$
$$\approx\frac{C+(c/\alpha)[1-(1-\alpha)e^{-\alpha\lambda s}]}{(1-\alpha)(s+1/\lambda)e^{-\alpha\lambda s}}g. \qquad \textbf{(6)}$$

Letting $s_{c,C}$ be that value of s that minimizes the right-hand side of (6), we propose the heuristic policy, which, when a round is to begin when the current locked-in reduction is $g-x$, continues a round until the total reduction is at least $\min\{x, s_{c,C}\}$. It should be noted that $s_{c,C}$ is the value of s that minimizes the ratio of the expected cost of a round whose goal is to reach a reduction of at least s divided by the expected reduction of such a round; that is, $s_{c,C}$ minimizes

$$\frac{C+cE[N(s)]}{E[T(s)]}=\frac{C+(c/\alpha)[1-(1-\alpha)e^{-\alpha\lambda s}]}{(1-\alpha)(s+1/\lambda)e^{-\alpha\lambda s}}.$$

Example 2: Letting, as in Example 1, $\lambda=1, \alpha=0.1, g=100$ and taking $C=1$, we have the following table for the approximate values of $s_{c,1}, E[P_{s_{c,1}}(100)], E[R_{s_{c,1}}(100)], E[\mathrm{TC}(s_{c,1})]$, for a variety of values of c.

c	$s_{c,1}$	$E[P_{s_{c,1}}]$	$E[R_{s_{c,1}}]$	$E[\mathrm{TC}(s_{c,1})]$
1	2.843	123.97	39.07	163.04
0.5	3.945	130.43	33.97	99.19
0.33	4.652	134.91	31.92	76.89
0.25	5.165	138.28	30.82	65.39
0.10	6.697	149.26	28.80	43.73
0	9	168.07	27.87	27.87

4. MULTIPLE PATIENTS PER ROUND

Suppose now that every round begins with a group of m patients who are each initially given the procedure. After each group of patients are given the procedure, the scores of all the nonfailed patients are observed and we must then decide which of them should be given another treatment and which of them should be discontinued so as to lock in the total of their so far amassed scores for the round. Say that a subround is completed each time a group of patients is given the procedure. Although the problem of minimizing the expected number of rounds—or minimizing the expected total cost when there is a cost for each round as well as a cost each time the procedure is used—can be set up as a dynamic programming problem, the state of the system at the beginning of a subround would be a multidimensional vector $(x, r, y_1, \ldots, y_r)$, where

https://doi.org/10.1017/S0269964809990040 Published online by Cambridge University Press

x is the remaining locked-in reduction needed, r is the number of surviving patients in the round, and $y_1, y_2, \ldots, y_r$ are the total reductions so far of these r patients. Because a simple coupling argument can be used to show that it would never be optimal to continue with a patient whose total reduction so far is s while locking in a patient whose total reduction so far is t whenever $s > t$, it follows that the optimal decision in state $(x, r, y_1, \ldots, y_r)$ can be represented as one of the values $i, i = 0, \ldots, r$, with the interpretation that if decision i is made, then the procedure is given to the i of the r patients whose total reductions so far are smallest. However, a great deal of computational effort would be needed to solve this dynamic programming problem. Thus, we propose slight variants of the heuristic strategies of Section 2 (if one wants to minimize the expected number of rounds needed to obtain a total reduction of g) or of Section 3 (if there are costs both for rounds and procedures). In the former case, after each subround we propose that all scores should be locked in if doing so yields a cumulative locked- in score of at least g; if this is not the case, then an individual patient's total score should be locked in if it exceeds s^*. The patients whose scores are not locked in are once again given the procedure. In the latter case, where there are costs C for each round and c for each time the procedure is used, we propose using the same strategy as in the preceding case except that $s_{c,C}$ is substituted for s^*.

5. MAXIMIZING THE PROBABILITY OF REACHING THE GOAL WITH n PATIENTS

Suppose now that our objective is to maximize the probability of having a total reduction of g within n rounds when a single patient is treated in each round. One thing to note is that if additional reductions are needed to reach our goal when only two rounds remain, then neither of those rounds should be voluntarily ended before the goal is reached. Clearly, this is true for the final round. To see that it is also true for the next to last round, note that if we would stop that round when an additional reduction amount x is needed, then we will only be successful if the next round's reduction reaches x before an abandonment and this is the same as the probability that the additional reduction in the next to last round reaches x.

Now, suppose that there are presently k rounds to go ($n - k$ have already been used) and that so far we have a locked-in total reduction of $g - f$ (so an additional amount f is needed). Imagine that we will continue each of the next $k - 2$ rounds until the reduction in that round exceeds s, and if an additional goal reduction is still needed when only two rounds remain, then we continue each round until reaching the goal. Let X be the total reduction obtained in the initial $k - 2$ rounds. Additionally, let N be the number of those $k - 2$ rounds that amassed a positive reduction. Because each round will end with a reduction of at least s with probability

$$p(s) \equiv (1 - \alpha)e^{-\lambda\alpha s},$$

it follows that N is a binomial random variable with parameters $k - 2$ and $p(s)$. Given N, the total reduction in the $k - 2$ rounds is distributed as Ns plus the sum of N

https://doi.org/10.1017/S0269964809990040 Published online by Cambridge University Press

independent exponential random variables with rate λ (the exponentials represent the amounts by which the reduction in a successful round exceeds s.) Hence,

$$E[X|N] = sN + N/\lambda = (s + 1/\lambda)N$$

and

$$\text{Var}(X|N) = N/\lambda^2,$$

implying that

$$E[X] = (s + 1/\lambda)(k-2)p(s) \equiv \mu_{k-2}(s)$$

and, by the conditional variance formula (see [6]),

$$\text{Var}(X) = (s + 1/\lambda)^2(k-2)p(s)(1-p(s)) + (k-2)p(s)/\lambda^2 \equiv \sigma^2_{k-2}(s).$$

Now, we can express X as

$$X = \sum_{i=1}^{k-2} I_i(s + Y_i),$$

where I_i is an indicator variable for the event that round i is successful (i.e., does not end with a failure), Y_i is exponential with rate λ, and all of the random variables in the sum are independent. Consequently, for k not small, by the central limit theorem it follows that X approximately has a Normal distribution. Thus, if, when k rounds remain and an additional reduction of f is needed, we were to utilize the policy that continues each of the first $k-2$ rounds until the total score s is reached and then "goes for broke" during the final two stages, then the probability this strategy results in a success is

$$P(\text{success}) = E[P(\text{success}|X)] = E[h(f, X)] \approx E[h(f, W_{s,k-2})],$$

where $W_{s,k}$ is a normal random variable with mean $\mu_k(s)$ and variance $\sigma^2_k(s)$ and where

$$h(f,x) = \begin{cases} 1 & \text{if} \quad x \geq f \\ 1 - (1 - p(f-x))^2 & \text{if} \quad x < f. \end{cases}$$

Letting $s(k,f)$ be the value of s that maximizes $E[h(f, W_{s,k})]$, we propose the following heuristic policy: When k rounds remain and the still needed locked-in reduction amount is f, then the next round should be voluntarily ended when the total score amassed in the round is at least $s(k-2, f)$.

Remarks:

1. The values $s(k,f)$ only need be computed when needed. A simulation will be needed to determine them.

2. For Z a standard normal,

$$P(W_{s,k} > f) = P\left(Z > \frac{f - \mu_k(s)}{\sigma_k(s)}\right),$$

leading us to believe that $s(k,f)$ will be close to the value of s that minimizes $(f - \mu_k(s))/\sigma_k(s)$.

References

1. Centers for Disease Control and Prevention. (2007). National Diabetes Fact Sheet, 2007.
2. Clancy, D.E., Dismuke, C.E., Magruder, K.M., Simpson, K.N., & Bradford, D. (2008). Do diabetes group visits lead to lower medical care charges? *American Journal of Managed Care* 14(1): 39–44.
3. Clancy, D.E., Huang, P., Okonofua, E., Yeager, D., & Magruder, K.M. (2007). Group visits: Promoting adherence to diabetes guidelines. *Journal of General Internal Medicine* 22(5): 620–624.
4. Clancy, D.E., Yeager, D.E., Huang, P., & Magruder, K.M. (2007). Further evaluating the acceptability of group visits in an uninsured or inadequately insured patient population with uncontrolled type 2 diabetes. *Diabetes Education* 33(2): 309–314.
5. Ross, S.M. (1983). *Introduction to stochastic dynamic programming*. New York: Academic Press.
6. Ross, S.M. (1996). *Stochastic processes*. New York: Wiley.
7. Trento, M., Passera, P., Tomalino, M., Bajardi, M., Pomero, F., Allione, A., Vaccari, P., Molinatti, G.M., & Porta, M. (2001). Group visits improve metabolic control in type 2 diabetes: a 2-year follow-up. *Diabetes Care* 24(6): 995–1000.
8. Trento, M., Passera, P., Bajadi, M., Tomalino, M., Grassi, G., Borgo, E., Donnola, C., Cavallo, F., Bondonio, P., & Porta, M. (2002). Lifestyle intervention by group care prevents deterioration of Type II diabetes: A 4-year randomized controlled clinical trial. *Diabetologia* 45(9): 1231–1239.
9. Watkinson, M. (2004). Group visits improved concordance with American Diabetes Association practice guidelines in type 2 diabetes. *Evidence-Based Nursing* 7(2): 57.

https://doi.org/10.1017/S0269964809990040 Published online by Cambridge University Press