

Estimating Probability Distributions by Observing Betting Practices

Dr. C. Lynch

National University of Ireland, Galway
caroline.lynych@nuigalway.ie

Prof. D. Barry

University of Limerick
vpa@ul.ie

Abstract

A bookmaker takes bets on a two-horse race, attempting to minimise expected loss over all possible outcomes of the race. Profits are controlled by manipulation of customers' betting behaviour; in order to do this, we need some information about the probability distribution which describes how the customers will bet. We examine what information initial customers' betting behaviour provides about this probability distribution, and consider how to use this to estimate the probability distribution for remaining customers.

Keywords. EM Algorithm, bookmaker, horse race, Markov decision process.

1 Introduction

A bookie takes bets on a contest for which there are only two possible outcomes, which we will label as A and B . The bookie wishes to maximise his minimum expected profit purely by manipulation of customers' betting practice. A gambler enters the bookie's shop seeking to place a wager on this contest. Let p denote the gambler's probability that outcome A will occur. The bookie quotes odds of O_1 against outcome A and of O_2 against outcome B . This means that a winning wager of one unit on outcome A produces a return of $O_1 + 1$ while a winning wager of one unit on outcome B produces a return of $O_2 + 1$. Hence a wager on outcome A will be attractive to the gambler if

$$p(O_1 + 1) \geq 1$$

or equivalently

$$p \geq \frac{1}{O_1 + 1} = \theta_1.$$

Similarly a wager on outcome B will be attractive to the gambler if

$$1 - p \geq \frac{1}{O_2 + 1} = \theta_2$$

or equivalently

$$p \leq 1 - \theta_2.$$

The quantities θ_1 and θ_2 are called the bookie's quoted probabilities for outcome A and B respectively. Hence the strategy for an individual gambler is simple - he places a wager on any outcome for which his probability exceeds that quoted by the bookie.

It should be noted that the "quoted probabilities", θ_1 and θ_2 , described above, are not probabilities, in the sense that their sum will generally be greater than 1. In fact, they may more properly be described as upper probabilities, as defined in [13]. It can be shown that it is never to the advantage of the bookie to have these upper probabilities sum to less than 1, as in Lemma 2 of [14]. Thus, for the remainder of this paper, we shall assume that $\theta_1 + \theta_2 \geq 1$.

We also assume that the quoted odds, O_1 and O_2 , are positive. This follows naturally from the requirement that a wager of 1 unit leads to a return of $O_1 + 1$ on outcome A or $O_2 + 1$ on outcome B , and the customer is unlikely to wager more than the expected return. By the definitions of θ_1 and θ_2 , this means that θ_1 and θ_2 , in turn, are positive. These conditions on θ_1 and θ_2 ensure the coherence of the upper probabilities in this case.

We idealise the bookie's shop by assuming that the bookie sells two types of tickets - one which guarantees a return of one unit should outcome A occur and costs θ_1 , and one which guarantees a return of one unit should outcome B occur and costs θ_2 . This avoids sure loss, as the customer's only options are to bet on A or B , individually; to bet on both would ensure a loss for the customer, as he would be required to bet an amount $\theta_1 + \theta_2$, greater than his (guaranteed) return of 1. We also assume that the bookie knows, before opening the book, that N customers will consider a wager on the contest and that their probabilities p_1, p_2, \dots, p_N of outcome A occurring behave like a random sample from a probability

distribution. Finally, we assume that customers can buy at most one of each type of ticket and that the bookie is free to alter the quoted probabilities after each customer leaves.

The bookie seeks to manipulate customers' betting behaviour as best suits himself; this depends, however, on knowing the probability distribution from which the customers' probabilities p_1, p_2, \dots, p_N derive. After first considering the optimal procedure when the distribution is known, we will consider how the bookie may estimate this probability distribution using information derived from customers' betting practices.

2 Distribution Known

Assuming the distribution of customers' probabilities to be known, the optimal algorithm for the bookie to follow is the "Dynamic Programming" Algorithm, as described in Barry & Hartigan[2]. This iterative algorithm depends on knowledge of the customers' probability distribution, F , the number of customers left to come, n , and the current "state of the book", i.e. the amount of the bookie's profit on outcome A , denoted a , and on outcome B , denoted b , if the book was closed at that instant, i.e. no more bets were taken.

Assuming knowledge of these quantities, the algorithm is then given as follows;

$$R_n(a, b) = \frac{a + b}{2} + P_n(d)$$

where

$$P_n(d) = P_{n-1}(d) + \max_{\theta_1, \theta_2} \left\{ [1-F(\theta_1)] \left[\theta_1 - \frac{1}{2} + P_{n-1}(d-1) - P_{n-1}(d) \right] \right. \\ \left. + F(1-\theta_2) \left[\theta_2 - \frac{1}{2} + P_{n-1}(d+1) - P_{n-1}(d) \right] \right\}$$

with $P_0(d) = -\frac{|d|}{2}$ and $d = a - b$. Here, $R_n(a, b)$ denotes the expected value of the bookie's final minimum profit between both outcomes. The algorithm gives the bookie a method for deriving optimal quoted probabilities for the next customer, given the current state of the book, n customers left to go and F known.

The above equation for $P_n(d)$ describes how it depends on the previous value, $P_{n-1}(d)$, then adds a term, maximised over θ_1 and θ_2 , which describes the profit accruing if the customer bets on A , with probability $1 - F(\theta_1)$, and if the bet is on B - with probability $F(1 - \theta_2)$; the only two possible bets.

3 Strategy for Distribution Unknown

Having determined a strategy for F known, we must consider how to estimate F when it is unknown. We subdivide the interval $[0,1]$ into r subintervals of equal width - the choice of the value of r will be discussed in Section 4. We then estimate F by means of a histogram, with r intervals. For each of these r intervals, the height of the histogram will be determined by the probability assigned to that interval, π_j . This probability will be determined by the betting behaviour of the customers, as described hereafter. $F(\theta)$ may then be determined by the formula

$$F(\theta) = \begin{cases} r\theta\pi_1 & 0 \leq \theta \leq \frac{1}{r} \\ \pi_1 + (r\theta - 1)\pi_2 & \frac{1}{r} \leq \theta \leq \frac{2}{r} \\ \pi_1 + \pi_2 + (r\theta - 2)\pi_3 & \frac{2}{r} \leq \theta \leq \frac{3}{r} \\ \vdots & \vdots \\ \sum_{i=1}^j \pi_i + (r\theta - j)\pi_{j+1} & \frac{j}{r} \leq \theta \leq \frac{j+1}{r} \\ \vdots & \vdots \\ \sum_{i=1}^{r-1} \pi_i + (r\theta - r + 1)\pi_r & \frac{r-1}{r} \leq \theta \leq 1 \end{cases}$$

3.1 Estimation of F

This involves the EM Algorithm; we have an estimation, and a maximisation, step.

3.1.1 Estimation

Each customer's betting pattern gives us information about their value of p , as follows;

$$\begin{array}{ll} \text{Bet on Horse A} & \theta_1 \leq p \leq 1 \\ \text{Bet on Horse B} & 0 \leq p \leq 1 - \theta_2 \\ \text{No Bet} & 1 - \theta_2 < p < \theta_1 \end{array}$$

We denote the lower limit of the range in which p falls by a_1^k for customer k and the upper limit by a_2^k such that $a_1^k \leq a_2^k$. We also denote the lower and upper limit of each of the subintervals of $[0,1]$, I_j , by $[L_j, R_j]$, with $R_j = L_{j+1}$. We have an indicator function, X_{jk} , defined as follows;

$$X_{jk} = \begin{cases} 1 & p \in [L_j, R_j] \\ 0 & \text{otherwise} \end{cases}$$

In this case, the log likelihood function is given by

$$\ell = \sum_{k=1}^N \sum_{j=1}^r X_{jk} \log \pi_j.$$

Given the customer's behaviour, we have a range for the customer's probability - i.e. $a_1^k \leq p \leq a_2^k$. Let us call this information Y_k .

We will seek to maximize the Expected value of the log likelihood, given this information, i.e.

$$E(\ell|Y) = \sum_{k=1}^N \sum_{j=1}^r E(X_{jk}|Y_k) \log \pi_j.$$

$$\begin{aligned} E(X_{jk}|Y_k) &= P(X_{jk} = 1|Y_k) \\ &= P(p \in [L_j, R_j] | p \in [a_1^k, a_2^k]) \\ &= \frac{P(p \in [L_j, R_j] \cap [a_1^k, a_2^k])}{P(p \in [a_1^k, a_2^k])} \\ &= \frac{P(p \in [L_j, R_j] \cap [a_1^k, a_2^k])}{\sum_{i=1}^r P(p \in [L_i, R_i] \cap [a_1^k, a_2^k])} \end{aligned}$$

We have

$$P(p \in [L_j, R_j] \cap [a_1^k, a_2^k]) = \pi_j \times \frac{l_{jk}}{R_j - L_j}$$

where l_{jk} is the length of $[L_j, R_j] \cap [a_1^k, a_2^k]$ and is given by

$$l_{jk} = \begin{cases} 0 & R_j \leq a_1^k \\ R_j - a_1^k & L_j \leq a_1^k \leq R_j \leq a_2^k \\ a_2^k - a_1^k & L_j \leq a_1^k \leq a_2^k \leq R_j \\ R_j - L_j & a_1^k \leq L_j \leq R_j \leq a_2^k \\ a_2^k - L_j & a_1^k \leq L_j \leq a_2^k \leq R_j \\ 0 & L_j \geq a_2^k \end{cases}$$

3.1.2 Maximisation

Next, we seek to maximise the expected value of the log likelihood function. The Maximum Likelihood Estimate for π_j is given by

$$\hat{\pi}_j = \frac{\sum_{k=1}^N E(X_{jk}|Y_k)}{N}.$$

Each of the subintervals of $[0,1]$ was assigned an initial probability, π_j^1 . For simplicity, this initial probability was the same for each subinterval, assuming the Uniform distribution, so that, with r subintervals, the initial values of π_j^1 are given by

$$\sum_{j=1}^r \pi_j^1 = 1 \Rightarrow \pi_j^1 = \frac{1}{r}, \forall j.$$

This initial probability was then updated by observing each customer's behaviour.

As will be seen in this, and the next, subsection, we will now divide our customers into three groups; the very first will be used to initialise the information matrix, the second group will be used for the purpose of maximising the information we may obtain,

leaving us with the third and final group for maximising profit, once F has been satisfactorily estimated. One of the questions with which we will be concerned is how many customers should be allocated to each group.

3.2 Early Customers

For the first few customers, the odds are chosen so as to maximise the information obtained.

We derive the information matrix, I , using the formula

$$I_{ij} = E \left[- \frac{\partial^2 \ell}{\partial \hat{\pi}_i \partial \hat{\pi}_j} \right]$$

where ℓ is the log likelihood, defined as before.

As described in the previous section, we decided to divide the interval $[0,1]$ into a number of subintervals, each of which was assigned a probability, π_j , which was updated by observation of customers' behaviour.

As before, we may express the log likelihood function as

$$\begin{aligned} \ell &= \sum_{k=1}^N \sum_{j=1}^r X_{jk} \log \pi_j \\ &= \sum_{k=1}^N \left[\sum_{j=1}^{r-1} X_{jk} \log \pi_j + X_{rk} \log \left(1 - \sum_{j=1}^{r-1} \pi_j \right) \right] \\ &= \sum_{j=1}^{r-1} \left(\sum_{k=1}^N X_{jk} \right) \log \pi_j \\ &\quad + \left(\sum_{k=1}^N X_{rk} \right) \log \left(1 - \sum_{j=1}^{r-1} \pi_j \right) \end{aligned}$$

as $\pi_r = 1 - \sum_{j=1}^{r-1} \pi_j$.

Hence, we find that

$$\frac{\partial \ell}{\partial \pi_j} = \sum_{k=1}^N \left[\frac{X_{jk}}{\pi_j} - \frac{X_{rk}}{1 - \sum_{j=1}^{r-1} \pi_j} \right]$$

and

$$\frac{\partial^2 \ell}{\partial \pi_i \partial \pi_j} = - \frac{\sum_{k=1}^N X_{jk}}{\pi_j^2} \delta_{ij} - \frac{\sum_{k=1}^N X_{rk}}{(1 - \sum_{j=1}^{r-1} \pi_j)^2},$$

where $\delta_{ij} = \{1 \text{ if } i = j, 0 \text{ otherwise}\}$. Thus, we have

$$E \left[- \frac{\partial^2 \ell}{\partial \pi_i \partial \pi_j} \right] = \frac{N \delta_{ij}}{\pi_j} + \frac{N}{1 - \sum_{j=1}^{r-1} \pi_j}.$$

The entries are added for each successive customer.

Having calculated the information matrix, we use it to choose the odds for each of the customers before F

is determined. Firstly, both θ_1 and θ_2 are set at $\frac{1}{r}$, for convenience of programming. The information matrix is recalculated for each combination of θ_1 and θ_2 , each being incremented in steps of $\frac{1}{r}$. Finally, that combination of odds which maximises the determinant of the information matrix is used for the next customer, so long as it satisfies the condition $\theta_1 + \theta_2 \geq 1$. In practice, this condition was satisfied by every optimal combination of odds. This procedure is repeated for each of the customers in turn.

The optimal number of customers used to estimate F is found by inspection. This procedure is described later.

After each of these customers bets, our estimate of F is updated using the EM Algorithm, as described previously. Finally, we must initialise the information matrix.

3.3 Initialisation

3.3.1 Odds for the Initial Customers

In order to initialise the information matrix, the theta-values for the first few customers are chosen according to the following plan;

If we divide the interval $[0,1]$ into r equally-spaced subintervals, placing the theta-values on the divisions of these subintervals will give us precise information about the distribution of probability within these subintervals. The optimal value of r is found by inspection, and is described subsequently.

We do not need to set either a theta-value equal to 1, which guarantees no bets, or equal to 0, which guarantees a bet from any customer.

Bearing these points in mind, we set the theta-values for the first customer as

$$\theta_1 = \theta_2 = \frac{r-1}{r}$$

We then take each theta-value down by a value $\frac{1}{r}$ in turn for each of the next few customers.

3.3.2 No. of Customers in this Group

As customers bet on Horse A with probability $1 - F(\theta_1)$, the value of θ_1 will provide us with information about the probabilities of the subintervals above θ_1 . Thus, this value provides us with information about the subintervals at the upper end of the interval $[0,1]$.

Similarly, customers bet on Horse B with probability $F(1 - \theta_2)$. Thus, the value of θ_2 provides us with information about the probabilities of the subintervals below $1 - \theta_2$, and thus provides us with information about the subintervals at the lower end of the interval $[0,1]$.

So the theta-values for the very first customer tell us something about the probability in the first, and last, subintervals. Each successive customer's set of theta-values tells us about an additional subinterval. Finally, we only need information about $(r-1)$ subintervals, as we know that the probabilities sum to 1 in total. Altogether, this tells us that we need $r-2$ customers in the first group, to initialise the information matrix.

4

Choice of No. of Subintervals and No. of Customers to Use in Estimation

Firstly, as discussed previously, we divide the interval $[0,1]$ into r subintervals, to each of which is assigned a probability, so as to estimate F .

We now need to determine

1. the optimal number of subintervals, and also
2. the optimal number of customers, as a percentage of the total (assumed known), whose odds we should use in order to maximise the information matrix, as described in the previous section. This number is in addition to the $r-2$ customers used in the beginning to initialise the information matrix.

These were estimated simultaneously, by calculating profits for the same Dynamic Programming profit function for a variety of combinations of **(1)** numbers of subintervals and **(2)** numbers of customers used in estimation of F , and choosing the combination which proved best overall. The *state of the book* for a particular outcome denotes the bookie's profit if that outcome occurs. Let A_n denote the state of the book for outcome A, and B_n the state of the book for outcome B, when n of the N customers remain. In the strategy which our bookie uses, we have

$$\tilde{\theta}_1 \left(\frac{A_n - B_n}{n} \right) \text{ and } \tilde{\theta}_2 \left(\frac{A_n - B_n}{n} \right),$$

which are chosen to maximise

$$\min\{E[A_0|A_n = a, B_n = b], E[B_0|A_n = a, B_n = b]\}.$$

This gives the function to be maximised as

$$\min \left\{ \frac{d}{2} - n(1 - \theta_1)[1 - F(\theta_1)] + n\theta_2 F(1 - \theta_2), \right. \\ \left. - \frac{d}{2} + n\theta_1[1 - F(\theta_1)] - n(1 - \theta_2)F(1 - \theta_2) \right\},$$

where $d = a - b$. This is the objective function which was used in the simulation study, a summary of whose results follows. It assumes that the quoted probabilities remain constant for all n remaining customers, and calculates the final expected profit if A occurs as the income from those customers who bet an amount θ_2 on B , with probability $F(1 - \theta_2)$, less the outgoing return of 1 unit to those who bet an amount θ_1 on A , with probability $1 - F(\theta_1)$. A similar calculation determines the final expected profit if B occurs. The algorithm involves the calculation of the minimum of these two expected final profits, given the current state of the book.

It will be noted that this is a different algorithm to the optimal one discussed in Section 2; as discussed in Barry & Hartigan[2], this algorithm provides an easier method for calculation of the quoted probabilities, without excessive penalty in terms of the bookie's profit.

The measure of which combination of (1) and (2) proved best was provided by obtaining the mean profit, over fifty replications in each case, for each individual combination. The difference between each of these values and the maximum value over all combinations was then obtained for each distribution. This was repeated for each of $N = 100, 500$ and 1500 customers, and for each of five distributions; namely,

1. Uniform i.e. $F(\theta) = \theta$

2.

$$F2(\theta) = \begin{cases} 0 & 0 \leq \theta \leq \frac{1}{8} \\ 2\left(\theta - \frac{1}{8}\right) & \frac{1}{8} \leq \theta \leq \frac{5}{8} \\ 1 & \frac{5}{8} \leq \theta \leq 1 \end{cases}$$

3.

$$F3(\theta) = \begin{cases} 0 & 0 \leq \theta \leq \frac{1}{4} \\ 2\theta - \frac{1}{2} & \frac{1}{4} \leq \theta \leq \frac{3}{4} \\ 1 & \frac{3}{4} \leq \theta \leq 1 \end{cases}$$

4.

$$F4(\theta) = \begin{cases} 0 & 0 \leq \theta \leq \frac{3}{4} \\ 4\theta - 3 & \frac{3}{4} \leq \theta \leq 1 \end{cases}$$

5.

$$F5(\theta) = \begin{cases} 2\theta & 0 \leq \theta \leq \frac{1}{4} \\ \frac{1}{2} + 2\left(\theta - \frac{3}{4}\right) & \frac{1}{4} \leq \theta \leq \frac{3}{4} \\ \frac{1}{2} & \frac{3}{4} \leq \theta \leq 1 \end{cases}$$

We found the maximum difference, for each combination of number of subintervals and percentage of customers, over the five distributions. This represents the maximum loss per customer. Thus, we use the combination which provides the smallest value of maximum loss. The maximum loss per customer over all distributions is shown in the following tables.

Maximum Difference over Five Distributions

Intervals	% of Customers to Estimate F		
	0	1	2
1	0.0711	0.0715	0.0714
2	0.0358	0.0357	0.0356
3	0.0229	0.0216	0.0259
4	0.0259	0.0181	0.0195
5	0.03296	0.03624	0.03621
6	0.03305	0.03005	0.0295
7	0.03139	0.03453	0.03369
8	0.03531	0.03939	0.03728
10	0.03357	0.03355	0.0361

N=100

Intervals	% of Customers to Estimate F		
	0	1	2
1	0.015	0.015	0.015
2	0.009	0.009	0.009
3	0.0049	0.0054	0.0054
4	0.0035	0.0034	0.0036
5	0.0044	0.0057	0.0066
6	0.0049	0.0047	0.0051
7	0.005	0.005	0.0048
8	0.004	0.004	0.004
10	0.004	0.004	0.004

N=500

Intervals	% of Customers to Estimate F		
	0	1	2
1	0.005	0.0049	0.0049
2	0.0031	0.0031	0.0031
3	0.0016	0.0018	0.0018
4	0.0015	0.0015	0.0015
5	0.0015	0.0019	0.0022
6	0.0019	0.0019	0.0019
7	0.0019	0.0012	0.0017
8	0.0015	0.0015	0.0015
10	0.0015	0.0014	0.0014

N=1500

From these tables, we may see that the optimal % of customers for maximisation of the information matrix in all cases is 1%; higher percentages are not shown here, as they led to greater loss. We also see that the optimal number of subintervals into which to divide the interval $[0,1]$ is 4 for $N = 100$ and 500, and 7 intervals for $N = 1500$. The optimal combination is that which minimises the difference shown in the above tables.

We may further see from these tables, however, that the maximum difference over all distributions decreases, for each combination, as the total number of customers increases- demonstrating that, for larger numbers of customers, there is reduced loss in using a non-optimal combination.

5 Summary and Conclusions

In summary, the method described in this paper provides us with a means of estimating the overall distribution of customers' probabilities, based solely on the betting practices of relatively few initial customers, which provide us with interval estimates of these probabilities. This proves a highly useful tool when distributions are unknown.

Further work on this topic might include the examination of whether it is possible to incorporate an element of profit maximisation into the stage where F is being determined. Another obvious extension of the work is to the case where there are more than two possible outcomes; however, each extra outcome leads to multiple extra possibilities for the customer, who may bet on any individual outcome, or possibly on a combination of them. As well as leading to a much more complicated model, this gives rise to the possibility of incoherence, and to the incurrence of sure loss; care needs to be taken in this scenario.

Acknowledgements

This paper is based on part of the first author's PhD thesis, completed at the University of Limerick under the supervision of Prof. D. Barry, whose help and support have proven invaluable.

References

- [1] Aoki, M. (1967). *Optimization of Stochastic Systems*. Academic Press, New York
- [2] Barry, D. & Hartigan, J. A. (1996). "The Minimax Bookie." *J. Appl. Prob.* **33** 1093-1107
- [3] Bellman, R. (1957) a. *Dynamic Programming*. Princeton University Press, Princeton, NJ

- [4] Bellman, R. (1967) b. *Introduction to the Mathematical Theory of Control Processes*. Academic Press, New York
- [5] Blackwell, D. (1976). "The Stochastic Processes of Borel Gambling and Dynamic Programming." *Annals of Statistics* **4** 370-374
- [6] Dubins, L. E. & Savage, L. J. (1965). *How to Gamble if you Must*. McGraw-Hill, New York
- [7] Henery, R. J. (1984). "An Extreme-value Model for Predicting the Results of Horse Races." *Appl. Statist.* **33** 125-133
- [8] Henery, R. J. (1985). "On the Average Probability of losing Bets on Horses with given Starting Price Odds." *J. R. Statist. Soc. A* **148** 342-349
- [9] Hoerl, A. E. & Fallin, H. K. ((1974). "Reliability of subjective evaluations in a high incentive situation." *J. R. Statist. Soc. A* **137** 227-230
- [10] Plackett, R. L. (1975). "The Analysis of Permutations." *Appl. Statist.* **24** 193-202
- [11] Rieder, U. (1976). "On optimal policies and martingales in dynamic programming." *J. Appl. Prob.* **13** 507-518
- [12] Whittle, P. (1982). *Optimization Over Time: Dynamic Programming and Stochastic Control*. Academic Press, New York
- [13] Walley, P. (1991). *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall
- [14] Barry, D. & Lynch, C. (2006). "The Minimax Bookie: The Two-Horse Case." *Adv. in Appl. Prob.* **Vol. 38** No. 4