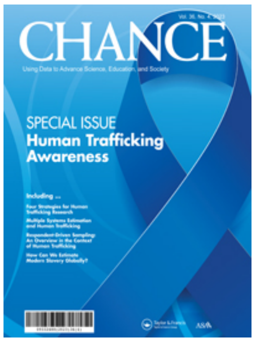

CHANCE

ISSN: (Print) (Online) Journal homepage: www.tandfonline.com/journals/ucha20

Respondent-Driven Sampling: An Overview in the Context of Human Trafficking

Jessica P. Kunke, Adam Visokay & Tyler H. McCormick

Respondent-Driven Sampling: An Overview in the Context of Human Trafficking

Jessica P. Kunke, Adam Visokay, and Tyler H. McCormick

Human trafficking is a global public health concern with widespread and long-lasting negative consequences. Understanding trafficking and estimating the number of people being trafficked is complicated by the stigma, sensationalization, and secrecy of trafficking.

Recent estimates of the number of people being trafficked worldwide range from 12.3 million to 45.8 million people. While prevalence estimation is just one of many research priorities in this field, constraining the prevalence estimates better is important for guiding policy decisions. A common sampling technique known as respondent-driven sampling (RDS) can be used to reach this population.

Human trafficking was legally defined in 2000, internationally by the Palermo Protocol adopted by the United Nations General Assembly, and within the U.S. by the Victims of Trafficking and Violence Protection Act of 2000. Generally, human trafficking is the use of force, fraud, or coercion to exploit one or more people through commercial sex or forced labor. Inducing a minor (someone under the age of 18 years old) into commercial sex is considered human trafficking regardless of the presence of force, fraud, or coercion.

There are many popular misconceptions. For instance, human trafficking is often confused with people-smuggling, in which people are moved consensually but illegally. By contrast, human trafficking can, but does not necessarily, involve movement, and requires the use of force, fraud, or coercion (Schroeder, et al. 2022). While human trafficking has often been defined and approached through the lens of criminal justice, it is increasingly recognized as a complex public health issue.

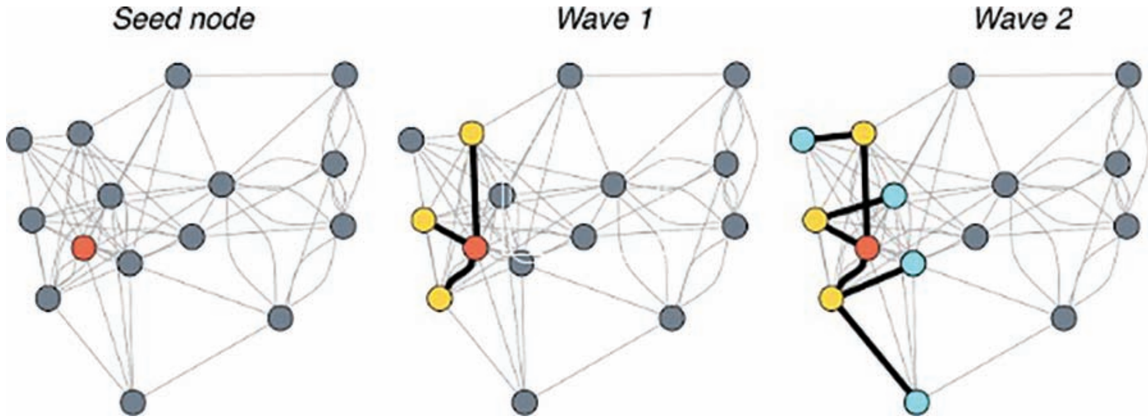

Figure 1. Each panel shows an entire example network. Nodes recruited in waves 1, 2, and 3 shown in red, yellow, and light blue, respectively. Gray nodes are never recruited during sampling waves, and bolded paths indicate directed recruitment links. Paths not bolded remain unobserved to researchers.

A major challenge in studying human trafficking stems from differences in definitions or in the operationalization of the same definition from one study to another. The Palermo Protocol explicitly mentions slavery and organ removal as forms of exploitation in its definition of trafficking, and some definitions include forced marriage, forced begging, or child soldiers.

Even under a single definition, what gets counted as trafficking on a case-by-case basis depends on popular conceptions of trafficking, which are shaped by racism, sexism, colonialism, and other systemic injustices. Gendered and racialized ideas of innocence and purity misinform the popular narrative about who is trafficked, why, and what they need. This focus also encourages sensationalism, which distracts from progress. Researchers, policymakers, law enforcement officials, service providers, and others who are central to identifying and combating human trafficking are subject to these same biases and misconceptions.

In addition to examining the complex question of what counts as trafficking, researchers have been working on developing effective, standardized statistical and sampling methods to understand the scope and nature of human trafficking. Many studies to date have used administrative data, but for various reasons, trafficking-related charges and prosecutions are thought to represent only a small, biased sample of existing trafficking cases.

For this reason, many other studies have collected fresh data, which raises other challenges.

Traditional survey methods assume that researchers have a sampling frame or a list of people in the general population that includes the target population of interest (in this case, people who are being trafficked) and that respondents will willingly identify whether they are being trafficked. In practice, accessing people in this population requires learning how to find them and building their trust. Individuals who are trafficked might not have autonomy over their movements, distrust officials, and not feel comfortable identifying themselves either because of stigma or fear of retribution. In this setting, traditional sampling techniques often do not reach trafficked individuals.

One strategy to reach individuals excluded from standard surveys involves leveraging the social networks of individuals in the group of interest. With these techniques, the researcher does not need to access a representative sample from the general population but, rather, can interact with a sample of people who are connected with members of the group of interest. These network-based methods, broadly, fall into two categories.

The first category does not involve interacting with members of the group directly. Strategies such as the network scale-up method (NSUM) ask individuals how many people they know in the target population; i.e., people who are experiencing trafficking or who have done so within some time period. Responses from the general population are then scaled with assumptions about how well the average fraction of the respondents' networks that consists of trafficked individuals extrapolates to the population.

These indirect methods have the advantage of not requiring respondents to identify themselves or specific other people as members of the group of interest. A disadvantage is that these methods are often limited to prevalence estimation rather than gaining additional insight into risk factors, experiences, or possible paths out of trafficking.

A second class of methods involves interacting directly with individuals in the group of interest. These methods fall under a general class of methods known as link-tracing or chain-referral designs because recruitment proceeds along links in the social network connecting of individuals who are victims of trafficking. Several related iterations of link-tracing designs have been proposed, so beginning with some terminology is helpful.

Chain-referral sampling is a general term for a method that "traces" respondents' networks as a means of recruitment.

Snowball sampling is a chain-referral method originally proposed as a way to learn about network features that starts with a probability sample of respondents and traces their networks.

Respondent-driven sampling (RDS) is sometimes called "non-probability snowball sampling" because it starts with a convenience sample and respondents choose network members to recruit. Researchers use RDS to estimate prevalence, understand characteristics of particular populations (e.g., the fraction of sex workers in an area who have been trafficked), or access members of a hard-to-reach group for an intervention.

RDS is an increasingly popular sampling and estimation strategy for human trafficking, but is not yet standardized in its application.

RDS Implementation and Assumptions

RDS can be used for three distinct purposes: (i) estimating characteristics of the group of individuals (e.g., the fraction of trafficking victims who are also minors); (ii) estimating a population size or prevalence; and (iii) accessing a representative sample of individuals for further study or intervention (e.g., to evaluate the effectiveness a particular type of outreach).

RDS relies on members of the group of interest to recruit other members into the study, thus leveraging the social connections and relationships of group members to increase participation.

The RDS process begins with a convenience sample of members of the group of interest—in this instance, people who have been trafficked within some period of time. These individuals are often known to researchers from previous studies or interventions, or have previously interacted with a public health infrastructure.

These initial recruits are known as **seeds**. The seed individuals are asked to recruit a particular number of additional individuals, also from the group of interest. Each of these recruits typically receives an incentive to participate in the study, known as the **primary incentive**.

Recruits are then asked to bring more individuals from the group of interest into the study. When a new recruit participates in the study, the person who recruited them receives a **secondary incentive**.

Each new recruitment cycle defines a **wave** of a **recruitment chain**—the complete set of individuals and their referral connections, including the original seeds.

There are several practical considerations when performing RDS. The order of recruitment can be consequential when analyzing RDS, so it is critical for researchers to keep track of the recruitment sequencing. Often, this is accomplished by passing out coupons with unique identifying numbers. Each participant gets a certain number of coupons containing their unique recruiter number, so when a new participant brings a coupon they were given, the study team knows who recruited them.

An additional consideration is the number of coupons available to each person. Allowing each person to recruit more people reduces the chances that a chain terminates in the early waves. For a given sample size, however, it also means that the study can incorporate fewer seeds and, thus, have chains originating in potentially different parts of the network.

Figure 1 shows the RDS process on a small example network, starting with a single seed (RDS typically begins with multiple seeds), selected as a convenience sample, denoted in red. The seed recruits, in this example, three additional participants in the first wave. These participants then, in turn, recruit additional participants in subsequent waves. The unbolded network edges in the figure are not observed.

This example also illustrates two challenges with the RDS procedure.

First, the sampling happens on top of an existing social network that is unknown to the researcher and difficult to recover from the RDS chain.

Second, the sampling process on that network is controlled by the respondents, not by the researcher, meaning that the choice of who is included in the study is up to the respondents and may or may not represent the population well or meet other desirable sample criteria.

Since the initial seeds for RDS are a convenience sample, they are not representative of the population. In an ideal world, however, subsequent recruitment waves would "move away" from the initial seeds in the social network, making the choice of initial seeds less and less consequential.

Under ideal circumstances, as the chains traverse the network, they will include respondents with heterogeneous characteristics, and the frequency of those characteristics in the sample will be roughly that of the population. If this happens, RDS behaves like a mathematical process known as a Markov Chain Monte Carlo.

This ideal behavior of RDS requires several assumptions.

First, RDS assumes that members of the population can be reached through their network, i.e., that they know one another reciprocally, interact frequently, are willing to recruit others, and have mobility. However, this assumption may not hold in the context of human trafficking. Restrictions on mobility, for example, may make it impossible for individuals to receive a coupon or to bring a coupon they receive to a research center and participate in the study.

Second, RDS assumes that respondents' network sizes are either known or accurately estimated. This assumption is necessary because the likelihood of being sampled depends on the respondent's network; a person with more contacts has more chances to be included.

Third, the sampling process has to continue through enough waves to mitigate the dependence on seeds. If there are too few waves, then the structure of the sample will be too closely related to the initial seeds. If there are substantial bottlenecks in the network, then the recruitment process can get "stuck" in one pocket and not explore the full extent of the graph. Particularly with small groups, the sample size can become close to the total population size.

Fourth, RDS assumes sampling is done with replacement, meaning that the study may recruit the same person more than once.

Fifth, RDS assumes that network connections are reciprocal—that person A is equally likely to refer person B as person B would be to refer person A.

Finally, RDS assumes that respondents recruit randomly from their contacts. Under this assumption, the only factor that affects how likely you are to be recruited is your number of contacts. In practice, though, recruitment may be highly preferential, or even if the recruitment is random, there may be selection bias in which people who receive coupons are more likely to participate in the study. Preferential referral can introduce bias. The recruitment process is likely to be based on several factors that are not visible to the researchers and therefore cannot be controlled for in a straightforward way.

RDS Estimation

As mentioned previously, researchers use RDS for a variety of estimation goals, such as estimating prevalence or a population fraction. Here, the focus is on estimating the population fraction. Readers interested in prevalence estimation can refer to Handcock, et al. (2014) or Crawford, et al. (2018).

A working example is conducting RDS to estimate the fraction of sex workers who have been trafficked. The researcher performs RDS on the population of sex workers (which is probably difficult to access with other sampling methods due to stigma, fear of prosecution, or other factors) and interviews each person recruited who indicates whether they have been trafficked.

The first estimator to consider would be to simply take the average. That is, take the number of sex workers recruited who report being trafficked and divide by the total number recruited.

This estimator would be biased because some people are more likely to be recruited than others. For example, people with more contacts have more chances to be included in the sample.

To compensate for this, a class of estimators called **Horvitz-Thompson estimators** (or sometimes generalized Horvitz-Thompson estimators) re-weights the average by the inverse of the likelihood that a person is included in the sample (Salganik and Heckathorn. 2004).

Respondents with fewer connections are less likely to be included in a referral chain, and thus have a lower inclusion probability. Therefore, the estimator gives their responses extra weight, proportional to how likely they are to be referred. In this case, the RDS estimator uses the inverse of each respondent's estimated degree—how many reciprocal ties they have to other members of the population of interest—as a correction factor for the estimate. Gile and Handcock (2010) provide a much more thorough discussion of these estimators and their properties.

Considerations in the Context of Trafficking

There are several considerations for successfully implementing RDS, particularly in the context of human trafficking. It is instructive to illustrate some of the difficulties that have arisen in the context of specific trafficking studies.

In their studies of sex workers in three different countries, Simic, et al. (2006) were unable to recruit enough study participants through RDS. They attributed this to several interrelated potential factors: lack of trust, social network structure, restricted movement, and inadequate study incentives.

Lack of trust. General mistrust of official agencies, combined with the sensitivity around sex worker status, tends to reduce participation. Even though the study team worked in advance to build trust and create community ties, tight control by brothels and police crackdowns increased potential participants' reluctance to identify themselves or others. Some participants did not want to reveal their own status as a sex worker to others by recruiting them, since the recruits would find out the eligibility criteria of the study when they were interviewed. Participants sometimes avoided recruiting particular people to avoid identifying them.

Simic, et al. (2006) suggest it may help to run the studies for a longer time to gain more trust, especially if the seed sampling and interviews occur at a location with ongoing services. However, in places where sex workers have little contact with local services, this may not be feasible.

Social network structure. RDS assumes dense, connected networks. By contrast, many of the sex workers in these communities were isolated, either due to restricted movement or because they worked independently and did not tend to reveal their status to others. In Serbia, Simic, et al. (2006) found that most sex workers worked independently and did not connect much with each other. They also found that street sex workers, organized sex workers, and independent sex workers tended not to connect with each other—street sex workers tended to be socially separated by ethnicity, sexuality, and other aspects of identity.

Restricted movement. In Montenegro, brothels were tightly controlled and sex workers were not allowed to leave the premises where they worked. Sex workers with restricted movement were unlikely to be able to receive a coupon or to go to the study locations to participate even if they received a coupon. This seems to have combined with intense policing practices, in retaliation for a recent police HIV infection, to hamper study recruiting.

Inadequate study incentives. The financial incentives provided were not high enough relative to the earnings sex workers could make and the opportunity cost of missing work to participate in the study. Sex workers were more interested in the HIV testing than the financial incentives.

If the incentive is too high and generally appealing beyond the target population, it can encourage people who are not trafficked to attempt to participate. This illustrates the importance of better identifying effective incentives before conducting a study.

Recent Advances and Future Directions

Extensions of RDS have been developed to address the challenges that arise when some of these critical assumptions are not met in practice: network sampling with memory (NSM) and randomized respondent-driven sampling (RRDS).

NSM is an application of RDS that builds on advancements in the mathematics and computer science literature about random walks on graphs proposed by Mouw and Verdery (2012). At a high level, the researcher supervises and strategically directs the recruitment process as it unfolds. This gives the researcher more control, ultimately yielding a more efficient sampling framework than can be attained with traditional RDS.

NSM begins with initial seeds from a convenience sample, all of whom provide a roster of their contacts known to be members of the target population in addition to answering the substantive interview questions chosen by the researcher.

NSM is then implemented as a two-step approach—the **Search** mode followed by the **List** mode. Search mode prioritizes **bridge** nodes—individuals who connect two or more clusters together—to sufficiently explore the network, while List mode ensures that nodes sampled early in the process are not over-represented in the sample.

Search mode takes the network information of respondents and uses the local topography to identify bridge nodes that connect unexplored portions of the network. These nodes are then given priority in the recruitment process. The researcher pre-specifies a threshold that triggers when the network has been sufficiently explored by Search mode.

After Search mode concludes, NSM proceeds to the List mode, which entails two steps: (1) keeping a list of all individuals in the revealed network and (2) sampling from that list with the same cumulative probability for each individual so new additions to the list receive priority.

One of the key advantages of NSM compared to RDS is improved efficiency in searching the network. Given high-quality network data collected from each respondent, the computation and processing costs associated with this method are small. However, collecting high-quality network data from human populations can be prohibitively expensive or logistically infeasible.

NSM might also be impractical in the context of human trafficking where estimates of each respondent's network degree can be highly variable. Real-time supervision and direction of the recruitment process requires additional time and effort, making NSM more challenging to implement in practice.

Respondents recruiting randomly from their contacts is an important assumption of RDS that is often violated in practice. Boudreau, et al. (2023) propose a cellphone-based variant of randomized respondent-driven sampling (RRDS) to address this challenge (see Figure 2).

The set-up resembles RDS in that researchers begin with a convenience sample for the initial seeds. From each seed, the researcher collects a list of phone numbers for their contacts believed to also be in the target population. The researcher

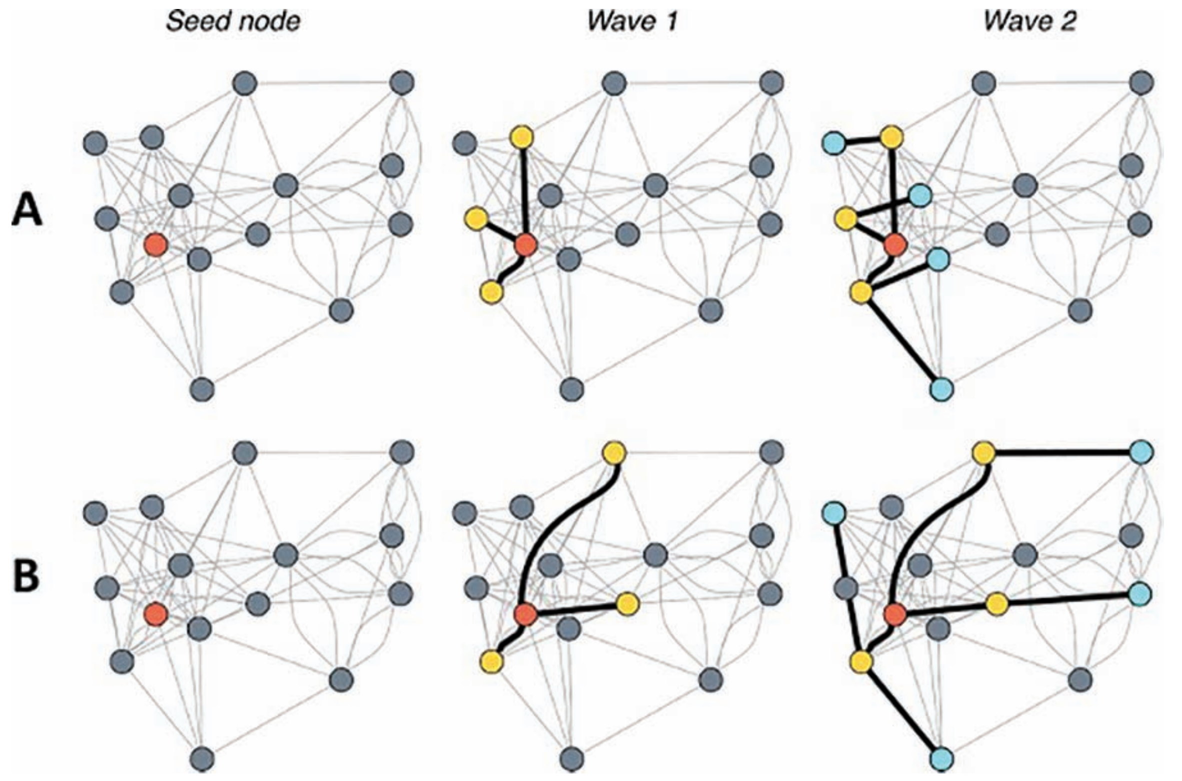

Figure 2. Each graph shows an entire example network. Nodes recruited in waves 1, 2, and 3 shown in red, yellow, and light blue, respectively. Gray nodes are never recruited, and bolded paths indicate directed recruitment links. Row A demonstrates process of generating traditional RDS recruitment tree. Row B shows how randomizing recruitment of respondents' contacts for each wave results in better coverage.

then chooses a random subset of the respondents' contacts, administers the survey, and collects a list of phone numbers of their contacts for the subsequent wave.

This process is then repeated until the desired sample size is reached.

RRDS has several advantages relative to traditional RDS. First, it provides more reliable randomization by introducing randomness at each wave of the recruitment process, making the method closer to the ideal RDS assumption of random selection among contacts. It also provides more control over recruitment.

Second, it is phone-based rather than venue-based. RRDS uses phone surveys and does not require in-person interview sites. This makes RRDS a useful option for applications like human trafficking where restricted mobility makes it difficult to recruit respondents in person. The costs (in time and effort) associated with administering the randomization are relatively low; it is as simple as drawing a simple random sample from each respondent's list of phone numbers before proceeding with the next wave of recruitment.

RRDS will not be well-suited to every context. The most obvious requirement is the need to work with a population where individuals have access to cellphones and the autonomy to use those phones as they wish. Individuals in the group of interest without access to a phone will be excluded. RRDS also relies on individuals having the phone numbers (or saving contacts) of other members of the group of interest.

Conclusion

Human trafficking is a complex, stigmatized, secretive, and constantly evolving phenomenon. Respondent-driven sampling offers several advantages over traditional survey methods for studying this population, but its success requires building trust with the relevant communities, providing well-informed and motivating incentives, and understanding and accounting for several aspects of the social network structure of the people being trafficked and their surrounding community. The literature provides some potential suggestions for addressing these concerns, as well as promising directions for future research.

See the online supplemental materials (*https://chance.amstat.org/2023/11/supp-material-36-4/*) for an annotated bibliography and additional concerns illustrated by other studies in the literature.

Further Reading

Boudreau, L., Heath, R., McCormick, T. H., and Visokay, A. 2023. Randomized respondent-driven sampling: A cellphone-based approach to sampling in disaster settings. Working Paper, University of Washington.

Crawford, F. W., Wu, J., and Heimer, R. 2018. Hidden population size estimation from respondent-driven sampling: a network approach. *Journal of the American Statistical Association* 113:755–766.

Gile, K. J. and Handcock, M. S. 2010. Respondent-driven sampling: an assessment of current methodology. *Sociological Methodology* 40:285–327.

Handcock, M. S., Gile, K. J., and Mar, C. M. 2014. Estimating hidden population size using respondent-driven sampling data. *Electronic Journal of Statistics* 8:1491.

Mouw, T., and Verdery, A. M. 2012. Network sampling with memory: A proposal for more efficient sampling from social networks. *Sociological Methodology* 42:206–256.

Salganik, M. J., and Heckathorn, D. D. 2004. Sampling and estimation in hidden populations using respondent-driven sampling. *Sociological Methodology* 34:193–240.

Schroeder, E., Edgemon, T. G., Aletraris, L., Kagotho, N., Clay-Warner, J., and Okech, D. 2022. A review of prevalence estimation methods for human trafficking populations. *Public Health Reports* 137:46S–52S.

Simic, M., Johnston, L. G., Platt, L., Baros, S., Andjelkovic, V., Novotny, T., and Rhodes, T. 2006. Exploring barriers to "respondent-driven sampling" in sex worker and drug-injecting sex worker populations in Eastern Europe. *Journal of Urban Health* 83:6–15.

Acknowledgment: The authors are supported by the National Institute of Mental Health of the NIH under Award Number DP2MH122405 and by the Center for Statistics and the Social Sciences at the University of Washington. The content is solely the responsibility of the authors and does not necessarily represent the official views of the NIH.

About the Authors

Jessica P. Kunke is a PhD student in statistics at the University of Washington. Her research interests include survey methodology, network models, and spatiotemporal statistics with environmental and social applications. She has an MS in statistics from the University of Chicago, for which she conducted her master's thesis at the Argonne National Laboratory, and teaches coding, math, and data management workshops for researchers in and outside of academia.

Tyler McCormick is a professor of statistics and sociology at the University of Washington, core faculty member in the Center for Statistics and the Social Sciences, and senior data science fellow at the eScience Institute who develops statistical models for inference and prediction in scientific settings where data are sparsely observed or measured with error. He holds a PhD in statistics with distinction from Columbia University and has received the NIH Director's New Innovator Award, NIH Career Development (K01) Award, Army Research Office Young Investigator Program Award, and a Google Faculty Research Award. During the 2019–2020 academic year, he was a visiting faculty researcher at Google People+AI Research (PAIR). He is the former editor of the *Journal of Computational and Graphical Statistics* (JCGS) and a Fellow of the ASA.

Adam Visokay is a PhD student in the Department of Sociology at the University of Washington, focused on networks, survey methods, NLP, and computational social science. He spent two years as a researcher at the Center for Statistics and the Social Sciences. Visokay has a BA in history and economics from the University of Virginia and an MA in economics from Syracuse University.