TSAS

canadian network for research on terrorism, security and society

Working Paper Series

No. 14-08

August 2014

A Framework for Estimating the Number of Extremists in Canada

Garth Davies

and

Stephanie Dawson

The Canadian Network for Research on Terrorism, Security, and Society

TSAS is supported as a national strategic initiative funded by **SSHRC** and **Public Safety Canada**, along with the following departments of the federal government:

- Canadian Security Intelligence Service (CSIS)
- Citizenship and Immigration Canada (CIC)
- Royal Canadian Mounted Police (RCMP)

TSAS also receives financial support from the University of British Columbia, Simon Fraser University, and the University of Waterloo.

Views expressed in this manuscript are those of the author(s) alone. For more information, contact the Co-directors of the Centre, Daniel Hiebert, Department of Geography, UBC (daniel.hiebert@ubc.ca) and Lorne Dawson, Department of Sociology and Legal Studies, University of Waterloo (ldawson@uwaterloo.ca).

Table of Contents

TSAS

canadian network for research on terrorism, security and society

Working Paper Series

A Framework for Estimating the Number of Extremists in Canada

Garth Davies

School of Criminology, Simon Fraser University

Stephanie Dawson

School of Criminology, Simon Fraser University

Introduction

Violent extremism and terrorism are among the most salient policy issues facing governments today. While Canadian authorities have made significant strides in combatting terrorism, there remains much that is unknown about the phenomenon. One of the most basic questions for which there is scant empirical data relates to numbers. Specifically, at present we lack any solid estimates of the number of violent extremists in Canada. The reasons for the paucity of information are clear. The very nature of terrorism, a socially stigmatized and illegal activity, necessitates secrecy. Consequently, terrorists go to great lengths to avoid detection. But having a reasonable idea of the size of the terrorist threat is a central policy consideration. For example, to that extent that various policy initiatives are aimed at preventing or reducing radicalization, knowing the number of terrorists provides a baseline against which these programs may be evaluated. Such knowledge would also allow for the better allocation and prioritization of limited resources. Finally, it would allow us to better understand the evolving nature of the threat posed by terrorism.

In the absence of sound estimates, disagreements over the number of violent extremists have been characterized by wild speculation. Following the events of September 11, this debate has focused almost solely on the number of "potential" Muslim terrorists. At the highest end of the spectrum, it has been suggested that up to 15 percent of Muslims have fundamentalist beliefs supportive of terrorism. Based on a worldwide population of between 1.2 and 1.5 billion Muslims, this would translate to at least 180 million, and as many as 225 million, terrorists-in-waiting. In a similar vein, radio and television personality Glenn Beck proposed in 2010 that the 10 percent of Muslims (according to his calculation, 160 million) were affiliated with terrorists. But if these numbers had any basis in reality, wouldn't the incidence of terrorism be much, much higher? In response to these astronomical estimates, Kurzman (2011, 7) asked "why we don't see terrorist attacks everywhere, every day". Depending on one's calculations, Kurzman places the number of potential Muslim terrorists between 10,000 and 15,000; still a large number, but microscopic compared with the estimates offered previously. In lieu of a convincing empirical framework, the debate over numbers is reduced to competing ideological and political positions.

While the explanation for the lack of data on the number of terrorists is reasonably straightforward, and a need to remedy this situation is pressing, the solution to the problem is less apparent. The objective of this report is to examine how the numbers of terrorists might be determined or estimated in a more systematic, or at least more methodologically compelling, way. Terrorists are not the only group of individuals that are challenging to study. In fact, terrorists share some of the same traits as other groups. For example, owing to their illicit, clandestine, and/or transitory nature, populations such as drug users, prostitutes, and the homeless are not amenable to traditional sampling methods. As a result, numerous techniques have been developed to sample and count "hidden populations." This study explores the range of "hidden population" techniques, and attempts to determine which, if any, might assist in better estimating the number of terrorists in Canada.

Hidden Populations

Prior to discussing the techniques available for estimating the number of extremists, it is necessary to clarify what the term "hidden populations" actually means. Hidden populations refer to small subpopulations of individuals who are unwilling to disclose themselves (Frank and Snijders 1994). More specifically, these groups are vulnerable or socially stigmatized and engage in largely low-visibility, and often illegal or illicit behaviours (Bloor, Leyland, Barnard, and McKeganey 1991; Johnston and Sabin 2010). Examples of these populations include, but are not limited to: the homeless, injection drug users, drug dealers, men who have sex with men, commercial sex workers, illegal immigrants, victims of human trafficking, and artists (Salganik and Heckathorn 2004). Given that extremists and terrorists engage in illegal activity and, as a result, operate in an extraordinarily clandestine manner, this group of individuals represents a prototypical example of a hidden population.

Studies on these hidden populations have raised a number of specific methodological concerns (Berk 2007; Faugier and Sargeant 1997; Salganik and Heckathorn 2004; Thompson 2012). First, hidden populations do not have a consensus-based sampling frame (i.e., a list of all members in the populations) from which to define and randomly sample these groups (Acharya 2007; Faugier and Sargeant 1997; Thompson 2012). This means that there is no known probability of selection. Simple random sampling methods, therefore, are not efficient methods to estimate categorical frequencies of hidden populations. Thus, traditional descending methodologies (i.e., those using representative sampling strategies drawn from highly standardized questionnaires to make inferences about a whole population) have limited utility in these situations (Berk 2007; Faugier and Sargeant 1997).

Second, it is often impractical to construct a sampling frame for these populations. Given their low visibility, dispersed hidden populations are difficult to find and sample (Faugier and Sargeant 1997; Goel and Salganik 2009; Lee 1993). For instance, many of these populations are "floating", (i.e., they have the unique quality of horizontal and vertical socio-geographic mobility), which prevents existing surveillance systems grounded in fixed community-based institutions, such as health care and police departments, from accounting for them (Faugier and Sargeant 1997; Lee 1993). Important segments of these populations, therefore, are often missed. To complicate matters further, even if members of the population are located, they may be unwilling to participate in surveillance data collection or to reveal their activities to an unknown interviewer (Faugier and Sargeant 1997; Goel and Salganik 2009). According to Faugier and Sargeant (1997), the more sensitive or threatening the phenomenon under study, the greater the potential for respondents to conceal their activities and, thus, the more difficult the sampling is likely to be. Participating in illegal or illicit activities that often carry harsh legal or social sanctions may deter members of hidden populations from cooperating in sampling procedures (Goel and Salganik 2009).

Finally, to achieve unbiased estimates of the prevalence of these hidden populations, large samples are required. A large sample provides sufficient data for an accurate estimation of what is a statistically rare event (Acharya 2007; Hendricks and Blanken 1992; Salganik and Heckathorn 2004). Compared to the general population, however, hidden populations are relatively small (Goel and Salganik 2009). In addition to being small in overall numbers, it is often difficult to distinguish between members of the hidden population and those of the general population;

there are seldom clear and/or obvious distinctions between members of these different groups (Acharya 2007; Goel and Salganik, 2009; Lee 1993). Drawing a large and representative sample of these populations, therefore, would be a labour-intensive, expensive and extremely challenging endeavour (Salganik and Heckathorn 2004; Medhi et al. 2012). Taken together, it is clear that obtaining data from these subpopulations and producing estimates is an extraordinarily difficult task.

In an attempt to address the issues associated with studying hidden populations, a number of non-probability sampling and estimation techniques have been developed. These approaches include link-tracing designs, targeted sampling, time-location sampling (TLS), capture-recapture techniques, fieldwork sampling, and multiplier methods. Each method has been applied to a number of distinct populations and offers a unique approach to estimating the prevalence of hidden populations. For ease of presentation, each technique will be outlined using the following format. First, the details concerning how the method works will be discussed. Examples of its use with specific hidden populations will then be examined. Finally, each section will conclude with a discussion of the strengths and limitations of the technique.

Link-Tracing Designs

Link-tracing is an umbrella term capturing a variety of sampling techniques, including snowball sampling, network sampling, respondent-driven sampling (RDS), and adaptive sampling. As a group, these techniques are meant to be used as both an informal means to reach, explore and describe a hidden population, as well as a formal method to make inferences about the population of interest (Atkinson and Flint 2001; Snijders1992). Premised on the idea that those best able to access members of hidden populations are their own peers, link-tracing methods seek to exploit the natural links that exist between individuals in a population (Heckathorn 1997; Atkinson and Flint 2001). Assuming that groups tend to be linked together by common social traits, such as their sexual partners or drug use (St. Clair and O'Connell 2012), link-tracing methods follow the links that tie together members of the population of interest. Using the resulting social networks of identified individuals, researchers are provided with an ever-expanding pool of potential contacts and a means to make inferences about the structure of these networks (Heckathorn 1997).

How It Works

Although conceptually straight-forward, understanding how link-tracing designs actually work requires more explanation. At the most basic level, all link-tracing designs operate by following the links between individuals in a hidden group using successive sampling waves (Frank and Snijders 1994). To begin, an initial sample of respondents must be selected (wave 0). This requires clearly outlining the standard for inclusion (i.e., defining the hidden population and the links that will be used to tie members of the population together), and finding a set of individuals ("units") who meet these selection criteria (Thompson 2012). After this initial sample of respondents is obtained, each individual in the sample is supposed to name other members of the target population (i.e., those who meet the inclusion criterion) with whom they are acquainted (Frank and Snijders 1994; Thompson 2012). Some or all of the links identified from the initially sampled units will then be traced to obtain a new set of population members (wave 1). Sampling will continue in this manner until a predetermined number of waves have been sampled, a predetermined sample size has been obtained, or until every unit obtained in a wave has already been sampled in a previous wave (Frank and Snijders 1994). The final sample will consist of the initial sample and all the waves successively found around it.

Once the final sample has been obtained, the networks are modeled as graphs. The graphs consist of a set of nodes, representing people (i.e., the members of the hidden population of interest), and a set of edges or arcs that represent the links[1] (i.e., social or spatial relationships) between the nodes (St. Clair and O'Connell 2012). Visually, the graph depicts a set of small circles representing the nodes, and lines or arrows representing the links (Thompson 2012). Revealing information about the characteristics of the nodes, as well as their networks, each graph represents the entire hidden population of interest (St. Clair and O'Connell 2012).

1It is important to note that the arcs represent different types of links depending on the type of link-tracing design used. In snowball sampling and respondent-driven sampling, arcs represent social links or interactions between individuals. In adaptive sampling, on the other hand, the arcs represent spatial relationships between neighbouring units (St. Clair and O'Connell 2012).

Culminating with this graphical representation, link-tracing designs can be used in an informal manner to qualitatively describe the identified network (Atkinson and Flint 2001; Frank and Snijders 1994). However, it is often the researcher's goal to use the data collected via these sampling strategies to make inferences about the hidden population (Atkinson and Flint 2001). But in order to make accurate inferences from the sample to the population, the zero sample must be considered a probability sample (Faugier and Sargeant 1997). However, given that link-tracing designs are classified as non-probability sampling techniques, using these methods to study hidden populations means that obtaining a full random or probability sample is not possible (Faugier and Sargeant 1997; Thompson 2012). The representativeness of these samples, therefore, must be assessed in a different manner. Relying on tentative maps (i.e., known information concerning the distribution of the target population, including types of individuals, times and places), a number of estimation methods have been developed to deal specifically with the non-probability samples created using link-tracing designs (Chow and Thompson 2003; Faugier and Sargeant 1997; Frank and Snijders 1994; Handock, Gile, and Mar 2012). For example, design-based estimates of population means and totals can be used with snowball samples. Treating the population size (N) as an unknown parameter[2], the number of vertices of the fixed arc indicators of a snowball sample can be estimated (Frank and Snijders 1994). In essence, information about the population size is fixed and probability plays a role only via the sampling procedure (i.e., the population size is drawn from the pattern in the sampling process) (Frank and Snijders 1994).

A variety of Bayesian models have also been utilized. For instance, successive sampling approximation creates Bayes estimates of the population size for respondent-driven samples. By exploiting the dependence in the respondent-driven sampling design, this method works by leveraging the information in the decreasing size of the sampled units over time to make inferences about population size (Handcock, Gile, and Mar 2012). Bayesian binary response models can be used to estimate a population proportion when using one-wave snowball or adaptive cluster sampling designs (Chow and Thompson 2003). For link-tracing designs allowing for subsampling within each wave, Markov chain Monte Carlo (MCMC) techniques have been used

2Treating the population size (N) as a unknown parameter requires a probability model for the observed data given (N), as well as a prior for N (Frank and Snijders 1994).

to compute Bayes estimates (Goel and Salganik 2009). Using this technique, a weighted sample mean is used to produce an unbiased estimate of the population mean[3] (Goel and Salganik 2009; Salganik and Heckathorn 2004). Lastly, estimates of hidden population size for link-tracing sampling designs can be obtained using model-based maximum likelihood estimation (Frank and Snijders 1994; Felix-Medina, Monjardin, and Aceve-Castro 2009). These estimates work by modeling the links obtained from random networks (Frank and Snijders 1994; Felix-Medina, Monjardin, and Aceve-Castro 2009).

All link-tracing designs operate using some variation of this basic process of sampling and estimation. However, there are some important and noteworthy differences between these techniques. To gain a more complete understanding of link-tracing methods, therefore, it is necessary to provide a separate discussion of the elements that are specific to each these techniques.

Snowball Sampling

Snowball sampling is used to estimate hidden populations by enlarging an initial, small sample of individuals through their reported contacts or relationships with others (Atkinson and Flint 2001; Frank and Snijders 1994; Kaplan, Korf, and Sterk 1987). This is accomplished in a series of steps. First, an initial number of hidden population members who have the desired characteristics are identified. This initial sample can be a simple random sample from the general population containing the hidden population as a subpopulation, or it can be obtained via site sampling (i.e., sampling certain areas where members of the hidden population are known to frequent) (Frank and Snijders 1994). These individuals are known as the "seeds" (Atkinson and Flint 2001; Thompson 2012). These seeds are then used to recruit similar participants in a multistage process. To start, the seeds' social networks are used to recruit new respondents. More specifically, by following a subsample of the network links, new respondents are selected. These respondents are then asked to provide information about other members of the subpopulation of interest. At this point, additional respondents will give the researcher the name of more subjects.

3MCMC techniques are often used with RDS. Based on the idea that the network is connected and every person is reachable from every other person, this method works off of two assumptions (Salganik and Heckathorn 2004). First, participants are assumed to recruit a single individual who was chosen uniformly at random from their network of contacts. Second, it is possible that participants are recruited into the sample multiple times (i.e., allows for sampling with replacement) (Goel and Salganik 2009).

This process of subject referrals and recruitment continues until a target sample size has been reached, or the sample has become saturated[4] (Atkinson and Flint 2001; Magnani, Sabin, Saidel, and Heckathorn 2005; Sadler et al. 2010).

The abovementioned process outlines the most basic form of snowball sampling. This technique, though, can also be adapted to better fit specific situations. For example, the random-walk design is a snowball sampling method used when respondents would be unable to divulge sensitive information about their peers (e.g., when federal guidelines for the protection of human rights prohibit sharing personal information) (Klovdahl 1989). In this instance, respondents are asked to list the people they know in the population under study, and then the researcher uses a randomized device to select the people to be included in the next wave of the sample. Essentially, only one of the links from any one node is randomly selected and followed to another node, and so forth (Klovdahl 1989). Ego networks, also known as one-layered snowball sampling designs, are used when the researchers can only obtain a single wave of respondents. This form of snowball sampling uses a convenience sample of respondents who report on their contacts and provide their knowledge about the ties that exist between these contacts (Frank and Snijder 1994). Regardless of the way this technique is implemented, however, snowball sampling works by using the information provided by the sampled units about themselves and about other units (Atkinson and Flint 2001; Sadler et al. 2010). Once the sample is finalized, estimates of the categorical frequencies of the hidden population can be obtained. This is accomplished by applying one of the previously described statistical methods (i.e., those that have been developed to deal with non-probability samples, such as those created by snowball sampling techniques) (Frank and Snijder 1994).

Network Sampling

Closely related to snowball sampling techniques, network sampling works by using each individual in the sample as a sampling node to generate the next subject, and repeats this procedure until the network is exhausted (Sudman and Blair 1999). Rather than accessing the target

4A sample becomes saturated when the information provided by new sample subgroup members does not differ from that obtained from previously interviewed respondents (Atkinson and Flint 2001).

population directly, however, network sampling uses a random or stratified random sample of the general population to report about members of the target population (Bernard et al. 2010; Sudman and Blair 1999). More specifically, a simple random sample or stratified random sample of units is selected from the general population, and all observation units (i.e., individuals in the hidden population) linked to any of the selected units are included in the sample (Sudman and Blair 1999). In order to use these samples to estimate the size of the hidden population of interest, two requirements must be fulfilled. First, there must be an informant within the sample who has the information needed to report about all network members. This requires that respondents know people with specific characteristics in a variety of populations (Sudman and Blair 1999). Second, because network sizes vary from person to person and the probability of being identified is directly proportional to network size, located members of the hidden population must be weighted inversely by their network size. As a result, estimates of network size are most accurate when network members are closely related (Sudman and Blair 1999).

Similar to snowball sampling, there are slight variations in the way the network sampling technique has been utilized. For instance, when the researcher's agenda requires that a structured instrument be administered to a large number of participants who inhabit different social networks within a selected area, the privileged access interviewer (PAI) method is used (Griffiths, Gossop, Powis, and Strang 1993). Using an interviewer who has contact or characteristics in common with the population of interest, this method enables the researcher to collect structured data outside of treatment or closed settings (Griffiths, Gossop, Powis, and Strang 1993). The network scale-up method (NSUM), on the other hand, is used to estimate the personal network size of the members of a random sample of a total population, and, using this information, to then estimate the number of members of a hidden subpopulation of the total population (Bernard et al. 2010). Resting on the assumption that people's social networks are, on average, representative of the general population in which they live and move, this method starts by estimating each respondent's personal network size (Bernard et al. 2010). This is accomplished by asking the respondent about the number of people he/she knows in various populations of known size, or by having each respondent enumerate the people they know in a list of specific relationship types or categories (e.g., family, friends, coworkers, etc.) (Bernard et al. 2010). Once the personal network sizes have been derived, these prevalence estimates are combined with known information about the size of the general population. This produces an estimate for the number of people in the population who engage in the target activity (Bernard et al. 2010). For example, if a sample of respondents knows an average of 300 people and two of these individuals engage in the activity of interest, then it is estimated that two out of every 300 people in the general population are part of the hidden population (Bernard et al. 2010). To improve the accuracy of the estimate, the responses from many different respondents should be combined (Bernard et al. 2010).

Respondent-Driven Sampling

Based on the major tenets of snowball sampling, respondent-driven sampling (RDS) is designed to estimate the proportion of a population with a specific characteristic (Goel and Salganik 2009; Heckathorn 1997; Heckathorn 2002). Applied in situations where the sampling population has an identifiable contact pattern[5], respondents are selected from friendship networks of existing members of the sample using a multistage process (Heckathorn 1997). To begin, the researcher selects a small number of "seeds"–the initial population members to participate in the study. These seeds are chosen based on some form of pre-existing contact with the study population, or because they have special attributes that ensure effective recruitment (Johnston and Sabin 2010). Regardless of the basis for their selection, all seeds will (1) have large social networks, (2) be respected by members of the hidden population of interest, and (3) be able to convince others to participate in the study and have interest in the study goals (Heckathorn 2002). In addition, the seeds should comprise key subpopulations to enable the sample to reach a stable composition and equilibrium with respect to the traits and characteristics upon which the research focuses in a timely manner (Goel and Salganik 2009). Once the seeds are selected, this forms wave 0 of the sample.

After participating in the study, the seeds are then asked to recruit others from the target population to participate. This requires that respondents know one another as being members

5A contact pattern refers to the activities that constitute membership in the population of interest that also provide a basis for creating connections among population members (Heckathorn 1997).

of the hidden population of interest (Johnston and Sabin 2010). During the recruitment process, each seed will provide a coupon to each person they recruit. To allow the researcher to trace the recruitment patterns in the population, each seed is given a fixed number of unique coupons (Goel and Salganik 2009; Salganik and Heckathorn 2004). Those successfully recruited from the people in wave 0 form wave 1. This process of existing sample members recruiting future sample members continues until the desired sample size is reached, or until the sample reaches equilibrium[6] (Acharya 2007; Goel and Salganik 2009; Magnani et al. 2005).

It is important to note that for the sampling process to continue, it is necessary that each participant recruit at least one new participant (Salganik and Heckathorn 2004). As such, the choice of number of recruits/coupons should be large enough to enable the recruitment process to continue even if some subjects choose not to recruit. Incentives, such as payment and gifts, are often used to encourage participation (Johnston and Sabin 2010). Furthermore, if the seeds and/or recruits are not drawn as desired, it may be preferable to create long chains using many sampling waves. This enables the researcher to explore parts of the network and more isolated segments of population that may have otherwise had zero chance of being included in the sample (Salganik and Heckathorn 2004).

Once the sample has been selected, estimates about the social network can be made. Rather than directly estimating from the sample to the population, though, RDS uses an indirect method. According to Heckathorn (2002), the sample is first used to make estimates about the social network connecting the population. This information about the social network is then used to derive the proportion of the population in different groups. To estimate from the social network to the population, a reciprocal model is used to calculate the number of cross-group ties between the groups in the population (Salganik and Heckathorn 2004). These cross-group friendships are created based on behaviour; when one respondent recruits another, this behavioural link represents a network link that can be verified by asking the recruitee to characterize the degree of their relationship to the recruiter (i.e., friend or acquaintance). By emphasizing the difference between the various levels of friendship, this method brings together information about the char-

6Equilibrium refers to a state in which the sample composition does not change during subsequent cycles of recruitment (Magnani et al. 2005).

acteristics of both the nodes and the network, and provides estimates of the proportion of the population with a specific trait (Salganik and Heckathorn 2004).

In order to make these estimates, two pieces of information are required. First, the self-reported degree of each person in the sample must be known (Thompson and Collins 2002; Thompson and Frank 2000; Thompson and Seber 1996). This requires the researcher to account for the connections between the recruiters and their recruits, as well as the size of the social network of each participant. Second, the coupon's unique identifier must be recorded. The coupon links each sample member to the person who recruited them, thus providing the pattern of recruitment (Magnani et al. 2005; Salganik and Heckathorn 2004; Thompson and Seber 1996; Thompson and Frank 2000).

In addition to having these pieces of information, in order to use RDS to produce asymptotically unbiased estimates about the hidden population, certain assumptions must be met (Goel and Salganik 2009). For example, the network of the hidden population is assumed to form one connected component; there is a path between every person and every other person in the sample. Within the sample, all respondents must have received and used one coupon, and the recruitment of new respondents was done randomly (Salganik and Heckathorn 2004). Finally, the seeds are assumed to be sampled with replacement[7] and drawn with a probability proportional to their degree[8]. This enables the researcher to assume that each unit has a known and constant probability of selection, which allows for the recovery of population proportions using knowledge of the network structure that connects the population (Salganik and Heckathorn 2004).

Adaptive Sampling

Contrary to the other link-tracing methods, which focus on identifying networks based on social relationships, adaptive sampling techniques were developed to take advantage of spatial patterns in the population (Seber and Salehi 2013; Thompson 1990; Thompson 2012). Used to obtain more precise measures of population abundance or density for a given sample size, adaptive sampling strategies are designed for situations where data are being gathered from popula-

7 Although sampling can be done without replacement, assuming replacement eliminates changes in the population as the sampling progresses, and, thus, simplifies the estimation procedure (Heckathorn 2002; Thompson and Seber 1996).

8This means that a person with ten friends is twice as likely to be a seed as a person with only five friends (Salganik and Heckathorn 2004).

tions located in specific places (Thompson 1992). For example, these techniques can be used with populations of animals and plants that have aggregation tendencies resulting from schooling, flocking, dispersal patterns and/or environmental patchiness (Thompson 2012). To obtain these samples, adaptive sampling works by adapting the sampling pattern to what happens to be discovered at each stage of the sampling process (Seber and Salehi 2013). Basically, the procedure for selecting the sites or units to be included in the sample largely depends on the values of the variable of interest that are observed during the survey process (Seber and Salehi 2013; Thompson 2012).

Adaptive sampling, much like snowball and network sampling, includes a variety of sampling techniques. One such method is adaptive cluster sampling, which is useful when the organisms of interest are rare and highly clustered in their geographic distribution (Seber and Salehi 2013). This technique works by locating a group of organisms and then sampling a complete cluster. More specifically, adaptive cluster sampling begins with an initial sample of quadrants, plots, or areas selected by simple random sampling with or without replacement. Once one of these selected quadrants contains the organism of interest, new quadrants located in the vicinity of the original quadrant are then added to the sample (Magnani et al. 2005; Thompson 1992; Thompson 2012). This process continues until a pre-determined area has been covered or a certain sample size has been reached. After the cluster has been finalized, the mean abundance of the organism in the sample quadrants can be calculated (Seber and Salehi 2013). Prior to engaging in this sampling process, though, the researcher needs to complete three steps. First, the researcher must have defined a neighbourhood where sampling might be carried out such as a quadrant or plot. Second, a condition for choosing when to sample the neighbourhood and who will be part of the cluster needs to have been outlined. Finally, the researcher had to have chosen a sampling method for selecting the initial locations (Seber and Salehi 2013).

Another form of adaptive cluster sampling is hierarchically ascending classification[9]. This technique is useful when a researcher needs to successively construct a more general class or type based on similarities between the descriptive variables being used to analyze the hidden

9Hierarchically ascending classification is also referred to as cluster analysis in some of the literature (van Meter 1990).

population (van Meter 1990). Starting with the entire population, hierarchically ascending classification works by successively dividing this population to obtain a descending hierarchy of levels and classes. Specifically, working up from the zero level, where each individual is in its own unique class, successive levels of classification can be constructed by grouping similar individuals into the same class. Focusing on the resemblance between individuals to form groups, this method creates polythetic classes–groups composed of individuals sharing a set of characteristics that are used to define the class (van Meter 1990). For example, while the first class consists of individuals who are most similar to one another, the higher levels will have fewer members resembling each other (van Meter 1990). Each one of the resulting levels will correspond to a particular and distinct classification of the entire population of interest. Establishing the statistical probability or likelihood of the occurrence of each particular classification is the final procedure for this particular technique (van Meter 1990).

Closely linked to hierarchical ascending classification is the technique known as cross-classification. Similar to all clustering techniques, cross-classification analyses the similarities between variables for all individuals in the population (van Meter 1990). Taking this one step further, however, this method also assesses the similarities between individuals for all of the variables used to study the hidden population (Kaplan et al. 1985; van Meter 1990). These similarities are then combined and crossed to form cohesive blocks of the individuals and the variables (van Meter 1990).

The final adaptive sampling method, disproportionate stratification, is used when rare populations are more heavily concentrated in certain parts of the general population (Kalton 2001). In these situations, the researcher will have prior knowledge about how a population is dispersed, which allows for the population to be divided into distinct strata or areas (Seber and Salehi 2013). Once the population is stratified, the sampling procedure begins. First, an initial stratified sample is selected from the population. When the value of the variable of interest for any unit is observed to satisfy a specified condition, additional units can be drawn from the neighbourhood of that unit and included in the sample (Thompson 2012). Using simple random sampling within the strata, the optimum sampling fraction in each stratum is based on the prevalence of the rare population in the given stratum (Kalton 2001; Seber and Salehi 2013). Higher sampling fractions, therefore, are used in the strata containing greater concentrations (Seber and Salehi 2013; Thompson 2012).

Populations Studied Using Link-Tracing Designs

With the exception of the adaptive sampling methods, which have largely been applied to plant and animal populations[10], link-tracing designs have been used to study a wide range of hidden human populations. One of the most frequently studied groups is the population of individuals at high-risk for HIV. Using high risk behaviours such as drug use, men having sex with men, or commercial sex work as the link, different link-tracing designs have been utilized to survey at-risk HIV individuals and provide estimates of their population size (e.g., Handcock, Gile, and Mar 2012; Johnston, Sabin, Hien, and Huong 2006; Kajubi et al. 2008; Kendall et al. 2008; Magnani, Sabin, Saidel, and Heckathorn 2005; Paniotto, Petrenko, and Kupriyanov 2009; Pollack, and Schiltz 1988; St. Clair, and O'Connell 2012). For example, using data from the Colorado Springs HIV study, St. Clair and O'Connell (2012) used snowball sampling to survey high-risk communities. Starting with a baseline group of 595 individuals and following the network links defined by social, sexual and drug relationships, they discovered that 160 members from the baseline group were part of the hidden population. Handcock, Gile, and Mar (2012), on the other hand, used RDS to estimate the number of those most at risk for HIV in El Salvador in 2010. Using five seeds of both female sex workers and men who have sex with men (MSM), and nine successive waves of sampling, the authors were able to estimate that a total of 3.5 percent of the El Salvadorian population belonged to the hidden population of at-risk HIV individuals.

Interested in comparing a variety of sampling methods, including snowball sampling, respondent-driven sampling and time-location sampling (TLS)[11], Kendall et al., (2008) examined three samples of men who have sex with men (MSM) in Fortaleza, Brazil. The snowball sample of 127 MSMs was collected by locating MSMs in a variety of venues (e.g., bars, cinemas, streets, and beaches), intercepting and interviewing certain individuals, and then asking them to identify

10 Adaptive sampling has been used to study rare trees (Roesch,1993; Acharyal, Bhattarai, de Gier, and Stein 2000), sediment load in rivers (Arabkhedri, Lai, Noor-Akma, and Mohamad-Roslan 2010), and seaweed (Goldberg, Hein, and Brown 2006).

11Time-location sampling will be discussed in detail later in this report.

others to be recruited and surveyed. Collecting the RDS sample (N = 406) required the researchers to use long-chain recruitment. In this case, this meant that all members of the population who participated in the study referred two other members of the target population. This continued until the sample reached equilibrium (Kendall et al. 2008). Finally, the 274 individuals captured via the TLS method were obtained using ethnographic mapping of different venues and targeted selection of members of the MSM who frequented these venues (Kendall et al. 2008).

Focusing on comparing the samples' composition of specific characteristics, including socio-economic status (SES) and risk behaviours to one another and to known AIDS cases in the general population, Kendall et al. (2008), found that RDS was able to produce a sample with a wider inclusion of lower SES and education levels, and higher condom use at last sex than the snowball and time-location samples. RDS was also found to achieve a sample size more quickly and at a lower cost. Based on their results, the researchers caution that different sampling procedures may be better suited to different populations and settings (Kendall et al. 2008).

In addition to studying hidden HIV populations, link-tracing designs have also been used to analyze populations of drug users. Using snowball sampling techniques, researchers have been able to produce estimates of the prevalence and distribution patterns of cocaine use (Avico et al. 1988) and heroin use (Kaplan, Korf, and Sterk 1987) in different European cities. For example, using an initial interview sample of 153 active cocaine users, Avico et al. (1988), discovered differences in the prevalence in cocaine use in two European cities, Munich and Rotterdam. While the prevalence in cocaine use among artists and actors was 4 percent higher in Rotterdam, Munich had higher prevalence rates in the restaurant/café milieu and white-collar workforce (by 4 and 7 percent respectively) (Avico et al. 1988). These results show that snowball sampling can be used to provide internationally comparable data.

Within the population of drug users, respondent-driven sampling has been aimed at enumerating the specific subpopulations of drug injectors (e.g., Heckathorn 1997, 2002; Platt et al. 2006). The results of Heckathorn's (1997) study of 277 active drug injectors in Connecticut revealed that RDS provides a framework for producing unbiased population estimates. Specifically, by using a sampling process that continues through enough waves to reach equilibrium, the final RDS sample composition will be independent of the initial subjects (Heckathorn 1997).

Although link-tracing designs have largely been used to study these high risk and at-risk hidden populations, they have also been used with other hidden populations. For example, the network scale-up method was used to assess choking in children (Snidero, et al. 2007). Acharya (2007) used the snowball sampling technique to locate and interview trafficked women and female prostitutes in Mexico City, while Browne (2005) examined the ability of snowball sampling to recruit non-heterosexual women via their social networks. Finally, Salganik and Heckathorn (2004) used respondent-driven sampling to estimate the number of jazz musicians in San Francisco. In sum, researchers have been able to use different link-tracing designs to study a variety of interesting hidden populations.

Advantages of Link-Tracing Designs

Taken together, these studies highlight a number of advantages to using link-tracing designs to study hidden populations. The most noteworthy strength of link-tracing sampling designs is their usefulness in situations where there are limited official records or data sources available for capturing members of the hidden population (Faugier and Sargeant 1997). Given the low-visibility, and thus, detectability of many of the target population members, there are seldom adequate and readily available lists for these populations (Faugier and Sargeant 1997). This is particularly true in instances where the focus of the study is on an extremely sensitive issue or rare population property, such as non-heterosexual sexual relationships (Browne 2005; Faugier and Sargeant 1997; Salganik and Heckathorn 2004). By taking advantage of accessible research settings such as public places including bars, parks, and streets, and directly engaging the target population, link-tracing sampling designs can assemble a relatively large and diverse participant group (Avico et al. 1988; Faugier and Sargeant 1997; Sadler et al. 2010). Not only does this produce a thick body of rich and descriptive information about the respondents and their network, it provides sufficiently systematized data–a large and representative sample of the target population–to permit formal analyses and comparisons of rare social phenomena (Acharya 2007; van Meter 1990).

Regarding their efficiency, link-tracing sampling methods have been found to be cheaper, quicker, and more accurate than other sampling methods (Acharya 2007; Atkinson and Flint 2001). In exploring a statistically rare event, for instance, household or other population samples require enormous samples to yield a sufficient number of subjects in the target population. Link-tracing sampling designs, on the other hand, have the potential to produce a rapidly growing database (Salganik and Heckathorn 2004). Thus, for a given amount of resources, link-tracing techniques provide researchers with more study sites and larger sample sizes than other methods (Salganik and Heckathorn 2004). Furthermore, requiring only a basic knowledge of the population to find a person to act as an initial contact, link-tracing designs use fewer resources (Acharya 2007). Utilizing strict criteria for inclusion, such as unique coupons and specific quotas for social characteristics (e.g., age or sex), these methods ensure that each sampling unit only appears in the sample once (Acharya 2007; Frank and Snijders 1994). This reduces the chance for duplication of respondents, and minimizes the biases associated with over-representation of those participants with large networks (Johnston and Sabin 2010). Finally, given that the number of sampling waves and traced links are allowed to vary, these techniques can be easily adapted to fit the population of interest (Atkinson and Flint 2001). As the characteristics, behaviours, and social structures of hidden populations tend to change over time, flexibility is an important asset. For example, with the Internet continually to providing evolving capacities for social networking, there are new mediums for link-tracing sampling methods to explore (Sadler et al. 2010).

In addition to these global advantages, each link-tracing method has its own unique strengths. Snowball sampling techniques, for instance, are efficient and economical ways to find cases that may otherwise be difficult or impossible to locate or contact (Faugier and Sargeant 1997). Specifically, by using members of the target population to recruit others, these methods allow researchers to obtain respondents when some degree of trust is required to initiate contact (Atkinson and Flint 2001). This is especially useful in situations where identifying members of the target population requires information some people may be reluctant to divulge. For example, participants may be unwilling to admit that they have a sexually transmitted infection (Sadler et al. 2010). Moreover, this technique provides a viable channel for reaching small subpopulations that are highly nested within the general population (Atkinson and Flint 2001; Sadler et al. 2010). For example, in diverse communities composed of a variety of subpopulations, it is often difficult for the researchers to distinguish between members of the different subgroups and identify those who are part of the target population (Sadler et al. 2010).

In addition to providing a viable means of obtaining a diverse array of representative samples from sensitive subpopulations, snowball sampling is highly regarded as a multipurpose technique (Atkinson and Flint 2001). Capable of providing detailed, descriptive information, snowball sampling is useful in qualitatively exploring hidden populations (Kaplan, Korf, and Sterk,1987). In addition, when coupled with appropriate estimation methods, this method enables the researcher to make inferences about social networks in areas involving sensitive, deviant, or illegal issues (Atkinson and Flint 2001; Faugier and Sargeant 1997; Kaplan, Korf, and Sterk 1987).

Similar to snowball sampling, network sampling allows researchers to contact a large number of participants who share a set of predetermined characteristics (Bernard et al. 2010; Griffiths et al. 1993). The advantage to using network sampling techniques, though, is that they do not require the respondents to be part of the target population's network (Bernard et al. 2010; Griffiths et al. 1993). This is useful in situations where the researcher's knowledge of the operational structures and social characteristics of the target population is limited, and/or when the researcher does not have any pre-existing contacts with access to the population of interest (Faugier and Sargeant 1997).

Designed for use in the real world, respondent-driven sampling has become recognized as the most viable option for rigorous sampling of hard-to-reach populations (Jonhston and Sabin 2010). Given that RDS relies on respondents to recruit their peers, this method does not require exhaustive mapping to construct sampling frames (Magnani et al. 2010). By circumventing this need for the researcher to spend time looking for named recruitees, RDS effectively minimizes the amount of funding, time, and/or cooperation required by subjects (Salganik and Heckathorn 2004). Related to the reduced onus on subjects, RDS has been successful in obtaining high levels of subject participation (Salganik and Heckathorn 2004). Utilizing incentives such as payment gifts as rewards, and unique coupons to prevent respondents from having to divulge sensitive information, RDS encourages respondents to recruit and new subjects to participate (Johnston & Sabin, 2010; Salganik, and Heckathorn 2004).

Statistically speaking, RDS has been found to produce asymptotically unbiased estimates about hidden populations (Salganik and Heckathorn 2004). This is the result of RDS's long-chain recruitment design. Geared toward expanding the recruitment across many interconnected networks, RDS increases the reach of the sample into less visible pockets of the target population and, at the same time, minimizes the influence of the initial seeds on the final sample's composition (Kendall et al. 2008; Magnani et al. 2010). According to Heckathorn (2002), sample composition becomes independent of the seeds when the sample distribution on key variables remains stable within 2 percent of the equilibrium distribution. Therefore, regardless of how the seeds are selected whether at random or using individuals known to the researcher, RDS will produce unbiased estimates pertaining to the network and overall population (Salganik and Heckathorn 2004). In addition to being used on its own to provide estimates about the target population, the data gathered during the RDS process can also be combined with existing institutional data to estimate the size of the hidden population (Salganik and Heckathorn 2004). In short, by using longer recruitment chains, strict recruitment limits (i.e., a fixed number of unique coupons), and collecting data used to statistically adjust for the biases inherent in how individuals sharing similar characteristics are networked and likely recruit each other, RDS lends statistical rigour to more conventional snowball sampling techniques (Handcock, Gile, and Mar 2012; Kendall et al., 2008; Magnani et al. 2010).

Contrary to the snowball, network, and respondent-driven sampling techniques, the adaptive sampling strategies provide fewer benefits for studying human hidden populations. These methods are most efficient for sampling plant and animal populations that are spatially clustered together (Seber and Salehi 2013; Thompson 2012). In these situations adaptive sampling can be cost effective, as less screening is required in areas containing higher concentrations of members of the rare population (Thompson 2012). Furthermore, when the prevalence of the hidden population is much higher than average in certain geographical strata and these strata contain a substantial proportional of this population, using adaptive sampling strategies will result in notable gains in precision (Thompson 2012).

Disadvantages of Link-Tracing Designs

Link-tracing designs, both as a group and individually, have a number of positive features making them useful for studying hard-to-reach populations. However, there are also a number of drawbacks to using these techniques. First and foremost, these methods can be extremely labour intensive. The researcher must actively develop and control the sample's initiation, progress, and conclusion (Faugier and Sargeant 1997). At the outset, finding respondents and starting referral chains can be difficult. Given the moral, legal and/or social sensitivities surrounding the behaviours in question, many hidden populations have low visibility, and, thus, pose a serious challenge for locating and contacting potential respondents (Faugier and Sargeant 1997; Sadler et al. 2010). In order to identify the initial respondents, then, the researcher must have extensive knowledge pertaining to the distribution of the hidden population (Atkinson and Flint 2001). Further, once the seeds have been found and the sampling process has begun, the researcher needs to control the types of chains and number of cases in any given chain. During this stage, the researcher needs to be able to verify the eligibility of potential respondents and ensure that all sample members are actually members of the target population (Salganik and Heckathorn 2004). Again, due to the social stigma and/or threat of criminal prosecution faced by members of the hidden population, this becomes a difficult task. Respondents may be unwilling to provide reliable and consistent information about their colleagues in the population (Salganik and Heckathorn 2004). Further, there may be certain social and physical barriers, such as ethnicity, occupation, and location of residence that reduce the likelihood that respondents know other people in the hidden population (Bernard et al. 2010). In such cases, respondents would be unable to provide information about specific members in the population. Relying on this restricted information, therefore, may result in missed respondents and/or significant overlap in the final sample (Frank and Snijders 1994). To address this issue, the researcher must employ a non-threatening procedure to keep track of respondents (e.g., the unique coupons used in RDS), as well as consistently monitor the referral chains and data quality (Faugier and Sargeant 1997; Sadler et al. 2010).

In addition to being labour intensive, link-tracing designs have been criticized for their high risk for sampling and selection biases (Heckathorn 2002). For instance, because initial seeds are generally known to researchers rather than being a simple random sample from the population, they cannot be considered to be representative of entire target population; they will have much higher degrees (i.e., a larger network size) than other members of the hidden population (Heckathorn 2002). Unless the level of degree is accounted for, the bias introduced during the seed selection phase may be compounded as the sampling process continues (Salganik and Heckathorn, 2004; van Meter 1990). Currently, RDS is the only technique that produces estimates that account for individuals' degree, and, thus eliminate the problems associated with selection bias (Salganik and Heckathorn 2004). Using snowball sampling, though, this issue can be minimized. For example, by modeling a stochastic process that permits the calculation of sample weights that are inversely proportional to the probabilities of selection, unbiased estimates can be generated (van Meter 1990).

The use of non-random referrals is also problematic. Given that referrals from all respondents are based on their subjective perceptions of others, some individuals will have a greater likelihood of being targeted than others (Atkinson and Flint 2001; Faugier and Sargeant 1997; Heckathorn 2002). For example, subjects tend to refer those with whom they have social ties, such as family, friends, and other associates (Heckathorn 2002). As a result, recruitment patterns will often reflect the recruiter's affiliations. Given that the sample is generally not representative of the entire target population, drawing inferences from the data collected using link-tracing designs can be challenging (Atkinson and Flint 2001; St. Clair and O'Connell 2012; Salganik and Heckathorn 2004). With equal selection probability design-based estimators producing biased results, these situations require the use of specific non-probability estimators (Heckathorn 2002; Salganik and Heckathorn 2004; Snijders 1992). Moreover, even with the appropriate statistical tools, link-tracing designs can only provide estimates about the structure of the network connecting the target population (Snijders 1992). It is not appropriate to use this data to make direct inferences about the population (Salganik and Heckathorn 2004).

Another design flaw concerns the link-tracing methods' over-emphasis on the cohesiveness of social networks (van Meter 1990). Because these techniques only select individuals with inter-connections, members of the population who are highly atomized and isolated will be missed (Salganik and Heckathorn 2004; van Meter 1990). As a result, these samples will be biased toward the inclusion of individuals with many inter-relationships (Faugier and Sargeant 1997; Heckathorn 2002; Salganik and Heckathorn 2004). This is especially the case in snowball samples, which provide researchers with a higher proportion of respondents who share similar characteristics with their initial contact. In such instances, the final sample will be over-represented by the characteristics of those individuals with more social connections and under-represented by the characteristics of sample members with a smaller social network (Johnston and Sabin 2010). Essentially, link-tracing methods are not appropriate for groups who do not have frequent contact with other members of the target population (Salganik and Heckathorn 2004).

In addition to these general limitations, there are a number of disadvantages that are unique to specific link-tracing techniques. For example, because the composition of snowball samples is dependent on the choice of seeds and short recruitment chains, snowball sampling has been criticized for lacking validity in representation (Kendall et al. 2008). In addition, snowball sampling often results in an incomplete sample. With a final wave requiring members to only mention individuals that have been previously mentioned, snowball samples often miss certain segments of the target population (Frank and Snijders 1994).

Network sampling techniques, on the other hand, face challenges because of their assumption that the network of social contacts in the general population is essentially random (Bernard et al. 2010). In many situations this is an unreasonable assumption. For instance, in the context of behaviours that increase the risk of HIV, it is unlikely that these characteristics are randomly distributed in the general population (Bernard et al. 2010).

Issues with respondent-driven sampling relate to its dependence on the community structure of the hidden population and the recruitment procedure (Goel and Salganik 2009). RDS estimates are based on a single, long run of chain. In practice, however, some sample members will not recruit others, and, thus, a chain can terminate (Goel and Salganik 2009). In order to ensure the chain continues, respondents are encouraged to recruit multiple members. This created increased dependence between participants, which, in turn, increases the variance of the RDS estimates (Goel and Salganik 2009).

That the use of adaptive sampling has largely been restricted to studying plant and animal populations (Thompson 2012) largely reflects the fact that there are significant difficulties in using these techniques to examine hidden human populations. The major issue with adaptive sampling is that these methods can only be used in situations where the rare population is geographically clustered (Sudman and Blair 1999). The efficiency of these methods is fully dependent on the rarity of clustering of the population (Magnani et al. 2005). In addition, because these techniques require at least two "passes" over the population area, it is necessary that all group members be accessible during the period of data collection (Magnani et al. 2005; Seber and Salehi 2013). This means that the population must be relatively stable and readily available. When studying human populations, however, this requirement is often unrealistic. Human populations tend to be highly mobile, and, thus, there are likely to be changes in the geographic distribution of these rare populations over time (Magnani et al. 2005; Seber and Salehi 2013; Thompson 2012).

Targeted Sampling

Building off of snowball sampling techniques, researchers have developed a small number of outreach methods to attract a sample of people from a hard-to-reach population (Salganik and Hecakthorn 2004; Watters and Biernacki 1989). Targeted sampling, which is also referred to as "street outreach," usually involves sending fieldworkers out into the streets to locate and recruit members of a hidden population (Acharya 2007). This method allows for sampling to take place in a variety of settings, including both open (e.g., a street corner, public park, bar, etc.) and closed (e.g., institutions) settings (Salganik and Heckathorn 2004).

How It Works

Targeted sampling is a purposeful and systematic method made up of two distinct components. First, targeted sampling requires the researcher to create controlled lists of specified populations within a geographical district (Watters and Biernacki 1989). This is usually accomplished using an initial ethnographic assessment, which is aimed at identifying the various networks and/or subgroups that might exist in a given setting (Magnani et al. 2005). The chosen area will be the venue where the members of the population might be found with high probabilities. Although this may include public places such as bars or parks, target population members are often sampled in an institutional setting. For example, injection drug users may be sampled in the building where they attend their drug rehabilitation program (Salganik and Heckathorn 2004). This provides the researcher with an easily accessible and readily available sample of the target population (Magnani et al. 2005).

The second element involves developing a detailed plan to recruit adequate numbers of cases within the target population and area searched (Watters and Biernacki 1989). In some instances, the identified population members are treated as sampling strata and quota samples are chosen within each stratum using systematic sampling (Salganik and Heckathorn 2004). The most common approach, though, is to ask the members of the target population who belong to the sampled site to nominate other members of the population (Felix-Medina, Monjardin, and Aceves-Castro 2009). With this latter approach, a maximum likelihood estimator of the size of the population can be obtained by calculating the probability that a person is nominated by any person in a particular sampled site. This nomination probability is either homogeneous (i.e., the probability depends on the site), or heterogeneous (i.e., the probability depends on the person being nominated) (Felix-Medina, Monjardin , and Aceves-Castro 2009).

Populations Studied Using Targeted Sampling

Although recognized as an alternative to link-tracing designs, targeted sampling has not been widely used to study hidden human populations. Thus far, its utility has largely been limited to HIV behavioural and biological surveillance research (Watters and Biernacki 1989). Watters and Biernacki (1989) developed the targeted sampling technique as a scientific method to monitor the spread of the HIV infection among drug users and their sexual partners in the San Francisco area. To develop their sample, the researchers used four stages: (1) initial mapping, (2) ethnographic mapping, (3) developing an initial target plan for each district, and (4) sampling and revising the target plan. In the initial mapping phase, the researchers defined the districts in which to conduct the research. This was accomplished using geographical maps and directly observing various city neighbourhoods for easily identifiable signs of drug transactions and drug use (especially intravenous drug use of heroin, cocaine, and methamphetamine). The neighbourhoods with the highest concentrations of drug use, the Tenderloin and Mission districts, were used for sampling (Watters and Biernacki 1989).

Once the study areas were finalized, the ethnographic mapping was used to uncover and analyze the social organization of target groups existing within the selected districts. This mapping generated social networks characterized in terms of drug use profiles, social customs, sexual relationships, geographic locations of groups and members, etc. (Watters and Biernacki 1989). The major subgroups found were black, male heterosexual heroin addicts, white methamphetamine-injecting transsexuals, and drag queens.

After identifying the major sub-groups within each district, targets were identified, and recruitment strategies were developed. The researchers used an outreach worker from the AIDS prevention program to access the target population. Once the target population was made aware of the study, the researchers relied more on a passive referral mechanism to obtain new sample members (Watters and Biernacki 1989).

Advantages and Disadvantages of Targeted Sampling

Despite its limited use, there are some advantages to using targeted sampling to study hidden populations. Being a flexible and interactive technique, targeted sampling may be a useful method for sampling hidden populations in urban settings (Watters and Biernacki 1989). It can provide researchers with valuable information pertaining to the social, behavioural, and cultural characteristics of different subgroups, as well as the geographic distribution of the hidden population (Salganik and Heckathorn 2004). Overall, this method offers a cohesive and systematic approach that can help researchers study health or social problems that exist among members of hidden populations (Watters and Biernacki 1989).

Although it has potential, targeted sampling's actual utility is quite limited. The time and resources required to develop and implement a targeted sampling strategy, make it an extremely labour intensive technique (Acharya 2007; Magnani et al. 2005). In addition, because it is often difficult to obtain samples in public places, this technique often does not provide access to large samples of non-institutional members of a population (Acharya 2007). Moreover, there are a number of sources of bias in samples resulting from targeted sampling. The first source of bias stems from the fact that this method does not obtain random samples. Members of the population who congregate in specific places, including institutional settings, represent non-random segments of the target population (Acharya 2007; Magnani et al. 2005). Not having an equal probability of selection introduces selection bias. This means that accurate estimates cannot be obtained and the results cannot be generalized to the entire target population (Salganik and Heckathorn 2004). Adding to this selection bias are the problems associated with the location and time of day utilized during the sampling procedure. The sampling location can have a large impact on the composition of the resultant sample. Given that different segments of the population may gather in different settings, the sample will be biased towards those who congregate in the chosen sampling location (Watters and Biernacki 1989). For instance, if the target location is a public park, then all population members who do not congregate in this area are certain to be missed in the sampling process (Salganik and Heckathorn 2004). In addition, it is unlikely that male drug users will be found in women's treatment facilities (Watters and Biernacki 1989). Similarly, the time of day can impact who is obtained in the sample. For instance, if data collection occurs during the day, the sample will not include those population members who are active at night. This may be the case with street workers and drug users (Watters and Biernacki 1989).

Time-Location Sampling (TLS)

Time-location sampling, often referred to as time-space sampling, is a more refined alternative to targeted sampling. It is a method used to sample individuals who visit specific locations, including museums, libraries, and polling places at known times (Kalton 2001). Conceptually, TLS works by creating a comprehensive map of the areas and time periods where the target population can be found, and then randomly selecting specific venue-day-time (VDT) periods to conduct recruitment (Ferreira et al. 2008). To accomplish this, TLS uses a two-stage design. First, primary sampling units are constructed as combinations of locations and time segments when the location is open. These become the units that are sampled with probabilities proportional to size, with careful stratification by location and time (Kalton 2001; Sudman and Blair 1999). Second, a form of systematic sampling is employed to select the visitors entering or exiting the location (Kalton 2001; Karon and Wejnert 2012; Magnani et al. 2005).

How It Works

Although conceptually straightforward, a more concrete explanation of this technique is required. The first step for TLS involves ethnographic fieldwork (Magnani et al. 2005). At this stage, the researcher surveys specific geographic areas to construct a sampling frame; this identifies the members of the target population, as well as when they gather at certain locations. After the ethnographic fieldwork has been exhausted, the specific venue-day-time segments are chosen (Magnani et al. 2005). These become the primary sampling units. It is important to note that these units are randomly selected (Ferreira et al. 2008). Once the sampling units are determined, the last stage involves sampling the individuals who enter or leave the selected location (Kalton 2001; Magnani et al. 2005).

A variety of statistical methods, such as cluster and regression analyses, can be used with the data obtained from TLS samples to provide estimates about the hidden population (Valleroy et al. 2000). To obtain accurate estimates, however, selection probability weights must be used. These weights control for each participant's probability of being sampled (i.e., the number of individuals recruited compared to the overall attendance of the venue at the time of recruitment) (Ferreira et al. 2008; Valleroy et al. 2000).

Populations Studied Using Time-Location Sampling

Similar to targeted sampling, TLS has primarily been used to survey HIV risk and illness in the specific locations where the affected individuals congregate (e.g., Ferreira et al. 2008; Karon and Wejnert 2012; Valleroy et al. 2000). For example, Valleroy and colleagues (2000) used The Young Men's Survey, which was conducted in seven U.S. cities between 1994 and 1998, to obtain estimates of the prevalence of HIV infection and associated risks in the population of adolescent and young adult men who have sex with men (MSM) in the United States. In this study, the sample, which consisted of 3,492 MSM's aged fifteen to twenty-two years old, was created using a venue-based sampling procedure. The procedure involved first identifying the public venues frequented by young MSMs. These included street locations, dance clubs, bars, businesses, social organizations, and health clubs (Valleroy et al. 2000). Next, the researchers determined the days of the week and times of day when these venues were most saturated. In each location, one four-hour period was chosen for sampling. During this period, young men were approached, and, if the researcher found them to be eligible, they were included in the sample and asked to consent to an interview and HIV testing (Valleroy et al. 2000). Using the data collected during this process, the researchers were able to determine that the overall prevalence of HIV infection was high (7.2 percent).

Ferreira and his co-researchers (2008) used the TLS technique to assess HIV related risk behaviour in male truck drivers at the crossroads of two of the major highways used for trucking in northeast Brazil. By mapping the venues where truck drivers congregate (i.e., truck stops), as well as the times of day these individuals could be found to frequent these locations, the researchers were able to obtain a sample of 686 male truck drivers. Within this sample, it was determined that 21.3 percent had sex while on the road, 12.3 percent had sex with commercial sex workers, and 1.8 percent had sex with hitch-hikers (Ferreira et al. 2008). Essentially, the data obtained from this sample allowed Ferreira and colleagues (2008) to evaluate the extent to which these truck drivers engage in behaviours known to increase the risk for HIV infection.

Advantages of Time-Location Sampling

Evidenced by its utility in surveying HIV risk, time-location sampling offers an efficient way to systematically sample hidden populations that congregate in specific locations (Ferreira et al. 2008; Valleroy et al. 2000). This holds true for highly mobile populations as well (Ferreira et al. 2008). Further, by mapping the universe of venues where the target population can be found to frequent, randomly selecting the day, time and location for recruitment, and systematically selecting participants from the venue, TLS is able to closely approximate probability sampling (Kendall et al. 2008). More specifically, by using this randomization process, the probability of selection can be determined (Ferreira et al. 2008). This probability is calculated based on the expected sample yield at the location and number of members of the target population that are intercepted and interviewed (Kalton 2001; Sudman and Blair 1999). In addition to reducing selection bias, because the place, day, and time are sampled with a known probability, time-location samples can be used to make statistical inferences about the hidden population that attends the identified locations (Acharya 2007; Salganik and Heckathorn 2004).

Disadvantages of Time-Location Sampling

Although TLS offers a number of advantages for studying geographically clustered hidden populations, it also has a number of limitations. For instance, TLS's reliance on specific locations to conduct all sampling can be problematic. Not all places are easily accessible. For example, private venues, including homes, are often not accessible (Kanouse et al. 1999). Furthermore, when there are safety concerns in a particular location, the venues in this area are usually not selected (Salganik and Heckathorn 2004). Due to the prohibitive cost, venues with low expected sample yields are also not sampled. The choice of venue, therefore, is often limited (Salganik and Heckathorn, 2004).

A related problem concerns the samples that are collected at the chosen locations. The sample will only include the population that frequents the selected venues (Kendall et al. 2008; Magnani et al. 2010). At issue is the fact that the place-attending population may differ from the true population in a variety of unknown ways (Salganik and Heckathorn 2004). For instance, drug injection users who appear in public places may differ from those who only frequent private venues. This means that there may be an unknown, and potentially substantial, amount of bias present in the estimates (Salganik and Heckathorn 2004). In addition to missing segments of the target population, the sample may be compromised by multiplicities. Individuals will often make multiple visits to the target location during the survey period. Thus, there are increased selection probabilities associated with multiple visits that must be taken into consideration when developing the probability sample weights (Kalton 2001). Further, because respondents are often unable to accurately recall past visits, estimating these multiplicities can be challenging (Kalton 2001).

Capture-Recapture Method

Originally developed in biology, capture-recapture designs are increasingly being used in studies of human populations to provide estimates of the size of hidden populations (Bouchard 2007). At its most basic level, capture-recapture works by capturing, counting, and marking a sample, and then re-capturing that same sample at a later point and noting the overlap between these two captured samples (Hope, Hickman, and Tilling 2005). By analyzing the overlap between these two samples, the researcher is able to estimate the probability of detection (Hope, Hickman, and Tilling 2005; Thompson 2012). Essentially, capture-recapture is an indirect method that generates prevalence estimates based on the degree of overlap between two or more separate samples in hidden populations (Choi and Comiskey 2003; Comiskey and Barry 2001).

How It Works

Capture-recapture works by identifying a relatively important recurring pattern in an observed population (e.g., re-arrest), and using this to infer the proportion of this population that is active, but unobserved in the data (Bouchard 2007). This works by equating the unknown proportion of the population recorded in a list to the known proportion of one or more subpopulations also recorded in that list (Cormack 1999). In its most rudimentary form, the process for using capture-recapture data to estimate the total number of individuals in a population is as follows. First, an initial sample is obtained, and the individuals in that sample are marked (Thompson 2012). Independent of this first sample, a second sample is then obtained and it is noted how many the marked individuals appear in this sample. If this second sample is representative of the population as a whole, the sample proportion of marked individuals should be approximately the same as the proportion of marked individuals in the population (Thompson 2012). It is from this relationship that the total number of individuals in the population can be estimated.

For estimating hidden human populations specifically, the process requires comparing two or more lists for overlap or by gathering original data in the population through the identification of individuals (Larson, Stevens, and Wardlaw 1994; Williams and Cheal 2002; Thompson 2012). These observations can be made simultaneously using two sources that represent approximately the same population, or they can come from the same source at two different time points. Most often, capture-recapture with human populations will involve collating information on individuals who have come into contact with various agencies or data sources (Larson, Stevens, and Wardlaw 1994; Williams and Cheal 2002; Thompson 2012). For example, while one list may be obtained from the census data, a second list may be derived from a follow-up survey (Hope, Hickman, and Tilling 2005).

Although capture-recapture always operates using this basic procedure, this method can be used in a variety of different forms. Single mark-release is the standard capture-recapture method; it relies on two lists and uses a single sampling occasion (Thompson 2012). In situations where multiple sampling occasions are desired or required, the researcher must use the several mark-release method. Additional capture-recapture variations include mark-recapture with closed (i.e., no change in the population during the study period) or open (i.e., allows for births, deaths, immigration, and emigration) populations, and covariate capture-recapture, which adds covariates into the model to adjust for population heterogeneity (Hope, Hickman, and Tilling 2005; Thompson 2012).

Once the form of capture-recapture has been chosen and the data collected, it is possible to fit statistical models to the observed data to generate estimates of the unobserved population (Hope, Hickman, and Tilling, 2005). In order to obtain the best estimate of the size of the hidden population, however, the correct statistical techniques must be used. Specifically, the estimation method must be appropriate for the type of data (i.e., it must be compatible with the type of capture-recapture method chosen) (Thompson 2012). There are a variety of different statistical methods available. The most commonly used, and basic technique is the ratio estimation method. Creating a known to unknown ratio of the population, this method works best with the capture-recapture methods that use only two data sources (i.e., standard capture-recapture approaches) (Hay and Gannon 2006). In order to use this method to estimate the size of the population, the researcher needs to know the number of persons observed at the first count, the number of persons observed at the second count, and the number of persons observed at both counts (Hay and Gannon, 2006). The accuracy of the estimates produced by this method is completely dependent on whether or not the three assumptions for this model were met. These assumptions are: (1) behaviourally homogenous members, (2) a constant encounter rate, and (3) independence between samples (Hays and Gannon 2006). While the first assumption is rather straightforward, the second two require further explanation.

A constant encounter rate means that the probability for an individual to be observed and re-observed must be held constant during the observation period (Bouchard 2007; Larson, Stevens, and Wardlaw 1994). Essentially, this means that the overall numbers in the population should not be different at the time each sample is taken (Choi and Comiskey 2003; Hay and Gannon 2006). Meeting this assumption often requires that the population under study is closed–for example, there is no movement in or out of the population in the period over which the population is being studied (Choi and Comiskey 2003; Hay and Gannon 2006; Larson, Stevens, and Wardlaw 1990; Rossmo and Routledge 1990). Given that humans tend to be highly mobile, dynamic, and change their behavioural patterns, this assumption is often violated when studying hidden human populations. As such, this is regarded as the least stringent of the three assumptions (Hay and Gannon 2006).

The final assumption, independence between samples, is the most crucial. It means that an observation in one sample should not have any effect on the observation of an individual in a second, or subsequent samples (Hay and Gannon 2006). More specifically, this assumption mandates that, at any point in time, the number of encounters with an individual does not depend on the number of encounters already recorded in the data (Rossmo and Routledge 1990). Meeting this assumption requires researchers to account for all individuals who are present in more than one source (Hay and Gannon 2006). This can be accomplished by creating lists in a systematic and consistent manner (Larson, Stevens, and Wardlaw 1994).

In situations where the data do not meet the assumptions for the ratio estimation method, other statistical methods can be used to obtain estimates of the size of the hidden population. For instance, truncated Poisson models–also known as homogeneous Poisson models and Zelterman's truncated Poisson estimator–are useful in situations where the researcher has only a single data source. Based on the idea that the population consists of some unknown number of individuals that can be filled by equivalently active, known individuals, this method provides an estimate of the total size of the population equalling the number of individuals ever observed plus the number of individuals never observed in the data source (Roberts and Brewer 2006; Rossmo and Routledge 1990). One caveat to using this method is that it assumes the hidden population of interest is closed. Therefore, it cannot account for changes in the population (Bouchard 2007; Collins and Wilson 1990).

Finally, in instances where the researchers have multiple samples of three or more, researchers can choose between a number of statistical methods. For instance, the researcher could opt to use loglinear modeling. This method works by expressing the logarithm of the expectation of the counts in different combinations of lists as sums and differences of parameters (Cormack 1999). Depending on their samples, the researcher could also select one of the varieties of Poisson models that could include the heterogeneous Poisson model and Poisson-based regression model. While the heterogeneous model provides a mixture of two Poisson populations, the Poisson-based regression models consider a series of covariates in fitting an estimation curve (Collins abd Wilson 1990; Hope, Hickman, and Tilling 2005).

Populations Studied Using Capture-Recapture

Capture-recapture methods have been used extensively in the field of substance abuse to estimate the prevalence of drug users in a variety of communities worldwide (e.g., Calkins and Aktan 2000; Choi and Comiskey 2003; Comiskey and Barry 2001; Domingo et al. 1995; Hay and Gannon 2006; Hope, Hickman, and Tilling 2005; Hickman et al. 1999; Hutchinson, Bird, Taylor, and Goldberg 2006; Hser 1993; Larson, Stevens, and Wardlaw 1994; Mastro, Kitayaporn, Weniger, and Vanichseni 1994). For example, Hay and Gannon (2006) utilized a capture-recapture method to estimate the prevalence of illicit opiate use for thirty-two local government areas in Scotland. Using data from the Scottish Drug Misuse Database pertaining to a sample of 22,795 individual drug users, the researchers estimated the prevalence of the problem drug use and determined that this value corresponded to an estimated hidden population number representing 2 percent of the general population in Scotland between the ages of 15 and 54 (Hay and Gannon 2006).

Using a different variation of the capture-recapture method, Hope, Hickman, and Tilling (2005) estimated the prevalence of crack-cocaine use in twelve London boroughs between 2000 and 2001. In this case, the researchers applied the covariate capture-recapture technique to three data sources containing information on subjects reporting crack-cocaine use. To add an element of control to the model, the individuals in the sample were matched based on three covariates: (1) age group (fifteen to twenty-nine years and thirty to forty-four years), (2) gender, and (3) opiate use. Using a sample of 4117 individuals aged 15-44, the model estimated that, out of a total of 21,000 crack-cocaine users, 16,855 were unobserved (i.e., part of a hidden population). This represented a prevalence of 1.5 percent of crack-cocaine use (Hope, Hickman, and Tilling 2005).

This work has also been expanded to include populations of drug dealers (e.g., Bouchard and Tremblay 2005) and cultivators (e.g., Bouchard 2007). Bouchard (2007) estimated the size of the marijuana cultivation industry in Quebec, Canada using a truncated Poisson estimator capture-recapture model. Using arrest data, co-offending patterns, and seizure data collected from a sample of 10,204 different offenders, Bouchard (2007) was able to assess the risk of arrest for different types of marijuana growers, as well as the risk of detection for the cultivation sites. According to his findings, by combining the total number of growers arrested and the proportion of those offenders re-arrested for a specific offence, it could be determined that there were almost twice as many soil-based growers as compared to hydroponic growers (Bouchard 2007). In addition, this model established that hydroponic growers were at a lower risk of arrest. Lastly, by looking at risk of detection, Bouchard (2007) was able to estimate the prevalence of cultivation sites. Specifically, it was determined that the risk of detection increased with the size of the cultivation site (Bouchard 2007).

Moving away from populations of substance abusers, researchers have begun to use capture-recapture methods to estimate the prevalence of other criminal and non-criminal populations. For instance, in addition to analyzing general populations of offenders (e.g., Greene and Stollmack 1981), capture-recapture methods have been used to estimate criminal populations, including burglars (Riccio and Finkelstein 1985), car thieves (Collins and Wilson 1990), migrating fugitives (Rossmo and Routlege 1990), and prostitutes (Rossmo and Routlege 1990) and their clientele (Roberts and Brewer 2006). Moving away from the sphere of crime, this method has been used to estimate the prevalence of homelessness in specific locations (e.g., Williams and Cheal 2002). For instance, using a two-sample approach by taking two samples from the same location, Williams and Cheal (2002) were able to estimate the mean number of homeless people in Plymouth at 446 and Torbay at 325.

Advantages of Capture-Recapture

Studies using the capture-recapture methods highlight a number of strengths for using this approach to study hidden populations. First, because this method is not highly technical, it is appealing and accessible to researchers who are not capture-recapture specialists (Roberts and Brewer 2006). In addition to be relatively easy to use, research has shown that capture-recapture methods provide valid measures for various hidden populations (Bouchard 2007). For instance, by allowing researchers to use criminal population estimates as a relevant denominator to calculate different types of arrest risk, capture-recapture methods can be used to assess the number of arrests for the number of offenders at risk (Bouchard 2007). Moreover, it can be used to assess detection. This is beneficial in situations where the individual or activity of interest is unlikely to be formally detected (Bouchard 2007). Capture-recapture methods can also be used to actually calculate the odds of capture (Bouchard and Tremblay 2005). For example, odds of arrest can be derived by defining a population of "susceptibles" (i.e., a population of individuals involved in a particular line of activity, such as prostitutes) who have yet to be arrested but who nevertheless have the same characteristics as those who have been arrested (Bouchard and Tremblay 2005). The odds of arrest (i.e., the number of arrests given the number of susceptibles) estimated by capture-recapture techniques work by providing an indication of the likelihood that the pool of ready-to-act offenders participating in a given criminal activity will be arrested for a related offence (Bouchard and Tremblay 2005).

In addition to providing reliable measures, capture-recapture methods enable researchers to uncover important macro-level patterns. For instance, capture-recapture methods can be used to measure prevalence in the general population, and these values can be compared with survey and other qualitative approaches to assess the prevalence of certain behaviours and characteristics (Roberts and Brewer 2006). Allowing for an integrative approach, capture-recapture data can also be combined with survey or qualitative research to test assumptions about the population (Williams and Cheal 2002).

Another benefit of capture-recapture methods is that they offer a means to overcome the problems of undercounting and double counting that tend to plague other techniques (Williams and Cheal 2002; Roberts and Brewer 2006). For instance, with covariate capture-recapture the population can be stratified to account for variability within the population (Hay and Gannon 2006). Moreover, by using multiple agencies and data sources, capture-recapture methods make it less likely that people would be unaccounted for (Williams and Cheal 2002).

Capture-recapture methods are also regarded as being robust methods for counting hidden populations (Rossmo and Routledge 1990). Because data is collected over time, capture-recapture methods allow researchers to take into account processes that include changes in population (Rossmo and Routledge 1990). In addition, it also affords individuals sufficient time to be observed and re-observed while actively engaging in the activity of interest (Bouchard 2007). Moreover, capture-recapture methods can be modified to calculate different probabilities for capture and re-capture (Roberts and Brewer 2006). For example, according to Roberts and Brewer (2006), by employing a deterrence or escalation effect, a re-arrest probability can differ from an initial arrest probability. By using a moving average, capture-recapture methods can also estimate populations that have slow and small movements of individuals in and out of the sample (Bouchard 2007). This specific element allows this method to capitalize on the mobility that members of human hidden populations have within their own social space. For example, injection drug users seeking treatment are known to move freely between service providers (Larson, Stevens, and Wardlaw 1994).

Disadvantages of Capture-Recapture H2

Like all techniques used to study hidden populations, there are a number of disadvantages to using capture-recapture methods. Despite being useful for studying a variety of criminal and non-criminal hidden populations, capture-recapture methods have limited utility in situations where the population of interest contains a relatively large number of members who have a very small expected encounter rate (Rossmo and Routledge 1990). In addition, it is difficult to use these methods in areas where few agencies would be in contact with the population, such as those living in rural areas (Hay and Gannon 2006). As most capture-recapture methods require two or more valid, representative, and independent samples of a population, it is hard to implement these techniques in regions where data sources are limited (Bernard et al. 2010; Choi and Comiskey 2003).

Another major area of concern is recording. The validity and reliability of capture-recapture methods rest on how consistently and accurately the observations (i.e., tagging) are made (Cormack 1999). Because humans cannot be physically tagged, researchers need to use other methods to uniquely identify which individuals were recruited into more than one sample (Bernard et al. 2010; Williamd and Cheal 2002). Specifically, these chosen markers must be systematically applied and recorded to enable researchers to accurately match individuals between lists (Cormack 1999). To ensure the recording of this data is accurate, the population must be defined in the same way at all times (Cormack 1999; Larson, Stevens, and Wardlaw 1994). However, because there are different ways to record information, there are often variations and inconsistencies in source data across different agencies. This is even more of an issue when the data being used originate in different countries (Hay and Gannon 2006). In situations where the data have been inaccurately recorded, there may be misclassifications of observations. This can lead to false positives and/or negatives, which means there will likely be data overlap (Mastro et al. 1994).

Related to these recording issues are concerns regarding data sources. The range of individuals obtained via the capture-recapture process depends entirely on the range of population members appearing in the official data. As a result, the population estimates will only reflect the underlying population whose members are likely to be part of the sample. To ensure that the target population is adequately represented, therefore, the data sources must be as complete and accurate as possible (Collins and Wilson 1990; Larson, Stevens, and Wardlaw 1994). Many external data sources, however, may be poor indicators of overall incidence of certain offences. For example, police data generally underestimates the true incidence of crime (Collins and Wilson 1990).

There are also a number of methodological issues that require consideration. Specifically, it is often difficult to meet the assumptions of this technique with many hidden human populations. For instance, as these populations tend to be very diverse in terms of lifestyle, the no heterogeneity assumption can be highly problematic (Hay and Gannon 2006; Roberts and Brewer 2006; Rossmo and Routlege 1990). In addition, it is often difficult to meet the assumption of a constant encounter rate. Specifically, obtaining a truly closed population, which requires no immigration, emigration, births, or deaths, presents a challenge for populations that are mobile and engage in risky behaviours (Collins and Wilson 1990). For example, because no city is ever truly closed, neighbourhoods will likely experience varying levels of inflow and outflow of target population members (Collins and Wilson 1990; Mastro et al. 1994; Williams and Cheal 2002). Moreover, individuals within any given population will likely modify their behaviour over time. In criminal populations, for instance, while some offenders will consistently go in and out of offending at different points in time, others will retire or mature out of crime (Collins and Wilson 1990; Roberts and Brewer 2006). It is almost certain, therefore, that different individuals engage in the behaviour of interest with different frequency and pattern. Thus, the probability of capture is likely to be different from that associated with recapture (Mastro et al. 1994).

The independence assumption is also difficult to satisfy. There are many different factors that may influence the chances of an individual appearing in one or more sources, including, for example, the type of area where they live, the provision of services in the area, and the individual's age or gender (Hay and Gannon 2006). Eliminating the possibility that the occurrence of one event will not affect the probability of a subsequent event is, therefore, challenging (Williams and Cheal 2002). It often requires the use of three or more overlapping samples to produce robust results free of dependence (Larson, Stevens, and Wardlaw 1994).

Lastly, because capture-recapture is an observationally inductive method, it is often not possible to verify if the quality of the observed data justifies the assumptions underlying the statistical model (Hay and Gannon 2006; Williams and Cheal 2002). There is no independent means to assess how good the data actually is. This makes it difficult to formally test whether or not a person in the original sample would still satisfy the criteria for inclusion at a later date (Hay and Gannon 2006). As a result, researchers are often not able to validate the parametric estimates achieved using capture-recapture techniques. Essentially, there is currently no way to determine whether or not the population estimates are actually on the mark (Bouchard and Tremblay 2005).

Fieldwork Method

How It Works

The fieldwork method is a relatively new variation of the capture-recapture technique (Berk, Kriegler, and Ylvisaker 2008; Bloor, Leyland, Barnard, and McKeganey 1991). Also referred to as street counts, fieldwork involves compiling a list of the areas that contain the population of interest. Lists can include public places, such as bars and health clubs, institutional settings, treatment centres, private dwellings, and temporary residential facilities (Berk 2007). Once the lists are completed, enumerators need to be sent out to each sample area to identify and count the number of individuals in the population of interest. For instance, within each sampled tract, the number of homeless individuals must be located and counted (Berk 2007). During this sampling process, the researcher must be able to differentiate between new and repeat fieldwork contacts (Berk, Kriegler, and Ylvisaker 2008). This is the crux of the method, as the increasing ratio of repeat to new contacts will form the basis for modeling the total population (Bloor, Leyland, Barnard, and McKeganey 1991). Specifically, the size of the overall population is modeled from records of the increasing ration of new to repeat fieldwork contacts (Bloor, Leyland, Barnard, and McKeganey 1991).

Populations Studied Using the Fieldwork Method

The fieldwork method is relatively new, and, thus, its use as a means to study hidden populations has been rather limited. The few studies that have capitalized on this technique have used it to provide estimates of populations of homeless individuals (Berk, Kriegler, and Ylvisaker 2008), prostitutes, and drug users (Bloor, Leyland, Barnard, and McKeganey 1991). In their study, Berk, Kriegler, and Ylvisaker (2008) used census tracks within the boundaries of Los Angeles County to count the homeless in this area. The 211 tracts anticipated to have large numbers of homeless individuals were selected, as well as a stratified random sampling of an additional

299 census tracts (Berk, Kriegler, and Ylvisaker 2008). Fieldwork was then undertaken, whereby enumerators were paired with homeless individuals who were hired to act as guides. Street counts were provided for each of the selected and sampled tracts, and shelter counts were added to contribute to the total count of homeless persons (Berk, Kriegler, and Ylvisaker 2008). Using this data, Berk, Kriegler, and Ylvisaker (2008) estimated that there were nearly 65,000 homeless individuals in the entire Los Angeles County.

Bloor and colleagues (1991) used the fieldwork method to estimate the prevalence of female street-working prostitution in Glasgow, as well as the proportion of these prostitutes who were injecting drugs. Targeting Glasgow's red light district, the researchers informally interviewed 208 female street-working prostitutes over the course of thirty-two fieldwork visits. During each visit, a record was made as to whether each woman contacted was a new or repeat contact, and whether or not she was an injecting drug user (Bloor, Leyland, Barnard, and McKeganey 1991). To model the behaviour of the target population, three assumptions were made: (1) the two groups of injecting and non-injecting prostitutes are independent; (2) the population of each group is fixed; and (3) the observations are independent meaning the behaviour patterns do no vary for certain days of the week, or times of night (Bloor, Leyland, Barnard, and McKeganey 1991). According to their results, by distinguishing between new and repeat fieldwork contacts, the contactable population of Glasgow street prostitutes during the period of observation could be estimated to be 304, with 172 being drug-injectors, and 132 non-injectors (Bloor, Leyland, Barnard, and McKeganey 1991).

Advantages and Disadvantages of the Fieldwork Method

Given that the fieldwork method has, thus far, received little attention in the research realm, the strengths and limitations of this approach for studying hidden populations are not fully known. At present, it is clear that the fieldwork method may be an economical alternative to other techniques (Berk, Kriegler, and Ylvisaker 2008; Bloor, Leyland, Barnard, and McKeganey 1991). Because it can be easily combined with existing outreach service programs, it reduces the number of resources required to obtain a sufficient sample of the target population (Berk, Kriegler, and Ylvisaker 2008; Bloor, Leyland, Barnard, and McKeganey 1991). Further, by using information gathered directly from the population of interest, this method does rely on the partial population information provided by official data sources (Berk 2007). However, this reliance on the population to provide the data comes at a cost. In order to build up a sufficient record of new and repeat contacts, this technique requires researchers to have a priori information pertaining to the range of sites where members of the population gather. In addition, conducting informal interviews with population members necessitates the use of strict time-sampling procedures (Bloor, Leyland, Barnard, and McKeganey 1991). This method, therefore, is not very flexible. Taken together, although useful in providing estimations of the prevalence of certain risky behaviours and hidden populations, this method's utility is restricted to populations with known geographical distribution (Bloor, Leyland, Barnard, and McKeganey 1991; Berk 2007).

Multiplier Method

How It Works

Another capture-recapture variant is the multiplier method. Conceptually, the multiplier method seeks to estimate an unknown population from a known population that is more reliable or easily estimated by using the information on the ratio of the known to unknown population (Hope, Hickman, and Tilling 2005; Medhi, Mahanta, Akoijam, and Adhikary 2012). This is accomplished by using information from two independent data sources that overlap in a known way. The first data source is usually an institution or intervention program. Representing the population that the target population would come into contact with, this becomes the benchmark population (Medhi, Mahanta, Akoijam, and Adhikary 2012). The second source of data is a survey of the target population itself. The members of the target population must all be included in both data sources. After the data sources have been acquired, estimates can be derived by multiplying the numbers of persons who officially attended or accessed selected institutes or services over a specific time frame by the inverse of the proportion of the population who claim to have attended to accessed these same services over the same time period (Hope, Hickman, and Tilling

2005; Medhi, Mahanta, Akoijam, and Adhikary 2012). It is assumed that the estimates of the proportions will be unbiased (Hope, Hickman, and Tilling 2005).

Essentially, there are two necessary elements in this method. First, there must be reliable benchmark data representing the known portion of the target population who are part of the first data source. Second, there must be a multiplier that tells the researcher how many more members of the target population are not part of the first data source (Medhi, Mahanta, Akoijam, and Adhikary 2012).

Populations Studied Using the Multiplier Method

Although used less frequently than capture-recapture methods, the multiplier method has been applied in a number of instances to provide estimates of the size of hidden populations. Medhi and colleagues (2012), for instance, used this technique to estimate the size of the intravenous drug user (IDU) population in five districts in India. The sample consisted of males aged eighteen years and older who belonged to a specific district and who had injected drugs for non-medical purposes at least once in the last six months. The benchmark data for each district was the number of IDUs receiving services from different on-going rapid intervention and care (RIAC) centres. The second source of data consisted of the separate anonymous interviews of the IDUs. Applying a multiplier, Medhi and colleagues (2012) found that the estimated sizes of the IDU population in the five districts were 7353 in Imphal West, 5806 in Imphal East, 3816 in Thoubal, 2615 in Churachandpur, and 2137 in Bishenpur.

Kinnell (1988; quoted in Medhi, Mahanta, Akoijam, and Adhikary 2012) used the multiplier method to look at female sex workers, their clients, and risk of HIV infection in Birmingham. Her approach consisted of asking all women whom she contacted (N = 258) whether or not they had been arrested for a prostitution-related offence within the past year. To create a multiplier, she calculated the ratio between the number of her contacts who had been arrested and the number of arrests recorded by the local police (Kinnell 1988, as quoted in Medhi, Mahanta, Akoijam, and Adhikary 2012). According to the multiplier, the original 258 contacts could be multiplied by a figure of up to 1200 female sex workers. After adjusting for differences in potential arrest rates between regular and occasional street workers and non-street workers engaging in prostitution, the final prevalence of sex workers was estimated to be 2000 women (Kinnell 1988, quoted in Medhi, Mahanta, Akoijam, and Adhikary 2012).

Finally, the multiplier method has been successfully used to estimate the prevalence of homeless persons (e.g., James 1991; Hudson 1997). In using this method, James (1991) was able to determine that if local administrators of soup kitchens estimate that they are able to serve only 25 percent of the homeless population in their given catchment area, a count of the number of meals served in a given day can be multiplied by four to obtain the estimate of the number of homeless individuals in a the given catchment area on that day in question. Hudson (1997), on the other hand, used counts of the number of homeless persons in specific census tracts and regressed them on known features of these tracts. The resulting regression coefficients were then used as multipliers. For example, for every 10 percent increase in the number of vacant dwellings in a census tract, Hudson (1997) determined that the number of homeless in that tract would increase by eight individuals.

Advantages and Disadvantages of the Multiplier Method

The appeal of the multiplier approach stems from the fact that it is a mathematically simple and straightforward approach (Medhi, Mahanta, Akoijam, and Adhikary 2012). By allowing for captures to be obtained via enumerations or convenience samples, the multiplier method moves away from strict sampling designs (Handcock, Gile, and Mar 2012) which adds to the accessibility of this method. Despite these advantageous features, however, researchers must use caution when applying this approach.

The whole method depends on the accuracy of the multiplier, which, in turn, varies with the quality of the captures data (James 1991). High quality record keeping is essential for benchmark data to provide an accurate estimation of population size. The researcher must try to eliminate data duplication and the inclusion of people outside of the target population. (Medhi, Mahanta, Akoijam, and Adhikary 2012). Moreover, in addition to obtaining accurate official data, the researcher must collect reliable information about exposure to intervention or changes in behaviour from respondents. For instance, they need to be able to ensure that the target population is not providing false information about names, addresses, and other pertinent factors (Berk 2007; Medhi, Mahanta, Akoijam, and Adhikary 2012). This can be challenging. Not only may there be issues with respondents' recall bias, some of the self-reported information may be subject to the interviewer's bias or socially desirable responses (Medhi, Mahanta, Akoijam, and Adhikary 2012).

Another issue concerns the method's requirement that the relationship between the two data sources is constant. This means that the model must be stable over time (Berk 2007). In order to meet this criterion, the time reference period must be clear and the same in both data sources. Specifically, both the target population and benchmark data must represent the same constraints. Interviewing IDUs about their drug use in the last six months, for example, would also require the benchmark data consider usage during this same six month time period (Medhi, Mahanta, Akoijam, and Adhikary 2012).

Finally, the multiplier method has been criticized for its inability to produce estimates that are generalizable to the complete target population (Berk 2007). Although the use of mapping as a means to obtain respondents helps to prevent self-selection biases or volunteerism on the part of respondents, it does not guarantee that the entire target population has the opportunity to be sampled (Medhi, Mahanta, Akoijam, and Adhikary 2012). Sampling the whole population would require that the bulk of this population congregates in an identifiable location, which is rarely the case with hidden populations. The resulting estimates, therefore, cannot be applied to geographical areas beyond those in which the model was built (Berk 2007).

General Issues

As evidenced from the above discussion, all of the methods currently available for studying hidden populations have their own strengths and limitations. In addition to these individual advantages and disadvantages, there are four general problems that characterize this entire group of techniques. The first issue involves defining the problem, and, thus, the population of interest (Williams and Cheal 2002). As many issues are dynamic, people can move in and out of situations over time. Regarding homelessness, for instance, a person may be homeless at one point but not at another (Williams and Cheal 2002). Identifying stable characteristics of such populations, therefore, becomes a challenge. Furthermore, even if a researcher can successfully develop a working definition, this definition may not match how the population members define themselves (Williams and Cheal 2002). The features and characteristics chosen by the researcher may not always represent the defining problem for the population members themselves. Adding to this problem is the fact that respondents may be reluctant to admit to participating in covert activities. Thus, even if respondents meet the inclusion criteria, they may self-select out of the sample (Faugier and Sargeant 1997).

Another issue stems from the fact that the information obtained from both formal and informal data sources is often incomplete (Faugier and Sargeant 1997). In addition to respondents being reluctant to divulge sensitive information, official sources usually include only those members of the hidden population with relatively high encounter rates. Potentially missing certain segments of the population, it is rare that any data source will contain a full spectrum of the types of individuals involved in the activity of interest (Faugier and Sargeant 1997). Without reliable means to assess the impact of these missing cases, none of the techniques provide a means to systematically use the information collected to make inferences about the entire target population. All methods potentially lack generalizability (Faugier and Sargeant 1997).

Applications to Extremist Populations

The techniques available to study hidden populations provide a number of different options for estimating the number of terrorists or extremists in Canada. In addition to being characterized as a socially stigmatized population that engages in covert illicit and illegal activities, extremists have a number of unique qualities that present additional challenges for researchers. In order to determine which technique would be best suited to studying this group, therefore, the unique features characterizing this population must first be considered.

Prior to discussing the issues, however, it is important to note that the term “extremist” is being used in the broadest sense. In this instance, it refers to an individual holding radical beliefs and attitudes, and whose actions are of a character far removed from ordinary. More specifically, these individuals will hold extreme political, religious, or other ideological views, and they will have engaged in politically motivated acts of terror, or have the potential to do so (Martin 2013). As such, this definition encompasses actual terrorists, as well as those individuals who have the potential to engage in acts of terrorism.

The first issue facing researchers who wish to study the population of extremists is defining the target behaviour, and, thus the population. Currently, there is no single universally accepted definition of terrorism. International law does not provide a coherent definition, and all countries utilize their own conceptualization of this term in their legal framework (Young 2006). In Canada, for instance, s. 83.01 of the *Criminal Code,* defines terrorism as an act committed "in whole or in part for a political, religious, or ideological purpose, objective or cause" with "the intension of intimidating the public" with "regard to its security, including its economic security, or compelling a person, a government or a domestic or international organization to do or to refrain from doing any act." Activities recognized as being criminal within this context include, but are not limited to, endangering a person's life, risks posed to the health and safety of the public, death and bodily harm with the use of violence, significant property damage, and interference or disruption of essential systems, facilities, or services (Department of Justice 2013). In contrast, in the United States, terrorism is defined as consisting of activities that "involve acts dangerous to human life that are a violation of the criminal laws of the United States or of any State… intended to intimidate or coerce a civilian population, influence the policy of a government by intimidation, or… affect the conduct of a government by mass destruction, assassination, or kidnapping" (Department of Justice 2013, para 3). Lastly, according to the British government, terrorism refers to "the use and threat of action designed to influence the government or to intimidate the public or a section of the public and made for the purpose of advancing political, religious or ideological causes" (Department of Justice, 2013, para 3). Even though the different agencies identify common features of terrorism, including political motivation and the threat or use of violence, there lacks consistency in their definitions. Without a clear and cogent definition of the activities and behaviours that constitute terrorism, it may be difficult for agencies to accurately capture the target population. As a result, it may be challenging for researchers to obtain accurate official data on extremist populations.

In addition to the lack of consistency, all available definitions of terrorism are both cumbersome and vague at the same time. Specifically, in addition to requiring that members of this population possess specific motivational factors and intensions, these definitions outline a number of activities in which these individuals could be engaged. The relative importance of each of these elements, however, is not clear. Determining the defining features of terrorism, therefore, is not an easy task. Not only are researchers presented with the challenge of clearly defining the population, but also determining their criteria for inclusion in the sample becomes difficult. Given that the inclusion criteria shape the final sample, researchers must use caution when deciding which features and characteristics to choose. As mentioned previously, to obtain accurate data on the hidden population, the definition must coincide with how the population members would define themselves (Williams and Cheal 2002).

Provided that the researcher is able to clearly define the target group and inclusion criteria, the next challenge would be obtaining the cooperation of members of extremist populations. Since its inception, the word "terrorism" has had a very negative connotation (Hoffman 2006; Weinbery 2005). For example, terrorists are characteristically viewed as evil-doers, or madmen (Post 1990). Based on their association to terrorism, terms including "extremist," "radical," and "fanatic" have also taken on these pejorative meanings. Given the extreme stigma associated with these terms, respondents may be reluctant or unwilling to participate in the study.

Complicating matters further is the fact that, terrorism and terrorism-related behaviours are criminal offences. According to the Canadian *Criminal Code* and the Anti-terrorism Act (S.C. 2001 c. 41), individuals caught engaging in activities that endanger a person's life, pose a risk to the health and safety of the public, cause death or bodily harm with the use of violence, result in significant property damage, and/or interfere or disrupt the functioning of essential systems, facilities or services, may face criminal prosecution. The potential for criminal prosecution may deter individuals from participating, as well as from identifying other members of the population. To increase chances of obtaining respondents, therefore, researchers would need to use methods that would eliminate the need for respondents to divulge sensitive information. For instance, the unique coupon method used in respondent-driven sampling may provide a solution.

In addition to these definitional issues, researchers must remain aware of the challenges presented by a number of the features of terrorism. The use or threat of violence is one such feature (Hoffman 2006). What makes this unique is the level of violence used by terrorists. For example, in the infamous terrorist attack on September 11, 2001, members of the Al-Qaeda terrorist organization drove two planes into the Twin Towers in New York City resulting in the death of over 2,000 people, as well as millions of dollars in property damage. Essentially, with terrorism, the potential for large-scale destruction of property and loss of life for individuals is high. Consequently, once individuals who have committed or are deemed likely to commit an act of terrorism are identified and captured, it is unlikely that they will ever be released. This may present a problem for the techniques requiring multiple captures from different sources, such as capture-recapture methods and multiplier methods.

Another unique feature requiring consideration is the clandestine nature of terrorism. Although all hidden populations can be characterized as operating in a covert manner so as to conceal their illegal or illicit behaviours (Bloor, Leyland, Barnard, and McKeganey 1991; Johnston and Sabin 2010), extremists take this degree of secrecy to another level. Because terrorists are fighting in an asymmetrical war, their success lies in the element of surprise (Weinbery 2005). As a result, it is imperative that terrorists conceal any activities related to planning, organizing, and carrying out an attack (Gibbs 1989). The furtive nature of terrorism extends beyond the non-violent and violent actions taken by terrorists to include the individual's lifestyle. To reduce their chances of being detected, extremists work hard to maintain the appearance of a normal life. Many have ordinary jobs and live with their families (Gibbs 1989). Taken together, the highly secretive features of terrorism may significantly decrease the likelihood that members of this population will come into contact with agencies. Potentially large segments of this population, therefore, will be missed by these agencies. This poses a significant problem for techniques requiring the use of one or more official information data sources (e.g., capture-recapture and multiplier methods).

This issue of detection is compounded by the fact that terrorism is a global phenomenon. Acts of terrorism can be committed within a country's own borders–domestic terrorism–or outside–international terrorism. This means that terrorists can operate within their own country, or anywhere else in the world. For example, an extremist who currently lives in Canada may subsequently move, or travel to carry out an attack in a different country. Essentially, there are no physical boundaries or borders to enclose terrorists. This potential for mobility makes it almost impossible to use techniques that require the populations to be closed and/or geographically clustered (i.e., targeted sampling, time-location sampling, adaptive sampling, etc.).

A final concern involves the dynamic nature of terrorism. Terrorism has traditionally been conceptualized as a group phenomenon, a collective and organized activity. Recent incidents, such as the 2011 Norway attack, however, have brought a different form of terrorism to the forefront: lone wolf or small group terrorism (Spaaij 2012). Operating as individual or highly atomized cells, these forms of terrorism eliminate the cohesive network structure that was inherent in group-based terrorism (Spaaij 2012). This presents a challenge for sampling methods that require links between members of the population, such as snowball sampling and other link-tracing designs.

Despite the changes in group-structure, however, there is evidence to suggest that links between extremists still exist. According to Spaaij (2012), even those individuals perceived to be acting alone have some level of connection with others. For example, both Andres Breivick and Timothy McVeigh were believed to have had an accomplice (Spaaij 2012). Furthermore, with increasing technological advancements, there are new mechanisms available to connect isolated and atomized individuals with other members of the target population. Specifically, the Internet has a number of forums, such as blogs, that provide individuals with the opportunity to join virtual networks (Pantucci 2011). Using the Internet as the medium, link-tracing designs would be able to trace the networks of population members who are connecting online.

Taken together, although there would be challenges to using any of the methods currently available for estimating hidden populations, it appears that some methods would be better suited than others for studying the population of extremists in Canada. Given the potential difficulties associated with the information obtained by agencies, methods relying on official data sources (e.g., capture-recapture methods and the multiplier method) are less desirable. Further, given the potential for population mobility, it is nearly impossible to use the methods that require the population be stable and highly clustered in specific locations (i.e., adaptive sampling techniques, targeted sampling, and time-location sampling). This leaves the link-tracing designs. Among these techniques, it appears that respondent-driven sampling would be the most appropriate fit. Because it is designed for use in the real-world, respondent-driven sampling provides a means to account for a number of the challenges presented by extremist populations. Specifically, by eliminating the need for exhaustive mapping to construct sampling frames, as well as providing anonymized methods of recruitment (i.e., unique coupons), RDS provides a way for researchers to access the extremely vulnerable and often reclusive population of extremists.

Application in the Canadian Context

Having determined that respondent-driven sampling is the most viable technique available for estimating the number of extremists, it is now necessary to explore exactly how this method would be applied to study this population. Using Canada as a case study, the procedure for implementing a RDS approach to count the number of violent extremists will be outlined. Acknowledging the inherent difficulties associated with studying such hidden populations, the challenges researchers may face at each stage of the sampling and estimating process will also be highlighted.

Prior to initiating the sampling process, the population must be defined. As previously stated, an extremist refers to an individual holding extreme political, religious, or other ideological views who has, or is likely to, engage in politically motivated acts of terror (Martin 2013). Narrowing our focus to the Canadian context and using s. 83.01 of the Canadian *Criminal Code* as a guide, politically motivated acts of terror are those acts committed "in whole or in part for a political, religious, or ideological purpose, objective or cause" with "the intension of intimidating the public" with "regard to its security, including its economic security, or compelling a person, a government or a domestic or international organization to do or to refrain from doing any act." Essentially, in order to identify the populations of extremists in Canada, it is necessary to first determine what their main ideological objectives may be. For instance, being a multicultural nation, Canada is home to people with various cultural, ethnic, and religious backgrounds. While still predominantly a Roman Catholic and Protestant nation, with increasing immigration from many Middle Eastern countries, Canada has seen a growth in Islam, Hinduism, Sikhism, and Buddhism (Statistics Canada, n.d.). Moreover, given its diverse geographical landscape, Canada is built on a number of resource-based industries, including forestry and hydro. Taken together, Canada has the potential to provide the foundation for a myriad of political, religious, and ideological causes. As a result, Canada may be hosting a variety of different groups of political, religious, and environmental extremists. Researchers, therefore, must be aware of the prominent and controversial issues that exist within Canada's borders at any given time.

Once the population has been defined and the study's inclusion criteria (e.g., specific political, religious, or ideological objective) have been outlined, RDS requires the researcher to locate, and then reach out to members of this target population. Contacting members of the Canadian extremist population, however, will require strategy. Provided that extremists can have many different causes and objectives, these individuals may belong to a variety of different groups. Furthermore, being divided by their motivational factors and ideological frameworks, it is unlikely that these different networks of extremists will overlap. For instance, religiously-motivated extremists such as the Canadian-based Al-Qaeda-inspired groups have markedly different aims and desires from environmentally-motivated extremists, including the Squamish Five. Understanding that extremists will likely be segregated into separate groups, it will be necessary for the researcher to identify members belonging to each group. Regardless of the type of group, though, the process of locating and contacting these group members presents the researcher with numerous challenges.

As previously stated, the very nature of terrorism, a socially stigmatized and illegal activity, necessitates secrecy. Terrorists, therefore, go to extreme lengths to maintain anonymity and avoid detection. As such, extremists will not be found operating in an open or obvious manner.

For instance, it is unlikely that groups of extremists will be found congregating in public settings such as religious institutions and schools. Essentially, this feature of terrorism reduces the possibility that researchers will have physical access to this population. As a result, researchers will have to rely on less direct means to locate and trace the networks of members of the extremist populations in Canada. The Internet is one such option. There are countless websites, chat rooms, and blogs devoted to the espousal of specific causes. As such, the Internet provides individual extremists with the opportunity to join virtual networks and anonymously communicate with others who share the same views and beliefs. Being open to the public, the Internet also provides researchers with a method to access those extremists who are connected by these virtual networks. In theory, researchers should be able to use the Internet to join different virtual extremist communities, communicate with the group members, and identify suitable participants to act as seeds in the RDS recruitment process. In reality, however, this process raises a number of challenges and concerns. For instance, although accessing a website or blog may be straightforward, engaging with group members, and finding willing study participants may be more difficult. For instance, due to ethical issues concerning transparency, the researcher will need to identify him or herself, as well as the purpose of the study. Given the negative connotation associated with "terrorism" and the potentially severe consequences for engaging in terrorist activity, many members of the target population may be reluctant to participate in the study. More specifically, those members with the most extreme views and violent ideations (those most likely to engage in actual acts of terrorism) may be the least likely to identify themselves and cooperate with the researcher. In addition to individual members refusing to participate, in some instances, the group may decide to remove the researcher from their site to avoid any further contact and detection. Essentially, the researcher may have difficulty obtaining cooperation from influential group members, and, thus, must be cautious in their initial approach. In order to gain the trust of respondents, the researcher will likely need to spend a significant amount of time building a strong rapport with the group members.

A connected issue concerns the longevity of these Internet sites. In an effort to avoid detection and keep their operations and communications as covert as possible, extremist groups may change their URLs and websites over time. In addition, due to Internet policing and reporting, many of these sites may be taken down once discovered; those sites and groups espousing violence or the threat of violence are usually the first to be removed. Changes to existing sites and the removal of others make it extremely difficult for researchers to track group members and all of their networks. To use the Internet effectively, then, the researcher must remain up-to-date with the current and active websites, chat rooms, and blogs used by the extremist groups of interest.

Lastly, even if the researcher is able to find the extremists' online communication forums, gain access to the group, and engage members to participate in the study, the highly anonymized mechanisms for communication make it hard for the researcher to identify the individual study participants. In the place of real names, individuals create "user names" or "accounts" for different sites. Although each user name is unique, because it can be anything the creator chooses, the user name does not necessarily provide identifiable information about the individual who created it. This makes it difficult for researchers to use any tangible incentives to encourage participation. Furthermore, there are no limits as to the number of accounts any one particular individual can create; the same individual could have multiple accounts for the same site. As a result, despite implementing the unique coupon method used in RDS, the researcher may not be able to account for repeat participants. Finally, given that individuals are masked by user names and only accessible in cyberspace, it is difficult for researchers to determine the actual, physical location of each participant. More specifically, it would be almost impossible to determine if the extremist was residing in Canada or another country. As such, the researcher would have to find a way to determine the probability that the extremist was in Canada as compared to anywhere else in the world. To create a precise estimate, then, the researcher would have to acknowledge the potential for repeat cases, as well as adjust for the percentage of the known population of extremists that would make up the Canadian extremist population.

In conclusion, although there remain a myriad of difficulties to overcome, RDS appears to offer a more systematic method for estimating the number of extremists in Canada. At this stage, though, the real value in RDS might lie with its ability to provide us with badly needed information as a starting point for further research. Through RDS we may be able to learn more about the networks that connect terrorist, setting the stage for future research.

References

Acharya, A.K. 2007. A methodological approach to study hidden populations: The case of trafficked women in Mexico City. *International Journal of Social Sciences & Humanities*, XVII (1):9-23.

Acharyal, B., G. Bhattarai, A. de Gier, and A. Stein. 2000. Systematic adaptive cluster sampling for the assessment of rare tree species in Nepal. *Forest Ecology and Management*, 137:65-73.

Anti-terrorism Act S.C. 2001, c. 41.

Arabkhedri, M., F.S. Lai, I, Noor-Akma, and M.K. Mohamad-Roslan. 2010. An application of adaptive cluster sampling for estimating total suspended sediment load. *Hydrology Research* 41:63-73.

Atkinson, R., and J. Flint. 2001. Accessing hidden and hard-to-reach populations: Snowball research strategies. *Social Research Update* (33).

Avico, U., C. Kaplan, D. Korczak, and K. van Meter. 1988 Cocaine epidemiology in three European community cities. A pilot study using a snowball sampling methodology. *Research Report of the Cocaine Steering Group.* Health Directorate Commission of the European Communities in Brussels, Belgium.

Berk, R. 2007. Some thoughts on estimating the size of hidden populations: The special case of forced labor. Report. Department of Statistics, Department of Criminology: University of Pennsylvania.

Berk, R., B. Kriegler, and D. Ylvisaker. 2008. Counting the homeless in Los Angeles County. *Probability and Statistics: Essays in Honor of David A. Freedman* (2):127-141.

Bernard, H.R., T. Hallett, A. Iovita, E.C. Johnsen, R. Lyerla, C. McCarty, M. Mahy, M.J. Salganik, T. Saliuk, O. Scutelniciuc, G.A. Shelley, P. Sirinirund, S. Weir, and D.F. Stroup. 2010. Counting hard-to-count populations: The network scale-up method for public health. *Sexually Transmitted Infection* 86 (Suppl 2): ii11-ii15.

Biernacki, P. 1986. *Pathways from hidden heroin addiction: Recovery without treatment*. Philadelphia: Temple University Press.

Bloor, M., A. Leyland, M. Barnard, and N. McKeganey. 1991. Estimating hidden populations: A new method of calculating the prevalence of drug-injecting and non-injecting female street prostitution. *British Journal of Addiction* 86:1477-1483.

Bouchard, M. 2007. A capture-recapture model to estimate the size of criminal populations and the risks of detefction in a marijuana cultivation industry. *Journal of Quantitative Criminology* 23:221-241.

Bouchard, M., and P. Tremblay, P. 2005. Risks of arrest across drug markets: A capture-recapture analysis of "hidden" dealer and user populations. *Journal of Drug Issues* 35 (4): 733-754.

Browne, K. 2005. Snowball sampling: Using social networks to research non-heterosexual women. *International Journal of Social Research Methodology* 8 (1): 47-60.

Calkins, R.F. and G.B. Aktan, G.B. 2000. Estimation of heroin prevalence in Michigan using capture-recapture and heroin problem index methods. *Journal of Drug Issues* 30:187-204.

Choi, Y.H. and C.M. Comiskey. 2003. Methods for providing first prevalence estimates of opiate use in Western Australia. *International Journal of Drug Policy* 14 (4): 297-305.

Chow, M., and S.k. Thompson. 2003. Estimation with link-tracing sampling designs – a Bayesian approach. *Survey Methodology* 29:197-205.

Collins, M.F., and R. M. Wilson 1990. Automobile theft: Estimating the size of the criminal population. *Journal of Quantitative Criminology* 6 (4): 395–409.

Comiskey, C.M., and J.M. Barry. 2001. A capture-recapture study of the prevalence and implication of opiate use in Dublin. *European Journal of Public Health* 11 (2): 198-200.

Cormack, R.M. 1999. Problems with using capture-recapture in epidemiology: An example of a measles epidemic. *Journal of Clinical Epidemiology* 52 (10): 909-914.

Criminal Code, R.S.C.1985, c. C-46, S. 83.01.

Department of Justice. 2013. *Memorializing the victims of terrorism: Definitions of terrorism and the Canadian context*. Accessed May 23, 2014. http://www.justice.gc.ca/eng/rp-pr/cj-jp/victim/rr09_6/p3.html (Department of Justice).

Domingo-Salvany, A., R.L. Hartnoll, A. Maguire, J.M. Suelves, and J.M. Anto. 1995. Use of capture-recapture to estimate the prevalence of opiate addiction in Barcelona, Spain, 1989. *Journal of Clinical Epidemiology* 52:909-914.

Faugier, J., and M. Sargeant, M. 1997. Sampling hard to reach populations. *Journal of Advanced Nursing* 26:790-797.

Felix-Medina, M.H., P.E. Monjardin, and A.N. Aceves-Castro. 2009. Link-tracing sampling: Estimating the size of a hidden population in presence of heterogeneous nomination probabilities. *Section on Survey Research Methods,* 4020-4033.

Ferreira, L.O.C., E.S. de Oliveira, H.F. Raymond, S.Y. Chen, and W. McFarland. 2008. Use of time-location sampling for systematic behavioral surveillance of truck drivers in Brazil. *AIDS and Behavior* 12 (1 Supplement):32-38.

Frank, O., and T. Snijders. 1994. Estimating the size of hidden populations using snowball sampling. *Journal of Official Statistics* 10:53-67.

Gibbs, J.P. 1989. Conceptualization of terrorism. *American Sociological Review* 54 (3): 329-340.

Goel, S., and M.J. Salganik. 2009. Respondent-driven sampling as Markov chain Monte Carlo. *Statistics in Medicine* 28:2202-2229.

Goldberg, N.A., J.N. Hein, J.N., and J.A. Brown. 2006. The application of adaptive cluster sampling for rare subtidal macroalgae. *Marine Biology* 151:1343-1348.

Greene, M.A., and S. Stollmack. 1981. Estimating the number of criminals. In *Models in uantitative criminology* edited by J.A. Fox. 1-24. New York: Academic Press.

Griffiths, P., M. Gossop, B. Powis, and J. Strang. 1993. Reaching hidden populations of drug users by privileged access interviewers: Methodological and practical issues. *Addiction* 88:1617-1626.

Gus, M. 2006. *Terrorism and homeland security*. Thousand Oaks, CA: Sage Publications Inc.

Gus, M. 2013. *Understanding terrorism*. Thousand Oaks, CA: Sage Publications Inc.

Handcock, M.S., Gile, K.J., Mar, C.M. (2012). Estimating hidden population size using respondent-driven sampling data. Retrieved from arXiv:1209.6241v1 [stat.ME].

Hay, G., and M. Gannon. 2006. Capture-recapture estimates of local and national prevalence of problem drug use in Scotland. *International Journal of Drug Policy* 17:203-210.

Heckathorn, D.D. 1997. Respondent-driven sampling: A new approach to the study of hidden populations. *Social Problems* 44(2): 174–199.

Heckathorn, D.D. (2002). Respondent driven sampling II: Deriving valid population estimates from chain-referral samples of hidden populations. *Sociological Problems* 49(Suppl. 1): 11-34.

Hendricks, V.M., and P. Blanken. 1992. Snowball sampling: Theoretical and practical considerations. In *Snowball sampling: A pilot study on cocaine use* edited by V.M., Hendricks, P. Blanken, and N. Adriaans. 17-35. Rotterdam: IVO.

Hickman, M., S. Cox, J. Harvey, S. Howes , M. Farrell, M. Frischer, M. Stimson, G. Taylor, and K. Tilling. 1999. Estimating the prevalence of problem drug use in inner London: A discussion of three capture-recapture studies. *Addiction* 94(11):653-1662.

Hoffman, B. 2006. *Inside terrorism*. New York, NY: Columbia University Press.

Hope, V.D., M. Hickman, and K. Tilling. 2005. Capturing crack cocaine use: estimating the prevalence of crack cocaine use in London using capture-recapture with covariates. *Addiction* 100:1701-1708.

Hser, Y.I. 1993. Population estimates of intravenous drug users and HIV infection in Los Angeles County. *International Journal of the Addictions* 28 (8): 695-709.

Hudson, C.G. 1997. Estimating homeless populations through structural equation modeling. *Housing Policy Debates* 8 (3): 631-647.

Hutchinson, S.J., S.M. Bird, A.Taylor, and D.J. Goldberg. 2006. Estimating the prevalence, incidence and cessation of injecting drug use in Glasgow 1960-2000: Combining expert opinion with capture-recapture prevalence data. *International Journal of Drug Policy* 17:29-34.

James, F.J. 1991. Counting homeless persons with surveys of users of services for the homeless. *Housing Policy Debates* 2(3): 733-753.

Johnston, L.G., K. Sabin, M.T. Hien, and P.T. Huong. 2006. Assessment of respondent driven sampling for recruiting female sex workers in two Vietnamese cities: Reaching the unseen sex worker. *Journal of Urban Health* 83 (7): 16–28.

Johnston, L.G., and K. Sabin. 2010. Sampling hard-to-reach populations with respondent driven sampling. *Methdological Innovations Online* 5 (2): 38-48.

Kajubi, P., M.R. Kamya, H.F. Raymond, S. Chen, G.W. Rutherford, J.S. Mandel, and W. McFarland. 2008. Gay and bisexual men in Kampala, Uganda. *AIDS and Behavior* 12 (3): 492–504.

Kalton, G. 2001. "Practical methods for sampling rare and mobile populations". Paper presented at Proceedings of the Annual Meeting of the American Statistical Association, Rockville, Maryland, August 2001.

Kanouse, D. E., S.H. Berry, N. Duan, J. Lever, S. Carson, J.F. Perlamn, and B. Levitan..1999. Drawing a probability sample of female street prostitutes in Los Angeles County. *Journal of Sex Research* 36 (1): 45-51.

Kaplan, C.D., D. Korf, and C. Sterk. 1987. Temporal and social contexts of heroin-using populations: An illustration of the snowball sampling technique. *The Journal of Nervous and Mental Disease* 175 (9): 566-574.

Karon, J.M., and C. Wejnert. 2012. Statistical methods for the analysis of time-location sampling data. *Journal of Urban Health* 89 (3): 565-586.

Kendall, C., L.R. Kerr, F.S. Gondim, R.C. Werneck, G.L. Macena, R.H.M. Pontes, M.K. Johnston, L.G. Sabin, and W. McFarland. 2008. An empirical comparison of respondent-driven sampling, time location sampling, and snowball sampling for behavioral surveillance in men who have sex with men, Fortaleza, Brazil. *AIDS Behavior* 12: S97-S104.

Klovdahl, A. 1989. Urban social networks: Some methodological problems and possibilities. In *The small world* edited by M. Kochen. 176-210. Norword, NJ: Ablex Publishing.

Larson, A., A. Stevens, and G. Wardlaw. 1994. Indirect estimates of 'hidden' populations: Capture-recapture methods to estimate the numbers of heroin users in the Australian Capital Territory. *Social Science and Medicine* 39 (6): 823-831.

Lee, R.M. 1993. *Doing research on sensitive topics*. London: Sage.

Magnani, R., K. Sabin, T. Saidel, and D. Heckathorn. 2005. Review of sampling hard-to-reach and hidden populations for HIV surveillance. *AIDS* 19 (supple 2): S67-S72.

Mastro, T.D., D. Kitayaporn , B.G. Weniger, and S. Vanichseni. 1994. Estimating the number of HIV infected injection drug users in Bangkok: A capture-recapture method. *American Journal of Public Health* 84 (7): 1094-1099.

Medhi, G.K., J. Mahanta, B.S. Akoijam, and R. Adhikary, R. 2012. Size estimation of injecting drug users (IDU) using multiplier method in five Districts of India. *Substance Abuse Treatment, Prevention, and Policy* 7:9-13.

Paniotto, V., T. Petrenko, and O. Kupriyanov. 2009. *Estimating the size of populations with high risk for HIV using the network scale-up method*. Ukraine: Kiev International Institute of Sociology.

Pantucci, R. 2011. *A typology of lone wolves: Preliminary analysis of lone Islamist terrorists*. London: ICSR.

Platt, L., M. Wall, T. Rhodes, A. Judd, M. Hickman, L.G. Johnston, A. Renton, N. Bobrova, and A. Sarang. 2006. Methods to recruit hard-to-reach groups: comparing two chain referral sampling methods of recruiting injection drug users across nine studies in Russia and Estonia. *Journal of Urban Health* 83 (7): 39–53.

Post, J. 1990. Terrorist psycho-logic: Terrorist behavior as a product of psychological forces. In *Origins of terrorism: Psychologies, ideologies, theologies, state of mind* edited by W. Reich. 25-40. New York, NY: Cambridge University Press.

Riccio, L.J., and R. Flinkestein. 1985. Using police arrest data to estimate the number of burglars operating in a suburban county. *Journal of Criminal Justice* 13:65–73.

Roberts, J.M., and D.D. Brewer. 2006. Estimating the prevalence of male clients of prostitute women in Vancouver with a simple capture–recapture method. *Journal of The Royal Statistical Society Series A* 169:745-756.

Roesch, F.A. Jr. 1993. Adaptive cluster sampling for forest inventories. *Forest Science* 39:655-669.

Rossmo, D.K., and R. Routledge. 1990. Estimating the size of criminal populations. *Journal of Quantitative Criminology* 6:293–314.

Sadler, G.R., H.C. Lee, F.S.H. Lim, and J. Fullerton. 2010. Recruitment of hard-to-reach population subgroups via adaptations of the snowball sampling strategy. *Nursing and Health Sciences* 12:369-374.

Salganik, M.J., and D.D. Heckathorn. 2004. Sampling and estimation in hidden populations using respondent-driven sampling. *Sociological Methodology* 34:193-239.

Seber, G.A.F., and M.M. Salehi, M.M. 2013. *Adaptive sampling designs: Inference for sparse and clustered populations*. New York, NY: Springer.

Snidero, S., B. Morra, R. Corradetti, and D. Gregori. 2007. Use of the scale-up methods in injury prevention research: an empirical assessment to the case of choking in children. *Social Networks* 29:527-538.

Snijders, T.A.B. 1992. Estimation on the basis of snowball samples: How to weight?" *Bulletin de Mithodologie Sociologique* 36:59-70.

Spaaij, R. 2012. U*nderstanding lone wolf terrorism: Global patterns, motivations and prevention*. New York, NY: Springer.

St. Clair, K., and D. O'Connell. 2012. A Bayesian model for estimating population means using a link-tracing sampling design. *Biometrics* 68:165-173.

Statistics Canada. (n.d.). *Overview: Canada still predominantly Roman Catholic and Protestant*. http://www12.statcan.ca/english/census01/Products/Analytic/companion/rel/canada.cfm.

Sudman, S., and E. Blair, E. 1999. Sampling in the twenty-first century. *Journal of the Academy of Marketing Science* 27 (2): 269-277.

Thompson, S.K. 1990. Adaptive cluster sampling. *Journal of the American Statistical Association* 85 (412): 1050-1059.

Thompson, S.K. 2012. *Sampling*. Somerset, NJ: Wiley.

Thompson, S.K., and L.M. Collins. 2002. Adaptive sampling in research on risk-related behaviours. *Drug and Alcohol Dependence* 68:S57-S67.

Thompson, S. K., and O. Frank. 2000. Model-based estimation with link-tracing sampling designs. *Survey Methodology* 26 (1): 87-98.

Thompson, S. K., and G.A.F. Seber. 1996. *Adaptive sampling.* New York: Wiley.

Valleroy, L.A., D. MacKellar, J. Karon, D.H. Rosen, W. McFarland, D. Shehan, S. Stoyanoff, M. Lalota, D. Celentano, D. Koblin, H. Thiende, M. Katz, L. Torian, and R. Janssen. (2000). HIV prevalence and associated risks in young men who have sex with men. *Journal of the American Medical Association* 284 (2): 198-204.

Van Meter, K. M. 1990. Methodological and design issues: Techniques for assessing the representatives of snowball samples. In *The collection and interpretation of data from hidden populations (NIDA Research Monograph, 98)* edited by E.Y. Lambert. 31-43. Rockville, MD: US Department of Health and Human Services.

Watters, J. K., and P. Biernacki, P. 1989. Targeted sampling: options for the study of hidden populations. *Social Problems* 36 (4): 16-30.

Weinbery, L. 2005. *Global terrorism: A beginner's guide*. Oxford: Oneworld Publications.

Williams, M., and B. Cheal. 2002. Can we measure homelessness? A critical evaluation of the method of 'capture-recapture'. *International Journal of Social Research Methodology* 5 (4): 313-331.

Young, R. 2006. Defining terrorism: The evolution of terrorism as a legal concept international law and its influence on definitions in domestic legislation. *Boston College International and Comparative Law Review* 29:23.