Randomized Respondent Driven Sampling: A Cellphone Based Approach

Adam Visokay

A thesis
submitted in partial fulfillment of the
requirements for the degree of

Master of Arts

University of Washington

2024

Committee:
Tyler H. McCormick
Zack W. Almquist

Program Authorized to Offer Degree:
Department of Sociology

University of Washington

Abstract

Randomized Respondent Driven Sampling: A Cellphone Based Approach

Adam Visokay

Chair of the Supervisory Committee:
Tyler H. McCormick
Department of Department of Sociology

Respondent-Driven Sampling (RDS) is a widely used method for accessing hidden populations when more traditional survey techniques may not be feasible. However, the reliance on non-random peer recruitment introduces substantial bias, particularly in the presence of homophily. This paper introduces Randomized Respondent-Driven Sampling (RRDS), a novel, cellphone-based adaptation that incorporates researcher-controlled randomization into the recruitment process. RRDS preserves the network-based advantages of RDS while mitigating selection bias by decoupling recruitment from respondent preferences. Through simulation on synthetic networks with high homophily and an empirical application among Bangladeshi garment workers during the COVID-19 pandemic, RRDS demonstrates superior performance in sample representativeness, recruitment efficiency, and convergence to population parameters. The empirical study also reveals gendered constraints in referral behavior, underscoring the importance of context-sensitive implementation. RRDS offers a scalable, remote-compatible alternative for sociological research in hard-to-reach populations, or in populations that are not traditionally hard to reach, but become temporarily inaccessible, such as the case of garment workers during the pandemic.

1 Introduction

Researchers studying known populations produce representative samples using probability sampling procedures drawn from complete sampling frames. These data serve as the foundation for empirical work across sociology, political science, economics, and public health. Sampling frames—such as voter registration lists, customer databases, census registries, or administrative records—provide comprehensive and detailed inventories of individuals within target populations. This enables researchers to weight their samples with appropriate inclusion probabilities, ensuring representativeness and facilitating valid population-level inferences. Such representative samples are the gold standard for statistical inference, though collecting data in this way is not always feasible. Sampling frames are often prohibitively expensive, increasingly difficult to maintain due to falling response rates, or otherwise impossible to construct in many settings.

These challenges are particularly acute when studying sensitive or hidden populations—such as unhoused individuals, injection drug users, sex workers, undocumented immigrants, or other marginalized groups—where no detailed census exists and traditional probability sampling becomes infeasible. Researchers in sociology, public health, demography, and allied disciplines who work with such populations frequently rely on Respondent Driven Sampling (RDS). RDS is a non-probability snowball sampling procedure that leverages social network structures to recruit participants in the absence of a sampling frame. Initially proposed by Heckathorn [1997], RDS has become a standard approach for accessing hard-to-reach populations worldwide.

Despite its widespread adoption, the principal challenge associated with RDS data lies in the inherent difficulty of quantifying recruitment bias, rendering subsequent statistical inference potentially unreliable. Furthermore, key assumptions that underlie RDS methodology frequently fail to hold in practice. For example, the assumption that respondents recruit randomly from their networks is often violated due to homophily in recruitment patterns, where participants tend to refer others with similar characteristics to themselves [Liu et al., 2012]. Other problematic assumptions include the requirement that respondents accurately report

their network size, that sampling occurs with replacement, and that the target population forms a single connected component [Gile and Handcock, 2010]. When these assumptions are violated, traditional RDS estimators can produce biased results, particularly in populations with strong within-group ties or when sampling a small fraction of the total population. This limitation is well-documented in the RDS literature [Gile and Handcock, 2010, Salganik and Heckathorn, 2004, Volz and Heckathorn, 2008], yet viable alternatives for accessing these populations have remained elusive. Furthermore, traditional survey methods face a broader crisis as response rates continue to decline globally[Meyer et al., 2015], exacerbating the challenges of obtaining representative samples even in settings where sampling frames exist.

Recent global events, particularly the COVID-19 pandemic, have further complicated data collection efforts, necessitating rapid adaptation in survey methodology. Physical distancing requirements, mobility restrictions, and heightened health concerns have disrupted traditional face-to-face interviewing approaches that RDS typically relies upon. These circumstances have accelerated the need for innovative, remote sampling methods that maintain methodological rigor while adapting to contemporary constraints.

In response to these challenges, I introduce Randomized Respondent Driven Sampling (RRDS), a cell phone-based approach that builds upon traditional RDS while addressing several of its key limitations. This method offers numerous advantages in the current survey environment: (1) it reduces physical contact requirements, making it particularly suitable for disaster settings and public health emergencies; (2) it leverages increasingly ubiquitous mobile phone technology, even among marginalized populations; (3) it introduces randomization procedures that help quantify and mitigate recruitment bias; and (4) it provides a scalable platform for reaching traditionally inaccessible populations while maintaining rigorous statistical properties. Drawing from both a rigorous simulation study and empirical data collected via cell phone from Bangladeshi garment workers during the COVID-19 pandemic [Boudreau et al., 2024], I demonstrate the effectiveness of this randomized recruitment approach.

The remainder of this paper is organized as follows: in Section 2 I review the theoretical

foundations and limitations of traditional RDS. In Section 3 I detail my proposed randomized recruitment methodology (RRDS) and in Section 4 I demonstrate performance with a simulation study. Section 5 describes the empirical data collection procedure and 6 details the results of the analysis. Finally, I discuss implications for survey research methodology and directions for future refinement of the RRDS approach in 7.

2 Respondent Driven Sampling

RDS is both a sampling strategy and statistical estimation procedure. RDS data are collected using a snowball mechanism where members of the target group leverage their relationships to other group members in order to increase overall participation. Responses are then weighted by the inverse of each respondent's degree to correct for differential recruitment probabilities throughout the network. The process begins with a convenience sample of individuals known to be in the population of interest. This initial sample is known as the *seed* sample; the individual recruits are referred to as *seeds*. Seeds are typically solicited with a primary incentive to participate in the study, normally a monetary payment or material reward. Then each seed is offered a secondary incentive to recruit some number of their peers into the study. The seeds' peers are offered the primary incentive to join the study and the secondary incentive to recruit their peers into the study. This recruitment cycle repeats, incorporating new recruits into the study with each wave, until the desired sample size is reached. This sampling is done with replacement, meaning the same individual could be recruited into the sample more than once if they are referred by multiple individuals.

It is important for the researcher to keep track of the order of recruitment for subsequent analysis. Each individual is assigned a unique identifier the first time they are recruited into the sample by the researcher so that the recruitment chain can be traced from the original seed through to the final respondent.

2.1 Theoretical Foundations and Key Assumptions

The statistical theory underlying RDS estimation rests on several critical assumptions that enable inference about the target population from non-probability samples [Heckathorn, 1997, Salganik and Heckathorn, 2004, Volz and Heckathorn, 2008]. These assumptions are:

1. **Respondents recruit randomly from their personal networks**: Perhaps the most consequential assumption of RDS is that participants select recruits at random from their eligible network connections. This assumption is critical for the Markov process model upon which RDS estimation depends, allowing researchers to model recruitment as a random walk on the social network. However, substantial evidence suggests this assumption is frequently violated due to homophily, convenience selection, or other social factors that produce non-random recruitment patterns [Liu et al., 2012]. I will address this in Section 3.

2. **Respondents can accurately report their personal network size**: RDS estimators require participants to report the number of eligible individuals in their personal network (their "degree"). These self-reported degrees are used to weight observations to account for differential recruitment probabilities. Inaccurate degree reporting—whether due to recall bias, social desirability, or definitional ambiguity—can significantly bias population estimates [Gile and Handcock, 2010].

3. **Sampling occurs with replacement**: Standard RDS estimation theory assumes sampling with replacement, meaning individuals could theoretically be recruited multiple times. In practice, most RDS studies implement sampling without replacement, prohibiting multiple recruitment of the same individual. This discrepancy becomes particularly problematic when sampling fractions are large relative to the size of the target population [**?**].

4. **The target population forms a single connected component**: RDS assumes that all members of the target population are connected, directly or indirectly, through a single social network. If the population contains disconnected subgroups, seeds from

all subgroups must be included for valid inference about the entire population.

5. **Recruitment reaches the full network**: RDS theory posits that after sufficient recruitment waves, the sample composition stabilizes regardless of seed selection, reaching a sampling equilibrium that reflects the target population. This requires chains to progress through multiple waves, typically six or more, to overcome the bias introduced by the non-random selection of seeds [Salganik and Heckathorn, 2004]. It also assumes that the target population is fully connected.

6. **Reciprocity of relationships**: The estimators assume that network ties are reciprocal—if Person A reports a relationship with Person B, then Person B would also report a relationship with Person A. Directional or asymmetric relationships can undermine this assumption.

When these assumptions are violated, as frequently occurs in real-world implementations, RDS estimators may produce biased population estimates with misleading confidence intervals [Gile and Handcock, 2010]. The degree of bias depends on the specific nature and extent of assumption violations, characteristics of the target population's social network, and properties of the variables being estimated.

Despite these limitations, RDS has become the predominant methodology for studying hidden populations due to its practical advantages in recruitment and its theoretical framework for statistical inference. However, the fragility of its underlying assumptions highlights the need for methodological innovations that can address these limitations while maintaining the practical advantages of peer-driven recruitment. I propose a practical implementation to address homophily in recruitment in the next Section.

3 Randomized-Respondent Driven Sampling

Randomized Respondent-Driven Sampling (RRDS) directly addresses one of the most problematic RDS assumption violations—non-random recruitment. By introducing an explicit randomization step controlled by the researcher, RRDS maintains the network-based

recruitment advantages while mitigating biases from homophily-driven selection. The RRDS procedure works as follows: begin with a set of initial seeds, as in traditional RDS. Each respondent provides contact information for multiple members of their social network who meet the study criteria (typically phone numbers for m contacts). Instead of allowing respondents to choose which contacts to recruit, we randomly select a subset (s) of these contacts to follow up with. This process repeats through multiple waves until the desired sample size is reached.

This approach keeps respondents engaged in identifying eligible participants from their networks but removes their ability to selectively recruit contacts based on similarity or other preferences. The randomization step serves as a barrier against homophily bias while preserving the network-based recruitment mechanism that makes RDS effective for reaching hidden populations. Figure 1 illustrates RRDS is particularly effective in settings where: The population has moderate to high network density, respondents are willing and able to provide multiple network contacts, and respondents can be incentivized to participate and refer others.

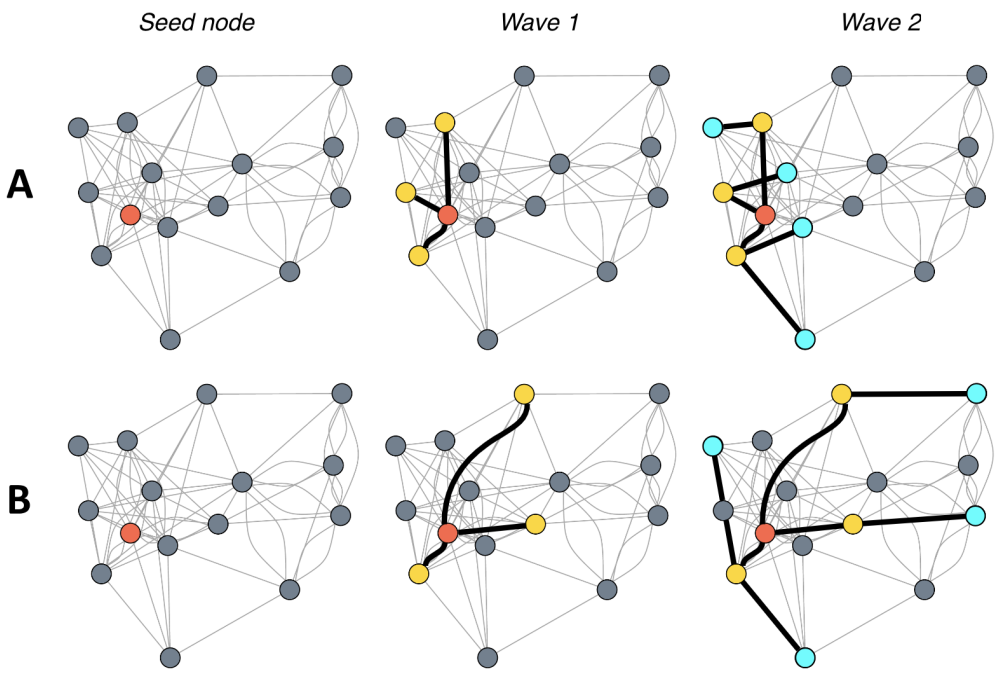

Figure 1: A) above illustrates tradtional RDS recruitment with homphily. When contacts are not referred at random, the recruitment chains can fail to achieve full coverage of the network. B) illustrates random referral by respondents, where wider coverage of the network can be achieved with the same number of recruitment waves.

For a phone-based respondent-driven survey, I suggest the following procedure. Assuming each respondent i has degree k in the target population: ask respondent i to provide phone numbers for m of their k contacts where $m < k$. Then, the researcher randomly recruits a subset s from the list of m contacts where $s < m$. Repeat these steps with each new respondent for each wave of recruitment until reaching the desired sample size. This randomization strategy is most effective when each respondent is both able and willing to provide a substantial list of contacts. The method is less effective in sparse networks (where respondents have few connections) or when respondents are reluctant to share contact information. As I will discuss in Section 6, a small number of referrals per respondent presents a major limit to the utility of this approach.

4 Simulation Study

To evaluate the effectiveness of the proposed Randomized Respondent-Driven Sampling (RRDS) approach compared to traditional RDS, I conducted a simulation study using a synthetic network with known population parameters. This controlled environment allows for direct comparison of how each method converges to known population values.

I generated a synthetic social network with 10,000 nodes representing individuals in a target population. The age distribution was set to be normally distributed with a mean of 41.5 years and standard deviation of 10 years, bounded between 18 and 65 years. Gender distribution was configured at 70% female and 30% male, implemented as a binary attribute. Network connectivity was established with an average degree of 2 connections per individual and a standard deviation of 2. A high homophily level of 0.9 was deliberately incorporated, meaning individuals have a strong tendency to connect with others sharing similar attributes in terms of gender and age. This high homophily parameter creates challenging conditions for sampling, reflecting real-world social networks where similar individuals tend to cluster together—a particularly important factor as it directly affects recruitment patterns in traditional RDS implementations.

4.1 Seed Selection

To create realistic conditions that demonstrate the resilience of both methods, I deliberately selected non-representative seeds, consisting of 76 initial participants – predominantly young men under 22 years old. This biased seed selection creates a significant initial deviation from the true population parameters of 41.5 years mean age and 70% female. Both sampling methods RDS and RRDS began with identical seeds to ensure fair comparison, allowing the recruitment process itself to be the only variable between the two approaches.

4.2 Sampling Methods Implementation

The traditional RDS algorithm was implemented to simulate homophily-driven recruitment patterns. For each participant in a given wave, the algorithm first identified all network

neighbors who had not yet been recruited into the sample. These eligible neighbors were then deliberately sorted by similarity to the recruiting participant, prioritizing same gender first and then closest age. This sorting step explicitly modeled the tendency of individuals to refer others similar to themselves. From these sorted neighbors, up to three were selected for recruitment into the next wave, always choosing the most similar eligible neighbors first. This process continued wave by wave until either reaching 12 waves or exhausting all possible recruits.

In contrast, the RRDS algorithm followed a different selection process while maintaining the same network structure and recruitment parameters. For each participant, all eligible neighbors (those not previously recruited) were identified. However, instead of sorting by similarity, the algorithm used random selection to choose up to three neighbors for recruitment into the next wave. This was implemented using R's sample() function without replacement, giving each eligible neighbor an equal probability of selection regardless of their attributes. This critical difference—random selection versus similarity-based selection—constitutes the core innovation of the RRDS method. Both algorithms used identical stopping criteria, proceeding through the same maximum number of waves and recruiting the same maximum number of participants per recruiter.

4.3 Performance Metrics

For both implementations, I tracked wave-by-wave statistics on recruitment success, demographic composition, and convergence to population parameters. After each wave, I calculated the number of new participants, their mean age, and proportion female. Cumulative statistics were also computed at each step to monitor overall sample growth and representativeness. These metrics allowed direct comparison of how quickly each method approached the known population parameters of 41.5 years mean age and 70% female distribution, as visualized in Figures 2 and 3.

4.4 Simulation Results

The results revealed substantial differences between the two methods. As shown in Figure 2, traditional RDS (panel a) showed persistent bias toward younger, male participants due to homophily-driven recruitment, with slow convergence toward population parameters across the 12 waves. The randomized approach (panel b) demonstrated more rapid demographic convergence and significantly larger recruitment success per wave, with points consistently larger in the RRDS visualization compared to the traditional RDS.

Figure 3 provides a compelling view of the cumulative effects of these recruitment differences. Traditional RDS implementation reached a final sample size of approximately 900 participants, while RRDS yielded a substantially larger sample of about 1,908 participants – more than double that of traditional RDS. As shown in panel (b), the RRDS approach not only produced larger samples but also achieved better approximation of the true population parameters, with the cumulative gender distribution and mean age more closely matching the dashed reference lines indicating population values.

The difference in recruitment efficiency is particularly striking, with RRDS yielding an average of approximately 172 participants per wave compared to just 70 for traditional RDS. This enhanced recruitment efficiency, coupled with improved representativeness, demonstrates the practical advantages of introducing randomization into the recruitment process for a population with high homophily. By mitigating the impact of initial seed bias and network homophily, RRDS provides a more robust sampling methodology for hidden populations while maintaining the network-based recruitment mechanism that makes RDS valuable for reaching these communities.

5 Empirical Data: Garment Workers in Bangladesh

The Bangladesh garment industry context offers an unusual opportunity to evaluate network-based sampling methodologies. Unlike most hidden populations studied with RDS, garment workers are typically accessible through traditional sampling methods. However, the COVID-19 pandemic created conditions that temporarily rendered this population difficult

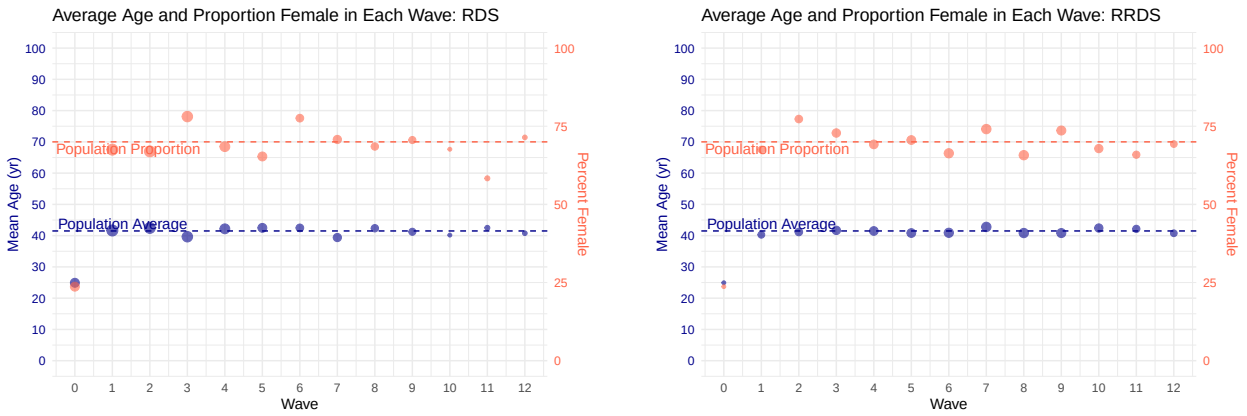

(a) Traditional RDS

(b) Randomized RDS

Figure 2: Point size corresponds to the number of respondents recruited into each wave. Both simulations begin with the same non-representative seed ($n = 76$). Randomized recruitment yields more recruits per wave on average ($\bar{n}_{RRDS} \approx 172$ vs $\bar{n}_{RDS} \approx 70$) and converges in fewer waves than traditional RDS.

to reach through conventional approaches, similar to traditionally hidden populations. This unique circumstance provides an ideal natural experiment for testing RRDS, as I have access to a representative baseline from pre-pandemic surveys against which to compare results to a "known" baseline.

The data collection process incorporated a parallel testing design, running both traditional RDS and RRDS arms simultaneously with the same starting conditions. Boudreau et al. began with seed participants drawn from two previous representative surveys of garment workers: a 2017 survey conducted by the Bangladesh Institute of Governance and Development (BIGD) that included 1,500 workers, and a smaller 2020 pilot survey. These seed samples were geographically representative of garment-producing areas in the Dhaka Division, which encompasses approximately 80% of Bangladesh's garment factories.

The RRDS methodology was implemented through a clear protocol. All seed participants were asked to provide contact information for multiple acquaintances who worked in the garment industry, with small monetary incentives for each referral provided. In the traditional RDS arm, respondents were explicitly asked which contacts they felt most com-

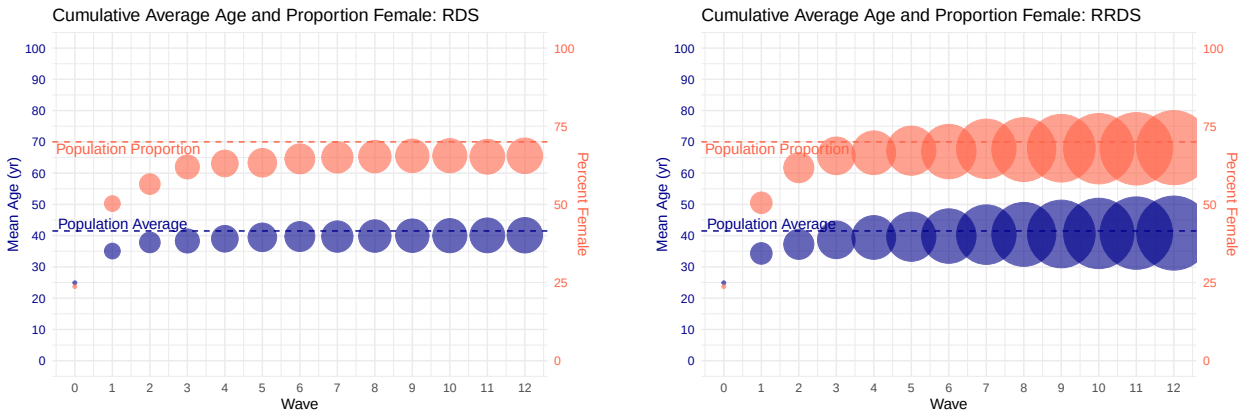

(a) Traditional RDS

(b) Randomized RDS

Figure 3: Cumulative sample from RRDS is larger than RDS ($N_{RRDS} \approx 1908$ vs $N_{RDS} \approx 900$) and better approximates the underlying network characteristics.

fortable referring, and these preferred contacts were followed up with. In the RRDS arm, contacts were randomly selected from the provided list without regard to the respondent's stated preferences.

The survey was conducted across five waves, beginning with the seed sample and proceeding through four subsequent waves of recruitment. Between the initial seed wave and the first full survey wave, an intermediate links survey was conducted that allowed tracking of referral patterns in detail. This structure enabled observation of how the demographic composition evolved differently between the traditional RDS and randomized RDS arms.

This empirical application revealed both strengths and challenges of implementing RRDS in real-world conditions. However, unexpected patterns in referral behavior were also encountered, particularly among female respondents who were less likely to provide multiple referrals. This was not anticipated and limits the generalizability of the conclusions, see Section 7 for a discussion about why females behaved differently than males in this context. This necessitated adaptive sampling protocols to maintain viable recruitment chains, highlighting the importance of flexibility when implementing novel sampling methodologies in field conditions. Ultimately, all of the results in section 6 focus only on male respondents because females referred less than two people on average – the modal number of referrals for

female respondents was zero.

Even despite these limitations, this study design provides empirical validation of the simulation findings while also revealing practical considerations for implementation. The ability to compare both methods against a known population distribution represents a significant methodological contribution, as most RDS studies lack a "ground truth" against which to evaluate their results. This real-world application demonstrates that RRDS can effectively address homophily-based recruitment bias while maintaining the network-based advantages of traditional RDS approaches.

6 Results

Demographic proportions for the male-only network-based samples were estimated using two different weighting procedures, the RDS Tree Bootstrap Method [Baraff et al., 2016] and traditional RDS weights. These estimates are compared against the representative estimates from the 2017 BIGD/LSE survey of 1500 garment workers [Kabeer et al., 2019] to assess the efficacy of the design. I compare estimates for proportion of native born workers, marital status, primary and secondary education, and presence of young children in Figures 4, 5, 6, 7, and 8. RDS Tree Bootstrap weighted estimates are shown in green and traditional RDS weights are shown in purple. RDS point estimates are shown with dots and RRDS estimates are shown with triangles.

Despite the limitations discussed in Section 5, these empirical results consistently demonstrate that RRDS provides improvements over traditional RDS across multiple demographic characteristics. For marriage status (Figure 5), RRDS yields substantially less biased estimates than traditional RDS. When traditional RDS weights are applied, the RRDS estimate is not statistically different from the true population parameter, while the RDS estimate shows significant deviation. This suggests that the randomization step effectively mitigates homophily-based selection bias for this characteristic.

Educational attainment shows more nuanced results. For primary school education (Fig-

ure **??**), the RRDS estimate using RDS Tree Bootstrap weights shows greater bias than the RDS estimate, though both methods retain statistical significance. With traditional RDS weights, both methods produce similarly biased point estimates that differ significantly from the baseline. For secondary education (Figure 7), both methods exhibit positive bias relative to the baseline, though RRDS demonstrates reduced uncertainty when RDS Tree Bootstrap weights are applied. With traditional RDS weights, the RRDS estimate shows less bias than the RDS estimate, though neither reaches statistical significance.

The proportion of respondents with children under 5 years old (Figure 8) reveals particularly strong homophily effects in traditional RDS. While the RRDS estimate using RDS Tree Bootstrap weights shows considerable uncertainty, the RRDS estimate with traditional RDS weights is unbiased and statistically significant. This demonstrates RRDS's effectiveness in counteracting homophily-driven recruitment patterns in family structure characteristics.

For the native-born proportion (Figure 4), both methods struggle to produce reliable estimates due to the small baseline proportion in the population and limited sample size of the network-based sampling. This highlights an important limitation: very rare population characteristics may remain difficult to estimate precisely even with methodological improvements to RDS.

Across all demographic dimensions examined, RRDS demonstrates either comparable or superior performance to traditional RDS, particularly when traditional RDS weights are applied. The consistency of these improvements across different demographic characteristics suggests that the randomization mechanism effectively addresses non-random selection bias while maintaining the practical advantages of network-based recruitment. These empirical results complement and reinforce the findings from the simulation in Section 4, providing real-world validation of RRDS's theoretical advantages.

7 Discussion

Randomized Respondent-Driven Sampling (RRDS) represents a methodological bridge between conventional probability sampling approaches and network-based recruitment methods typically used for hidden populations. By introducing researcher-controlled randomization into the recruitment process, RRDS preserves the advantages of network-based recruitment while significantly reducing the biases introduced by homophily and non-random selection [Salganik, 2012, Goel, 2010].

My simulation study provides compelling evidence that RRDS produces larger, more representative samples with faster convergence to population parameters compared to traditional RDS. However, the real-world application with Bangladeshi garment workers revealed important contextual challenges, particularly related to gender dynamics that affected referral patterns.

A particularly striking finding was the significantly lower rate at which female respondents provided multiple referrals compared to their male counterparts. This gender disparity in referral behavior has important implications for implementing RRDS in contexts with strong gender hierarchies. The Bangladesh garment industry presents a complex gender landscape worth exploring to understand these patterns. Despite women comprising approximately 60-80% of the garment workforce [Kabeer et al., 2019], their social position outside the workplace remains constrained by traditional gender norms.

Several factors likely contributed to women's reduced willingness or ability to provide referrals. First, phone ownership patterns in Bangladesh exhibit strong gender disparities, with many households sharing a single phone that is primarily controlled by male family members. Even when women work outside the home, their communication technologies are often monitored or controlled by husbands or fathers [Boudreau et al., 2024]. This technological constraint directly impacts women's ability to maintain independent social networks that can be mobilized for survey recruitment.

Second, cultural norms regarding appropriate social connections for women may limit

their network diversity. While male garment workers frequently socialize with colleagues outside work hours, female workers often return directly home after shifts due to safety concerns and family responsibilities. This constrains women's opportunities to form and maintain the broad occupational networks that facilitate successful referral chains in RDS methodologies.

Third, the post-Rana Plaza environment has heightened sensitivities around garment industry working conditions [Kabeer et al., 2019]. Female workers, who typically occupy more vulnerable positions in factory hierarchies, may hesitate to refer colleagues into a survey addressing potentially sensitive topics out of fear of workplace repercussions. As one female respondent noted during recruitment, "I know other workers, but I don't want to create problems for them."

These gender dynamics present important methodological considerations for implementing RRDS in contexts with similar gender hierarchies. Future applications should develop gender-sensitive recruitment protocols that might include different incentive structures for male and female participants, privacy assurances tailored to gender-specific concerns, and potentially modified referral requirements that acknowledge women's constraints in providing multiple contacts.

Beyond addressing gender dynamics, future research should explore optimal incentive structures to maximize referral compliance, examine the performance of RRDS across different network topologies, and develop specialized variance estimators that account for the unique features of randomized network recruitment. Additionally, expanding the approach to incorporate digital recruitment channels may further enhance its applicability in increasingly connected populations, while remaining attentive to gender-based digital divides. In summary, RRDS represents a valuable addition to the methodological toolkit for researchers working with difficult-to-sample populations [Heckathorn and Cameron, 2017, Sarah Raifman, 2022], offering a pragmatic compromise between statistical rigor and practical feasibility in challenging research environments. However, its implementation must be attentive to contextual factors, particularly gender dynamics, that shape referral patterns and ultimately

determine sample composition.

8 Figures

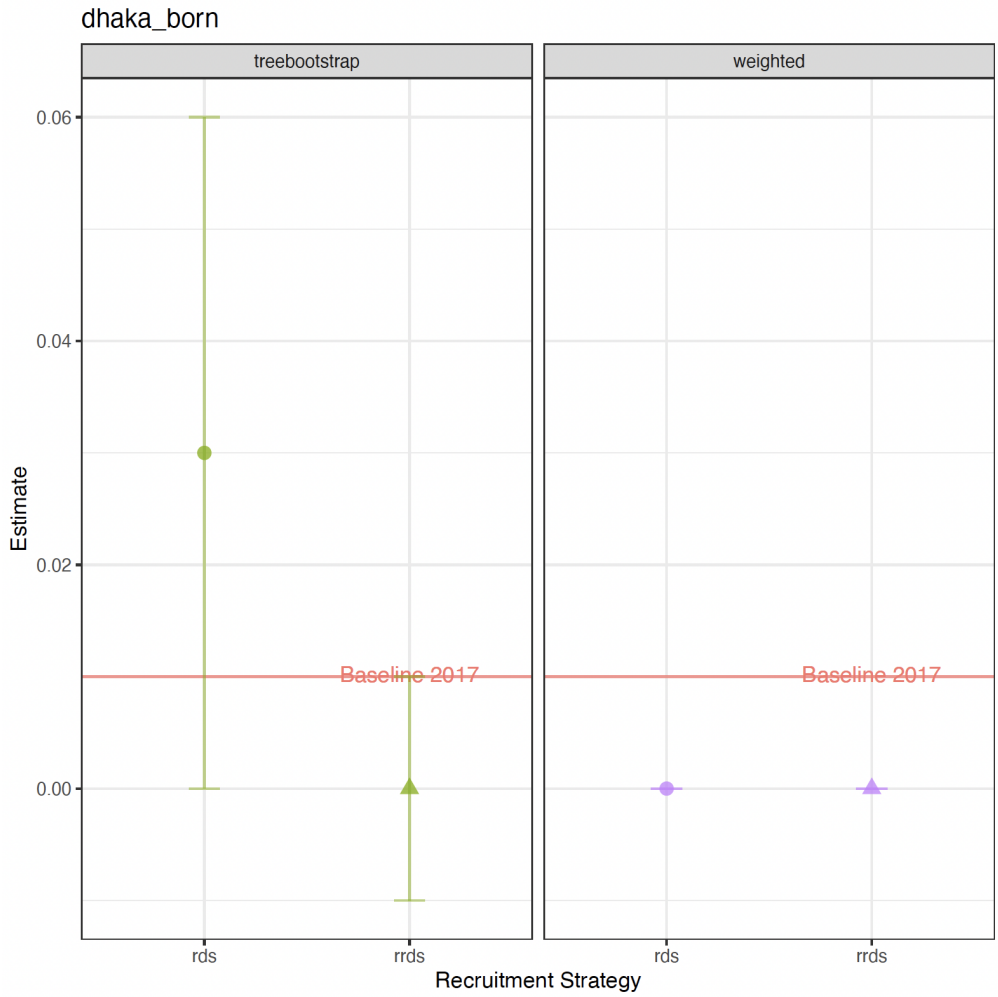

Figure 4: Proportion of respondents who are native born to the Dhaka region across recruitment waves. The horizontal dashed line represents the population parameter from the representative 2017 survey. This is a very small proportion of the population, and due to the limited sample size of from the network-based sampling it is not possible to recover a reliable estimate of this quantity with either method.

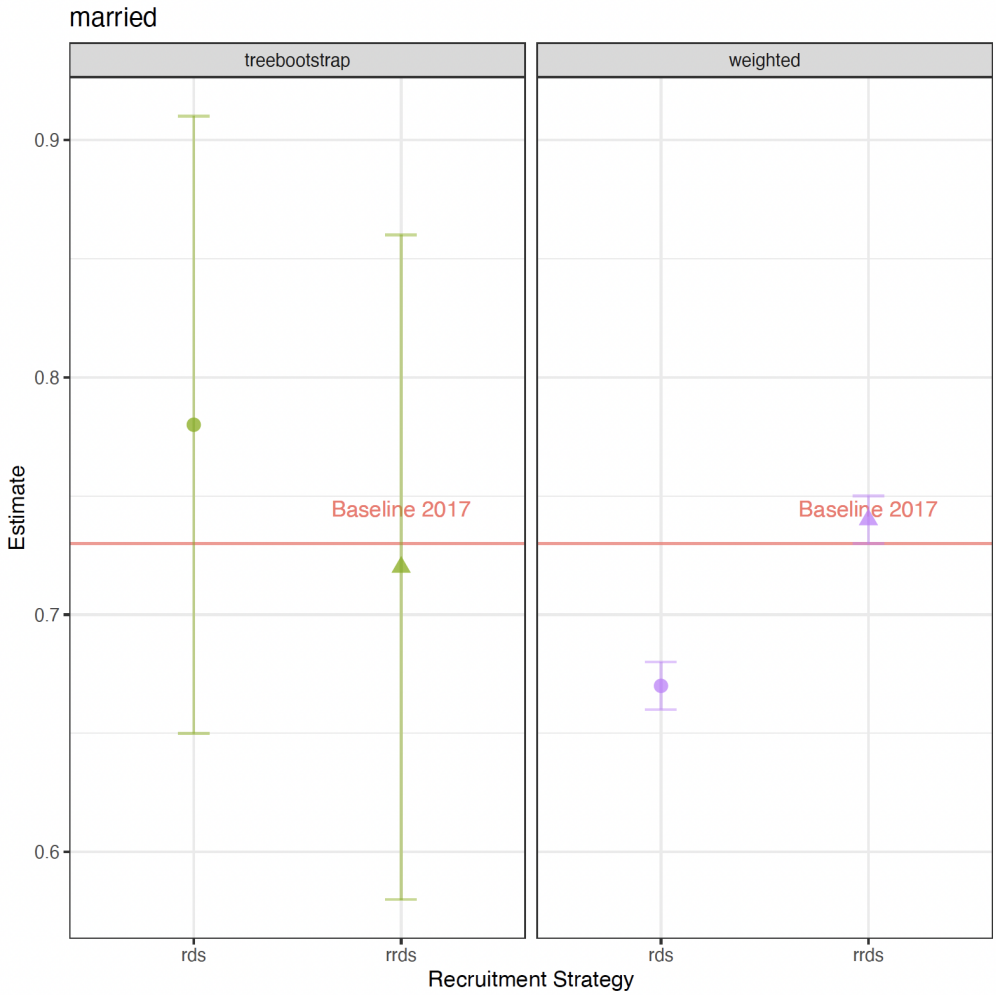

Figure 5: Proportion of married respondents across recruitment waves. RRDS point estimates are far less biased than traditional RDS. With traditional RDS weights, the RRDS estimate is not statistically different from the true population parameter, while the RDS estimate is.

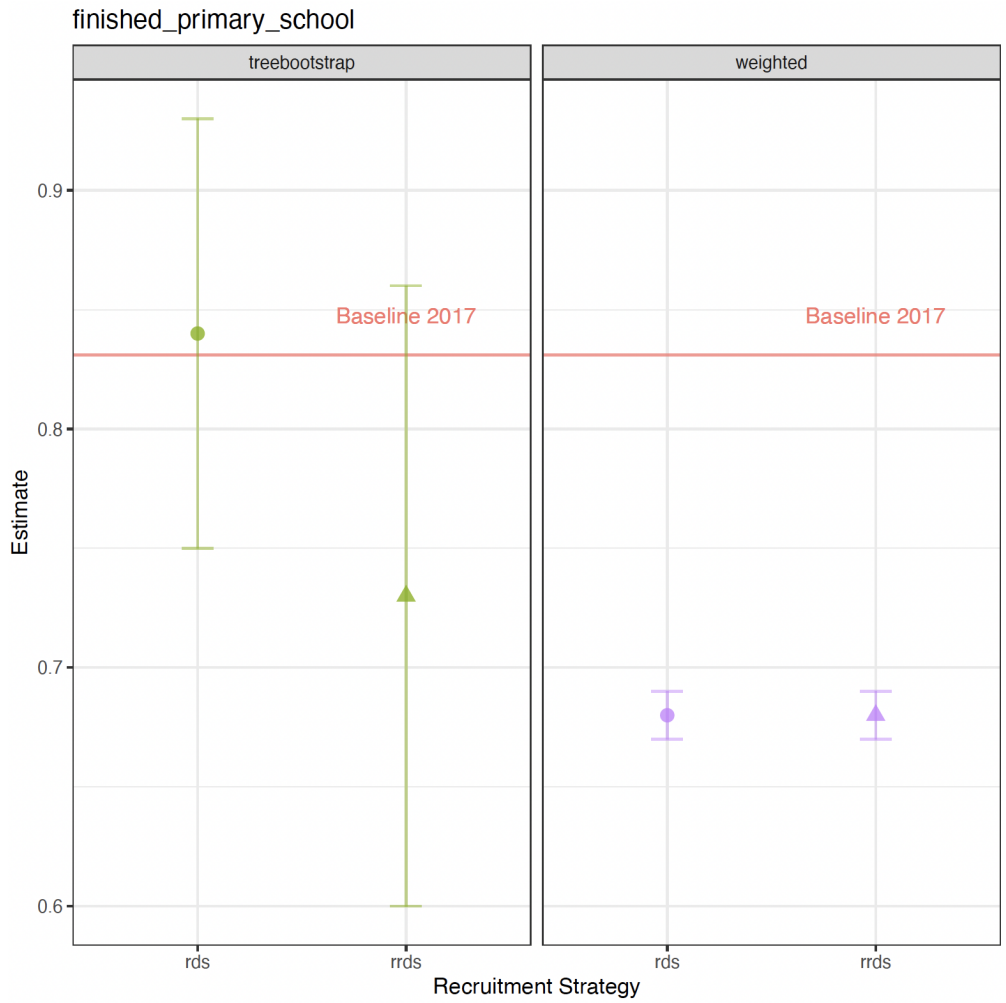

Figure 6: Proportion of respondents with primary school education across recruitment waves. The RRDS estimate with RDS Tree Bootstrap weights is more biased than the RDS estimate, but not enough to lose statistical signficance. For the traditional RDS weights, both methods give similarly biased point estimates that are statistically different from baseline.

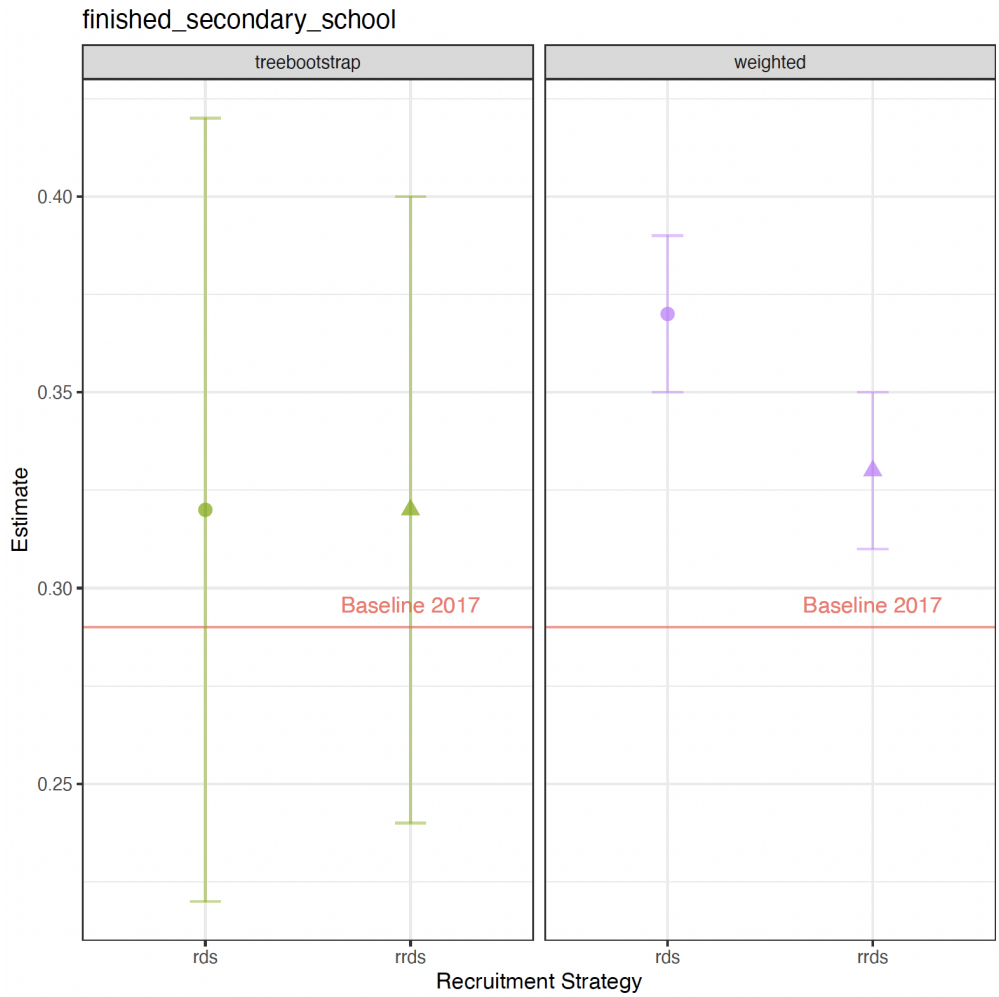

Figure 7: Proportion of respondents with secondary school education across recruitment waves. Both methods are positively biased relative to baseline, though the RRDS estimate has less uncertainty than the RDS estimate when RDS Tree Bootstrap weights are used. For traditional RDS weights, the RRDS estimate is less biased than the RDS estimate, though neither are statistically significant.

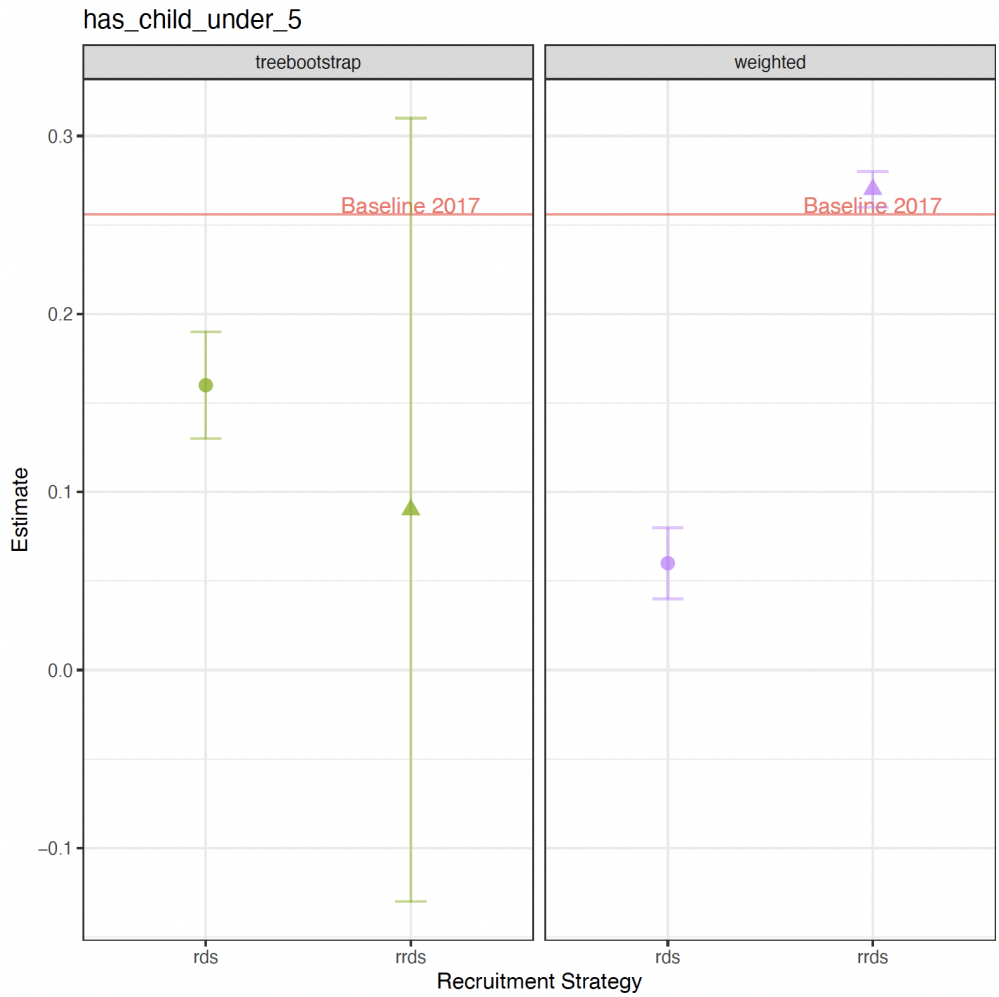

Figure 8: Proportion of respondents with children under 5 years old currently living at home across recruitment waves. This demographic characteristic shows particularly strong homophily effects in traditional RDS. We find considerable uncertainty around the RRDS estimate with RDS Tree Bootstrap weights. However, with traditional RDS weights, the RRDS estimate is unbiased and statistically signficant.

References

Aaron J. Baraff, Tyler H. McCormick, and Adrian E. Raftery. Estimating uncertainty in respondent-driven sampling using a tree bootstrap method. *Proceedings of the National Academy of Sciences of the United States of America*, 113(51):14668–14673, December 2016. doi: 10.1073/pnas.1617258113. URL https://doi.org/10.1073/pnas.1617258113.

Laura Boudreau, Rachel Heath, and Tyler H. McCormick. Migrants, experience, and working conditions in Bangladeshi garment factories. *Journal of Economic Behavior & Organization*, 219:196–213, March 2024. doi: 10.1016/j.jebo.2023.02.013. URL https://doi.org/10.1016/j.jebo.2023.02.013.

Krista J Gile and Mark S Handcock. Respondent-driven sampling: An assessment of current methodology. *Sociological Methodology*, 40(1):285–327, 2010. doi: 10.1111/j.1467-9531.2010.01223.x.

Salganik M. J. Goel, S. Assessing respondent-driven sampling. *Proceedings of the National Academy of Sciences of the United States of America*, 2010. doi: https://doi.org/10.1073/pnas.1000261107.

Douglas D. Heckathorn. Respondent-Driven Sampling: A New Approach to the Study of Hidden Populations*. *Social Problems*, 44(2):174–199, 05 1997. ISSN 0037-7791. doi: 10.2307/3096941. URL https://doi.org/10.2307/3096941.

Douglas D. Heckathorn and Christopher J. Cameron. Network Sampling: From Snowball and Multiplicity to Respondent-Driven Sampling. *Annual Review of Sociology*, 2017. doi: https://doi.org/10.1146/annurev-soc-060116-053556.

Naila Kabeer, Lopita Haq, and Munshi Sulaiman. Multi-stakeholder initiatives in bangladesh after rana plaza: Global norms and workers' perspectives. Working Paper Series 19-193, London School of Economics and Political Science (LSE), Department of International Development, London, 2019.

Hongjie Liu, Jianhua Li, Toan Ha, and Jian Li. Assessment of random recruitment assumption in respondent-driven sampling in egocentric network data. *Social Networking*, 1(2): 13–21, 2012. doi: 10.4236/sn.2012.12002.

Bruce D Meyer, Wallace KC Mok, and James X Sullivan. Household surveys in crisis. *Journal of Economic Perspectives*, 29(4):199–226, 2015. doi: 10.1257/jep.29.4.199.

Matthew J. Salganik. Respondent-driven sampling in the real world. *Epidemiology*, 2012. URL `https://www.princeton.edu/~mjs3/salganik_commentary_2012.pdf`.

Matthew J Salganik and Douglas D Heckathorn. Sampling and estimation in hidden populations using respondent-driven sampling. *Sociological Methodology*, 34(1):193–240, 2004. doi: 10.1111/j.0081-1750.2004.00152.x.

Jean C. Digitale Yea-Hung Chen Meghan D. Morris Sarah Raifman, Michelle A. DeVost. Respondent-Driven Sampling: a Sampling Method for Hard-to-Reach Populations and Beyond. *Current Epidemiol Reports 9, 38–47 (2022)*, 2022. doi: https://doi.org/10.1007/s40471-022-00287-8.

Erik Volz and Douglas D Heckathorn. Probability based estimation theory for respondent driven sampling. *Journal of Official Statistics*, 24(1):79–97, 2008.

A Appendix

Here I share further detail about the several waves of data collection. The data is all freely available here, on the personal website of the original author, Professor Rachel Heath.

A.1 Wave 1 - Seed Sample

Seed for referral chains were collected from two prior surveys:

1) Administered in 2017, 1500 workers were surveyed by the Bangladesh Institute of Governance and Development (BIGD) and London School of Economics [Kabeer et al., 2019].

2) Administered in 2020, 60 workers were surveyed in a pilot survey by Boudreau and Heath.

These two samples include both the names and phone numbers of respondents and, when originally recruited, were geographically representative of garment-producing areas in the Dhaka Division of Bangladesh, home to 80% of the country's garment factories.

The data from the 2017 BIGD survey were stratified by cells defined by location (evenly split between Ashulia, Gazipur, Narayanganj, Mirpur, and Hemayetpur), gender, and above or below median experience (for one's given gender-location group).

A.2 Intermediate Links

All wave 1 seed participants were asked to recruit their contacts using the following script:

> *We would like to survey more garment workers to learn about their experiences. We are interested in calling workers who you have ever spoken to on the phone, that is, whose numbers are in your recent call log. We will provide you BDT 10 for each worker (i.e., a name and phone number) that you give us, for up to 10 workers total. We will then use a lottery to choose one or more of them to call. They may be selected to answer a much shorter survey than the one you just*

received, which will take about 10 minutes, and receive BDT 50 as a thank - you for their time. We would let them know that you provided a referral. Would you be willing to provide the names and phone numbers of family members and/or friends who work in the garments sector who may be willing to participate?

Additionally, all wave 1 respondents - and subsequent referrals given by these respondents - were asked the following questions:

- If the respondent listed 2 or more names:

  Out of the workers you just listed, whom would you feel most comfortable referring?

- If the respondent listed 3 or more names:

  Whom would you feel next most comfortable referring?

- If the respondent listed 4 or more names:

  Whom would you feel next most comfortable referring?

Referrals from these wave 1 seed respondents were randomly assigned to either "randomized" (RRDS) or "traditional" (RDS) recruitment groups according to the following protocols.

1) The first round of the intermediate links survey began on November 26, 2020. In the RRDS group, 3 contacts (or however many contacts were given, if less than 3) were picked at random to be followed and asked for referrals of their own. In the RDS group, three preferred contacts (or however many contacts were given, if less than 3) were followed.

2) On December 12, 2020, we began following all referrals of women, regardless of treatment group.

3) From December 21, 2020 until we finished the intermediate links survey on January

16, 2021, we began following up to four referrals per respondent, using the following algorithm:

a. From respondents in the RRDS group, we selected four referrals randomly.

b. For respondents in the RDS group, we first selected one referral randomly from all referrals that they gave (both the preferred and others). Then, if the randomly selected referral was one of the three preferred referrals, we just follow the 3 preferred referrals, as we have been doing already. If the randomly selected referral was not one of the three preferred referrals, we follow that random referral and the 3 preferred referrals, for up to four referrals total.

These updates to the sampling procedure were made to adjust for the behavior of respondents. Early on we observed that women were exceedingly unlikely to provide referrals despite, resulting in the expiration of those recruitment chains. As we will discuss in section 5, this was an unforeseen challenge that we did our best to overcome.

In total, there were 30 waves of the intermediate links survey, where the selected referrals from wave n constitute the sample for referrals in wave $n + 1$.

A.3 Wave 2a

Wave 2a consisted of a random selection of 25% of the 2017 sample, as well as selections from intermediate links, selected in the following way:

1) 25% of factories were selected, and 3 workers per factory were selected.

2) Up to 10 workers from wave 10 on-wards from each factory that agreed to be surveyed in the Bangladesh Garment Manufacturers and Exporters Association (BGMEA) matched survey conducted by Laura Boudreau beginning in January 2021. If a BGMEA factory appears after wave 10 but has fewer than 3 post wave 10, allow up to 3 workers using respondents from waves before 10.

3) All respondents not currently employed in garment factories were selected.

4) All women from wave 10 on-wards were selected , up to a max of 15 per factory.

5) 50% of workers with 3 or fewer years of experience from random seeds in wave 10 on-wards were selected.

6) For those who have left the garment sector, select up to 3 workers from that same factory if they don't already have at least 3 workers sampled.

7) There was a cap of 5 men per factory (that wasn't matched to BGMEA). So randomly throw out men if more than this number was selected.

A.4 Wave 2b

Among referrals given by the respondents to wave 2a:

1) All referrals not currently employed in garment factories were selected.

2) All referrals of women were selected, unless they lead to greater than 15 workers per factory in expectation (assuming a 50% probability that a selected respondent leads to a completed survey, based roughly on the intermediate links surveys).

3) All referrals of newly agreed upon factories in the BGMEA survey.

4) Add men until 3 per factory in expectation are selected. Choose from among men with 3 or fewer years experience if available.

5) 20 men with more than 3 years experience were randomly selected.

A.5 Wave 2c

Among referrals given by the respondents to wave 2b:

1) Select referrals in factories in the BGMEA survey who currently have less than 10 successful surveys. Select women to yield 10 successful surveys in expectation, again assuming at 50% success rate.

2) Select all women unless their factory has more than 20 current surveys.

3) Select 100 men randomly.

Additional selections were made from the intermediate links:

1) In factories in the BIGD survey that did not yield 10 successful surveys so far, select workers to yield 10 successful surveys in expectation.

2) Select women whose factories have 5 or fewer surveys currently.

A.6 Wave 2d

Among referrals given by the respondents to wave 2c:

1) Select all referrals from those in BGMEA matched factories.

2) Choose 60 new factories at random from factories not matched to BGMEA survey. Choose all women from these factories and 25% of men in these factories.

3) Select all women who have left the sector.

4) Select 25% of men who have left the sector.