Journal of Official Statistics
Vol. 6. No. 4, 1990, pp. 421–435

Developing an Optimal Call Scheduling Strategy for a Telephone Survey

Betsy S. Greenberg[1] *and S. Lynne Stokes*[1]

Abstract: We use a Markov decision process to minimize the expected number of telephone calls needed to make contacts for a random digit dialing survey. Our states include information about the history of calls made to a telephone number and the action to be selected is the time of the next call attempt. Transition probabilities are estimated using a polytomous logistic regression model with data collected by the Bureau of the Census during a test of telephone interviewing for the Current Population Survey. We find that the optimal strategy for that survey should reduce the number of call attempts required and improve the nonresponse rate.

Key words: Random digit dial; logistic regression; Markov decision processes.

1. Introduction

Many large U.S. government survey organizations collect data by telephone because of the potential for improvements in timeliness and cost over personal visit interviewing. The high telephone ownership rate in the United States has lowered concerns about coverage bias. Development of computer assisted telephone interviewing systems (CATI) has increased data processing efficiency.

A variety of sample designs are used for these surveys, such as sampling from lists developed from directories or Waksberg-Mitofsky schemes (Waksberg 1978), which produce self-weighting samples of households. The steps of the data collection process are the same regardless of the design used. The status of the number must first be ascertained (nonworking, business or other nonresidential, or residence), the household contacted, and then the interview completed. All three steps are completed in one call for some sample units, while many calls are required for others. The first step is sometimes difficult to complete because, for example, "out-of-service" recordings are not used by some telephone companies. The second step requires attempting to contact at a time when an eligible respondent or at least a knowledgeable person (who can suggest a convenient appointment time) is available. And finally, the third step requires winning a respondent's cooperation.

Successful completion of the first step requires good information about how telephone companies identify nonworking numbers (such as with a nonworking number recording) and how to obtain information from them for numbers that are not identified. It also requires a well-designed

[1] Management Science and Information Systems Department, CBA 5.202, Business School, University of Texas at Austin, Austin, TX 78712-1175, U.S.A.
Acknowledgements: We thank John Mote and Anthony Brooke for their help in using GAMS.

series of questions to distinguish businesses and other nonresidential numbers from residential numbers. Completing the third step, once a contact has been made, requires a good questionnaire and a well-trained and experienced interviewing staff. Good survey organizations are well aware of these needs and make attempts to satisfy them.

The goal of reaching an uncontacted case is more problematic. In face to face interviews, information from neighbors and appearance of the sample unit or neighborhood are frequently used by interviewers for scheduling the next call. In telephone surveys, only the calling history of the case is available for use in scheduling contact attempts. There is no strategy known to be best for minimizing the effort of reaching households in telephone surveys. In fact, the details of such a strategy would be likely to vary from one survey to the next, depending on such factors as characteristics of the target population, the sample design, and even the staffing capabilities of the survey organization.

Weeks (1988) points out that the effort in making the phone calls typically consumes a significant portion of a survey's available resources. He writes (p. 420) that "calling protocols are still dominated by intuition and folklore" and that serious methodological research in these areas would make a significant contribution. The research that has been done on the question of call scheduling, for both face to face and telephone surveys, has concentrated nearly exclusively on the goal of finding the best time and day for making the first attempt (Falthzik 1972; Weeks, Jones, Folsom, and Benrud 1980) or by grouping together all calls (Vigderhous 1981). But most call attempts made in telephone surveys are not first attempts. So, an efficient call scheduling strategy for subsequent attempts would provide a potential for improving efficiency in data collection.

Attempts to improve the timing of callbacks have been reported. Warde (1989) tried to determine optimal waiting times for callbacks based on the last outcome only. Weeks, Kulka, and Pierson (1987) present evidence suggesting that second calls (in addition to first calls) should ideally be scheduled on weekday evenings or on weekends. They also mention that other factors such as the availability of qualified interviewers and the capacity of the interviewing facility need to be considered. Kulka and Weeks (1988) examined call outcomes across a series of three calls for the purpose of developing optimal calling protocols. Stokes and Greenberg (1990) developed a method for assigning calling priorities to uncontacted numbers with the goal of maximizing the number of contacts or interviews during a shift. This scheme attempts to use the calling resources in the best way possible, and has the advantage of being easy to implement. However, it does not consider the importance of quickly identifying nonresidential numbers, nor does it give high priority to phone numbers that are difficult to reach.

In this paper a method is described for the development of a strategy which has the broad objective of minimizing the total number of calls (both first and subsequent ones) required to contact a sample of households and at the same time, to reduce the proportion never contacted for a surveying period of fixed time duration. The strategy consists of a set of rules for scheduling the time of the next call to each unit, based on that unit's calling history. Using this method, it is easy to include staffing constraints in the strategy.

Section 2 of this paper describes the methodology. In Section 3 we illustrate an application of this methodology. We describe our Markov decision model, the factors that we found useful for estimating the parameters needed, and the solution we obtained from our optimization. Recommendations and conclusions follow in Section 4.

2. Methodology

In a survey, when an interview attempt has been unsuccessful because the unit to be interviewed was not contacted, the time of the next attempt to contact that unit must be selected. In telephone surveys, the only information available to aid in this decision is the amount of time remaining for the survey to be completed, the history of calls to the sample unit, and their outcomes. The small number and objectivity of the available information suggest that a generalizable method for improving the efficiency of callback scheduling may be possible to develop. This section describes a Markov decision process and how a process of this type could be used to minimize the number of phone calls required to contact a survey population. In Section 3 the use of this method is illustrated with a model developed to find the optimal strategy for a specific telephone survey.

2.1. *Markov decision process*

We assume that a process observed at discrete time points is in one of a number of possible states. After observing the state of the process, an action must be selected from the set of possible actions. If the process is in state i and action a is selected, then we assume that

i. a reward $R(i, a)$ is earned, and

ii. the new state of the system is chosen according to the transition probabilities $P_{ij}(a)$.

This implies that both the costs and the transition probabilities are functions only of the last state and the subsequent action. In order to choose actions, we must follow some policy. If a stationary policy is employed, i.e., if the action taken at time t depends only on the state at time t, then the sequence of states forms a Markov chain and the process is called a Markov decision process.

A Markov decision process can be used to model the process of contacting units in a telephone survey. That state of the process includes information about the time and outcome of previous calls to the unit and the time remaining in the survey. Based on the state, an action is selected, i.e., when, if ever, to make the next call. Following each attempt, the state of the process is updated. By including all of the information pertaining to the previous call attempts in the state description, the Markovian assumption will be satisfied.

Each time a case is completed, for example, because an interview was completed or the survey period ended, we consider that the Markov decision process starts over. The objective of the decision process is to maximize the long run average reward earned per unit time. Since each call attempt changes the state of the Markov process, the optimal policy will maximize the average reward earned per call. Rewards are assigned so that the highest reward is given for the most desirable outcome, such as an interview. The lowest reward is given for the most undesirable outcome, such as completing the survey period without making a contact. The optimal policy gives, for each state, the appropriate action that will result in a maximum average return.

Optimal policies do not necessarily exist for all Markov decision problems with expected reward criteria. However, if there is a state, call it 0, and $\beta > 0$ such that $P_{i0}(a) \geqslant \beta$ for all i and a, then an optimal stationary policy will always exist (Ross 1983). In a telephone survey there is always a positive probability that the next call will result in an interview. By calling the resulting state after an interview state 0, it is clear that an optimal stationary policy is guaranteed to exist for this problem.

The stationary policy is allowed to be "randomized" in that actions are selected according to a probability distribution. If π_{ia} denotes the steady-state probability of being in state i and choosing action a, then the expected reward is $\Sigma_i \Sigma_a R(i, a)\pi_{ia}$. Therefore the problem of finding the optimal policy can be written as

$$\text{maximize} \sum_i \sum_a R(i, a)\pi_{ia}$$

$$\text{subject to } \pi_{ia} \geqslant 0, \text{ for all } i, a, \quad (1)$$

$$\sum_i \sum_a \pi_{ia} = 1, \text{ and} \quad (2)$$

$$\sum_a \pi_{ja} = \sum_i \sum_a P_{ij}(a)\pi_{ia} \quad \text{for all } j. \quad (3)$$

Constraints (1) and (2) are obvious. The first requires that the probability of being in a particular state and taking a particular action is not negative. Constraint (2) requires that the sum of all these probabilities is equal to one. Constraint (3) follows because the left-hand side is the probability of being in state j and the right-hand side is the same probability computed by conditioning on the state and action chosen one stage earlier. This maximization problem is a linear program (Manne 1960) and can be solved using any simplex code. Although the model allows for a randomized strategy, the actual solution will not be random. For each state, only one state/action pair will have a positive probability indicating that this is the optimal action.

White (1985, 1988) surveyed applications of Markov decision processes and found that despite a surge of work in the computational field and extension of Markov decision process theory, real applications were surprisingly rare. Reasons suggested for the lack of applications included the difficulty in obtaining transition probabilities and the sheer size of the problems.

The size of the optimization problem is determined by the number of possible states and state/action pairs used. The number of state/action pairs determines the number of variables and the number of possible states determines the number of constraints in the linear program. Each state/action pair requires one parameter estimate (the transition probability) for each of the possible outcomes.

The more information about the call history included in the state of the system, the better the Markov model will represent reality. However, this will also require a large number of possible states and a correspondingly large number of probability estimates. Obviously the quality of the optimization model and the optimal policy obtained reflects the quality of the estimates used to define the model. Therefore, it is important to carefully consider the trade-off between model accuracy and accuracy of the estimated parameters by only including information in the states that is actually found to be predictive of future outcomes. We next describe an approach that can be used to find a parsimonious set of model states and to estimate the required transition probabilities.

2.2. *Logistic regression*

In order to develop an optimal strategy, it is necessary to select model states and estimate probabilities for the possible outcomes of each state/action pair. To do this, one could arbitrarily select states and proceed to estimate the probabilities directly from the raw call history data of a similar survey; i.e., $P_{ij}(a)$ could be estimated as the proportion of calls in state i which were observed to move to state j after action a was taken. The problem with this direct approach is that there will be a very large number of transition probabilities to be estimated for even the simplest set of state/action pairs. As a simple example, suppose the states are defined only by the number of calls previously made to the number and the outcome of the most recent call, and that the possible actions are restricted to only two: call back during the next possible day or evening shift. Suppose that only four outcomes are possible (busy signal, unanswered ring, business or other nonresidential, and interview) and that a maximum of 10 calls are allowed to any telephone number. Since there are only two possibilities for previous outcome (busy signal and unanswered ring), it can be seen that there are still 160 transition probabilities (10 values for number of previous calls $\times 2$ values for previous outcome $\times 4$ values for next outcome $\times 2$ possible actions) to be estimated. Most researchers, ourselves included, do not have a call history data set large enough to support estimation of the number of parameters that would be required for even this moderately detailed specification of state/action pairs.

One solution to this problem is to approximate the transition probabilities using a model with a manageable number of parameters fit to the call history data. A logistic regression model is one possible choice. The categories to be predicted by the model are the possible call outcomes and the explanatory variables are the characteristics of the call history and selected action. Since more than two call outcomes are possible, a polytomous logistic regression model is needed. If there were N possible outcomes, this model assumes that the outcome for a case having call history characteristics represented by the vector $\boldsymbol{X}$ is multinomial with P_i, the probability of outcome i satisfying

$$\frac{P_i}{P_N} = e^{\boldsymbol{X}\boldsymbol{\beta}_i} \text{ and } P_1 + \ldots + P_N = 1$$

where $\boldsymbol{\beta}_i$ is a vector of parameters to be estimated for each of the $N - 1$ possible outcomes. (The Nth outcome is the complement of the others.)

For example, for the simple set of state/action pairs described above, the $N = 4$ outcomes are busy signal, unanswered ring, business of other nonresidential, and interview. There are three explanatory variables, two of which are class variables. They are the number of previous attempts, last outcome (busy or unanswered ring), and the next attempt time (day or evening). If there were no interactions between explanatory variables and the logit were linear in the number of previous attempts, then there would be 12 parameters to be estimated to fit this model: 4 parameters (1 intercept + 1 coefficient for previous attempt + 1 indicator coefficient for last outcome + 1 indicator coefficient for next attempt time) for each of $N - 1 = 3$ outcomes.

Once $\boldsymbol{\beta}_i$, the vectors of regression coefficients, are estimated, the model can be used to predict the probability of each outcome for a unit having a specified call history $\boldsymbol{X}$. Such predictions are smoother than ones made directly from the data, some of which will be based on few observations. They will also be subject to less sampling variability, since far fewer parameters will be estimated. However, since they may be subject to substantial bias resulting from a poor fit to the data, careful checks of model fit, such as examination of residual plots, should be made (see, for example, McCullagh and Nelder 1983).

A further advantage for the modeling approach to estimating transition probabilities is that it can aid in the selection of state/action pairs. Hypothesis tests help in identification of which characteristics of call history and actions are predictive of call outcome. These tests are better suited to identifying the effects than ones performed directly on observed outcomes, since tests on regression coefficients provide a way of separating the effect of one characteristic from that of the others; in other words, it allows for testing the amount of predictive ability for one characteristic while holding the others constant.

2.3. *Model sensitivity to parameter estimates*

When an optimal strategy is implemented, we expect that the actual improvement will be somewhat less than predicted by the model because the model selects optimal actions based on estimates rather than true probabilities. If we consider the linear program defined above, we see that equation (2) requires that the decision variables sum to 1. The effect of this constraint is that the objective function will be very flat. Therefore, the optimal value found by the model should be relatively insensitive to errors in estimation of the transition probabilities. In contrast, the optimal policy may in some instances be very sensitive to errors in the estimates. For example, if in a particular state the probability of an outcome with a high reward is about the same for two different possible actions, then the model will select the action with the higher estimated probability. If the actual probabilities are close, then even for very accurate estimates, it is likely that the wrong action will be selected. Fortunately, this should not matter, because the two actions are nearly equal with respect to the expected reward.

3. Illustration

To illustrate the method, this section develops a calling strategy for a survey conducted by the U.S. Bureau of the Census. The target population is the same as that of the Current Population Survey, which is that of noninstitutionalized civilians in the United States. The survey has a planned sample size of 600 households to be collected over a two-week period. The strategy was developed using data collected by the Census Bureau during an experimental random digit dial telephone survey (RDD-I). The data includes call histories from two replicates of the survey.

A Waksberg-Mitofsky design was used for the survey, so there are nonworking, business, and other nonresidential telephone numbers among the numbers called. When a nonresidential telephone number is identified, it is replaced with a new randomly selected number. It is important to identify these numbers quickly so that they can be replaced and a residential respondent can be interviewed before the end of the survey. Calls to nonworking numbers were excluded from our analysis because we felt that no strategy specifying calling times could improve the efficiency of identifying such cases. Business and other nonresidential numbers were not excluded because the data showed that the time at which these numbers are called does affect how quickly they can be identified.

3.1. *Estimates of transition probabilities*

In order to develop a call scheduling strategy, it is first necessary to determine the important factors that are predictive of future outcomes. These are the factors that will be included in the state description and used to estimate probabilities for the possible outcomes of a call. The data are from a survey identical to the one we are designing and are therefore ideal for determining the important predictive factors, for estimating probabilities, and comparing the strategy obtained from the optimization model.

The data base used to make these estimates contained information about the time of day, day of the week, and outcome of each call made during the production phase of interviewing. This included information on 4,196 calls made to 1,474 telephone numbers, about 1,200 of which were residences. The factors that affected probabilities of a contact or interview and a method for estimating the probabilities were described in Stokes and Greenberg (1990). Those results are summarized in this section, along with a description of those factors affecting other possible outcomes.

The outcome of each call to a telephone number in our sample was classified into one of six categories: business or other nonresidential, busy, unanswered ring, other noncontact (such as fast busy or silence), interview, or contact. A call was classified as an interview only if the entire questionnaire was completed. Every other case in which a residential respondent answered the telephone, such as a partial interview, an appointment to call back, or even a refusal, was classified as a contact.

The possible actions making up the strategy are the day and time at which a call to a given telephone number will be made. The choices for time were described by the shift: day (9:00 a.m.–5:00 p.m. weekdays), evening (5:00 p.m.–9:00 p.m. weekdays), or Saturday (10:00 a.m.–4:00 p.m.). The data did not support a finer classification of these time shifts. That is, morning and afternoon did not show a statistically significant difference, nor did early and late evening. Data were not available for call attempts made outside of these shifts.

In order to find the optimal strategy, estimates of probabilities of each call outcome for any action and call history pair were needed. In the case of first calls there is no previous history, so the only possible explanatory variable is the time of the call. Estimates of these probabilities for first calls were made from the 1,474 first calls by assuming a simple multinomial model. These estimates showed that the probability of a successful outcome (either an interview or contact) was highest for the evening. This observation has previously been made by several authors for a variety of target populations (e.g., Kulka and Weeks 1988; Vigderhous 1981; Warde 1989).

Estimates of the probabilities for subsequent calls to a number were made from a polytomous logistic regression model built from the 2,722 additional calls made to the telephone numbers which were not contacted in the first attempt, the same data base used for fitting a model for prediction of the probability of interview or contact in Stokes and Greenberg (1990).

The dependent variable in the model is the categorical variable "outcome" which has the six possible values described above. The explanatory variables found to be useful for prediction of outcome probabilities are basically the same as those found in Stokes and Greenberg to predict the probability of the one outcome interview or contact. The effect from each of these factors will be discussed briefly.

1. Shift during which the call was made.
   As with first calls, we found that subsequent calls made in the evening have higher probability of resulting in a contact or interview. The shift of a call attempt is indicated using the variables DAY SHIFT and EVENING SHIFT. (The third level is indicated when both of these variables equal zero.)

2. Number of previous attempts to the number.
   The probability of an interview or contact decreases steadily as number of previous attempts (ATTEMPTS) increases, while the probability of an unanswered ring increases rapidly.

3. Timing of previous unsuccessful calls.
   Next, the timing of previous unsuccessful

calls has an effect on the likely outcome. When a previous weekday attempt made in the day had resulted in an unanswered ring, indicated by the variable *D*, the probability of a contact or interview was reduced for subsequent daytime attempts. Similarly, a previously unanswered evening call, indicated by the variable *E*, reduced the probability of a successful outcome on subsequent evening calls. These effects are included in our model as interactions between the variables *D* and *E* and specific shifts. A similar observation could not be confirmed from the data for Saturday calls.

4. Elapsed time since the previous call.
Longer waiting times between calls increased the probability of interview or contact outcomes. Two variables are used in the model to describe waiting time. One is the number of days since the last attempt (LAGDAY) and the other is an indicator of whether or not the previous call was within two hours (IMM).

5. Previous outcome.
The previous outcome received for a telephone number is related to its current outcome. There are three possible previous outcomes: busy, unanswered ring, and other noncontact. Of these three, busy has the highest probability for a subsequent contact or interview outcome, while unanswered ring has the lowest probability. This remains true regardless of the elapsed time since the previous call; that is, there is no interaction between elapsed time and previous outcome. Also, repeated unsuccessful outcomes are extremely common; that is, previous outcome of busy increases the probability of a busy outcome, with similar effects for unanswered ring and other noncontact. This effect is included in the model by the presence of two indicator variables, LAGRING and LAGBUSY, which describe the previous outcome.

In addition, interactions among several pairs of variables were considered; no other interactions were significant. The logistic regression model required estimation of 50 parameters. In contrast, if we had estimated the probabilities directly from the data without using a model, we would have to estimate 2,592 parameters from only 2,722 observations. The model was fit by maximum likelihood. The analysis of variance table describing the fit of this model to the data is shown in Table 1. A strong predictive ability is shown for each of the variables described above. The fit of the model was

Table 1. Analysis of variance table for logistic regression model

Source	Degrees of freedom	Chi-square	P-value
INTERCEPT	5	26.11	0.0001
DAY SHIFT	5	49.69	0.0001
EVENING SHIFT	5	12.63	0.0271
ATTEMPTS	5	69.26	0.0001
LAGDAY	5	11.46	0.0429
IMM	5	16.16	0.0064
LAGBUSY	5	43.71	0.0001
LAGRING	5	130.70	0.0001
DAY SHIFT × D	5	81.69	0.0001
EVENING SHIFT × E	5	28.68	0.0001

evaluated by examining plots of weighted residuals for all six categories against predicted values and the variables LAGDAY and ATTEMPTS. The only trend apparent in our plots was that the model was slightly more likely to underestimate the probability of contact and interview when those probabilities were small than when they were large. The number of such cases was small enough that we do not consider this to be a problem.

While our model appears to fit well, the variance of parameter estimates would be smaller if more data had been available. With a larger data set, it is possible that other factors, interactions, or finer classifications of the time shifts may have been statistically significant. In addition, if more data were available, it would be possible to develop separate models for each call so that the predictive factors could be different for different values of the variable ATTEMPTS.

3.2. The optimization model

A Markov decision process is used to model the problem of selecting an optimal calling strategy for working telephone numbers. The state of the system is described by the following information about the history of calls to a particular phone number:

1. whether or not this number was ever called and received an unanswered ring during a weekday day shift (D),
2. whether or not this number was ever called and received an unanswered ring during an evening shift (E),
3. the outcome (PREVOUT) and time (TIME) of the previous call, and
4. the number of attempts (ATTEMPTS) already made to this phone number.

After observation of the state, an action, i.e., the time of the next call is chosen. Based only on the state and the action chosen, the probability distribution for the next outcome is determined. For each state, only certain actions are feasible, as the time of the next call must be after the time of the previous call. This is why we include TIME in the state description even though it is not used as a variable predictive of outcome. The model state combined with the selected action determines the values for all of the variables used in the logistic regression model described in the previous section. Therefore the transition probabilities for this optimization model can be estimated from the RDD-I data using the logistic regression model. The action selected and the outcome obtained determine the next state and the reward earned.

When the outcome of a call is a business or other nonresidential number, the probability distribution for the next outcome is determined by assuming that a brand new phone number is being called. When a call is unsuccessful, i.e., when the outcome is an unanswered ring, busy, or other noncontact, the probability distribution for the next state is determined by the call history to that phone number as described by the new state. When a call results in an interview or a contact without an interview, a reward is earned, no further calls are made to that number and the process starts over. When the survey period ends before a contact is made a negative reward (penalty) is earned as the process starts over. The objective of the Markov decision process is to maximize the long run average reward earned.

The objective of this study is to minimize the number of phone calls required to contact a household. However, an interview is certainly better than a refusal or even an appointment to call again. Contacts at certain times of the day, regardless of the skill of the interviewer, are more likely to result in an immediate interview, while contacts at other times are more likely to result in refusals or appointments. In the RDD-I data, approximately 54% of first contacts that were made on the weekend resulted in immediate interviews, while only 50% of first contacts made in the evening and 40% of first contacts in the day resulted in immediate interviews. By making contacts at better times, the additional work required to complete the case can be reduced. Therefore a smaller reward is assigned when a call results in a contact without an immediate interview.

In the RDD-I data there were 866 interviews and a total of 6,227 calls to working numbers. Therefore, the strategy used in selecting times to call resulted in 0.1391 interviews/call. Of the interviews, 492 were made at the first contact and the remaining required additional calls. A total of 613 phone numbers had contacts that did not result in immediate interviews. In the Markov decision model the process starts over with a new phone number after a contact. Therefore the model maximizes the expected reward per call for calls made up to and including the first call in which a residential respondent answered the phone. In the RDD-I data this included 4,196 calls. If a reward of 1 is assigned for interviews, then the reward, r, for contacts must be such that

$$\frac{492 + r613}{4196} = 0.1391.$$

Therefore a reward of 0.1493 is assigned each time a call results in a contact without an interview.

If by the end of the two week survey period a phone number has not been distinguished as residential or not, the phone number is assumed to be a residence and counted as a nonresponse. In the model, when the outcome of the last call of the survey is unsuccessful, the call is counted as a final noncontact. In addition to minimizing the number of phone calls required to contact households, it is also very desirable to reduce bias by decreasing this final noncontact rate. One way to do this is to assign a penalty (negative reward) for this outcome. Another way this noncontact rate could be reduced is by setting a constraint on the proportion of calls that result in a final noncontact. To illustrate the first approach, a penalty of 1 was assigned for a final noncontact.

We have selected a reward structure that depended only on the state of the system and not on the action. Therefore, our objective function could be written as maximize $\Sigma_i \Sigma_a R(i)\pi_{ia}$. However, if the cost of calls in certain shifts (such as evenings and weekends) was more expensive because of overtime costs, then it would be appropriate to have a different reward structure that depends on the action taken.

In our state description, D and E may each take on 2 possible values (either 0 or 1). The variable ATTEMPTS takes on any positive integer value and the outcome of the previous call (PREVOUT) takes on 6 possible values (business or other nonresidential, contact, interview, unanswered ring, busy signal, or other noncontact). The time of the previous call (TIME) could be any one of 22 possible values (10 day shifts, 10 evenings, and 2 Saturdays). In addition, we allow two phone calls to the same number in the same shift and so the previous call may be either the first or second in the particular shift, which means that the time of the previous call may take on 44 different values. Fortunately, not every possible combination of the variables, D, E, ATTEMPTS, PREVOUT, and TIME is possible. For example, when the outcome of the previous call was a business or other nonresidential, a new phone number is selected so the new state must have $D = 0$, $E = 0$, and

ATTEMPTS = 1. The number of possible states was further reduced by requiring that ATTEMPTS = 4 on the fourth and subsequent attempts. In all, 982 states are allowed in our model.

The number of variables required in the linear program is equal to the number of possible state/action pairs. If all feasible actions are allowed, there are, on average, 11 possible choices for each action and nearly 11,000 state/action pairs. Each state/action pair requires 6 parameter estimates (the transition probabilities for the possible outcomes). The number of variables and parameters required was reduced by allowing the choice of significantly fewer actions. Following an unsuccessful call (with outcomes of ring, busy, or other noncontact), we chose to require another attempt to be made by the end of the next weekday. This is an example of a managerial decision that is easily incorporated into the model. In all, we allow 2,763 state/action pairs.

3.3. The solution

The linear programming formulation described in the previous section was specified with algebraic equations using GAMS, the General Algebraic Modeling System (Brooke, Kendrick, and Meeraus 1988) and the optimal solution to the linear program was found using MINOS. GAMS greatly simplified the effort required to generate and solve a linear program with approximately 1,000 variables and 3,000 constraints. In this section, the optimal strategy obtained from the optimization model is discussed.

The strategy obtained from our model is only applicable to a two week household survey similar to the RDD-I telephone survey. The strategy obtained depends on some of the modeling decisions we have described above. For example, the only feasible actions allowed in our model were calls made within the next weekday, so the resulting optimal policy obtained from the model only includes those actions. The quality of the resulting policy obtained from the model also depends on the quality of the data we had available for prediction. For example, our data showed that the probability of a successful outcome on a first call (either an interview or contact) was the lowest for the Saturday shift. This is not consistent with the findings of Weeks et al. (1987) and Kulka and Weeks (1988). This discrepancy may have occurred because our data base included very few Saturday calls. This would affect the optimal policy obtained from the model. In spite of model and data limitations, the policy obtained from our model illustrates the possible improvements from our approach. We describe the resulting policy and improvements below.

Although the optimization model allows a randomized strategy, the solution obtained is not random. For each state, only one state/action pair has positive probability indicating that this is the optimal action. In addition to indicating the optimal actions, the solution to the linear program predicts the fraction of calls with particular outcomes. In the solution of the demonstration model, the fraction of calls with predicted outcomes interview, contact (without interview), and final noncontacts are 0.1641, 0.1732, and 0.0124, respectively. If 600 households are to be interviewed, this would translate to 282 interviews, 297 contacts, and 21 phone numbers that would not have been identified as residential or not.

The optimal strategy obtained from the model could be summarized as follows: Make the first call during the day and all additional calls in the evening until the phone number is identified as being residential or not. If the outcome of the first call is a busy signal or other noncontact, then following an unanswered ring in the evening shift, another attempt should be made to make a contact during the day. Unsuccessful call attempts made in the day are followed by an attempt in the evening of the next day. This strategy is different from that obtained by Kulka and Weeks (1988), possibly because our data base included business and other nonresidential telephone numbers which are most easily identified by calls during the day.

In general, there should not be two calls made to the same number in the same day and certainly not in the same shift. Exceptions to this rule of thumb occur toward the end of the survey and before the first weekend. For example, after identifying a nonresidential number, a new phone number is randomly selected. Calls to new numbers are usually made during the day shift; i.e., if a nonresidential number is identified during an evening shift, the first call to a newly selected number should be made in the next day shift. However, on Friday evenings, Saturdays, as well as Wednesday and Thursday evening of the second week of the survey, calls to newly selected numbers should be made immediately. Similarly, following unsuccessful calls, subsequent attempts should be made earlier as the time allowed to complete the survey runs out. It is not until the last weekday of the survey that it is optimal to make two calls in the same shift.

Following an unsuccessful daytime call, an evening call should be made. Typically, these calls are made the next day of the survey. However, beginning on Wednesday of the second week of the survey, it is sometimes better to make the next call the same evening instead of waiting and possibly running out of time. Whether or not a telephone number is called twice in one day seems to depend on how difficult it is to contact a number given its call history. If the probability of reaching a number is moderately low, the number should be called sooner, so that more calls to that number will be possible. However, if the call history indicates that the chance of reaching the phone number is extremely low, the number should not be called twice in one day. This may be because it is expected that both calls will be unsuccessful.

The solution to the linear program also allows calculation of the proportion of call attempts that should be made in each shift. This information allows planning for the number of interviewers required for each shift. The optimal solution in the demonstration model requires that about 42%, 51%, and 6% of calls be made in the day, evening, and Saturday, respectively. More notable, however, is the fact the model solution specifies that 39% of all calls should be made in the day shift on the first day of the survey. Almost no calls should be made until the evening of the second day of the survey when 19% of all calls should be made. On subsequent days, significantly fewer calls should be made in each shift. As Weeks et al. (1987) point out, in practice, this schedule may not be practical to implement. Because of a limited number of CATI calling stations and interviewers, it will probably be required to spread the calls more evenly over the survey period.

Constraints on the availability of interviewers and calling facilities could be easily included in the optimization model by adding constraints limiting the proportion of calls that are allowed in each shift. To illustrate this, we added the constraint that at most 15% of the calls could be made in any one shift. Our original model only allowed the first day or evening as possible choices for calls at the beginning of the survey. To make the new linear program feasible, we allow first calls during the first two days and evenings. The optimal policy obtained for this model is a randomized strategy. The new strategy requires that first calls are now

Table 2. Comparison of optimization model results with RDD-I data

	RDD-I	Model 1	Model 2
Interviews	492	563	557
Contacts	613	594	604
Business	277	205	229
Non-contact	96	43	39
Total calls	4,196	3,432	3,465
Score	0.139	0.190	0.187

made in each of these shifts. If the optimal policy is followed, 33% of calls will be made in the day, 60% in the evening, and 7% on Saturday.

We find that the optimal policy for subsequent calls is very similar to the one described above. For telephone numbers that are first called in the evening, the optimal strategy is to follow an outcome of unanswered ring in the evening by a call attempt in the day. All subsequent calls should be in the evening unless the call in the day had an outcome of busy or other noncontact. In this case, the model predicts that it will be necessary to make additional attempts in the day. Otherwise, all additional attempts should be made in the evening. With the added constraint limiting the proportion of calls made in each shift, the effect of planning ahead for the end of the survey is seen earlier. Starting on Tuesday of the second week, some phone numbers are called twice in one day. This apparently happens earlier because calls to some phone numbers were started later.

In Table 2, the predicted results of the optimal policy obtained for the two models are compared with the RDD-I data. The results are scaled to represent 1,200 households, the planned sample size for the two replicates of the RDD-I survey, so that the numbers could be compared directly with the actual data. The score computed for each is the number of interviews and weighted contacts divided by the number of call attempts. The first model represents a 36% improvement in this score over the actual data. Specifically, 18% fewer calls should be made resulting in more interviews and fewer phone numbers that were never reached by the end of the survey. The addition of the constraints limiting the number of calls in a shift resulted in a slightly worse performance, but still one that is significantly better than the RDD-I data.

4. Recommendations and Conclusions

The optimal strategy obtained from the Markov decision model does not always take the action that will result in the highest probability of an interview or contact on the next call. Instead, this model is maximizing the average reward earned on all of the calls. For example, on first calls, the probability of success is higher if the call is made in the evening rather than the day. However, the optimal strategy requires that first calls be made in the day (unless time is running out in the survey). While this strategy does not result in the highest immediate reward, it is clearly "planning ahead" for the possibility that additional calls will be required. The planning in the strategy is also seen clearly as time runs out in the survey. Often a calling time will be selected that does not have the highest probability of success. This is required to allow for a sufficient number of calls to reduce the probability of a final noncontact at the end of the survey period.

The results from the two optimization models show that the calling strategy used to collect the RDD-I data could have been significantly improved. The number of calls required and the number of phone numbers that were never contacted are both significantly lower when an optimal strategy is used. The improvement is due to several differences between the way the RDD-I data was collected and the optimal strategy. In the optimal strategy, a given phone number is not (except near the end of the survey) called twice in one day and once an unanswered ring has been obtained during the day hours, the phone number is not called again during the day. During the actual survey, these rules were frequently violated. By maximizing an average reward criterion, the optimal policy plans ahead for the possibility of subsequent calls at the end of the survey period. When an optimal strategy is implemented, we expect that the actual improvement will be somewhat less than that predicted by the model because the model selects optimal actions based on estimates rather than true probabilities.

The optimal solution to the Markov decision model also allows prediction of the number of call attempts that will be made in each shift. This information will allow planning for the number of interviewers required for each shift. The Markov decision model will be very useful for staff planning. In practice, during a particular shift, the interviewers may not complete all the calls that should be made, or they may complete the required calls and have time at the end of the shift to call additional numbers. This means that the optimal strategy will not be precisely followed.

A priority rule could be incorporated to simplify this implementation problem. First priority should be given to calling those numbers which are to be called based on the strategy of the Markov decision model. Since it is possible that only some of these numbers will actually be called, they should also be ranked to determine which numbers should be called first. For example, highest priority could be given, as suggested in Stokes and Greenberg (1990), to those with the highest probability of a contact or interview. Since it is also possible that there will be time to call additional numbers, the remaining numbers should also be priority ranked. Combining the results of the Markov decision model with a priority rule will result in a tool that is useful for planning a survey as well as the day to day operation of the interview facility.

Another possible use of the methodology is to evaluate the currently used calling strategy. Since optimal expected rewards can be estimated from the model, a comparison with actual rewards provides information about whether or not change in the current calling rules are likely to produce large gains.

The methodology we describe in this paper could be used to find optimal calling strategies for other telephone surveys. The states and allowable actions would reflect the specifics of the particular survey. For example, the survey may be conducted over a different length of time or may allow different choices for the times of repeat attempts. The constraints on the proportion of calls should reflect the specifics of the resources available for the survey. The results from several different models could be used to compare different survey alternatives such as a longer survey period versus a shorter one with more interviewers working simultaneously.

The transition probabilities used in the Markov decision model should reflect the target population that is being surveyed. These probabilities should be estimated either using case management data collected in a survey with a similar target population or preferably from data collected in an experiment for that purpose. Use of a logistic regression model has the advantage of allowing identification of predictive characteristics of call histories and also gives a method for estimating transition probabilities from a limited amount of data.

5. References

Brooke, A., Kendrick, D., and Meeraus, A. (1988). GAMS: A User's Guide. Redwood City, California: The Scientific Press.

Falthzik, A.M. (1972). When to Make Telephone Interviews. Journal of Marketing Research, 9, 451–452.

Kulka, R.A. and Weeks, M.F. (1988). Toward the Development of Optimal Calling Protocols for Telephone Surveys: A Conditional Probabilities Approach. Journal of Official Statistics, 4, 319–332.

Manne A. (1960). Linear Programming and Sequential Decisions. Management Science, 6, 259–267.

McCullagh, P. and Nelder, J.A. (1983). Generalized Linear Models. London: Chapman and Hall.

Ross, S.M. (1983). Introduction to Stochastic Dynamic Programming, New York: Academic Press.

Stokes, S.L. and Greenberg, B.S. (1990). A Priority System to Improve Callback Success in Telephone Surveys. Proceedings of the Survey Research Methods Section, American Statistical Association.

Vigderhous, G. (1981). Scheduling Telephone Interviews. Public Opinion Quarterly, 45, 250–259.

Waksberg, J. (1978). Sampling Methods for Random Digit Dialing. Journal of the American Statistical Association, 73, 40–46.

Warde, W. (1989). Contacting the Farm Operator by Telephone: A Re-examination. U.S. Department of Agriculture, Research Report Number SRB-89-07, Washington, D.C.

Weeks, M.F. (1988). Call Scheduling with CATI: Current Capabilities and Methods. In Telephone Survey Methodology, edited by R.M. Groves, P.P. Biemer, L.E. Lyberg, J.T. Massey, W.L. Nicholls, II, and J. Waksberg. New York: John Wiley & Sons.

Weeks, M.F., Jones, B.L., Folsom, R.E. Jr., and Benrud, C.H. (1980). Optimal Times to Contact Sample Households. Public Opinion Quarterly, 44, 101–114.

Weeks, M.F., Kulka, R.A., and Pierson, S.A. (1987). Optimal Call Scheduling for a Telephone Survey. Public Opinion Quarterly, 51, 540–549.

White, D.J. (1985). Real Applications of Markov Decision Processes. Interfaces, 15, 73–83.

White, D.J. (1988). Further Real Applications of Markov Decision Processes. Interfaces, 18, 55–61.

Received August 1989
Revised May 1990