Public Opinion Quarterly, Vol. 81, No. 1, Spring 2017, pp. 144–156

THE LOW RESPONSE SCORE (LRS)
A METRIC TO LOCATE, PREDICT, AND MANAGE HARD-TO-SURVEY POPULATIONS

CHANDRA ERDMAN*
NANCY BATES

Abstract In 2012, the US Census Bureau posed a challenge under the America COMPETES Act, an act designed to improve the competitiveness of the United States by investing in innovation through research and development. The Census Bureau contracted Kaggle.com to host and manage a worldwide competition to develop the best statistical model to predict 2010 Census mail return rates. The Census Bureau provided competitors with a block group-level database consisting of housing, demographic, and socioeconomic variables derived from the 2010 Census, five-year American Community Survey estimates, and 2010 Census operational data. The Census Bureau then challenged teams to use these data (and other publicly available data) to construct the models. One goal of the challenge was to leverage winning models as inputs to a new model-based hard-to-count (HTC) score, a metric to stratify and target geographic areas according to propensity to self-respond in sample surveys and censuses. All contest winners employed data-mining and machine-learning techniques to predict mail-return rates. This made the models relatively hard to interpret (when compared with the Census Bureau's original HTC score) and impossible to directly translate to a new HTC score. Nonetheless, the winning models contained insights toward building a new model-based score using variables from the database. This paper describes the original algorithm-based HTC score, insights gained from the Census Return Rate Challenge, and the model underlying a new HTC score.

Chandra Erdman is a senior solutions consultant at Google in Mountain View, CA, USA. Nancy Bates is a senior researcher for survey methodology at the US Census Bureau, Washington, DC, USA. The authors thank Barbara O'Hare, Julie Coombs, Kathleen Kephart, Travis Pape, Gregg Robinson, and Antonio Bruce for helping develop the updated Planning Database with the LRS; Suzanne McArdle for creating the LRS map; and finally, Robert Groves for sponsoring the Kaggle crowdsource competition. *Address correspondence to Chandra Erdman, Google Inc., 1600 Amphitheatre Parkway, Mountain View, CA 94043, USA; e-mail: chantastic@google.com.

doi:10.1093/poq/nfw040 Advance Access publication December 8, 2016

Introduction

Although declining response rates remain a concern, social scientists still depend on sample surveys and censuses as primary data sources. However, subpopulations of interests are often the same segments least likely to participate (Feskens et al. 2007; Hu, Link, and Mokdad 2010; Smith 2014). In the monograph *Hard-to-Survey Populations*, Tourangeau (2014) classifies such populations according to the survey life cycle. Tourangeau argues that there are important nuances that typify the broader group—a population may be hard to *identify*, hard to *reach*, hard to *persuade*, hard to *interview*, or perhaps all of these. How one goes about classifying a hard-to-count population guides everything from sampling to questionnaire design, contact and recruitment strategy, mode choice, and language of interview (Lyberg et al. 2014; Stoop 2014).

Having an overall "indicator" of hard-to-survey areas is useful for many reasons. First, it can be used to stratify and oversample hard-to-survey areas for purposes of experimentation or targeting. Second, it may be used to staff and plan for nonresponse follow-up. Third, in the context of a large communications campaign (such as the 2000 and 2010 US Census campaigns), it can be used to allocate resources for advertising and community partnership activities.

Both the UK Office for National Statistics (ONS) and the US Census Bureau have developed metrics reflecting overall degree of difficulty by geographic area. For the 2011 Census, the UK ONS produced a hard-to-count index that categorized areas according to predicted propensity to respond. The index was used in allocating field resources and in sampling for the census coverage survey and coverage estimation procedures (Abbott and Compton 2014). The Office for National Statistics (ONS 2009) employed an area-level logistic-regression model to predict response propensities using variables including the proportion of persons claiming Income Support/Jobseeker's Allowance, proportion of young persons, proportion of persons not "White British," relative housing price, and area density. The model predicted nonresponse rates for standard statistical areas, which were then sorted and grouped into five categories. This index (used in conjunction with anecdotal evidence) allowed ONS to identify and prioritize hard-to-count groups, including young adults, ethnic minorities, low-income households, illegal immigrants, and persons with more than one residence.

For the US Census Bureau, Bruce, Robinson, and Sanders (2001) developed a summary score identifying areas that are difficult to enumerate—the "hard-to-count" or HTC score. Based on ethnographic research regarding barriers to enumeration (De la Puente 1993), the authors selected twelve variables for inclusion in the score. The variables reflect reasons people are missed in censuses and included both housing variables (e.g., percentage of vacant houses, percentage of housing units without a phone, and percentage of multi-unit structures) and socio-demographic/economic indicators (e.g., the percentage of people below poverty, percentage of linguistically isolated households, and percentage of renter households). To compute the HTC score, each of the twelve variables was sorted individually across tracts[1] from high to low, with high values indicating greater potential for difficulty. An integer score from 0 to 11 was then assigned to each tract, for each variable, according to where the tract fell along the 50th–100th percentiles of the variable. For example, tracts that had below-median proportions of renter households were assigned a score of 0 for this variable; tracts with proportions of renter households between the median and 55th percentile were assigned a score of 1, and so on. Finally, the integer scores were summed across the twelve variables for a final score ranging from 0 to 132.

This HTC score has been useful both in planning for the 2010 Census and in managing daily operations of many national surveys conducted by the US Census Bureau. However, the original score has a few limitations. First, the score is only available at the tract level, and there can be a great deal of variation in the characteristics of neighborhoods within tracts. Second, by design, the score only included variables that are negatively associated with survey response rates, and excluded variables describing race and ethnicity. Third, each of the twelve HTC variables were arbitrarily given equal weight in the score.

To overcome these limitations, we describe a new, publicly available model-based hard-to-count score developed by the US Census Bureau and based in part on results from a crowdsource challenge sponsored by the agency.

The US Census Bureau Return Rate Challenge

In 2012, the US Census Bureau launched a nationwide prize competition dubbed the Census Return Rate Challenge. The challenge encouraged individuals and teams to compete for prize money[2] in predicting 2010 Census mail-return rates. The objective of the contest was to create a statistical model to accurately predict 2010 Census mail-return rates for small geographic areas, namely, census block groups.[3] Nationwide, 79.3 percent of households that received a 2010 Census mail questionnaire completed it and mailed it back prior to non-response follow-up. However, the level of mail return varied greatly by geography. The Census challenge asked participants to model

1. Census tracts are relatively permanent subdivisions of counties that have between 1,200 and 8,000 residents and an optimum size of 4,000 people.

2. A total of $25,000 in prize money was awarded under the challenge: $14,000, $7,500, and $2,500, respectively, for first-, second-, and third-place model winners, and $1,000 for a data-visualization winner.

3. Census blocks are statistical areas bounded by features such as streets, roads, streams, and railroad tracks. Census block groups are contiguous clusters of blocks that subdivide census tracts.

these variations using variables found in the block group-level 2012 Census Planning Database (PDB).

The PDB is a publicly available database provided at both the tract and block group-levels, which contains a wide range of housing and person characteristics compiled from the 2010 Census and the American Community Survey (ACS). The ACS was implemented by the US Census Bureau as a replacement for the long-form decennial census. It uses a series of monthly samples to produce annually updated estimates for Census tracts and block groups.[4] The survey samples about 3.54 million addresses each year. ACS data are collected by self-response via the Internet and mailed paper questionnaires, and during nonresponse follow-up via computer-assisted telephone interviewing and computer-assisted personal interviewing.

Kaggle.com, a company that hosts predictive modeling competitions, was contracted to manage and judge the competition. The contractor withheld a portion of the PDB for validation, to confirm models with the lowest population-weighted mean squared error (MSE). Submissions were validated and updated daily via the Kaggle website "leader-board." The census challenge proved very popular, with 244 teams and individuals entering the contest.

The competition was the Census Bureau's first use of crowdsourcing as a solution to a technical problem, and the agency hoped to draw upon the winning model as a means to produce a model-based hard-to-count metric. The intent was to develop a new score that was replicable, publicly available, easy to interpret and use in the field, and consistent across various levels of geography, in particular, census block groups and tracts.

The Winning Challenge Model

At the conclusion of the challenge, a software developer was awarded the top monetary prize for his predictions that yielded a population-weighted MSE of 2.60. An examination of the top three challenge models revealed some commonalities. The three top contestants all used ensemble methods (gradient boosting or random forests) that fall under the heading of Machine Learning. These are the same methods used in a large number of statistical modeling competitions, including the million-dollar Netflix Prize competition (Koren 2009). These methods generate a multitude of alternative models for prediction or classification of a given data set, fitting model after model in an effort to minimize a loss function (in this case, population-weighted MSE). This results in an ensemble of weakly predictive models that together yield highly accurate predictions.

4. Detailed descriptions of ACS sampling methods, response, and coverage rates can be found at http://www.census.gov/programs-surveys/acs/methodology.html.

Winning contestants also used hundreds of predictors (the top model contained over 300), including many variables from public datasets external to the PDB. The use of external data meant that the winning model could not be directly applied as a predictive model for the new score. However, examination of the winning model predictors in rank order of relative influence proved enlightening and, in fact, confirmed previous research around census mail-return behavior. For example, the single most influential predictor in the winning model was the percentage of renter households in a block group. Previous research has noted wide variation in census participation between homeowners and renters as far back as the 1990 Census. Research by Word (1997) documents that renters are much less likely to mail back a census questionnaire than homeowners. The percentage of renters in a geographic area is also one of the twelve variables used in the Census Bureau's original HTC score.

Developing an Updated Hard-to-Count Score

We ranked covariates in the winning challenge model by their relative influence (Friedman 2001). Figure 1 displays this statistic for the 50 most influential predictors. As noted earlier and reflected in the figure, the most influential variable is the percentage of renters in a block group. This variable is followed by the percentage of people aged 18–24, and the percentage of households headed by unmarried females. The key features of this plot are: (1) relative influence decreases sharply over the first several variables; (2) smaller drops occur over the next twenty variables; and (3) beyond this point, the relative influence of the predictors is small. For this reason, we pursued a model using the 25 most influential variables to develop a new HTC metric.

Because we wanted a model that is easy to interpret, we examined individual ordinary least squares (OLS) regression models in lieu of the ensembles of regression trees that won the challenge. Moreover, because we wanted model predictors associated with actionable strategies, we restricted our attention to predictors that fit this criterion. For example, we included variables such as the percentage of persons aged 65+ and percentage of households with children because variables like these can guide design decisions such as mode choices, targeted messaging, and the like. Alternatively, we excluded a few variables such as nearest neighboring block group return rates and margins of error for various ACS estimates. These were good predictors of mail-return rates, but cannot be used to guide the development of strategies for increasing response rates. Furthermore, neighboring return rates can only be used to predict return rates in the next census *after* the next census.

Table 1 displays coefficients of the "Top 25 Challenge" model,[5] at the block group and tract levels, estimated for both levels of geography so that the results

5. In the end, we removed the variable "Ages < 5," as all other age groups are accounted for in the model, and this group is very closely related to the more influential variable, "Households with related child < 6." We then included the next most influential variable, median household income.

Figure 1. Relative Influence of the 50 Most Influential Variables in the Winning Challenge Model.

may be included with both versions of the PDB. Because we want to predict areas that are hard to enumerate, we used 100 minus return rate—"non-return rate"—as the dependent variable. Most predictors are highly significant at both levels of geography, and all are significant at at least one level of geography. Seven of the twelve variables from the original HTC score are found in the model.[6] Other variables include length of residence, presence of young children, and married-couple households—variables that describe the "place attachment" construct found in urban sociology literature (Brown, Perkins, and Brown 2003). This construct figures into theories explaining behaviors such as level of civic engagement, voting, and even participation in surveys (Guterbock, Hubbard, and Hoilan 2006). From table 1, we see that (given all other covariates) the presence of renters, vacant units, and persons aged 18–24 in a block group are all positively associated with low response. That is, the greater the presence of these characteristics in a block group, the lower the self-response rates. Alternatively, presence of persons aged 65+, married couples, and persons with a college degree are negatively associated with low response.

6. Original HTC variables that were not included in the model are percentages of households with public assistance income, unemployed persons, crowded units, linguistically isolated households, and households with no phone service.

Table 1. Low-Response Model Summaries at the Block Group and Tract Levels

	Block group n = 217,417			Tract n = 72,763		
Variable	Coef.	Z-value	Sig.	Coef.	Z-value	Sig.
Intercept	10.29	12.49	***	16.61	10.56	***
Renter-occupied units	1.08	50.57	***	0.95	23.51	***
Ages 18–24	0.64	21.53	***	0.47	9.57	***
Female head, no husband	0.58	17.26	***	0.33	5.37	***
Non-Hispanic White	–0.77	–38.76	***	–0.87	–26.17	***
Ages 65+	–1.21	–39.61	***	–1.29	–24.31	***
Related child < 6	0.46	15.82	***	0.08	1.38	
Males	0.09	20.43	***	0.04	4.91	***
Married family households	–0.12	–37.43	***	–0.14	–25.46	***
Ages 25–44	–0.06	–1.74		0.11	2.08	*
Vacant units	1.08	52.74	***	0.91	25.11	***
College graduates	–0.32	–17.33	***	–0.53	–12.62	***
Median household income	0.24	4.62	***	0.34	2.88	**
Ages 45–64	–0.08	–2.54	*	–0.16	–2.69	**
Persons per household	3.44	13.19	***	3.30	6.78	***
Moved in 2005–2009	0.09	7.19	***	0.13	4.38	***
Hispanic	0.41	24.45	***	0.52	18.23	***
Single-unit structures	–0.52	–53.11	***	–0.56	–27.32	***
Population density	–0.40	–41.93	***	–0.46	–29.25	***
Below poverty	0.11	9.95	***	0.26	9.56	***
Different housing unit 1 year ago	–0.12	–11.09	***	–0.35	–12.55	***
Ages 5–17	0.17	4.30	***	0.24	3.24	**
Black	–0.04	–2.69	**	0.01	0.24	
Single-person households	–0.24	–5.19	***	–0.35	–4.26	***
Not high school graduate	–0.06	–4.84	***	–0.19	–6.75	***
Median house value	0.71	25.56	***	0.78	14.69	***

NOTE.—All variables are percentages unless otherwise indicated. Most variables are square- root, log or logit transformed (see the appendix for details). Block group R-squared: 56.10; Tract R-squared: 55.25.

*$p < .05$; **$p < .01$; ***$p < .001$.

With this model, we propose a new census challenge-based score—each geography's score is simply the fitted value from the OLS regression. We refer to this new metric as the low response score, or LRS. To illustrate this new model-based score, figure 2 displays quintiles of the LRS and actual 2010 Census mail-return rates for block groups in the District of Columbia. With few exceptions, hard-to-count areas have the highest Low Response Scores (and darker shading), and the easiest-to-enumerate areas have the lowest scores (and lighter shading).

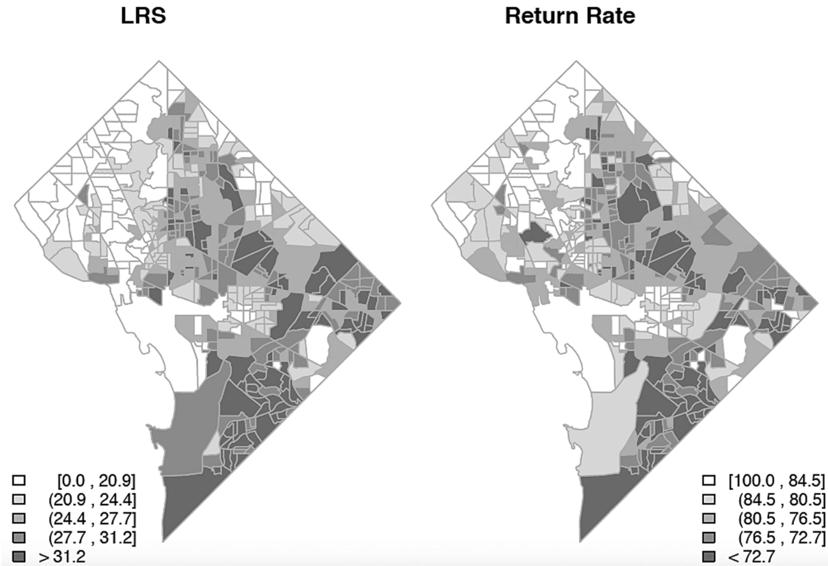

Figure 2. Quintiles of the LRS and 2010 Census Mail Return Rates for Block Groups in the District of Columbia.

Discussion

In this article, we describe the application of results from a crowdsourcing project to produce the best predictive model of 2010 Census mail-return rates. An overall goal of the project was to encourage new approaches to understanding self-response behavior in the decennial census. A secondary goal was to use the winning model as a means for updating the US Census Bureau's HTC score. The HTC score was a metric pioneered by Census Bureau demographers over two decades ago that delineates areas of the country according to difficulty of enumeration.

The winning challenge model had many predictor variables—342 in all. After close examination, we excluded many variables when constructing a new HTC model because they came from sources external to the Census Bureau's public database offered for the challenge, had little actionable meaning regarding hard-to-count populations, or had low predictive power. Still, our final model containing only 25 variables was highly predictive of mail response at the block group level (R-squared = 56.10). Our model relies heavily on the highest rank-order predictors from the winning census challenge model, which includes a majority of the twelve variables used in the original HTC score. The variables that constitute our score are a robust set that can inform many aspects of a census.

Using block groups from three neighborhoods in the District of Columbia as an example, we see how areas with a similarly high LRS can have very different characteristics (figure 3). The LRS is 32.5 for Columbia Heights, 37.4 for Trinidad, and 37.8 for Anacostia—all above the 90th percentile (and correspondingly, the actual mail-return rates for these areas are in the lowest decile). We can use the publicly available PDB[7] to further explore these areas. All three neighborhoods have high proportions of renters (82 percent or higher) but are otherwise quite dissimilar. The Columbia Heights block group comprises 45 percent Hispanic residents, where 33 percent speak a language other than English. Seventy-four percent of the households are multi-unit structures, 52 percent are non-family households, and 50 percent of the householders moved in within the past five years. When conducting the census, this block group could benefit from in-language advertising and Spanish-language forms. Alternatively, the Anacostia block group consists of 98 percent Black residents, of whom 46 percent are below poverty level and 89 percent are single-unit homes. This block group has a low percentage of non-family households (15 percent), and only 21

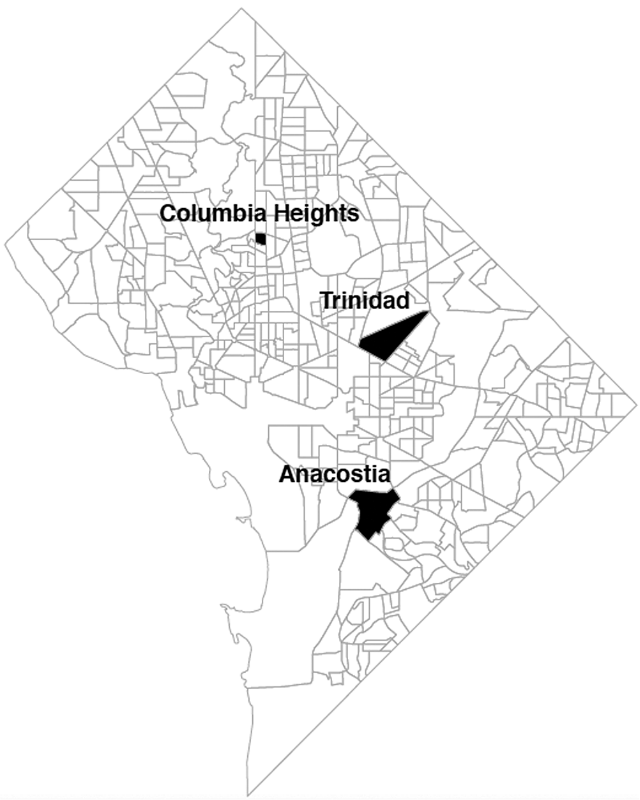

Figure 3. Three HTC Block Groups in the District of Columbia.

7. http://www.census.gov/research/data/planning_database/

percent of householders moved in within the past five years. A successful technique for engaging this area might be community partnerships because the area is characterized by families who have lived in the area for a longer time. Finally, the Trinidad neighborhood skews younger and mobile, with well over one-third (37 percent) aged 18–24 and 59 percent of householders having moved to the address within five years. The area is also characterized by poverty and vacant units (one-third are below poverty, and just over one-quarter of the houses are vacant). The racial/ethnic composition is varied, with 55 percent Black householders and 31 percent White. These characteristics suggest a neighborhood in high transition with in-moving younger households along with the more traditionally hard-to-count. For this area, we would encourage Internet response for younger households while leveraging community leaders to engage and motivate longer-time residents. These differentiated outcomes are particularly useful when trying to plan a census or survey that includes large-scale interviewer hiring, a targeted communication campaign, delivery of in-language forms, and/or development of community partnership programs.

The LRS is now publicly available on both the block group and tract-level PDBs, so users can perform similar exercises for their own neighborhoods and communities (see online appendix for a national map of the LRS). We suggest users not be constrained by the 25 variables that ended up in the final model. Instead, we recommend using the LRS to first pinpoint harder-to-count areas, then using the full set of variables found in the PDB to develop tailored methods designed to boost cooperation.

While we believe the new LRS will be useful to census and survey planners, we note an important caveat to the score—namely, the response metric predicted by the model is based on a single mode of self-response (mail). For the next census in 2020, the Internet is expected to be the majority self-response mode, with great effort made to promote and encourage its use. The survey literature suggests that, in mixed-mode approaches, different segments of the population adopt the Internet response mode to different degrees (Datta, Walsh, and Terrell 2002; Link and Mokdad 2006; Nicolaas et al. 2014).

Internet response behavior can be illustrated with the ACS. When mail was the only self-response option, a population segment coined the "Single Unattached Mobiles" (Bates and Mulry 2011) had below-average 2009 mail-return rates (42.5 percent). This segment skews toward single, young people living in urban multi-units who rent and move frequently. Beginning January 2013, the ACS added Internet as means of self-response. ACS return rates from 2013 that reflect both mail and Internet mode indicate that this segment prefers Internet (64 percent of total self-response was by Internet). Additionally, when both modes are offered, the overall self-response rate was actually higher for this group than four years prior (45.5 percent; see Baumgardner et al. 2014).

Because the dependent variable used to build the LRS does not account for Internet response, our score is not as accurate as it could be and some groups and geographical areas may be characterized as harder to count than they really are. However, in 2013, in addition to adding the Internet as a response option, the ACS also added a question about household Internet subscription. When ACS Internet data become available at low levels of geography, we can revisit our LRS model to adjust both the dependent variable (to include additional modes of self-response) and predictor variables (to include degree of Internet penetration). In the meantime, the LRS is available on the most recent versions of the publicly available tract- and block group-level PDBs and will be updated and appended to future iterations.

Appendix

List of Transformations

Variable	Transformation
Renter-occupied units	Logit
Ages 18–24	Square root
Female head, no husband	Square root
Non-Hispanic White	Logit
Ages 65+	Square root
Related child < 6	Square root
Ages 25–44	Square root
Vacant units	Log
College graduates	Logit
Median household income	Log
Ages 45–64	Square root
Persons per household	Log
Moved in 2005–2009	Square root
Hispanic	Logit
Single-unit structures	Logit
Population density	Log
Below poverty	Square root
Different housing unit 1 year ago	Square root
Ages 5–17	Square root
Black	Logit
Single-person households	Square root
Not high school graduate	Square root
Median house value	Log
Public assistance	Square root
Unemployed (ages 16+)	Logit
Crowded units	Square root
Linguistically isolated households	Square root
No phone service	Square root

NOTE.—IN order to obtain fitted values, a small amount was added to each variable before each log or logit transformation.

Supplementary Materials

Supplementary materials are freely available online at http://poq.oxfordjournals.org/

References

Abbott, Owen, and Garnett Compton. 2014. "Counting and Estimating the Hard-to-Survey Populations in the 2011 Census." In *Hard-to-Survey Populations*, edited by Roger Tourangeau, Brad Edwards, Timothy P. Johnson, Kirk M. Wolter, and Nancy Bates, 58–81. Cambridge, UK: Cambridge University Press.

Bates, Nancy, and Mary H. Mulry. 2011. "Using a Geographic Segmentation to Understand, Predict, and Plan for Census and Survey Mail Nonresponse." *Journal of Official Statistics* 27:601–18.

Baumgardner, Stephanie, Deborah Griffin, Rachel Horwitz, and Elizabeth Nichols. 2014. "Who Benefitted from the Addition of an Internet Response Option in the American Community Survey." Paper presented at the US Bureau of the Census Survey Methods Seminar Series, Washington, DC.

Brown, Barbara, Douglas D. Perkins, and Graham Brown. 2003. "Place Attachment in a Revitalizing Neighborhood: Individual and Block Levels of Analysis." *Journal of Environmental Psychology* 23:259–71.

Bruce, Antonio, J. Gregory Robinson, and Monique V. Sanders. 2001. "Hard-to-Count Scores and Broad Demographic Groups Associated with Patterns of Response Rates in Census 2000." Paper presented at the Joint Statistical Meetings of the American Statistical Association, Atlanta, Georgia, USA.

Datta, Pratim, Ken Walsh, and Dek Terrell. 2002. "The Impact of Demographics on Choice of Survey Modes: Demographics Distinctiveness Between Web-Based and Telephone-Based Survey Respondents." *Communications of the Association for Information Systems* 9:223–40.

De la Puente, Manuel. 1993. "Why Are People Missed or Erroneously Included by the Census: A Summary of Ethnographic Coverage Reports." Paper presented at the Research Conference on Undercounted Ethnic Populations, Richmond, VA, USA.

Feskens, Remco, Joop Hox, Gerty Lensvelt-Mulders, and Hans Schmeets. 2007. "Nonresponse Among Ethnic Minorities: A Multivariate Analysis." *Journal of Official Statistics* 23:387–408.

Friedman, Jerome H. 2001. "Greedy Function Approximation: A Gradient Boosting Machine." *Annals of Statistics* 25:1198–232.

Guterbock, Thomas M., Ryan A. Hubbard, and Laura M. Hoilan. 2006. "Community Attachment as a Predictor of Survey Nonresponse." Unpublished paper.

Hu, S. Sean, Michael W. Link, and Ali H. Mokdad. 2010. "Reaching Linguistically Isolated People: Findings from a Telephone Survey Using Real-Time Interpreters." *Field Methods* 22:39–56.

Koren, Yehuda. 2009. "The BellKor Solution to the Netflix Grand Prize." Available at http://www.netflixprize.com/assets/GrandPrize2009_BPC_BellKor.pdf.

Link, Michael W., and Ali Mokdad. 2006. "Can Web and Mail Survey Modes Improve Participation in an RDD-Based National Health Surveillance?" *Journal of Official Statistics* 22:293–312.

Lyberg, Lars, Mathew Stange, Janet Harkness, Peter Mohler, Beth-Ellen Pennell, and Lilli Japec, 2014. "A Review of Quality Issues Associated with Studying Hard-to-Survey Populations." In *Hard-to-Survey Populations*, edited by Roger Tourangeau, Brad Edwards, Timothy P. Johnson, Kirk M. Wolter, and Nancy Bates, 82–108. Cambridge: Cambridge University Press.

Nicolaas, Gerry, Lisa Calderwood, Peter Lynn, and Caroline Roberts. 2014. "Web Surveys for the General Population: How, Why and When?" Available at http://www.natcenweb.co.uk/genpop-web/ documents/GenPopWeb-final-report.pdf.

Office for National Statistics. 2009. “Predicting Patterns of Household Nonresponse in the 2011 Census.” Available at www.ons.gov.uk/ons/guide-method/census/2011/the-2011-census/processing-the-information/statistical-methodology/predicting-patterns-of-household-non-response-in-the-2011-census.pdf.

Smith, Tom W. 2014. “Hard-to-Survey Populations in Comparative Perspective.” In *Hard-to-Survey Populations*, edited by Roger Tourangeau, Brad Edwards, Timothy P. Johnson, Kirk M. Wolter, and Nancy Bates, 21–36. Cambridge: Cambridge University Press.

Stoop, Ineke. 2014. “Representing the Populations: What General Social Surveys Can Learn from Surveys Among Specific Groups.” In *Hard-to-Survey Populations*, edited by Roger Tourangeau, Brad Edwards, Timothy P. Johnson, Kirk M. Wolter, and Nancy Bates, 225–44. Cambridge, UK: Cambridge University Press.

Tourangeau, Roger. 2014. “Defining Hard-to-Survey Populations.” In *Hard-to-Survey Populations*, edited by Roger Tourangeau, Brad Edwards, Timothy P. Johnson, Kirk M. Wolter, and Nancy Bates. Cambridge: Cambridge University Press, 3–20.

Word, David L. 1997. *“Who Responds/Who Doesn’t? Analyzing Variation in Mail Response Rates During the 1990 Census.”* US Bureau of the Census Population Division Working Paper Series, Washington, DC.