Evaluating the Effect of Within-Household Subsampling on the Precision of Crime Victimization Rates

Vincent G. Iannacchione and Bonnie E. Shook-Sa

July 2013

RTI Press

About the Authors

Vincent G. Iannacchione, MS, is a senior research statistician at RTI International in Washington, DC.

Bonnie E. Shook-Sa, MAS, is a research statistician at RTI specializing in survey sampling and analysis.

Suggested Citation

Evaluating the Effect of Within-Household Subsampling on the Precision of Crime Victimization Rates

Vincent G. Iannacchione and Bonnie E. Shook-Sa

Contents

Abstract

The decision to select a subsample of eligible members of a sampled household is influenced by a number of factors including burden on the household, data quality, cost, and the sampling variance of survey estimates. Design effects quantify the influence of a complex sampling design on the variance of survey estimates. Selecting a subsample of eligible persons within a sampled household can have counteracting impacts on design effects. On one hand, subsampling increases the design effects attributable to unequal weighting. On the other hand, subsampling could reduce the design effects attributable to clustering because the potential intra-household correlation among respondents in the same household may be reduced or eliminated. If the reduction in correlation is greater than the increase caused by unequal weighting, subsampling can achieve the same sampling variance as selecting all eligible household members, with less cost and burden. We present the results of a simulation study that evaluates the design effects associated with subsampling household members on personal victimization rates based on the 2008 National Crime Victimization Survey, which selected all persons 12 and older in a sampled household.

Introduction

In a complex sample design, the estimated sampling variance associated with selecting all members of a household is more accurate than random subsampling, and subsampling two persons per household produces more accurate variance estimates than selecting just one person. With this in mind, a survey researcher might ask, "Are there situations where within-household subsampling can produce nearly the same variance estimates as selecting all household members?" *Design effects* (Kish, 1965) can be used to answer this question by "deconstructing" the sampling variance in a way that identifies the factors that drive the variance when implementing within-household subsampling.

In addition to the size of a sample, the sampling variance of a survey estimate is affected by stratification, clustering, and unequal weighting caused by differential selection and response rates. Design effects quantify how the complex design of a sample affects the variances of survey estimates. Specifically, moving from a sample design that selects all eligible persons per household to a design that subsamples eligible persons within households could have counteracting impacts on the design effects of survey estimates. On one hand, the design effects could increase as a result of unequal selection probabilities associated with subsampling persons from multi-person households. On the other hand, subsampling could reduce the design effects because it reduces the potential correlation among outcomes reported by members of the same household. If the reduction in correlation is greater than the increase caused by unequal weighting, subsampling can achieve the same sampling variance as selecting all eligible household members—with less cost and burden.

We present the results of a simulation study based on the 2008 National Crime Victimization Survey (NCVS) that evaluates how subsampling within NCVS households affects the design effects of personal victimization rates. We begin with a brief description of the current NCVS sampling and weighting methodology. We then describe how we selected the simulation samples that we used to estimate the design effects that may accompany

selecting one or two persons per household. We estimate the number of households that would be needed to equalize the precision of the current "all person" sample design and that of a one-person or two-person design. We conclude by discussing the advantages and disadvantages of within-household subsampling.

Methods

Summary of the NCVS Sampling Design

The NCVS is a survey of the US civilian noninstitutionalized population that focuses on personal and property crimes. Within each selected household, a screening questionnaire is administered to all persons aged 12 and older to determine whether they were victims of personal crimes during the previous 6 months. Personal crimes are those committed against individuals and include rape/sexual assault, robbery, assault, and personal theft. In addition, a *household* screener is administered to a single household member (the household respondent) and is used to report property crimes against the household, such as burglary, theft, and motor vehicle theft (US Department of Justice, 2008). These data are used to estimate yearly victimization rates[1] and changes in victimization rates from year to year.

The NCVS is a panel survey in which each sampled household or household equivalent[2] is interviewed once every 6 months over a 3-year period for a total of seven interviews. Interviews one and five are conducted face-to-face, and the remaining interviews are conducted by telephone. Currently, everyone 12 years of age or older in a sampled household is asked if her or she was victim of a crime that occurred during the previous 6 months.

In response to a recommendation from the National Research Council (2008), the Bureau of Justice Statistics, which administers the NCVS, is considering restricting the sample to one eligible person from

---

[1] Victimization rates are the estimated number of victimizations per 1,000 persons (for personal crimes) or households (for property crimes).

[2] Group quarters are living quarters where residents share common facilities or receive formally authorized care. For the NCVS, group quarters are divided into clusters of four expected persons. These clusters are referred to as *household equivalents*.

each sampled household. This change in the NCVS sample design could have far-reaching and unknown implications for response rates, survey costs, and survey estimates and their associated sampling variances.

From a data quality standpoint, a design of one person per household is preferable to either the current "all persons" design or a design of two persons per household because either may be subject to certain biases for intra-familial crimes such as domestic violence. For example, under the current design a husband who has been interviewed may tell his wife not to report domestic violence. The resulting underreporting could be reduced with a design of one person per household (assuming privacy is maintained in the interview setting) if the wife were selected because she would be the only person interviewed from the household. These potential biases cannot be measured in a simulation study. Therefore, this analysis is restricted to the effects that subsampling eligible persons within households has on the precision of resulting survey estimates.

The NCVS uses a stratified, four-stage sampling design to estimate crime victimization rates for the national civilian noninstitutionalized population age 12 and older (US Census Bureau, 2009). At the first stage, primary sampling units (PSUs) are demographic areas consisting of large metropolitan areas, counties, or groups of adjacent counties. Large PSUs are included in the sample automatically, and each is assigned its own self-representing stratum. The remaining non-self-representing PSUs are combined into strata by grouping PSUs with similar geographic and demographic characteristics.

At the second stage, each selected PSU is divided into segments (clusters of about four households each), and a systematic sample of segments is selected. At the third stage, all households in a sampled segment are selected, and at the fourth stage, all persons aged 12 and older are selected from each sampled household. This type of sampling design enables the selection of a self-weighting probability sample of eligible persons. That is, prior to any weighting adjustments for nonresponse or noncoverage, each eligible person has

the same design weight, which is the inverse of the overall probability of being selected.

Because of the complex sampling design used for the NCVS, the usual sample variance that assumes simple random sampling needs to be multiplied by the design effect to approximate the sample variance associated with the complex design. Kish (1987) proposed a production model of the overall design effect ($DEFF_T$) as the product of two components: $DEFF_C$, which is attributable to clustering and is dependent on the within-cluster sample sizes, and $DEFF_W$, which is attributable to differential sampling rates (or unequal weighting). That is,

$$DEFF_T = DEFF_C \times DEFF_W \,.$$

Gabler, Haeder, and Lahiri (1999) provide a model-based justification for using Kish's formula.

In the four-stage design used for the NCVS, the design effect attributable to clustering[3] can be approximated as:

$$DEFF_C = 1 + (\bar{b}_1 - 1)\rho_1 + (\bar{b}_2 - 1)\rho_2 + (\bar{b}_3 - 1)\rho_3$$

where $\bar{b}_1$ is the average number of sampled persons per PSU, and ρ_1 is the intracluster correlation that measures the homogeneity of the characteristic being measured for persons within the PSUs. Similarly, $\bar{b}_2$ is the average number of sampled persons per segment, and ρ_2 is the intracluster correlation for persons within segments. Finally, $\bar{b}_3$ is the average number of sampled persons per household, and ρ_3 is the intracluster correlation for persons within households.

The design effect attributable to differential sampling rates and weighting adjustments for nonresponse or noncoverage (Kish, 1965) can be expressed as:

$$DEFF_W = \frac{n\Sigma W_i^{\,2}}{(\Sigma W_i)^2}$$

where W_i is the analysis weight assigned to respondent i.

Typically, within-household subsampling would cause $DEFF_C$ to be smaller than that associated with the current NCVS design because the third component ($[b_3 - 1]\rho_3$) would either be reduced with a two-person-per-household design or even eliminated with a one-person-per-household design. Conversely, $DEFF_W$ would be greater because of

---

[3] See Hanson, Hurwitz, and Madow (1953, p. 401) for more details.

the unequal weighting caused by the selection of a subsample from each multi-respondent household. The combined effect on $DEFF_T$ would depend on the relative decrease of $DEFF_C$ to the increase in $DEFF_W$. By simulating the selection of a subsample of respondents from each multi-respondent household, we can estimate the combined effects of $DEFF_C$ and $DEFF_W$ on $DEFF_T$.

Selection of the Simulation Samples

The 2008 NCVS public-use database contains survey data for 88,700 respondents[4] in 48,111 unique households in 2008 (US Department of Justice, 2008). As one can determine from Table 1, some 69,007 respondents (78 percent) were in 28,418 households with two or more respondents. The remaining 19,693 respondents (22 percent) were in single-respondent households. Selecting one respondent from each of the 28,418 multi-person households would reduce the total number of respondents to 48,111, which is 54 percent of the 88,700 NCVS respondents in 2008. Selecting two respondents from each of the 8,211 households with three or more respondents would reduce the total number of respondents to 76,529 (86 percent of the 88,700 NCVS respondents in 2008).

To account for the demographic fluctuations associated with subsampling within a household, we selected two sets of 1,000 replicated samples from the NCVS public-use database for the 2008 collection year. For the one-respondent simulation, we selected

one respondent at random from each of the 28,418 households with two or more respondents and combined this subsample with data from the 19,693 single-respondent households. For the two-respondent simulation, we selected two respondents from each of 8,211 households with three or more respondents and combined this subsample with data from the 39,900 single- and two-respondent households. Although some respondents completed the screening interview twice during 2008, each respondent's probability of selection was independent of the number of times he or she responded.

We extracted the 2008 collection year interviews for all respondents selected during subsampling. We then adjusted the weights of subsampled respondents in multi-respondent households to reflect the respondents' within-household probabilities of selection. For each simulation sample, we post-stratified the weights assigned to subsampled respondents so that the gender, race, and age characteristics of the subsamples would match the same control totals as those used in the full NCVS sample. The subsampling and weighting process is outlined for the one-person subsample in Figure 1.

Table 1. Distribution of responding households and interview respondents in the 2008 National Crime Victimization Survey by number of respondents per household, N = 88,700

Number of Interview Respondents per Household	Responding Households[1]		Interview Respondents[2]	
	Number	**Percent**	**Number**	**Percent**
One	19,693	41%	19,693	22%
Two	20,207	42%	40,414	46%
Three or more	8,211	17%	28,593	32%
Overall	**48,111**	**100%**	**88,700**	**100%**

[1] NCVS households with one or more interview respondents during 2008.

[2] Persons who completed a NCVS screening interview during one or more quarters in 2008.

[4] A respondent is a person who completed the screening interview during one or two quarters in 2008.

Figure 1. Overview of the sampling, weighting, and estimation process for the one-person-per-household replicate samples
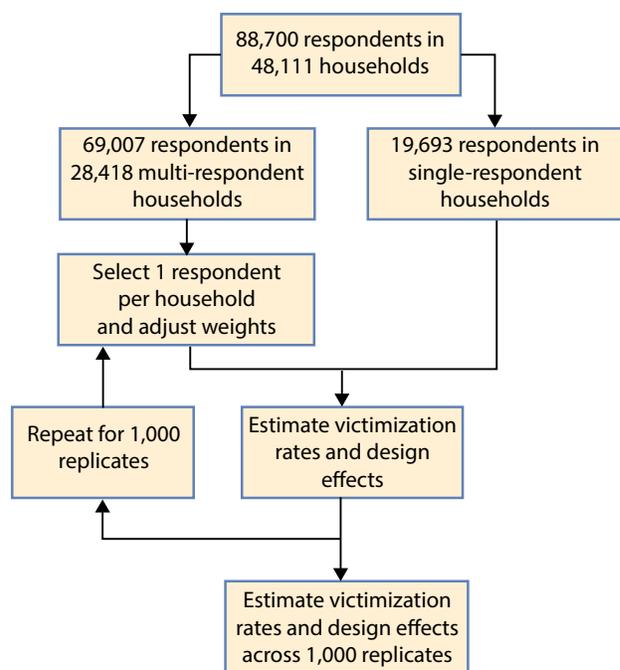
An analogous process was used for the two-person subsample.

Results

We calculated the design effects of four key victimization rates for the designs with one and two persons per household using the adjusted weights and the replicate subsamples and for the existing design using the unadjusted weights and the complete sample. We computed variance estimates using SUDAAN® software (RTI International, 2008) with the pseudo-stratum code and the half-sample code as described in the 2008 NCVS codebook. We obtained overall design effects of victimization rates for the one- and two-person designs by averaging

across the 1,000 replicate samples. We then used the resulting design effects of victimization rates to compare the precision of estimates obtained from the one- and two-person designs to estimates calculated under the existing design. The median design effects across the four crimes are shown for key demographic domains in Table 2.

As expected, the median $DEFF_C$ is lower for the one- and two-person samples than for the full sample because the sample with one person per household eliminates the intra-household correlation and the sample with two persons per household reduces it. However, the opposite is true for $DEFF_W$. As Figure 2 shows, the unequal probabilities of selection for persons in multi-respondent households causes

Table 2. Median design effects for victimization rates: full sample vs. two persons vs. one person per household samples

Domain	$DEFF_C$[1]			$DEFF_W$[2]			$DEFF_T$[3]		
	Full Sample	Two Persons[4]	One Person[4]	Full Sample	Two Persons[4]	One Person[4]	Full Sample	Two Persons[4]	One Person[4]
Overall	1.26	1.23	1.16	1.09	1.24	1.53	1.38	1.52	1.77
Gender									
Male	1.33	1.30	1.21	1.09	1.23	1.50	1.45	1.61	1.82
Female	1.19	1.13	1.07	1.10	1.24	1.55	1.31	1.39	1.66
Race									
White Only	1.19	1.12	1.05	1.08	1.23	1.52	1.28	1.38	1.59
Black Only	1.39	1.27	1.10	1.14	1.27	1.61	1.58	1.61	1.78
Other	1.20	1.09	1.06	1.10	1.20	1.43	1.31	1.31	1.51
Hispanic Origin									
Hispanic	1.02	0.95	0.94	1.08	1.27	1.56	1.11	1.21	1.46
Non-Hispanic[5]	1.22	1.23	1.17	1.09	1.23	1.52	1.34	1.51	1.79
Age									
12–15	1.22	1.21	1.17	1.07	1.14	1.14	1.30	1.38	1.34
16–19	1.17	1.11	1.06	1.05	1.20	1.34	1.24	1.33	1.43
20–24	1.16	1.04	0.94	1.10	1.23	1.48	1.27	1.28	1.39
25–34	1.30	1.21	1.17	1.10	1.18	1.43	1.43	1.43	1.67
35–49	0.99	0.93	0.88	1.05	1.15	1.40	1.04	1.08	1.23
50–64	1.09	1.07	1.00	1.06	1.11	1.33	1.16	1.19	1.32
65+	1.16	1.13	1.15	1.05	1.08	1.21	1.22	1.23	1.39

Note: Table refers to median design effects for victimization rates associated with rape/sexual assault, robbery, assault, and personal theft.

[1] Design effect attributable to clustering.

[2] Design effect attributable to unequal weighting.

[3] Overall design effect is the product of $DEFF_C$ and $DEFF_W$.

[4] Median design effects for the one- and two-person-per-household samples are averaged across the 1,000 replicate samples.

[5] Includes 184 persons with ethnicity unknown.

Figure 2. Comparison of median design effects attributable to clustering ($DEFF_C$) and unequal weighting ($DEFF_W$) for the current "All persons per household" NCVS sample and the within-household subsamples
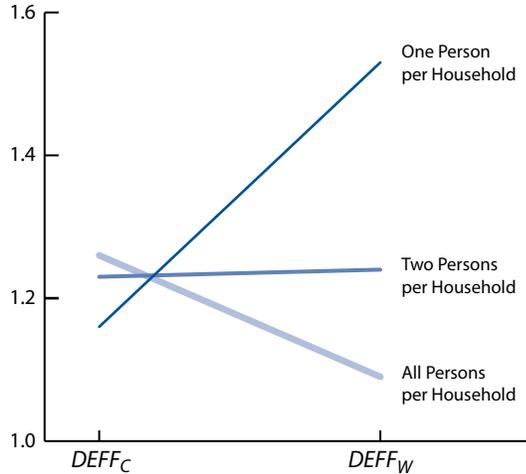
$DEFF_W$ to be highest for the one-person sample, the next highest for the two-person sample, and the lowest for the full sample. When combined, the loss in precision attributable to unequal weighting outweighs the gains in precision from eliminating or reducing within-household clustering.

To determine the stability of the design effects using one or two persons per household, we calculated the simulation variance and coefficient of variation (CV) for each estimate. The simulation variance and CV of a design effect are defined as follows:

$$Var(deff) = \frac{1}{R-1} \sum_{r=1}^{R} (deff_r - \overline{deff})^2$$

where $\overline{deff} = \frac{1}{R} \sum_{r=1}^{R} deff_r$ and R is the number of simulation samples ($R = 1{,}000$), and

$$CV(deff) = \frac{\sqrt{Var(deff)}}{\overline{deff}}$$

The CVs associated with the three types of design effects are presented in Table 3 for both the one- and two-person-per-household designs. The table shows that $DEFF_C$ is more variable than $DEFF_W$ and accounts for most of the variability in the overall design effect. The stability of the CVs indicates that our conclusions about increased design effects are not subject to excessive random variation.

Table 3. Coefficients of variation for median design effects associated with victimization rates

	One Person per HH	Two Persons per HH
$DEFF_C$[1]	0.20	0.11
$DEFF_W$[2]	0.02	0.01
$DEFF_T$[3]	0.20	0.11

Note: Table refers to median design effects for victimization rates associated with rape/sexual assault, robbery, assault, and personal theft.

[1] Design effect attributable to clustering.

[2] Design effect attributable to unequal weighting.

[3] Product of $DEFF_C$ and $DEFF_W$.

Discussion

The results of the simulations indicate that subsampling either one or two eligible persons from each multi-person household selected for the NCVS is likely to significantly increase the design effects of the victimization rates. Increased design effects would cause either increased costs associated with sampling more households to maintain the current precision of victimization rate estimates, or a loss in precision of victimization rate estimates.

The nominal sample sizes for the existing sample and the sample with one person per household could be equalized by enrolling an additional sample of 88,700 − 48,111 = 40,589 households. Simply equalizing the nominal sample sizes, however, does not consider the increased design effects that are associated with a sample of one or two persons per household. Specifically, selecting one person per household would require one respondent to be enrolled from each of 123,898 households to achieve the same precision as the existing 2008 sample victimization rate estimates. This represents an increase of nearly 75,800 participating households over the 48,111 households achieved in the 2008 NCVS. A design with two persons per household would require at least 18,796 additional participating households to equalize the current precision. Although the cost savings associated with interviewing a subsample of persons in a multi-person household would offset at least part of the increased cost needed to enroll additional households, it is reasonable to assume that additional resources would be needed to equalize the precision

of a within-household subsample with that of the current design.

Our results indicating increased design effects with a design of one person per household are at odds with those reported by Groves and Heeringa (2006). Their empirical study compared the relative sampling variance associated with selecting one adult in a two-person household to selecting both adults and was conducted as part of the National Comorbidity Survey–Replication (NCS-R). The study found that the addition of a second adult respondent in eligible two-person households increased the average sampling variance associated with prevalence rates for mental health diagnoses by 10 to 15 percent. Unlike our simulation study, which included all households regardless of size, the NCS-R empirical study was restricted to households with two or more persons. Therefore, it excluded single-adult households, which account for approximately 22 percent of both the NCVS and the NCS R target populations. In general, persons living in single-person households will have a much higher selection probability than those living in multi-person households.

Our simulation study indicates that the unequal weighting that results when multi-person households are combined with one-person households more than offsets any reduction in design effects caused by the lack of intra-household correlation in a one-person per household selection. However, three important caveats are associated with this analysis.

1. The simulation assumes that the response propensities of NCVS sample members are not significantly affected by within-household subsampling. However, the survey literature (e.g., Sharp & Frankel, 1983) suggests that the size of the survey request (intention to interview everyone aged 12 or older in a household versus a subsample) may affect response rates (i.e., the greater the burden, the lower the participation rate).

2. Attempting to interview everyone in a household may result in privacy concerns that cause deliberate concealment of one or more household members (Martin, 1999). In addition, a positive (or negative) interview experience for one household member may help to gain (or discourage) the cooperation of the other household members. This group dynamic would not apply to a single-respondent design.

3. The results of this simulation are specific to the NCVS survey design and the types of estimates analyzed (i.e., personal victimization rates). Design effects attributable to clustering are outcome- and design-specific, so the homogeneity patterns for crime victimization may differ from other social phenomena.

Despite these limitations, this research provides an estimate of the loss in statistical precision that would result if the NCVS were to transition to selecting one or two persons per household. Although within-household subsampling would reduce the burden on individual households, the resulting increase in design effects would lead either to higher costs associated with selecting significantly more households or to a loss in statistical precision of NCVS survey estimates.

References

Gabler, S., Haeder, S., & Lahiri, P. (1999). A model based justification of Kish's formula for design effects for weighting and clustering. *Survey Methodology, 25*, 105–106.

Groves, R., & 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.

Hanson, M., Hurwitz, W., & Madow, W. (1953). *Sample survey methods and theory, volume I*. New York: Wiley & Sons.

Kish, L. (1965). *Survey Sampling*. New York: John Wiley & Sons.

Kish, L. (1987). Weighting in Deft2. *The Survey Statistician*. The Newsletter of the International Association of Survey Statisticians, June.

Martin, E. (1999). Who knows who lives here? Within-household disagreements as a source of survey coverage error. *Public Opinion Quarterly* 63(2), 220–236.

National Research Council. (2008). *Surveying victims: Options for conducting the National Crime Victimization Survey*. Panel to review the programs of the Bureau of Justice Statistics. Robert M. Groves and Daniel L. Cork, eds. Committee on National Statistics and Committee on Law and Justice, Division of Behavioral and Social Sciences and Education. Washington, DC: The National Academies Press.

RTI International. (2008). *SUDAAN® language manual, release 10.0*. Research Triangle Park, NC: RTI International.

Sharp, L., & Frankel, J. (1983). Respondent burden: A test of some common assumptions. *Public Opinion Quarterly, 47*(1), 36–53.

US Census Bureau. (2009, Oct 22). *General description of the 2000 sample design of the National Crime Victimization Survey* (NCVS00SDE01). Internal Census Bureau memorandum from Kathlene Garland to Stephen Ash, Chief, Victimization and Expenditures Branch, Demographic Statistical Methods Division.

US Department of Justice, Bureau of Justice Statistics. (2008). *National Crime Victimization Survey, 2008* [Collection Year Record-Type Files] [Computer file]. ICPSR25461-v2. Ann Arbor, MI: Inter-university Consortium for Political and Social Research [distributor], 2009-09-11. doi:10.3886/ICPSR25461

RTI International is an independent, nonprofit research organization dedicated to improving the human condition by turning knowledge into practice. RTI offers innovative research and technical solutions to governments and businesses worldwide in the areas of health and pharmaceuticals, education and training, surveys and statistics, advanced technology, international development, economic and social policy, energy and the environment, and laboratory and chemistry services.

The RTI Press complements traditional publication outlets by providing another way for RTI researchers to disseminate the knowledge they generate. This PDF document is offered as a public service of RTI International.