



2013 AMERICAN COMMUNITY SURVEY RESEARCH AND EVALUATION REPORT MEMORANDUM SERIES
#ACS13-RER-18

DSSD 2013 AMERICAN COMMUNITY SURVEY MEMORANDUM SERIES #ACS13-MP-01

MEMORANDUM FOR ACS Research and Evaluation Steering Committee

From: Patrick J. Cantwell (*signed 12/20/2013*)
Chief, Decennial Statistical Studies Division

Prepared by: John Chesnut
ACS Methods Branch
Decennial Statistical Studies Division

Subject: Model-Based Mode of Data Collection Switching from Internet to Mail
in the American Community Survey

Attached is the final American Community Survey Research and Evaluation report on Model-Based Mode of Data Collection Switching from Internet to Mail in the American Community Survey. This report describes the results of research exploring the feasibility and benefits of applying adaptive survey design methods to switch ACS sample cases from Internet to mail.

If you have any questions please contact John Chesnut at (301)-763-6732 or Jennifer Tancreto at (301)-763-4250.

Attachment: Final report on Model-Based Mode of Data Collection Switching from Internet to Mail in the American Community Survey

cc:

ACS Research and Evaluation Work Group
Todd Hughes (ACSO)
Andrew Roberts
Michael Bentley (DSSD)
Mary Davis
Joan Hill
Alfred Meier
Doug Olson
Joseph Powers
Courtney Reiser

Daniel Sommers
Michael Starsinic
Jennifer Tancreto
Mary Frances Zelenak
Dawn Nelson (CSM)
Gina Walejko
Chandra Erdman (CSRM)
Barbara Ohare (DIR)

Model-Based Mode of Data Collection Switching from Internet to Mail in the American Community Survey

FINAL REPORT

John Chesnut

Decennial Statistical
Studies Division

TABLE OF CONTENTS

EXECUTIVE SUMMARY	iii
1. BACKGROUND	1
2. METHODOLOGY	2
2.1 Adaptive Design	2
2.2 Developing an Adaptive Design Process using the ACS Internet Test Samples	2
3. LIMITATIONS	2
4. RESULTS.....	3
4.1 Augmenting the ACS Sampling Frame	3
4.2 Predicting Daily Internet Response with Discrete-Time Survival Analysis	4
4.3 Creating the Group-Based Mode Switch Rules	8
4.4 Mode Switch Assignments	8
4.5 Adaptive versus Non-Adaptive Approach	9
4.6 Applying the Maximization Principle.....	11
4.7 Validating the Adaptive Design	14
5. SUMMARY	14
References	16

EXECUTIVE SUMMARY

Purpose of the study

The former director of the Census Bureau, Dr. Robert Groves, emphasizes in his 2011 vision document, *The Future of Producing Social and Economic Statistical Information*, the importance of the Census Bureau remaining relevant by responding to the threat of competing organizations producing estimates cheaper and faster through passive data sources. More specifically, in addressing this issue in the context of mixed-mode surveys, he states, “Prior to contacting them [respondents], we will be ignorant about the mode-preferences of our sample units; we must be able to switch across modes in real-time during the data collection phase to produce timely estimates.” He went on to say that, “we need to mount mixed-mode surveys that have real-time, rule-based switching across modes based on statistical analysis of paradata.” Some of the potential benefits of switching modes in real-time include reducing lag time for following up with nonrespondents, reduced response burden by providing respondents with a convenient or preferred mode of response, and cost savings by maximizing the number of respondents willing to respond in the cheaper modes of data collection such as Internet and mail.

As a first step in addressing this vision, we explore the use of adaptive survey design methods in the ACS by linking administrative data to the 2011 ACS Internet Test samples and developing a model-based assignment of mode switch strategies, focusing only on switching Internet nonrespondents to mail. We conduct a simulation of the adaptive survey design process with the April Internet test sample and then attempt to validate our model-based approach using the November Internet test sample comparing outcomes in Internet response and length of time to Internet nonresponse follow-up.

Background

The goal of the 2011 ACS Internet Tests was to determine the feasibility of including a fourth mode of response in the ACS – an Internet response option. The Census Bureau conducted the first test in April 2011 and then repeated the test with modifications in November 2011. Our focus is to use those sample case members of the treatment groups that received a letter and instructions for completing the ACS online in lieu of a paper questionnaire. Depending on their assigned treatment, nonrespondents in the April test received a paper questionnaire two or three weeks later. Nonrespondents in the November test received a paper questionnaire two weeks later. For the purposes of our study, we restrict our time frame of interest to the two weeks following the mailing of the letter and instruction for completing the ACS online.

Summary

This study demonstrated the successful development and application of an adaptive survey design process for assigning mode switch strategies in an optimal manner that minimized the time to follow up with Internet nonresponse cases while controlling for cost and error. Simulating our mode switches for

varying assumptions on our controls for cost and error, we found that we could achieve gains in timeliness through the acceptance of controlled increases in predicted follow-up cost and/or through the cost neutral approach of allowing more cases not likely to respond via the Internet to switch to mail prior to data collection. In addition, we demonstrated that in order to identify the optimal cost-benefit trade-off for our application as well as others involving intangible benefits of timeliness, data quality, sampling variance, etc. a valuation step would be required for a successful adaptive design process. Furthermore, in applying our adaptive design framework to an independent sample we demonstrated its portability. More importantly, we developed a simple framework that could serve as a foundation for more complex designs.

1. BACKGROUND

The goal of the 2011 ACS Internet Tests was to determine the feasibility of including a fourth mode of response in the ACS – an Internet response option. The Census Bureau conducted the first test in April 2011 and then repeated the test with modifications in November 2011. The April test included a national sample of 120,000 housing unit addresses for the U.S. and 7,400 for Puerto Rico. The April Internet Test sample design was similar to that used in the production ACS, with the exception of an additional level of stratification and not using differential sampling rates across strata (Keathley, 2010). Using segmentation clustering results from the Census 2010 advertising campaign research (Bates and Mulry, 2008), the ACS Internet Test sampling frame was stratified into “targeted” and “not targeted” Census tracts (Tancreto et al., 2012). Census Bureau researchers classified any tract that was in the segmentation clusters for single, detached, mobile households, or advantaged homeowners as “targeted” and classified any tract in the other clusters as “not targeted.” They expected sample addresses in the targeted stratum to use the Internet response option at a higher rate than those in the not-targeted. The researchers allocated the Internet Test sample equally between the two strata. They also selected cases in “neighboring” quartets (the k^{th} address systematically selected and the next three addresses in the sort formed a quartet) and each address in a quartet was randomly assigned to one of the four treatment groups.

The main objective of the April 2011 Internet Test included studying whether giving households a choice to respond by mail or Internet versus Internet only (in lieu of the first questionnaire mailing) affected overall self-response relative to the current ACS methodology which uses only a paper questionnaire to collect data in the first month of data collection. An additional objective of the test was to answer the research question, given that a household has the choice to respond by Internet or mail, does the prominence of the message to respond via the Internet in the mailing materials impact self-administered response? Furthermore, given that a household has only the option to respond via the Internet, does giving the nonrespondents two weeks versus three weeks before sending the questionnaire package affect the self-administered response rate? The researchers met these objectives using a nested experimental design that included four treatment groups labeled: Choice-Prominent, Choice-Not Prominent, Push-Accelerated, Push-Regular.

Based on further results from the 2011 November Internet Test (Matthews et al., 2012) and cost savings analysis, the Push-Accelerated with a New Postcard methodology was recommended by the Census Bureau and approved by the Office of Management and Budget to be implemented in the 2013 ACS. This method was identical to the Push-Accelerated methodology in the April Test in that sample addresses in this treatment group received a letter and instructions for completing the ACS online in lieu of a paper questionnaire. Nonrespondents received a paper questionnaire about two weeks later. However, in the November Test the same nonrespondents received a new reminder postcard three days after receiving the questionnaire. This new reminder postcard increased the overall response rate (Internet and Mail combined) by 5 percentage points compared to the Push-Accelerated treatment in the April test, but most of the increase was attributed to mail response. The researchers anticipated this result.

In order to use sample cases from the Internet Tests that are presented with similar conditions as production ACS in 2013, we used only sample cases from two of the four April test treatment groups – Push-Regular and Push-Accelerated in our research. In addition, we used only cases from two of the five November test treatment groups – Push-Accelerated and Push-Accelerated with a New Reminder Postcard.” This resulted in 59,964 sample cases from the April test and 39,978 sample cases from the November test eligible for our study. To compensate for the difference in mailing schedules between the Push-Regular and the other treatment groups, we restricted our study of Internet responses to the first two weeks of data collection.

2. METHODOLOGY

2.1 Adaptive Design

We can describe the basic framework for an adaptive survey design process in three major components (Groves and Heeringa, 2006; Wagner, 2008; Schoeten et al., 2011). First, we pre-identify a set of design features or strategies affecting cost and errors of the survey estimates. Second, we identify indicators of the cost and error properties. We derive the indicators from paradata and administrative record data sources that we link to the survey’s sample cases. Third, monitoring the indicators, we alter or adapt the survey design features based on rules that control in an optimal manner the cost-error trade-offs.

2.2 Developing an Adaptive Design Process using the ACS Internet Test Samples

Applying this framework to the limited scope of our study, we identified our set of design strategies as the mode switch days eligible for switching sample cases from Internet to mail. Our cost indicator monitored the mail follow up cost for Internet nonrespondents. In addition, our error indicator monitored the level of Internet response and the contrast between Internet respondents and nonrespondents. These indicators were derived from a propensity model we developed using data from the April ACS Internet Test to predict the daily Internet response propensity for a given sample case, with linked auxiliary frame information serving as the model covariates. We then evaluated these daily indicators of cost and error and altered the assignment of mode switch days in such a way that we minimized the time to Internet nonresponse follow-up.

3. LIMITATIONS

Given that this study is not a pilot or field study, we were not able to evaluate the impact of our simulated adaptive design process in the initial phase of data collection on subsequent phases or modes of data collection.

We were not able to evaluate the cost-benefit trade-offs adequately in determining the mode switch rules that would result in a maximum net benefit since we were not able to express the intangible benefit of timeliness and the tangible cost of nonresponse follow-up in comparable units of value (e.g.,

dollar amounts). Accurately describing the value of timeliness would require a valuation study beyond the scope of this study.

Given that sample case addresses vary in their geographic proximity to the point of mailing origin, the length of time to response may not accurately reflect the true length of time to response.

4. RESULTS

4.1 Augmenting the ACS Sampling Frame

The first step in developing our adaptive design process was to identify sources of data from both administrative records and paradata that we could use to inform the process for switching households from Internet to mail. These data sources needed to demonstrate a relationship with Internet response propensity. To begin, we proposed using the paradata from the instrument that provides time stamp data on a sample case's first access to the Internet instrument. However, we found that 93.6 percent of cases that access the instrument result in an Internet response within 2 days or less. This indicated an almost perfect correlation between this paradata variable and Internet response causing a quasi-complete separation in the data, preventing us from finding a maximum likelihood estimate for this covariate in our daily Internet response propensity model. Not having to monitor paradata collected during data collection, we assigned mode switch days prior to data collection, representing a static as opposed to a dynamic adaptive survey design (Schouten et al., 2011).

Table 1 lists the administrative record data sources selected for this study and their associated variables. We augmented the ACS sample frame data with the proposed administrative record data using the Census Bureau's Master Address File Identifier (MAFID) as the linking variable with the exception of the National Telecommunications and Information Administration (NTIA) data. This linking was possible due to previous work performed by the Census Bureau's Center for Administrative Records Research and Applications (CARRA) that resulted in a unique MAFID associated with nearly every record contained in the proposed administrative data record sources. CARRA extended the process developed for creating the Statistical Administrative Records System (StARS) database (cf. Judson, 2000 and Farber and Leggieri, 2002 for more information on StARS) to link Census files to administrative record data sources. The NTIA data is available through the website, <http://www.broadbandmap.gov/data-download> (December 2011 release) at the Census block-level, therefore these data were linked to the ACS sample data at the Census block-level.

Table 1. Administrative Record Data Sources

Administrative Record Data Source	Variables	ACS Internet Test Sample Linked by MAFID (%)
2010 Census – Housing Unit Response Data File	self-administered questionnaire, language of interview or questionnaire, proxy respondent	96
2010 Census – Edited Household Data File	householder - age, race, and Hispanic origin; tenure; large household	96
2010 Census – Edited Person Data File	no spouse, not related	87
2010 Census – Unedited Operation Data File	mail enumeration area, response check-in-date	99
Master Address File	urban-rural	100
Info USA	do not call flag, high-tech household	85
United States Post Office – National Change of Address Database	change of address flag	100
National Telecommunications and Information Administration	broadband flag	96
Internal Revenue Service	1040 total income reported for 2010	66
2010 Census – Advertising	targeted stratum	100

For the most part, we were successful in linking the ACS sample to the administrative data sources, with the IRS 1040 data being the exception. Excluding the IRS result, the percent of ACS Internet Test cases linked ranged from 85 to 100 percent. Our attempt to link to the IRS 1040 data resulted in only 66 percent of the cases having data on household income.

Administrative record data not available for some of our ACS sample cases posed a missing data problem. Since a portion of those cases with missing IRS data likely represent cases that are late income tax filers or households that failed to file, we chose to create an additional income category of ‘income not reported’ to address the missing IRS data. Note that some cases with missing IRS data may also represent cases where the Census Bureau was not able to match to address information included in the IRS records (e.g., PO boxes and rural routes). To remedy the remaining missing data, we used a multi-stage imputation methodology, first applying a Markov chain Monte Carlo imputation method (Schafer, 1997) to impute enough missing data to create a monotone missing data pattern. Given the monotone missing pattern, we then applied a multiple imputation procedure (Rubin, 1987) using regression based imputation.

4.2 Predicting Daily Internet Response with Discrete-Time Survival Analysis

Initially, we explored the relationships between the proposed administrative record data variables and Internet response. We dropped those variables where we observed a lack of a significant relationship. These variables included the USPS change of address flag, do-not-call status flag, and the 2010 Census proxy response flag. Furthermore, to avoid multi-collinearity issues due to the high correlation, we combined the two variables 2010 Census check-in date of the earliest received census form and the 2010 Census self-administered response indicator, to create a composite variable.

Proceeding with a reduced set of variables from our initial results, our next step was to explore their relationship with time (in days) to Internet response. We accomplished this by using a discrete-time survival analysis modeling approach (Allison, 2010). To use this method, we right-censored sample cases that did not respond on or before day 14. Using this approach our objective was to arrive at an estimate

of daily Internet response propensity for each of the households in sample. First, we used a discrete time logistic model to model the hazard function h_{it} for sample case i at time t (i.e., the probability that sample case i responds via the Internet given that t days have lapsed):

$\log\left(\frac{h_{it}}{1-h_{it}}\right) = \alpha_t + \beta_1 x_{it1} + \dots + \beta_k x_{itk}$ where α_t is a time specific intercept, x_{it1}, \dots, x_{itk} are the covariates at time t for the i^{th} sample case, and β_1, \dots, β_k are the regression parameter estimates. Note that in our case the covariates were not time varying.

To ensure that we developed a model that performs well in modeling Internet response propensity as it relates to time in days, we varied our assumptions about this dependency. Table 2 shows the model calibration and discrimination results for varying assumptions about the dependence of the hazard on time in days. The Akaike’s Information Criterion (AIC) (Akaike, 1974) and the Schwarz’s Criterion (SC) (Schwarz, 1978) are similar to the model fit statistic, -2 times the log-likelihood, but include a penalty adjustment for including more parameters in the model. Since the unrestricted assumption for the variable DAY, where we assume the time dependent variable DAY is categorical, results in the smallest AIC and SC values, we conclude that this assumption provides the best model calibration. In addition to calibration, the area under the Receiver Operating Characteristic (ROC) curve provides a measure of the model’s ability to discriminate between cases that will and will not respond on a given day (Hosmer and Lemeshow, 2000). All of our models provide excellent discrimination, but the unrestricted model demonstrates the best discrimination capability. Overall, we conclude that the unrestricted model was the best performing model.

Table 2. Model Fit Results

Model	Akaike’s Information		
	Criterion (AIC)	Schwarz’s Criterion (SC)	Area Under the ROC Curve
Unrestricted – DAY	511,223,724	511,224,057	0.800
Linear – DAY	550,975,216	550,975,493	0.723
Quadratic – DAY	548,952,150	548,951,025	0.723
Logarithmic – DAY	550,704,122	550,704,399	0.726

Sources: April 2011 ACS Internet Test Data, 2010 Census, Master Address File, Info USA, NTIA Broadband Data, 2010 IRS 1040 Tax Return Data

Table 3 lists the regression parameter estimates and their respective hazard ratios controlling for all other variables for the unrestricted model. From these results, we observe that the parameter estimates for the hazard function, range from -5.67 to 0.29 with all parameters being significant at a significance level of 0.10. Reviewing the hazard ratios for each variable independently (controlling for all other variables), we observe an increased propensity for responding via the Internet for households with characteristics indicating census mail enumeration area geography, broadband service area, census-reported not-related household, census-reported owner-occupied household, and non-high tech household. The remaining ratios indicate a decreased propensity for the non-reference categories. We note that the hazard ratio result for the variable indicating a household’s interest in high tech products and services is contrary to what we would expect. We speculate that households with interest in high tech products and services may have a heightened sense of awareness of the security risk of providing detailed data, such as the data required for the ACS, over the Internet and prefer to use a ‘safer’ mode

of data collection such as the paper questionnaire. Alternatively, the data indicating a household's interest in high tech products may not be accurate.

Table 3. Discrete Time Logistic Model Parameter Estimates and Hazard Ratios

Variable	$\hat{\beta}$	$SE(\hat{\beta})$	Hazard Ratio
Intercept	-1.71	0.07	
<i>Census Advertising Segmentation Cluster</i>			
Not targeted	-0.23	0.02	0.79
Targeted	Reference		
<i>Census Type of Enumeration Area</i>			
Mail Enumeration Area	0.18	0.05	1.19
Non-Mail Enumeration Area	Reference		
<i>Broadband Status</i>			
Broadband Service	0.08	0.03	1.08
No Broadband Service	Reference		
<i>Urban Rural Status</i>			
Rural	-0.25	0.03	0.78
Urban	Reference		
<i>Census Householder Age</i>			
Age >65	-0.47	0.03	0.62
Age ≤65	Reference		
<i>Census Large Household Status</i>			
Large Household	-0.43	0.05	0.65
Non-Large Household	Reference		
<i>Language of Census Form/Interview</i>			
Non-English	-1.15	0.18	0.32
English	Reference		
<i>Census Minority Status</i>			
Minority	-0.41	0.02	0.66
Non-Minority	Reference		
<i>Census Household Composition 1</i>			
No Spouse	-0.28	0.02	0.76
Spouse	Reference		
<i>Census Household Composition 2</i>			
Not Related	0.17	0.04	1.18
Related	Reference		
<i>Census Tenure</i>			
Owned	0.19	0.03	1.21
Not Owned	Reference		
<i>High Tech Household (Info-USA)</i>			
Non-High Tech	0.29	0.03	1.33
High Tech	Reference		
<i>Census Check-in Date and Form Type</i>			
Late Check-In and Non-SAQ	-0.89	0.03	0.42
Late Check-In and SAQ	-0.65	0.05	0.52
Normal Check-In and SAQ/Non-SAQ	Reference		
<i>IRS Reported Household Income</i>			
Income Not Reported	-0.83	0.05	0.44
\$0 - \$10,000	-0.50	0.05	0.61
\$10,001 - \$15,000	-0.72	0.07	0.49
\$15,001 - \$25,000	-0.68	0.05	0.51
\$25,001 - \$35,000	-0.61	0.05	0.54
\$35,001 - \$50,000	-0.45	0.05	0.64
\$50,001 - \$75,000	-0.29	0.05	0.75
\$75,001 - \$200,000	-0.13	0.04	0.88
\$200,001 +	Reference		
<i>Time (in Days)</i>			
Day 1	-5.67	0.30	0.00
Day 2	-2.28	0.04	0.10
Day 3	-1.35	0.04	0.26
Day 4	-1.32	0.03	0.27
Day 5	-1.58	0.03	0.21
Day 6	-1.73	0.04	0.18
Day 7	-1.50	0.04	0.22
Day 8	-1.53	0.04	0.22
Day 9	-1.97	0.05	0.14
Day 10	-2.31	0.05	0.10
Day 11	-2.61	0.05	0.07
Day 12	-2.79	0.06	0.06
Day 13	-2.80	0.07	0.06
Day 14	Reference		


Note: All regression parameter estimates are significant at the 10 percent significance level.

Sources: April 2011 ACS Internet Test Data, 2010 Census, Master Address File, Info USA, NTIA Broadband Data, 2010 IRS 1040 Tax Return Data

Lee and Wang (2003) show that the hazard function, the survival function, and the probability density function used in survival analysis are mathematically equivalent. Therefore, from the hazard function estimates, we were able to derive the density function values or daily Internet response propensities for our sample cases.

To create our model-based mode switch groups we sorted or stratified our sample cases by their respective predicted daily Internet response propensities, resulting in 667 groups ranging in size from one sample case to 2,001 cases with an average size of 90 cases. The graph in Figure 1 illustrates probability density functions for a subset of our mode switch groups, but represents the range of daily propensities across all groups – ranging from those least likely to most likely to respond by Internet. From this graph, we observe that the effect due to day of offering the Internet response option persists across mode switch groups, but is attenuated by the other covariates that define the groups. Additionally, we observe a diminishing return in the predicted yield of Internet responses across groups at day 4 or 5.

Figure 1. Estimated Probability Density Function by Group



4.3 Creating the Group-Based Mode Switch Rules

If we were to base our mode switch rules strictly on the point where we observed a diminishing return in the predicted daily response yield for a given mode switch group, we would switch 28.9 percent of the cases at day 4 and 71.1 percent at day 5. However, we wanted to derive a set of mode switch rules using an optimization procedure such as integer programming (Ignizio and Cavalier, 1994) that also takes into account the group size, Internet nonresponse follow-up cost, the lag time to follow-up with Internet nonrespondents, and nonresponse bias. To account for our follow-up cost, we only considered fixed cost attributed to the mail follow-up such as printing, mailing, postage, and data capture. To account for nonresponse bias in our optimization problem, we used the predicted level of response and the Representativity Indicator (R-indicator) statistic as an indirect measure of the level of contrast between respondents and nonrespondents (cf. Schouten et al. 2009).

Given our available parameters, our goal was to find a solution set of group-based mode switch days that minimize the average group Internet nonresponse follow-up time as defined by

$\bar{y} = (\sum_{g=1}^G n_g \cdot M_g) / \sum_{g=1}^G n_g$ where M_g is the mode switch day and n_g is the sample size for group g , $M_g \in \{0, 1, \dots, 13\}$, constrained by

- (i) $K \sum_{g=1}^G (1 - \sum_{t=1}^{M_g} f_g(t)) \cdot n_g \leq (1 + \%increase) \cdot C_{baseline}$, where K is the cost per Internet nonresponse attributed to mail follow-up, $f_g(t)$ is the predicted response propensity for group g at time t , and $C_{baseline}$ is the baseline mail follow-up cost if we were to wait until day 13 to follow-up with all Internet nonrespondents
- (ii) $M_g \geq t_g^*$ where $f_g(t_g^*)$ is defined as the local maximum for the probability density function for group g at day t_g^* , i.e., $f_g(t_g^*) \geq f_g(t) \forall t \neq t_g^*, t \in \{1, \dots, 13\}$. Note that we relax the constraint, letting $M_g \geq 0$, in the case where the local maximum $f_g(t_g^*)$ fails to exceed a lower bound cutoff f_{LB_cutoff} , i.e., $f_g(t_g^*) \leq f_{LB_cutoff}$ (e.g., $f_{LB_cutoff} = 0.01$).
- (iii) $\hat{R} \geq \hat{R}_{baseline}$ where $\hat{R}_{baseline}$ is the level of sample representativity we would achieve if we were to wait until day 13 to follow-up with all Internet nonrespondents.

$$\hat{R} = 1 - 2\hat{S}(\bar{\rho}) = 1 - 2 \left[\frac{1}{\hat{N}-1} \cdot \left(\sum_{g=1}^G w_g (\hat{\rho}_g - \bar{\rho})^2 \right) \right]^{\frac{1}{2}}$$

$$\text{where } \hat{\rho}_g = \sum_{t=1}^{M_g} f_g(t), \bar{\rho} = \left(\sum_{g=1}^G w_g \sum_{t=1}^{M_g} f_g(t) \right) / \sum_{g=1}^G w_g, w_g = \sum_{i=1}^{n_g} w_i,$$

$$\hat{N} = \sum_{g=1}^G \sum_{i=1}^{n_g} w_i, \text{ and } w_i \text{ is the sample design weight.}$$

4.4 Mode Switch Assignments

Given the framework for assigning mode switch strategies, we now provide an example of the resulting solution that minimizes our objective function while meeting our constraints on cost and error. Table 4 shows an example scenario where we accept a controlled predicted cost increase of no larger than 2 percent and an R-indicator value that is equal to or greater than our baseline R-indicator of 0.79. In addition, we specify a response propensity cut-off that preserves the period of maximum response only for groups with a maximum daily response propensity of greater or equal to 0.01. For groups below this cut-off, we switch them prior to the start of data collection. Under these assumptions, the solution that minimizes our objective function while meeting our constraints requires that we switch 9.3 percent of our sample cases at day zero or in other words not offer the Internet response option to them. We switch a substantial portion of our cases, 29.8 percent, on day 10. Other notable days include switching 17.6 percent on day 11, 17.3 percent on day 9, 15.3 percent on day 8, and 8.9 percent on day 13. For the remaining balance of our cases, we switch on days 5, 7, and 12. Based on this allocation of mode switch days, we have an average time to nonresponse follow-up of 9.1 days compared to 13 days without mode switching.

Table 4. Example Mode Switch Assignment

Mode Switch Day	Group Count	Sample Unit Count	Percent
0	51	5,584	9.3
5	1	2	0.0
7	3	10	0.0
8	47	9,159	15.3
9	124	10,365	17.3
10	272	17,884	29.8
11	100	10,563	17.6
12	18	1,088	1.8
13	51	5,309	8.9
Total	667	59,964	100.0
Average = 9.1 days			
2% cost increase, max predicted propensity cut off = 0.01, R-indicator = 0.79, Internet Response Rate = 13.9% (on day 13)			

Sources: April 2011 ACS Internet Test Data, 2010 Census, Master Address File, Info USA, NTIA Broadband Data, 2010 IRS 1040 Tax Return Data


4.5 Adaptive vs. Non-Adaptive Design Approach

The previous example showed the effect of mode switching on timeliness for one set of assumptions, but we are interested in its effect for a range of assumptions. The graph in Figure 2 shows the impact of mode switching on the average reduction in the group-based length of time to Internet nonresponse follow-up as we change our assumptions. Our assumption changes entail increasing the predicted nonresponse follow-up cost in a controlled manner (0 to 15 percent) and increasing the lower bound for relaxing the constraint that preserves the data collection period where respondents for a given group are more likely to respond (0 to 0.025). Note that across our assumptions, we control the R-indicator value such that we maintain our baseline value. The general trend across the range of lower bound

assumptions appears as a logarithmic relationship between our controlled cost increases and number of days saved – a rapid rise in benefit as we accept a higher cost, followed by a diminishing return.

Reviewing the range of lower bound cut-offs for our group maximum daily propensities, we observe that as we increase the lower bound values for when we relax the constraint on preserving periods of maximum response, our timeliness improves substantially without an increase in the predicted cost. In other words, we achieve higher cost-neutral gains in our timeliness by switching more groups not likely to use the Internet response option at day zero. Groups with little to no likelihood of responding via the Internet have a cost-neutral impact on cost and can provide substantial contribution to improving timeliness. Note that we observe a diminishing return from raising this cutoff at about the 0.025 level.


Figure 2. Average Reduction in Time to Internet Nonresponse Follow-Up by Cost Increase



The next step in our study was to simulate the mode switch assignments for the range of assumptions previously outlined and observe the outcomes in Internet response. Note that our simulation did not account for nonresponse cases that may respond via the Internet after we switch them from Internet to mail. The graph in Figure 3 plots the relationship between the difference in the adaptive versus non-adaptive Internet response rates and the average reduction in time to nonresponse follow-up. We

observe that as we control for a more accelerated timing for Internet nonresponse follow-up, Internet response declines under the adaptive design due to the shortened window of opportunity for respondents to respond via the Internet. Additionally, we observe that we achieve gains in timeliness without compromising the response rate by increasing the lower bound cutoff for preserving periods of maximum response. These gains diminish at the 0.025 level.

Figure 3. Average Reduction in Time to Internet Nonresponse Follow-Up by Response Rate Outcome



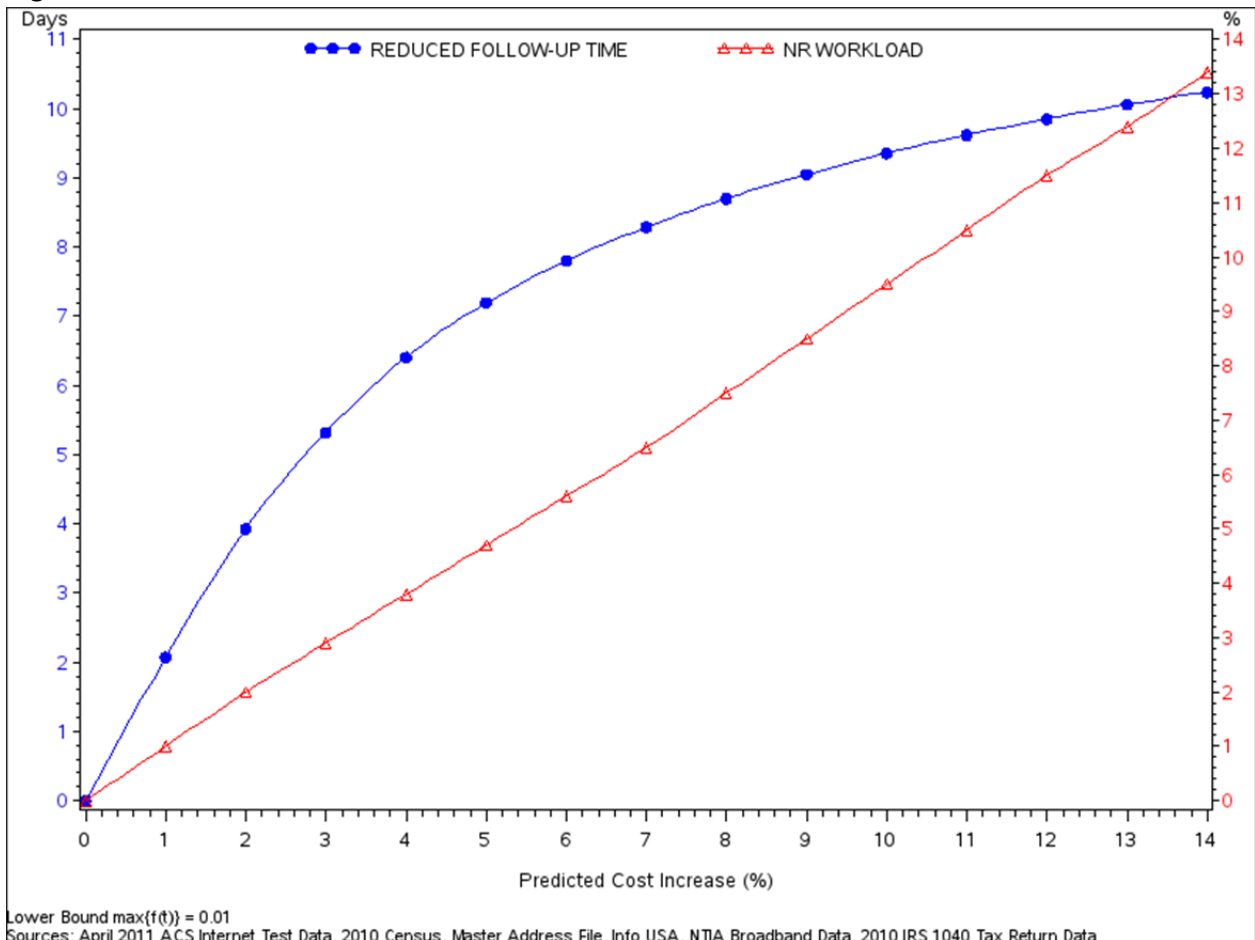
4.6 Applying the Maximization Principle

Previously, we illustrated the benefit in timeliness due to mode switching for a range of control values on predicted cost and response. These results demonstrated the nature of the relationship of improvements in timeliness and controlled increases in the nonresponse workload, however we do not have a clear understanding of the net benefit we achieve through mode switching, that is, the benefit we achieve after we subtract the cost of increasing the nonresponse workload. More importantly, of the possible cost and workload scenarios, which scenario would give us the highest return on our investment? Borrowing from the field of economics' principle of maximization, this process involves identifying the output level and price where an organization realizes a maximum profit or maximum net benefit. By definition, the maximum profit occurs at the point where the marginal revenue and marginal

cost equal each other (Samuelson and Marks, 2012). Applying this principal, our problem is analogous in that our acceptance of additional increases in the expected nonresponse workload represents our 'output' level or quantity produced. Furthermore, through simulating the mode switches, the actual decline in Internet responses represents our actual cost of accepting an increase in the expected nonresponse follow-up workload and the improvement in timeliness of the nonresponse follow-up represents our revenue or total benefit. Our only limitation in using this principal is that we do not have comparable valuations of the 'benefit' and 'cost,' for example dollar amounts. Therefore, for illustration purposes of how we might use this principle, we assume that a change in one day saved in our time to nonresponse follow-up is equivalent to a 1.3 percentage point change in our Internet nonresponse workload.

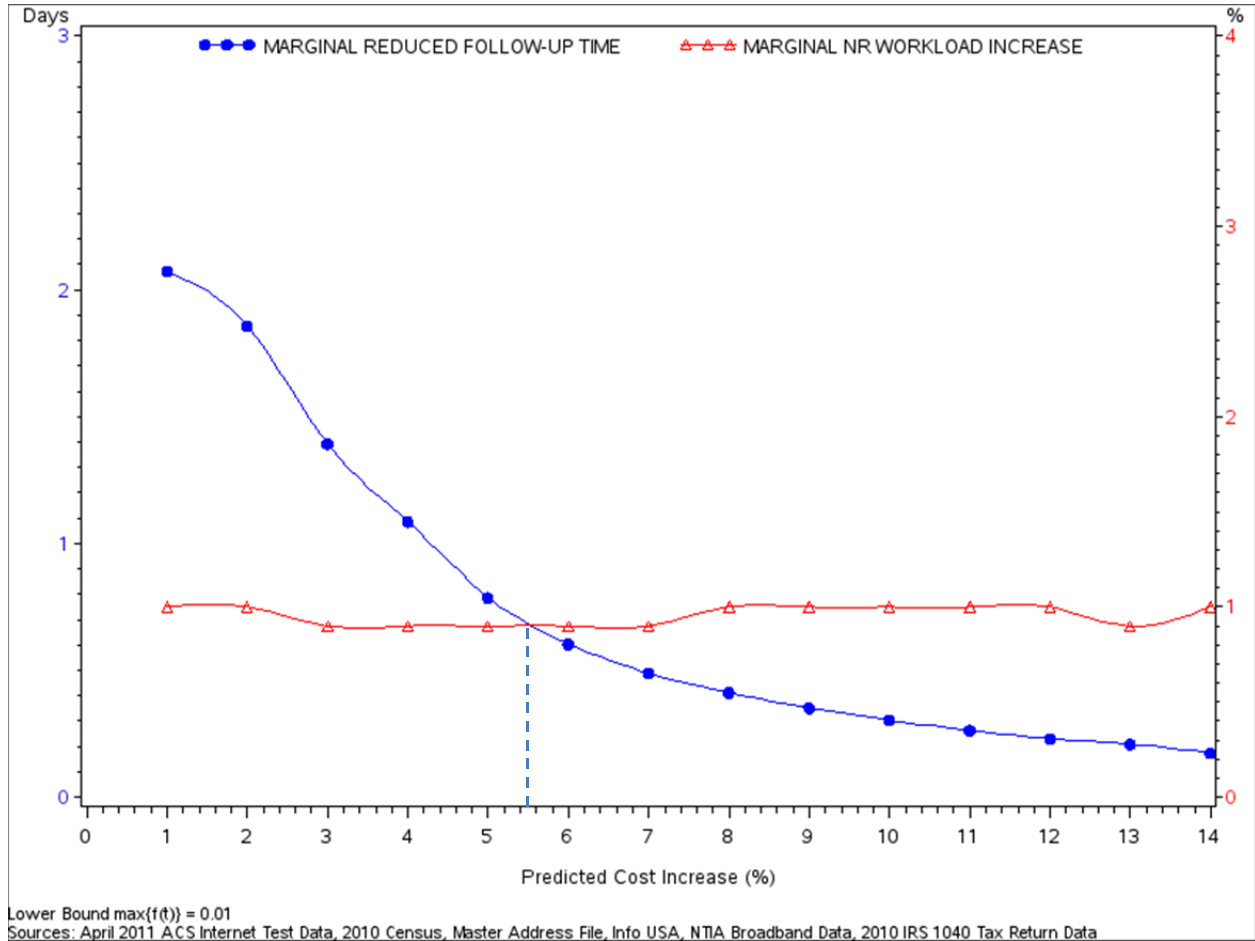
Figure 4 shows the total benefit curve in terms of days saved in time to nonresponse follow-up and the total cost curve in terms of the percent increase in the nonresponse workload. We observe that the distance between the benefit and cost curve (the net benefit) increases and then decreases as we accept a higher cost associated with Internet nonresponse follow-up. In other words, we observe an increasing net benefit followed by a diminishing net benefit as we accept higher costs up to a point where our costs begin to exceed the benefit.

Figure 4. Total Benefit and Total Cost



To identify the point where our marginal benefit and marginal cost are equal or the point at which we achieve a maximum net benefit, we need to derive the marginal benefit and cost curves. Figure 5 shows the marginal benefit curve, representing the change in time savings per unit increase in cost, and the marginal cost curve, representing the change in the actual nonresponse workload per unit increase in the predicted cost. We observe a decreasing trend in our marginal benefit curve and a horizontal trend in our marginal cost with an intersection between the two curves corresponding to a 5.45 percent increase in our predicted cost. Controlling our predicted cost to this level, we can expect the largest net benefit. As noted earlier, a requirement for applying this principle is that we be able to express tangible cost and intangible benefits in comparable units of value, preferably dollar amounts. Therefore, we point out that meaningful evaluation of cost-benefit trade-offs for our application as well as others involving intangible benefits of data quality, timeliness, sampling variance, etc. will require a valuation step for a successful adaptive design process.

Figure 5. Marginal Benefit and Marginal Cost



4.7 Validating the Adaptive Design

The final step in our study was to use the November 2011 Internet Test sample to validate the adaptive design approach we developed using the April 2011 Internet Test sample. This allowed us to gauge the portability of our adaptive design framework to other samples. To perform this validation, we followed the steps previously outlined in forming mode switch groups and assigning mode switch days. We then simulated the mode switches, comparing outcomes in response and reductions in time length to Internet nonresponse follow-up.

Applying the adaptive design we developed using the April Internet Test data, we first applied the model we developed to predict the daily Internet propensities to the November sample cases. Using the household-level regression parameters calculated from the April data, we inputted the covariate values associated with the November cases into our model to calculate the daily response propensities. Using the resulting propensities, we created the mode switch groups resulting in 631 groups (compared to 667 groups created from the April test sample). Applying the same integer programming method described previously, we assigned the mode switch days to the given mode switch groups such that we minimized the time to nonresponse follow-up while controlling for cost and error.

We first compared the Internet response rate outcomes for our range of assumptions. This resulted in 96 comparisons, subtracting the November response rate from the April response rate. For each comparison, we performed a two-sided t-test using a Bonferroni correction to control the family-wise error rate. Reviewing our test comparisons, 26 of our comparisons show evidence that differences exist between the April and November data at a significance level of 0.10, with an average difference of 0.7 and ranging from 0.5 to 0.9. Given these results, our assessment is that our model-based adaptive design appears to have produced response outcomes that closely resemble our results from the April sample.

Next, we compared the improvements in timeliness (in days) between the April and November samples. Note that our measure of reduced time to nonresponse follow-up is a non-inferential statistic. Therefore, any measureable difference between the April and November is the actual difference not an estimate. Based on our comparisons, the average difference between the April and November test samples in days saved in time to nonresponse follow-up is -0.028. In addition, the differences in days saved ranged from -0.044 to 0.002. These results again demonstrate that our design performed well in replicating the improvements in timeliness observed with the April sample.

5. SUMMARY

Using the ACS Internet Test sample, we developed and applied an adaptive design process for the initial phase of data collection in the ACS using administrative records. Augmenting the sample frame data with administrative record data, we successfully developed a discrete-time logistic model to predict household-level daily Internet propensities. Creating mode switch groups based on these propensities, we used integer programming to assign mode switch strategies in an optimal manner. These group-

based assignments minimized the length of time to follow-up with Internet nonrespondents while controlling cost and error. These cost and error controls included controlling increases in predicted nonresponse follow-up cost, maintaining a baseline level of sample representativity, and preserving periods of maximum response with the exception of groups not likely to respond via the Internet. We then successfully simulated switching Internet nonrespondents to mail based on their assigned group-based mode switch strategy. Performing this adaptive design process for a range of assumed values for our controls, we reviewed the outcomes in terms of benefit in improved timeliness in the data collection process and the impact on actual response in the Internet mode of data collection. In addition, we demonstrated the use of the maximization principle for identifying the specific percent increase in the predicted cost that leads to the maximum net benefit in improved timeliness in Internet nonresponse follow-up. However, lacking a comparable valuation measure between our benefit in timeliness and nonresponse cost, we were not able to explore this area of research in depth.

In addition to developing and simulating an adaptive design process under the April test data, we validated our approach using an independent sample, the November Internet Test data. Comparing the outcomes in timeliness and response to the April results, we concluded that the results were almost identical. This provides support for the portability of our process to other samples.

Due to our limited scope, our discussion of improvements in timeliness centered around controlled increases in cost. To address our concerns about the cost of improved timeliness, we point out cost-neutral gains in timeliness can be achieved by switching groups not likely to respond via the Internet prior to data collection. In addition, switching groups preferring a paper questionnaire sooner may reduce the burden or frustration associated with the wait time for receiving the mailed questionnaire. Furthermore, the time saved may benefit later modes in targeting difficult respondents. However, complicating the workflow may create difficulties in workload management. Future advancements in our mode switch process could offer respondents longer periods to respond in their preferred 'cheap' mode of response, thus maximizing response in the cheaper modes of data collection.

Given the exploratory nature of our research, we focused on a simple framework that may serve as a foundation for more complex designs. For example, we could include time varying covariates from paradata sources in the model. In addition, we could alter the constraints in the optimization step to control on other indicators of cost and error as well as change the objective function. In addition, we could include other survey design features affecting cost and error.

References

- Akaike, Hirotugu (1974), "A new look at the statistical model identification," *IEEE Transactions on Automatic Control* 19 (6): 716–723.
- Allison, P.D., (2010), *Survival Analysis Using SAS: A Practical Guide, Second Edition*. Cary, NC, SAS Institute, Inc.
- Bates, N. and Mulry, M. H. (2008), "Building a Segmentation Model to Target the 2010 Census Communications Campaign," 2008 JSM Proceedings of the Section on Survey Research Methods – AAPOR paper, pp 4065-4071.
- Farber, J.E. and Leggieri, C.A. (2002), "Building and Validating the Statistical Administrative Records System 1999," Administrative Records Research Memorandum Series, Washington, DC: Bureau of the Census.
- Groves, R. (2011), "The Future of Producing Social and Economic Statistical Information," Director's Blog, <http://directorsblog.blogs.census.gov/2011/09/>, U.S. Census Bureau.
- Groves, R. and Heeringa, S. (2006), "Responsive Design for Household Surveys: Tools for Actively Controlling Survey Errors and Costs," *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 169 (3): 439-457.
- Hosmer, D.W. and Lemeshow, S. (2000), *Applied Logistic Regression* (2nd ed.), New York, John Wiley & Sons.
- Ignizio, J. P. and Cavalier, T. M. (1994), *Linear Programming*. Englewood Cliffs, NJ, Prentice Hall.
- Judson, D. (2000), "The Statistical Administrative Records System: System Design, Successes, and Challenges," NISS/Telcordia Data Quality Workshop, Trenton, NJ.
- Keathley, D. (2010), "Specifications for Selecting the 2011 American Community Survey and Puerto Rico Community Survey Internet Test Samples," DSSD 2011 American Community Survey Memorandum Series #ACS11-S-9.
- Lee, E. and Wang, J. (2003), *Statistical Methods for Survival Data Analysis* (3rd ed.) Hoboken, NJ, John Wiley & Sons.
- Matthews, B., Davis, M.C., Tancreto, J., Zelenak, M.F., and Ruitter, M. (2012), "2011 American Community Survey Internet Tests: Results from Second Test in November 2011," 2012 American Community Survey Research and Evaluation Report Memorandum Series #ACS12-RER-21.
- Rubin, D. B. (1987), *Multiple Imputation for Nonresponse in Surveys*, New York, John Wiley & Sons.

Samuelson, W. and Marks, S. (2012), *Managerial Economics* (7th ed.), Hoboken, NJ, John Wiley & Sons.

Schafer, J. L. (1997), *Analysis of Incomplete Multivariate Data*, New York: Chapman and Hall.

Schouten, B., Calinescu, M. and Luiten, M. (2011), "Optimizing Quality of Response through Adaptive Survey Designs," Discussion Paper, Statistics Netherlands.

Schouten, B., Cobben, F. and Bethlehem, J. (2009), "Indicators for the Representativeness of Survey Response," *Survey Methodology* 35 (1): 101-113.

Schwarz, Gideon E. (1978). "Estimating the dimension of a model". *Annals of Statistics* 6 (2): 461–464.

Tancreto, J., Zelenak, M. F., Davis, M.C., Ruiter, M., and Matthews, B. (2012), "2011 American Community Survey Internet Tests: Results from First Test in April 2011," 2012 American Community Survey Research and Evaluation Report Memorandum Series #ACS12-RER-13.

Wagner, J. (2008), "Adaptive Survey Design to Reduce Nonresponse Bias", PhD thesis, University of Michigan.