Multiobjective Optimal Allocation for Hard to Reach Populations

Benjamin Phillips, Stanislav Kolenikov, and Chintan Turakhia, Abt SRBI

Introduction

We outline the development of the sampling scheme for a forthcoming national survey of American Jews. The survey will receive considerable scrutiny as the last major national surveys of this population were conducted in the early 2000s (Kotler-Berkowitz et al. 2004; Mayer, Kosmin, and Keysar 2002). These surveys occasioned considerable wrangling regarding methodology and the accuracy of estimates. To this day, there is little agreement about the size or characteristics of this population, creating an environment in which new surveys will receive heavy scrutiny.

American Jews are challenging population to survey because of their low incidence (c. 1.9% of the adult population) and the absence of official statistics on population size and characteristics. Although American Jews are a hard to reach population due to their low incidence, some characteristics of American Jews make it somewhat easier to survey them. The population is geographically concentrated (see, e.g., Figure 2) and some information is available from sample surveys and nonsurvey auxiliary data.

Our approach to sample development involved a wide variety of techniques from different corners of applied statistics, mathematics, and quantitative social sciences that are brought to bear on the interconnected aspects of this problem, rather a single dramatic advancement in any single area. This paper addresses the development of the database used to support estimates, the small area estimate approach used to develop estimates of Jewish incidence for stratification, the development of design effects estimates encompassing the effects of cell phone/landline allocation, and the approach used to develop optimal allocation for the multiple objectives of the study, described below.

Design Objectives and Tradeoffs

The study has four substantive objectives, presented in descending order of importance:

1. Estimate characteristics of the Jewish population (minimum *n*=2,000).
2. Estimate characteristics of the Orthodox Jewish population (minimum *n*=200).
3. Estimate the size of the adult Jewish population.
4. Estimate characteristics of the “Russian Jewish” population (minimum *n*=200). This is defined as consisting of migrants from the former Soviet Union and their descendants.

These objectives involved significant trade-offs, as we shall show later. Estimation of Jewish population characteristics ideally would involve drawing an *epsem* sample of Jews. Estimates of the Orthodox population would be best served by a much more tightly focused sample, given the greater geographic concentration of the Orthodox population, particularly in New York City and surrounding areas. Estimates of the size of the Jewish population would be most accurate with an allocation focused on low incidence areas, due to the high degree of variance therein. Finally, a sample of Russian Jews, like that of the Orthodox, would benefit from a more concentrated sample.

The central component of our design efforts is an Excel file that integrates information from a wide variety of sources. A schematic view of the layout, contents, and functions performed by the Excel file is shown in Figure 1.

Figure 1. Flow of information in design optimization

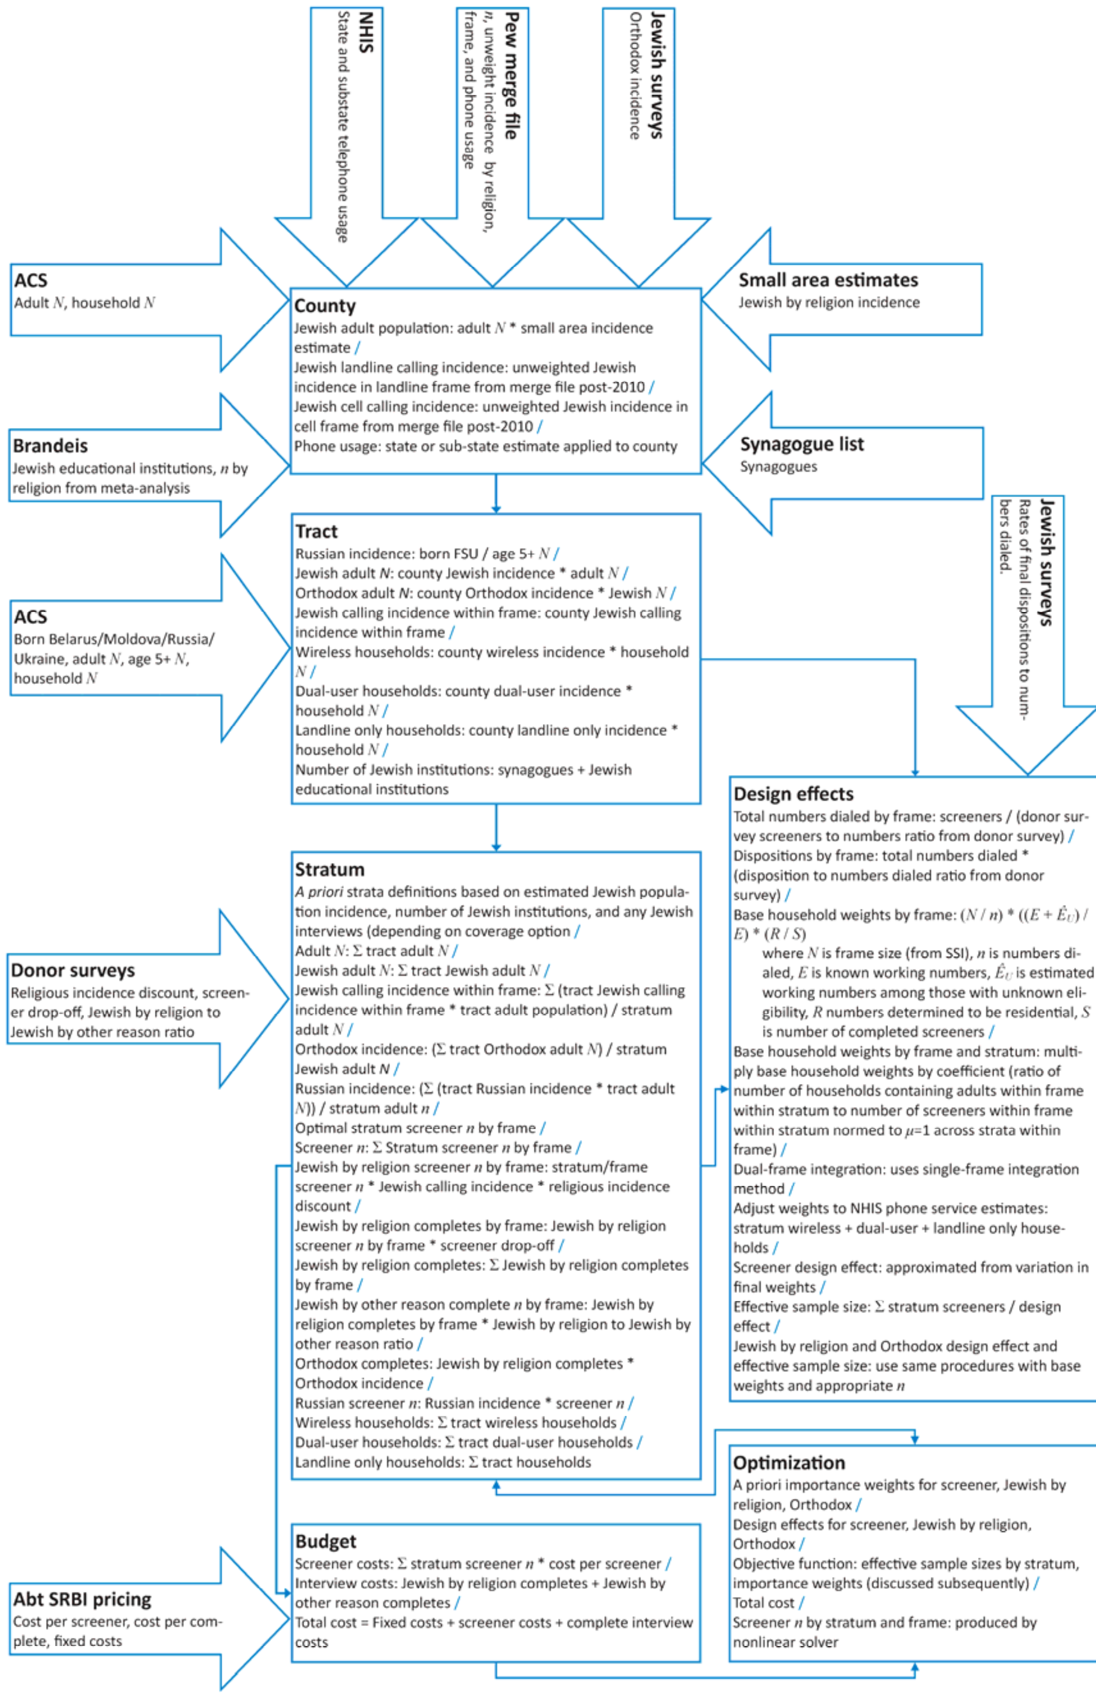

Small Area Estimates

To achieve a highly flexible design for a tightly geographically clustered population like the Jewish one, accurate estimates at a sufficiently small level of geographic aggregation that would allow sample targeting are required. No such estimates existed, and we had to compute our own estimates. For practical reasons of data availability, we focused at the estimates at the county level. The following data sources were used:

1. Pew merge file: a public use file of the completed interviews from 58 surveys fielded from May 2000 through March 2012 (*n*=302,789).
2. ICPSR (2008) county data, combined from a variety of sources. The full data set contains 470 variables for 3,097 counties and county equivalents for the time period from 2000 to 2007, roughly matching the first half of the Pew merge file. We used the following variables: gender by age composition (6 categories), race composition (4 categories), internal and international migration, housing stress, low education, low employment, non-metro recreation, retirement destination, urbanicity (9 categories), economic activity (6 categories), income level and structure (5 sources).
3. List of Jewish educational organizations from JData.com (a project of the Cohen Center for Modern Jewish Studies at Brandeis University), with addresses geocoded to the county level by Abt SRBI.
4. A commercially acquired list of synagogues, likewise geocoded to the county level.
5. A commercially acquired incidence of Jewish names by county, based on an ethnic surname algorithm of a sample provider.

The modern approach to estimates for small geographic areas involves use of statistical models to predict the variable of interest (such as the incidence of a rare characteristic, like Jewish religion/origin). Model-based small area estimation (SAE) proceeds in the following steps.

1. An appropriate regression model (linear for continuous response, logistic for binary response, Poisson for counts, etc.) is first formulated for the response of interest. It would use variables available for all sampled areas and all sampled units.

2. This model is fitted to the existing survey data, with sampling weights if available and necessary.

3. Predictions from the model (called *synthetic estimators*) are obtained.

4. If the area level data is available from the survey data, *direct survey estimators* (weighted means, rates, proportions) are calculated, along with their estimated variances.

5. The synthetic estimators are combined with the direct estimators so as to minimize the mean squared error of the resulting *composite estimator*.

The estimates of the population incidence of Jewish population have been obtained from a small area model that combined county-level explanatory variables, county-level random effects, and individual level outcomes of endorsing Jewish-by-religion response in the religious affiliation questions in earlier studies conducted by Pew Forum for Religion and Public Life. We used a Fay-Herriott-type area-level mixed model with a binary response (Rao 2003):

$$\ln \frac{p_{ij}}{1-p_{ij}} = x_i' \beta + \nu_i \tag{1}$$

where p_{ij} is the probability of the jth respondent in county i being Jewish, x_i are the county-level data, β are regression coefficients, and v_i are area-specific random effects. The evidence of Jewish presence variables were used in the form of the square root of the number of synagogues and educational organizations per 10,000 population, separately for Orthodox and non-Orthodox denominations, for a total of four variables. This transformation was chosen to ensure approximately linearity of the relation between this transformed variable and the empirical best predictions (EBPs) of Jewish incidence obtained at earlier stages.

Small area model estimation was performed using a combination of the `lme4` R package (Bates et. al. 2012, Pinheiro and Bates 2000) with Laplace approximation (one integration point) to quickly produce initial estimates and Stata's `gllamm` package (Rabe-Hesketh and Skrondal 2004) to produce the final estimates that account for the complex survey weights in Pew merge file. Effective sample size weighting (Pfeffermann et. al. 1998) has been used within clusters at level 1; unit weights were used at level 2. Composite estimates (EBPs, Jiang and Lahiri 2001) were computed as

$$\hat{h}_i = \exp(x_i'\hat{\beta}) \frac{\mathbf{E}\exp\{(t_i\{y\}+1)\,\hat{\sigma}\xi-(t_i\{1\}+1)\ln[1+\exp(x_i'\hat{\beta}+\hat{\sigma}\xi)]\}}{\mathbf{E}\exp\{t_i\{y\}\,\hat{\sigma}\xi-t_i\{1\}\ln[1+\exp(x_i'\hat{\beta}+\hat{\sigma}\xi)]\}} \quad (2)$$

where $t_i\{y\}$ is the weighted sum of the responses, $t_i\{1\}$ is the effective sample size, $\hat{\beta}$ are the regression (1) coefficients, $\hat{\sigma}$ is the standard deviation of the are random effects, and ξ is the standard normal variate with respect to which the expectations are taken. SAE models and the various county-level predictors they utilized are summarized in Table 1. Area level R^2 was computed as one minus the ratio of the estimated random effect variance in a given model to that of the intercept only model. The final full model with all predictors was used to produce the estimates (2). The restricted models were fit originally, and as the new data were coming into the project, the county-level predictors had been expanding.

Table 1. Summary of small area estimation models

	Intercept Only	County Demographics	County Demographics + Jewish names	County Demographics + Synagogues	County Demographics + Synagogues + Names
County demographics		$\chi^2(39) = 891.31$	$\chi^2(39) = 856.21$	$\chi^2(39) = 274.14$	$\chi^2(39) = 324.91$
Jewish names			$\chi^2(1) = 123.46$		$\chi^2(1) = 35.62$
Synagogues + educational institutions				$\chi^2(4) = 172.29$	$\chi^2(4) = 84.27$
Random effect variance	1.251	0.286	0.122	0.100	0.073
Area level R^2	0%	77.1%	90.2%	92.0%	94.1%

Table 2 provides the distribution of the population Jewish incidence. At the highest end of the spectrum, the 9 highest incidence counties are home to 22% of the adult Jewish population and slightly more than 3% of the general adult population. At the low end of the spectrum, the 364 lowest incidence counties are home to only 0.06% of the adult Jewish population and about 1.5% of the total population. The last two columns provide input into the decision on establishing the cut-off in the sample design. E.g., if 0.5% cut-off will be used, about 5% of the adult Jewish population will be excluded from the sampling frame. This, however, would allow us to concentrate the dialing effort in only 69% of the general population. The information is also presented graphically on Figure 3.

The proposed SAE model is effective in explaining the patterns of variation of the population incidence of Jews by religion. Not only does it provide higher quality estimates for sampling design decisions, but may also be able to generate some additional insights, e.g., small counties with unexpectedly high predictions of Jewish incidence may be worth further investigation. The incidence of Jewish names, although heavily undercounting the actual Jewish incidence, was slightly more informative than the presence of synagogues and educational organizations. The resulting EBPs demonstrated the expected behavior of reproducing the direct estimates in large counties (with 1,000+ interviews), on one hand, and the synthetic, model-only estimates in small counties (less than 100 interviews), on the other.

Table 2: Distribution of estimated incidence

Jewish incidence	Total adult population	Estimated Jewish population	# of counties	% general adult population	% Jewish population	Cumulative % adult population	Cumulative % Jewish population
0–0.1%	2,320,154	1,618	239	1.02	0.04	1.02	0.04
0.1–0.2%	17,004,524	2,7296	904	7.49	0.63	8.51	0.67
0.2–0.3%	20,977,007	51,659	746	9.23	1.20	17.74	1.87
0.3–0.4%	17,737,664	62,619	375	7.81	1.45	25.55	3.33
0.4–0.5%	15,625,851	70,040	231	6.88	1.63	32.43	4.95
0.5–0.6%	14,620,282	80,875	153	6.44	1.88	38.86	6.83
0.6–0.7%	9,089,521	58,587	78	4.00	1.36	42.87	8.19
0.7–0.8%	10,546,016	78,602	64	4.64	1.83	47.51	10.01
0.8–0.9%	8,461,281	70,239	47	3.72	1.63	51.23	11.65
0.9–1.0%	7,337,668	69,754	37	3.23	1.62	54.46	13.27
1.0–1.2%	9,836,442	107,070	46	4.33	2.49	58.79	15.75
1.2–1.4%	9,209,725	118,145	22	4.05	2.74	62.85	18.50
1.4–1.6%	6,122,651	91,350	26	2.70	2.12	65.54	20.62
1.6–1.8%	11,043,556	189,652	23	4.86	4.40	70.41	25.02
1.8–2.0%	3,785,621	72,990	12	1.67	1.69	72.07	26.72
2–3%	21,040,429	49,6590	37	9.26	11.53	81.33	38.25
3–4%	12,325,025	416,293	21	5.43	9.67	86.76	47.91
4–5%	10,526,230	473,401	10	4.63	10.99	91.39	58.91
5–7.5%	9,831,013	612,360	14	4.33	14.22	95.72	73.12
7.5–10%	2,786,563	256,187	5	1.23	5.95	96.95	79.07
10+%	6,933,023	901,192	7	3.05	20.93	100.00	100.00

Frames

The design uses a dual-frame landline/cell phone RDD sample. Economic efficiencies are obtained solely through the use of disproportionate sample allocation and exclusion of areas with very low Jewish incidence. There were a number of other frames that could have potentially yielded higher calling incidence, but their use would have been associated with mode effects, coverage issues, and even face validity. We wanted to minimize the potential for methodological criticism of the study, and decided to not use recontact sample from previous RDD surveys, ethnic names or list frame (e.g., membership lists of Jewish organizations) to draw the actual samples. The primary motivation for foregoing strategies that could have increased the study's effective incidence was the contested nature of Jewish population statistics. Use of ethnic names had, in particular, acquired a poor reputation. Moreover, the logistics of combining sub-national Jewish organizational lists would have proved time-consuming and expensive and the potential impact of recontact sample on design effects were of serious concern.

The set up of our design optimization toolkit allowed us to obtain the optimal allocation of the interviews between the landline and the cell phone frames as one of the optimization results. Depending on the tuning parameters and preferences built into the optimization procedure, the percentage of the screener interviews placed into the cell phone frame ranged between mid-30s to mid-40s, which appears reasonable and in line with practice of public opinion research studies circa 2012.

Figure 2. Map of estimated Jewish population incidence

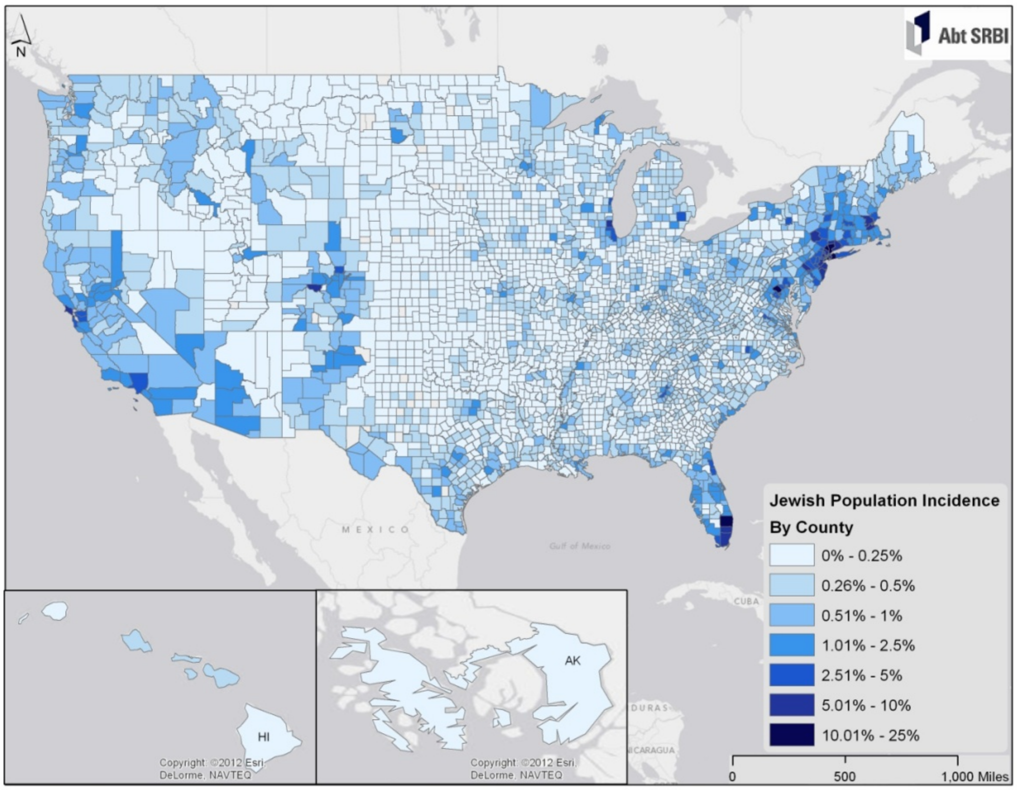

Coverage

As close to half of U.S. counties have little or no evidence of Jewish life, some level of under-coverage by design was a foregone conclusion given low Jewish incidence and the decision to exclude additional sampling frames that would have increased effective incidence (see Fig. 3). The operative question was what degree of under-coverage would be used. Had we high quality information on the effect of coverage on Jewish and Orthodox population estimates, use of MSE in the objective function to be optimized would have been desirable. We did not, however, have such information. Instead, we prepared estimates at three coverage levels: 95%, 97.5%, and 99% and considered trade-offs in a qualitative sense. The coverage options impacted strata definitions, as defined below.

Figure 3. Undercoverage of general and Jewish population by Jewish incidence cutoff

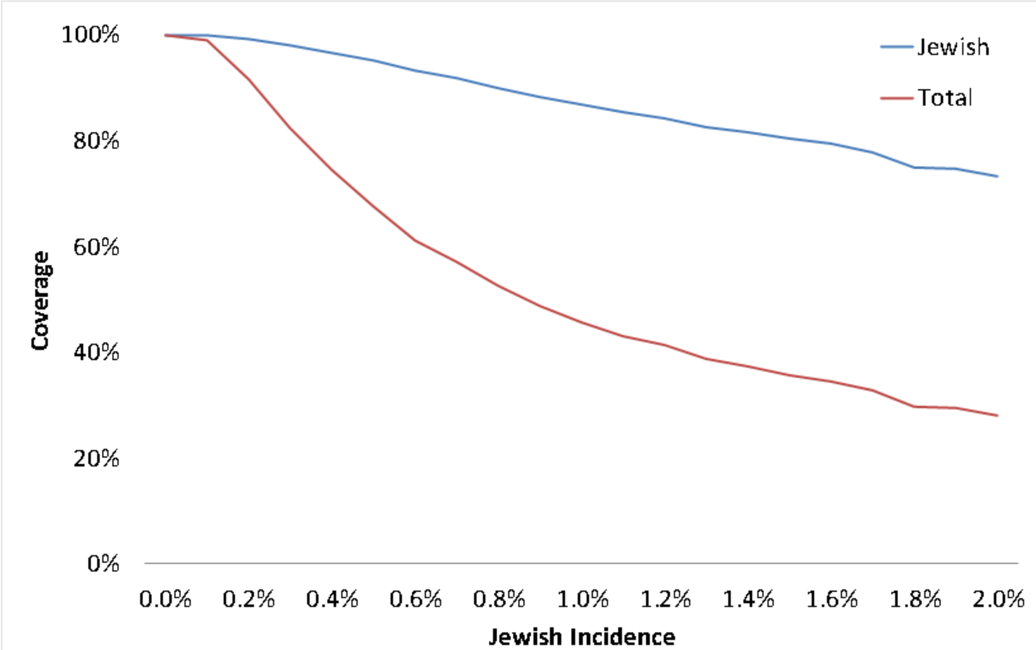

Strata Definitions

A total of five strata were originally defined by Jewish population incidence. As design progressed, additional strata were added for Orthodox and Russian Jews to provide sufficient flexibility in targeting these populations to meet design goals. Strata were defined as follows, where lower numbered strata have priority over higher numbered strata:

1. *Russian*. Born in former Soviet Union incidence of 10% or greater and tract in Orthodox, very high density, high density, medium density, low density, or very low density county.
2. *Orthodox*. Orthodox incidence of 35% or greater in Jewish population of county and county Jewish incidence of 5% or more.
3. *Very high density*. County Jewish incidence of 10% or greater.
4. *High density*. County Jewish incidence of 5% or greater.
5. *Medium density*. County Jewish incidence of 3% or greater.
6. *Low density*. County Jewish incidence of 1.5% or greater.
7. *Very low density*. Definitions varied by coverage option:
   - 95% coverage: County Jewish incidence of .5% or greater.
   - 97.5% coverage: County Jewish incidence of .35% or greater.
   - 99% coverage: County Jewish incidence of .25% or greater *or* 1+ Jewish by religion interviewee in county *or* 1+ Jewish institution (synagogue or educational institution) in county.

Optimization and Parameters

We calculated optimal allocations for different scenarios using Microsoft Excel's nonlinear solver. The Excel solver had the advantage of being able to use the Excel sheet used to hold other study information and avoided the need to recreate the existing algorithms and tables in an optimization package such as AMPL or MATLAB.

As there were multiple study objectives, it was necessary to combine these in some fashion, with the relative importance of each objective being specified through the use of importance weights. The objective function used was the Cobb-Douglas production function, a popular choice for multicriteria optimization in economics for problems of utility maximization, where the sample designer derives greater utility from greater accuracy (lower variances, greater effective sample sizes):

$$U = \frac{\tilde{n}^{\alpha}\tilde{n}^{\beta}\tilde{n}^{\gamma}}{1+\delta\sum_{h=1}^{L}\left(\ln\left(\frac{p_{cell}n_{h,cell}}{n_{h,LL}}\right)\right)^{4}} \tag{3}$$

where $\tilde{n}$ is the effective sample size for the Jewish by religion (a), Orthodox (b), and screeners (c) ; α, β, and γ are the non-negative importance weights of these outcomes ($\alpha + \beta + \gamma = 1$); p_{cell} is the proportion of screeners on the cell frame; $n_{h,LL}$ and $n_{h,cell}$ are the number of screeners in the landline and cell frames respectively in the hth stratum ($h = 1,2,\ldots,L$); and δ is a parameter that controls the degree to which landline and cell phone allocations in the hth stratum are penalized for diverging from p_{cell}.

The effective sample sizes took into account the additional sampling error associated with a dual-frame sample. These were calculated by simulating the weighting process for a dual-frame sample (see Figure 1, design effects panel for additional details). Final dispositions used to calculate the base weights were based on ratios of the numbers dialed to the various dispositions in a "donor" survey of similar design, in this case the fresh landline and cell phone RDD portions of a Pew survey of Asian Americans conducted by Abt SRBI. With the final dispositions estimated, weighting proceeded normally: base weights were created within each stratum and dual-users from each frame were composited using a single frame method (Lohr 2009). Resulting weights were then adjusted to match NHIS-derived estimates of telephone usage in the sampled areas. The weights were then applied to the screener, Jewish, and Orthodox parts of the sample to yield design effects using the $1 + cv^2$ approximation.

The weights corresponding to one of the α, β or γ optimization parameters equal to 1, and others, equal to 0, give the limits of how large the effective sample size can be for the Jewish by religion, Orthodox Jews, or screener interviews, respectively. Based on our communication with the client to solicit their preferences, we used the importance weights of .67 to estimates of Jewish population characteristics, .22 to estimates of Orthodox characteristics, and .11 to estimates of Jewish population size.

The nonlinear solver could vary the number of screeners within each frame X stratum cell. Allocation of screeners to each frame X stratum cell was constrained to be ≥ 10 in order to avoid undesirable solutions that allocated no sample to a given cell. Other constraints on the study were as follows. Study cost was constrained to be less than or equal to the study budget. Orthodox sample size was constrained to be ≥ 200. Screened Russian households were constrained to be ≥ 1,000 to have sufficient cases for analysis. This constraint was specified in numbers of screened households because there were no reliable data on the proportion of Russian households that were Jewish. The proportion of cell phone screeners was constrained to be between 10% and 90% to ensure that both landline and cell phone frames were represented on the design. We also introduced a constraint that the number of Jews interviewed be ≥ 2,500 based on original design specifications and ran models without a constraint.

Each optimization run would take between several seconds and 2–3 minutes when the allocation of counties and tracts to strata was fixed in the Excel file. With a dynamic allocation using Excel lookup and string matching functionality, the computing time increases by about one order of magnitude.

Some Specific Trade-Off Designs

Allocations for 95%, 97.5%, and 99% coverage options with and without a constraint for $n \geq 2{,}500$ Jewish interviews are shown in Tables 3 and 4. The design effects (i.e., the difference between the nominal and effective sample size) are due to widely ranging sampling fractions: higher sampling fractions are used in the more productive higher incidence strata.

One trade-off is coverage vs. sampling error. The higher the estimated coverage of the Jewish population was, the lower the effective sample size, due to the increased cost per completed case. Comparing 95% and 99% options, the cost of greater coverage is a loss of c. 300 in effective sample size with the $n \geq 2{,}500$ constraint and c. 200 without the constraint. Requiring a minimum of $n \geq 2{,}500$ is associated decreases in the amount of sample allocated to the low and very low density strata and increases in sample allocation to the Orthodox, very high, and high incidence strata. Allocation to the cell phone frame also declines. In spite of the greater nominal sample size with the $n \geq 2{,}500$ constraint, effective sample size declines between 148 for the 95% coverage option and 205 for the 99% option.

Discussion

We demonstrated how sample design for a complicated study of an H2R population can be approached. While our multiobjective optimization approach allows to answer some of the design questions, other potentially interesting questions remain unanswered, such as the number of interviews with individuals who consider themselves Jewish by reasons other than religion (e.g., Jewish parents or upbringing, with current religious affiliation other than Judaism), mostly for the reasons that this group has not been studied sufficiently well before, and it is impossible to construct similar models for them.

The study will be in the field in 2013. Large screener sizes would allow constructing reliable direct estimates of Jewish incidence that will be compared to the SAEs for the ultimate test of the small area model.

Acknowledgments

We are grateful to the Pew Forum on Religion & Public Life for their leadership in studying religion in the United States; we particularly wish to acknowledge helpful feedback from Alan Cooperman, Scott Keeter, and Greg Smith as well as the use of data from Pew surveys conducted since 2000. We wish to acknowledge Survey Sampling International for providing county level counts of Jewish ethnic names and the Cohen Center for Modern Jewish Studies at Brandeis University for supplying lists of Jewish day schools, supplementary schools, and early childhood education centers from JData.com and county counts of Jewish by religion adults from their ongoing meta-analysis project. We would also like to acknowledge insights from Courtney Kennedy and Mark Schulman.

Table 3. Sample allocation by coverage options, n=*2,500 constraint*

	95% coverage		**97.5% coverage**		**99% coverage**	
Stratum	Landline	Cell	Landline	Cell	Landline	Cell
Russian	3.5%	1.9%	3.6%	1.9%	3.7%	1.9%
Orthodox	4.1%	7.7%	3.8%	7.6%	3.6%	7.4%
Very high density	10.7%	12.1%	13.6%	10.9%	16.2%	10.2%
High density	17.3%	11.4%	16.2%	10.9%	15.1%	10.3%
Medium density	18.8%	20.0%	17.0%	19.2%	15.4%	18.2%
Low density	20.7%	21.0%	18.4%	19.9%	16.6%	18.7%
Very low density	24.9%	26.0%	27.4%	29.6%	29.4%	33.4%
Total	50,729	29,001	51,561	28,617	51,855	28,505
Percent from frame	63.6%	36.4%	64.3%	35.7%	64.5%	35.5%
Goal	*n*	Effective *n*	*n*	Effective *n*	*n*	Effective *n*
Jews by religion	2,500	1,427	2,500	1,259	2,500	1,118
Orthodox Jews	370	224	367	206	363	188
Screeners	79,730	57,749	80,178	55,491	80,360	54,112
Russian screeners	1,000		1,000		1,000	

Table 4. Sample allocation by coverage options, no sample size constraints

	95% coverage		**97.5% coverage**		**99% coverage**	
Stratum	Landline	Cell	Landline	Cell	Landline	Cell
Russian	3.7%	2.3%	3.8%	2.3%	3.9%	2.3%
Orthodox	2.7%	5.3%	2.5%	5.1%	2.3%	4.9%
Very high density	4.8%	8.2%	4.3%	7.8%	3.9%	7.4%
High density	12.7%	11.4%	11.7%	10.7%	10.9%	10.0%
Medium density	18.4%	20.3%	16.9%	19.2%	15.6%	18.0%
Low density	24.5%	22.9%	22.5%	21.4%	20.8%	19.9%
Very low density	33.3%	29.6%	38.3%	33.6%	42.6%	37.6%
Total	45,650	34,489	45,659	35,151	45,359	35,912
Percent from frame	57.0%	43.0%	56.5.3%	43.5%	55.8%	44.2%
Goal	*n*	Effective *n*	*n*	Effective *n*	*n*	Effective *n*
Jews by religion	2,045	1,575	1,940	1,440	1,845	1,323
Orthodox Jews	293	221	280	206	268	192
Screeners	80,140	66,444	80,811	65,755	81,271	65,585
Russian screeners	1,000		1,000		1,000	

References

Bates, D., M. Maechler, and B. Bolker. (2012). lme4: Linear Mixed-Effects Models Using S4 classes. R package version 0.999999-0 (published 2012-06-23).

Beck, P., J.B. Ukeles, and R. Miller. (2002). The Jewish Community Study of New York 2002: Geographic Profile. UJA-Federation of New York, New York.

Blumberg, S.J., J.V. Luke, N. Ganesh, M. Davern, and M.H. Boudreaux. (2012). Wireless Substitution: State-Level Estimates from the National Health Interview Survey, 2010-2011. National Health Statistics Reports. No. 61. National Center for Health Statistics, Hyattsville, MD.

Kotler-Berkowitz, L., S.M. Cohen, J. Ament, V. Klaff, F. Mott, and D. Peckerman-Neuman. (2004). The National Jewish Population Survey 2000-01: Strength, Challenge and Diversity in the American Jewish Population. United Jewish Communities, New York.

Jiang, J. and, P. Lahiri (2001). Empirical Best Prediction for Small Area Inference with Binary Data. Annals of the Institute of Statistical Mathematics 53(2): 217–243.

Lohr, S. L. (2009). Multiple-Frame Surveys, in D. Pfeffermann and C. R. Rao (eds), Handbook of Statistics: Volume 29A Sample Surveys: Design, Methods and Applications, pp. 71-88. Elsevier

Mayer, E., B. Kosmin, and A. Keysar. 2002. American Jewish Identity Survey: AJIS Report. Center for Jewish Studies, Graduate Center of the City University of New York, New York.

Pfeffermann, D., C. J. Skinner, D. J. Holmes, H. Goldstein, and J. Rasbash (1998). Weighting for Unequal Selection Probabilities in Multilevel Models. Journal of the Royal Statistical Society: Series B (Statistical Methodology) 60(1): 23-40.

Pinheiro, J.C., and D.M. Bates. (2000). Mixed-Effects Models in S and SPLUS. New YorkL Springer.

Rabe-Hesketh, S., A. Skrondal, and A. Pickles (2004). Generalized Multilevel Structural Equation Modeling. *Psychometrika* 69(2): 167-190.

Rao, J. N. K. (2003). *Small Area Estimation*. New York: Wiley.

Tighe, E., L. Saxe, and C. Kadushin. (2012). Cross Survey Analysis to Enumerate Religious Minorities in the United States. Paper presented at the Survey Methods for Hard to Reach Populations Conference, October 31, New Orleans, LA.

United Jewish Communities. (2003). National Jewish Population Survey [MRDF]. Waltham, MA: North American Jewish Data Bank [producer][distributor].