Measuring the measurement error:

A method to qualitatively validate survey data*

Christopher Blattman Julian Jamison Tricia Koroknay-Palicz
Katherine Rodrigues Margaret Sheridan[†]

January 11, 2016

Abstract

Empirical social science relies heavily on self-reported data, but subjects may misreport behaviors, especially sensitive ones such as crime or drug abuse. If a treatment influences survey responses, it biases causal estimates. We develop a validation technique that uses intensive qualitative work to assess survey misreporting and pilot it in a field experiment where subjects were assigned to receive cash, therapy, both, or neither. According to survey responses, both treatments reduced crime and other sensitive behaviors. Local researchers spent several days with a random subsample of subjects after surveys, building trust and obtaining verbal confirmation of four sensitive behaviors and two expenditures. In this instance, validation showed survey underreporting of most sensitive behaviors was low and uncorrelated with treatment, while expenditures were under-reported in the survey across all arms, but especially in the control group. We use these data to develop measurement error bounds on treatment effects estimated from surveys. [*JEL codes: O1, C93, I32, C81, K4*]

**Acknowledgements:* For comments we thank Neal Beck, Alex Coppock, Dan Corstange, Macartan Humphreys, Don Green, Cyrus Samii, Chris Udry, several anonymous referees, and participants at the NYU 2014 CESS conference. This study was funded by the National Science Foundation (SES-1317506), the World Bank's Learning on Gender and Conflict in Africa (LOGiCA) trust fund, the World Bank's Italian Children and Youth (CHYAO) trust fund, the UK Department for International Development (DFID) via the Institute for the Study of Labor (IZA), a Vanguard Charitable Trust, the American People through the United States Agency for International Development's (USAID) DCHA/CMM office, and the Robert Wood Johnson Health and Society Scholars Program at Harvard University (Cohort 5). The contents of this study are the sole responsibility of authors and do not necessarily reflect the views of their employers or any of these funding agencies or governments. Finally, for research assistance we thank Foday Bayoh Jr., Natalie Carlson, Camelia Dureng, Mathilde Emeriau, Yuequan Guo, Rufus Kapwolo, James Kollie, Rebecca Littman, Richard Peck, Patryk Perkowski, Colombine Peze-Heidsieck, Joe St. Clair, Joseph Sango Jr., Helen Smith, Abel Welwean, Prince Williams, and John Zayzay through Innovations for Poverty Action (IPA).

[†]*Blattman* (corresponding author): Columbia University SIPA, 420 W 118th St., New York, NY (chris-blattman@columbia.edu); *Jamison*: Global Insights Initiative, The World Bank, 1818 H St NW, Washington, DC (julison@gmail.com); *Koroknay-Palicz*: The World Bank, 1818 H St NW, Washington, DC (t.gonwa@gmail.com); *Rodrigues*: International Rescue Committee, Research Department, 122 East 42nd St., New York, NY (katherine.rodrigues@rescue.org); *Sheridan*: University of North Carolina at Chapel Hill, Clinical Psychology, sheridan.margaret@unc.edu.

1 Introduction

The trouble with many survey topics, whether it’s abortion, drug use, crime, domestic violence, or support for terrorism, is that people may not tell the truth. This makes survey data on any sensitive topic suspect. Even without incentives to misreport, self-reported data are often inaccurate. Studies show people even misreport their gender and education.¹ When measuring subjects that can embarrass or endanger the respondent, we worry that people underreport their attitudes or actions.²

When we are interested in the impact of a program or event, measurement error will also affect our ability to estimate unbiased causal effects. In dependent variables, random measurement error reduces precision but won’t bias estimates.³ Systematic reporting errors, however, generally bias causal estimates, especially when the measurement error is correlated with the treatment or exogenous event of interest. For instance, people who receive an anti-crime message or an addiction treatment might be more likely to respond that they are non-violent or drug free, both because it’s socially desirable and because of perceived experimenter demand (where participants conform to the expectations of the people who ran the program).

Researchers have come up with a number of ways to limit bias in self-reported data. In developed countries, it is common to use administrative data. For example, studies of crime-reduction programs (such as the one we study in this paper) often prefer arrest and incarceration records to self-reported crime (e.g. Deming, 2011). Such data are seldom available in developing countries, however. Moreover, arrest data have serious systematic measurement error problems of their own.⁴

Others use survey experiments and indirect questioning. In list experiments, respondents report the number of items they agree with on a list, which randomly includes or excludes a sensitive item.⁵ In endorsement experiments, respondents rate their support for actors expressing sensitive ideas (Bullock et al., 2011). These are valuable tools, albeit with limitations. They can be imprecise and require large samples, and they can be cumbersome when measuring an array of items. Survey experiments also rely on two key assumptions: that

¹See Asher (1974); Bound et al. (2001).

²For instance, Karlan and Zinman (2008) find that large numbers of borrowers do not report high-interest consumer loans, potentially because they feel embarrassed.

³See Asher (1974); Hausman (2001). This statement applies primarily to linear models.

⁴Arrests underreport true criminal behavior, and they require strong assumptions: that arrests are responses to crimes rather than statistical or other discrimination; and that the treatment doesn’t affect the likelihood of being arrested for a crime, by changing the location and observability of the crime for example.

⁵e.g. Raghavarao and Federer (1979). For recent applications see Blair and Imai (2012); Karlan and Zinman (2012); Jamison et al. (2013).

people do not lie when counting on a list or endorsing a person, and that the presence of sensitive items doesn't affect reporting of non-sensitive ones (Blair and Imai, 2012).

Finally, in some cases data are physically verifiable and researchers can use a little of what Freedman (1991) called “shoe leather” and simply verify behavior. For instance, in Mexico, the government sent administrators to audit self-reported asset data used to decide who was in or out of a cash transfer program and found underreporting of assets to increase eligibility (Martinelli and Parker, 2009).

This paper develops and field tests an alternative approach for testing the direction and degree of survey misreporting. It is intended to be useful when objective administrative data are not available, survey experiments are impractical, and direct physical verification is impossible. We pilot the approach on self-reported measures of crime, drug use, homelessness, gambling, and discretionary spending. In principle the method could be applied to other sensitive topics where objective assessments are difficult—intimate partner violence, prostitution, risky sex behaviors, participation in communal violence, voting behavior, sexual identity, stigmatized diseases, and so forth.

The approach is relatively simple. We use intense qualitative work—including in-depth participant observation, open-ended questioning, and efforts to build relationships and trust—to try to elicit more truthful answers from a random subsample of experimental subjects. We focus on a very small number of key behaviors, and over several days of trust-building and conversation, we try to elicit a direct admission or discussion of the behavior.

We then compare these qualitative findings to survey responses, and use the difference to estimate the direction, magnitude, and patterns of measurement error. It is effectively a shoe leather approach for difficult-to-verify, often covert behaviors. Like survey experiments, the method relies on the assumption that people are more truthful in this context than in a survey. The techniques we use—spending time with respondents, interacting in their natural environment, developing a rapport, and trying to attain “insider” status—are central techniques in qualitative and ethnographic research to obtain honest and valid responses (e.g. Wilson, 1977; Bryman, 2003).

This paper illustrates the approach, including when, where, and how it could be applied to other field experiments or other causal analysis using survey data. It also describes the patterns of reporting bias we observe in this particular crime-reduction study, upending the priors we held about the nature and direction of measurement error in these circumstances.

The study recruited a thousand destitute young men in the slums of Liberia's capital, Monrovia, with an emphasis on men involved in petty crime or drugs. The formal evaluation by Blattman et al. (2015) randomized two interventions designed to reduce crime

and violence: an 8-week program of group cognitive behavior therapy (CBT) to discourage impulsive, angry, and criminal behaviors; and an unconditional cash transfer of \$200.

Obviously, we should be wary of self-reported survey measures of illegal or immoral behavior, especially from a population suspicious of authority, some of whom make their living illicitly. We should be doubly concerned when one of the treatments (therapy) tried to persuade people away from “bad” behaviors, potentially triggering additional social desirability bias or the perception of experimenter demand among the treated. We can imagine any informational or behavioral intervention would raise similar concerns. List experiments were one option, but we found them difficult to implement with a largely uneducated, illiterate population that were selected in part for impulsive behavior.⁶ Thus we developed this alternative.

Of more than 4,000 endline surveys conducted over the study, we randomly selected roughly 7.3% and attempted to validate survey responses on just six behaviors. Within days of the survey, one of a small team of Liberian qualitative research staff (“validators”) would visit the respondent four times over ten days, each day spending several hours as a participant observer or in active conversation with the man, his peers, and community members. Validators sought a direct admission of the behavior after building trust and familiarity. In effect the method is a very intensive, relationship-based form of survey auditing, which cost roughly as much as a regular survey to implement.

Validators and the authors then coded an indicator for whether or not the respondent had engaged in each behavior in the two weeks prior to the survey (i.e. during the timeframe about which survey questions on recent behavior were asked). Beforehand, we deemed four behaviors “potentially sensitive”: marijuana use, thievery, gambling, and homelessness. Two others were common, non-sensitive behaviors that could be subject to recall bias or other forms of error: paying to watch movies in a video club, and paying to charge their mobile phone at a kiosk. We call these the “expenditure” measures.

This qualitative approach is not free from error: validators could still miss behaviors, make faulty inferences, or let suspicions of treatment status influence their judgment (among other things).⁷ These limits of participant observation are well-known (Power, 1989). But

⁶For instance, a list experiment read aloud would require many ideas to be held in mind, and we were concerned that answers would be correlated with cognitive abilities.

⁷For instance, as with the survey, conversations between validators and participants may have been influenced by social desirability bias or experimenter demand. Additionally, had the validation exercise relied on observation as the primary source of evidence and the presence of an observer prompted good behavior, we would have underestimated sensitive behaviors in the validation. People have been shown to increase hand-washing behavior, for example, when directly observed, suggesting a Hawthorne effect of observation (Ram et al., 2010). This kind of desirability bias could be greater in a treatment arm, and validators might not eliminate it. Even validators could be biased if they can glean a subject’s treatment status. Thus we cannot eliminate all measurement error correlated with treatment status through our approach.

these errors, we argue, are less likely to bias treatment effect estimates than the experimenter demand and social desirability bias we worried would cause underreporting in the survey. It comes down to the following proposition: that we can reduce the appearance of experimenter demand (plus other biases correlated with treatment) through four days building rapport and trust, and a focus on only six facts, in the context of what feels to the study participant like everyday conversation rather than a formal survey in which a stranger asks about the same six behaviors in a 300-question, 90-minute questionnaire.

This is the key assumption underlying the technique. It parallels the “no liars” and “no design effects” assumptions in list experiments. Like list experiments, the assumptions cannot be tested directly. But if we accept them, then by comparing survey data to the data collected by validators, we can assess the presence and degree of measurement error in the survey data, and its correlation with treatment assignment.

In this specific crime study, one of our main concerns was that men would under-report potentially sensitive behaviors due to social desirability bias, and that the therapy treatment might further increase social desirability related under-reporting, leading us to observe a treatment effect that in reality is merely a treatment-correlated increase in under-reporting. The validation, however, found no evidence this was the case. Survey-based reporting of potentially sensitive behaviors was quite high: at endline, 22% of men reported stealing in the past two weeks, and 48% admitted to marijuana use. For the four sensitive behaviors, survey responses and validated measures were identical about 80% of the time. Men reported slightly fewer sensitive behaviors in the survey than the validation, but this underreporting was driven mainly by the underreporting of gambling. To the extent there is survey underreporting of the sensitive behaviors, it is the opposite of what we anticipated: the group that received both cash and therapy were the least likely to under-report sensitive behaviors.

Another prior was that expenditure data would be less prone to social desirability bias or measurement error correlated with treatment. In contrast, we found that across all treatment arms expenditures seem to be underreported in the survey relative to validation. This underreporting of expenditures was largest in the control group.

One benefit of the validation is that it upended our priors about the nature and direction of measurement error. Another is that it affects our conclusions in the larger study. Using outcome data from the survey, Blattman et al. (2015) found that cash led to short run income and expenditure gains, which dissipated after a year. They also found that therapy reduced anti-social behaviors, such as crime, immediately and dramatically, but that this change persisted only if the men received therapy and cash. The validation exercise largely bolsters this core result. Indeed, it implies that the treatment effects based on survey data could underestimate the true effects by up to 20%. The validation calls into question, however, the

finding that cash led to short run income gains, given that underreporting of expenditures is greater in the control group.

There are several possible explanations for these patterns. Regarding the low level of underreporting of sensitive behaviors, the majority of the men in the study sample are part of a counterculture where drugs and crime are commonplace. Thus, they may be less likely to feel stigma around these behaviors than “normal” society members. Additionally, these men (and their entire counterculture) are already seen by “normal” society as pariahs, and thus, there may be little advantage to hiding such behaviors.

Regarding the higher level of sensitive behavior reporting in the treatment groups, especially therapy plus cash, it’s conceivable that therapy makes the men more accustomed and willing to discuss these sensitive issues openly with a member of the project, or that the control group wants to appear better behaved or more deserving of a future program.

Turning to expenditures, underreporting across all arms is consistent with simple recall bias in consumption surveys. Underreporting of gambling and expenditures, especially in the control group, is also consistent with control group members hoping to become eligible for a future program by appearing poorer or more deserving.

Altogether, these findings are crucial to the credibility of the study’s experimental estimates, in this case bolstering the claim that the therapy reduced crime and other anti-social behaviors, and moderating the claim that the cash transfer increased incomes. Perhaps more broadly, the findings also challenge conventional notions of the direction of measurement error.

It would be a mistake, however, to cite this paper as evidence that systematic measurement error of sensitive behaviors in high-risk populations is low; that behavioral treatments foster trust and reduce measurement error; or that low-salience expenditures are especially vulnerable to experimenter or recall bias. These are all plausibly true, but before we can generalize more validation needs to be done in more places. An important takeaway message is that, despite several years working with this and similar populations, including extensive quantitative surveys and qualitative interviews, our priors about the most important sources of measurement error were wrong.

We include a detailed description of our procedures to make it easier for other researchers to adapt and use the method. In principle, we think it is applicable to a wide range of risky or stigmatized sexual, health, and economic activities. The cost, in our case, was roughly 3% of the total evaluation budget, a modest amount given that measurement error in self-reported data was the key causal identification concern in the evaluation.

One analog to our approach is in psychology, where virtually every self-reported survey measure of mental health has been validated using structured clinical interviews (e.g. Spitzer

et al., 1999). Another is a recent surge of behavioral and other measures to validate survey data on violence, prejudice, and other troublesome outcomes. In addition to the list and endorsement experiments mentioned before, Scacco (2010) interviewed a random subsample of potential religious rioters behind a screen that shielded their identity, and Paluck and Green (2009) measure cooperation by the patterns of distribution of a group survey gift. Finally, business profits and consumption have also proven troublesome to measure and have been the subject of experimental measurement studies. de Mel et al. (2009) experimentally test alternative approaches to measuring microenterprise profits, and find (counterintuitively) that the least intensive methods yield the least biased estimates. Beegle et al. (2012) have experimentally tested various consumption measures against one another. These studies have proven important to the studies where income is the crucial outcome. Ours could prove as useful to interventions targeted at violence, crime, and other risky or stigmatized behaviors. One thing is certain: systematic measurement error is a large and largely unaddressed problem, calling for more such new tools and their refinement and replication.

2 Context and experimental design

In poor countries like Liberia, governments are especially fearful of urban young men and the possibility they will commit crimes, rioting, or election violence. We designed a study to test the economic and behavioral roots of crime and violence among high-risk men. Blattman et al. (2015) describes the study in full detail.

2.1 Full experimental sample

The study recruited 999 young adult men in five neighborhoods Monrovia, a city of roughly 1.5 million. The study sought out “hard-core street” men—men in their 20s and 30s who live in extreme poverty and may be involved in violence, drugs and crime. We recruited and implemented the study in three phases over two years, typically in different, distant neighborhoods (see Appendix A). Table 1 describes the study sample at baseline.

On average the men were age 25, had nearly eight years of schooling, earned about \$68 in the past month working 46 hours per week (mainly in low-skill labor and illicit work), and had \$34 saved. 38% were members of an armed group during the two civil wars that ravaged the country between 1989 and 2003. At baseline, 20% reported selling drugs, 44% reported daily marijuana use, 15% reported daily use of hard drugs, 53% reported stealing something in the past two weeks, and 24% reported they were homeless in the last two weeks.

Table 1: Description of the study sample (n=999)

Baseline covariate	Mean	Std. Err.	Baseline covariate	Mean	Std. Err.
Age	25.4	(4.86)	Average weekly work hours in:		
Married/living with partner	16%	(0.37)	Potentially illicit activities	13.6	(27.26)
# of women supported	0.5	(0.64)	Agricultural Labor	0.4	(3.69)
# children under 15	2.2	(3.17)	Low-skill wage labor	19.4	(28.85)
Muslim	10%	(0.30)	Low-skill business	11.5	(23.98)
Years of schooling	7.72	(3.29)	High-skill work	1.5	(7.59)
Literacy score (0-2)	1.2	(0.90)	Ex-combatant	38%	(0.49)
Math score (0-5)	2.8	(1.57)	Currently sleeping on the street	24%	(0.43)
Health index (0-6)	4.9	(1.38)	Times went hungry last week	1.26	(1.36)
Disabled	8%	(0.26)	Sells drugs	20%	(0.40)
Monthly cash earnings (USD)	68.30	(84.49)	Drinks alcohol	75%	(0.43)
Durable assets index, z-score	0.00	(1.00)	Uses marijuana daily	44%	(0.50)
Savings stock (USD)	0.34	(0.67)	Uses hard drugs daily	15%	(0.35)
Able to get a loan of \$300	11%	(0.31)	Stole in past two weeks	53%	(0.50)

Notes: Surveys were completed with all men, but there are a small number of missing baseline values per respondent. For purposes of regression analysis, these are imputed with the sample median to avoid losing the observation.

2.2 Intervention and experimental design

We designed, implemented, and evaluated two interventions—group cognitive behavior therapy and cash—in a factorial experimental design. We first randomly assigned half the sample to an offer of therapy. Therapy was completed within eight weeks. Following this, we held a second lottery for grants of \$200 with the full sample.⁸

Treatment 1: Cognitive behavior therapy and counseling The therapy was designed and implemented by a local non-profit organization, Network for Empowerment and Progressive Initiatives (NEPI) Liberia. The 8-week program had two main goals. The first was “transformation,” or the shift from the position (and self-identity) as an outcast living on the fringe of society to an economically- and socially-integrated member of mainstream society. The second goal was to shift men from present-oriented decision-making to future-oriented goals and behavior.

⁸None knew of the cash grant until after therapy was completed. Randomization was done through public draw in blocks of roughly 50. There is balance across treatment and control groups. 90% of all men assigned to the therapy attended at least six days of the therapy. Those who did not attend had slightly less schooling, slightly higher earnings and assets, and are less likely to use drugs or alcohol or steal. Thus it appears the highest risk young men were the most likely to attend. See Appendix A for details.

The approach and curriculum grew out of NEPI’s experience, but were largely grounded in cognitive behavioral therapy (CBT) theory and practice. Group-based CBT approaches have been validated, typically in US populations, to reduce substance abuse, criminality, and aggression.

Participants met three times a week in groups of about 20, for four hours at a time, led by two facilitators. The only compensation provided for attendance was a bowl of rice and simple stew. On alternate days when the group did not meet, the facilitators visited the men at their homes or work areas to provide individual advising and encouragement. Many of the facilitators who ran the group intervention and individual counseling were themselves ex-combatants or reformed street youth.

The CBT element of the program manifested itself in the emphasis on small practical changes each session, which are reinforced through encouragement and praise. These included reducing substance use and abuse, improving body cleanliness, improving the cleanliness of the area in which they lived, and managing their anger without resorting to violence. Facilitators also formally encouraged participants to engage with society in planned and unaccustomed ways.

Facilitators also taught skills around planning and goal setting to help participants enhance their future-oriented attitudes, anticipate potential setbacks, and build skills for dealing with adversity. Finally, throughout the eight weeks, facilitators articulated a set of mainstream social norms and encouraged them to adopt these norms.⁹

Treatment 2: Unconditional cash grant All men were eligible for a cash grant of \$200. The cash was both a treatment and also a measurement tool (to see whether spending patterns were affected by the therapy).¹⁰ The framing of the grant was minimalist—people were told that it was random, one-time, and unconditional.¹¹

⁹These include discouragement of crime, substance abuse, and interpersonal violence (encouraging instead the use of peaceful solutions to conflict). The program also encouraged good financial management, especially saving money, as an important aspect of future- and goal-oriented behavior.

¹⁰An international non-profit, Global Communities, conducted the cash distribution. These partners conducted all recruitment and program implementation to minimize the perceived connection between the research team and programs.

¹¹Prior to the lottery, the group merely received a short lecture (15-30 minutes) on how to safeguard the funds once received. Of those assigned to the cash grant, 98% received it.

2.3 Survey data collection

The research team, a Liberia branch of the non-profit research organization Innovations for Poverty Action (IPA), presented themselves as independent evaluators.¹²

We attempted to collect survey data from each recruit five times: at baseline prior to the intervention; at “short-run” endline surveys roughly 2 and 5 weeks after the cash transfers; and at two “long-run” endline surveys 12 and 13 months after the cash grants.¹³

Because the sample was exceptionally mobile and difficult to track over time, we took special measures to minimize attrition. At baseline we were clear about our desire to stay in touch. We took photos and signature samples, and collected as many as ten different ways to contact each respondent. We documented contact information for each respondent, including all the places they said they sometimes stay, plus contact information for the network of people around them who have a more stable location. Respondents were often on the run from the police or other people, and so their contacts might be uncomfortable speaking to enumerators and revealing the respondent’s location. Thus, after the baseline survey, we asked respondents to use the enumerator’s phone to call their most stable contact and introduce the enumerator and study and give permission. At each endline, enumerators would typically start with the phone numbers of the various contacts or respondent and try to arrange an appointment. Contacts received no financial incentive. If this strategy failed, the enumerators would begin visiting the various locations listed. A slight majority of respondents were found within a few hours. In other cases, all leads were cold and more extensive sleuthing and asking around the neighborhood was required. If someone had traveled or moved far away, enumerators either waited until they returned or traveled across the country to find them in person. On the upper tail, it could take three to four days of physical searching to find the hardest-to-locate people. Enumerators only stopped searching when all possible leads had been exhausted.

¹²They visually distinguished themselves from other organizations by wearing uniquely colored emerald green t-shirts and identification badges over the years of the study. The exception to this is the validators, who wore street clothes that helped them blend in with the study participants.

¹³The exception is the 100 men in the pilot phase, who had a single “short run” survey 3 weeks after the grant, and a pair of “medium-run” surveys at 5 and 7 months in addition to the 12- and 13-month surveys. We ran pairs of short-run and long run surveys because it allowed us to take two measures of relatively noisy outcomes with potentially low autocorrelation such as earnings, expenditures, criminal activity, drug use, and so forth. Taking multiple measurements at short intervals allows one to average out noise, increasing power (McKenzie, 2012). Each survey was roughly 90 minutes long, followed by roughly 90 minutes of interactive behavioral games and psychological tests. Liberian enumerators conducted face-to-face interviews in Liberian English using handheld electronic devices.

By making at least four attempts to track each man, we were able to track and survey around 93% of the target sample across all endline survey rounds. Attrition is not strongly correlated with baseline covariates or treatment assignment.¹⁴

3 Empirical strategy

To motivate the empirical tests, we outline a simple model of the effect of different forms of measurement error in outcomes in the context of an experimental intervention. We adapt the simple linear model from the Bound et al. (2001) review of measurement error for these illustrative purposes. We use the language of experiments throughout, but the same principles could be applied to observational causal inference using survey data. In this simple example, we suppose the true treatment effect specification is:

$$y^* = \alpha + \theta T + \varepsilon \quad (1)$$

where y^* is the true outcome and T is an indicator for assignment to treatment.¹⁵ The observed survey outcome y^s , however, measures the true outcome with both systematic and random error:

$$y^s = \delta^s y^* + \gamma^s T + \mu \quad (2)$$

where we assume the random error μ is uncorrelated with y^* , T and ε . Throughout this illustration, δ (which we take to be positive throughout) denotes systematic measurement error of the true outcome (such as underreporting due to social desirability bias) and γ indicates error associated with treatment only (as in the case of experimenter demand, for example).

Attempting to calculate treatment effects on y^* using only y^s , the researcher estimates the following potentially erroneous equation:

$$y^s = \hat{\alpha} + \hat{\theta} T + \hat{\varepsilon} \quad (3)$$

¹⁴A majority changed locations between each round, many changing sleeping places every few weeks or nights. We generally made at least four attempts to locate each person, in all corners of the country, including prison (to be interviewed only when released). See Appendix A for formal analysis of attrition. The joint significance of all covariates and treatment assignment for survey attrition has a p-value of .53. Attrition is also roughly one percentage point lower in the treatment groups (not statistically significant).

¹⁵Bound et al. (2001) consider a continuous covariate X rather than indicator T . They also assume that other right-hand side variables are measured without error and have been partialled out. We ignore other covariates in this simple example, but the basic intuitions would hold with them present.

By substituting equation 1 into 2 and comparing to 3, we can see that the researcher estimates the treatment effect $\hat{\theta} = \delta^s \theta + \gamma^s$, and the bias from the true treatment effect θ is:

$$E(\hat{\theta} - \theta) = (\delta^s - 1)\theta + \gamma^s \quad (4)$$

There are three main cases to consider:

- $\delta^s = 1$ and $\gamma^s = 0$ is the special case of classical (random) measurement error involving μ only;
- $0 < \delta^s < 1$ and $\gamma^s = 0$ is the case where the survey measure systematically underreports the true outcome (but underreporting is uncorrelated with treatment status), in which case under-reporting would bias the estimated treatment effect towards the null, and over-reporting (if instead $\delta > 1$) away from the null; and
- $\gamma^s > 0$, which is the more worrisome case in which case we mistake measurement error (such as experimenter demand) for a treatment effect.

Now imagine we can collect validation data, y^ν , for a random subsample:

$$y^\nu = \delta^\nu y^* + \gamma^\nu T + \eta \quad (5)$$

where η is uncorrelated with T , y^* , ε , and μ . We define the difference in the survey and validation measures as:

$$y^\Delta \equiv y^s - y^\nu = (\delta^s - \delta^\nu)y^* + (\gamma^s - \gamma^\nu)T + \mu - \eta \quad (6)$$

The key assumption in this paper is that the measurement error in the survey and validation data are in the same direction and that validation data correspond more closely to y^* than survey data. That is:

$$0 \leq |\delta^\nu - 1| < |\delta^s - 1| \quad (7)$$

$$0 \leq |\gamma^\nu| < |\gamma^s| \quad (8)$$

$$(\delta^\nu - 1) * (\delta^s - 1) \geq 0 \quad (9)$$

$$\gamma^s * \gamma^\nu \geq 0 \quad (10)$$

If assumptions 7 through 10 hold, then y^Δ is a proxy for over-reporting (under-reporting if negative). The approach described in this paper is only suitable for validation techniques that

meet these assumptions. Note that in practice one cannot test them formally for clandestine or otherwise hidden behaviors, and the assumptions must be argued based on context and quality of the process. In much the same way, randomized response and list experiments rely on the assumption of less lying and no design effects, and instrumental variables estimates rely on the exclusion restriction.

If assumptions 7 and 8 hold, however, it means we can identify the direction and approximate magnitude of systematic survey error from the sample mean of y^Δ and assess whether the survey error is correlated with treatment by estimating the treatment regression:

$$y^\Delta = \alpha^\Delta + \theta^\Delta T + \zeta \quad (11)$$

where, since there is a treatment indicator in y^* , $\theta^\Delta = (\delta^s - \delta^\nu)\theta + \gamma^s - \gamma^\nu$.

As the validated measure approaches the true outcome measure, then θ^Δ approaches the value of the treatment effect bias described in equation 4. That is, as $\delta^\nu \rightarrow 1$ and $\gamma^\nu \rightarrow 0$ then $\theta^\Delta \rightarrow E(\hat{\theta} - \theta)$. The main focus of our analysis will be to calculate y^Δ in equation 6 and estimate θ^Δ from equation 11.

This formalization draws attention to several important caveats associated with any validation technique of this nature:

1. Identification of the bias $E(\hat{\theta} - \theta)$ hinges entirely on the credibility of the validation method and measure. The assumption of lower systematic measurement error is generally untestable and is a judgment call based on the nature and quality of the process.
2. Validation data cannot help us to separately identify the bias arising from general systematic error δ apart from treatment-specific error γ , except in the case where we are willing to make an a priori assumption about one of them, such as that $\gamma = 0$ (i.e. no “John Henry” effects or other forms of experimenter demand). In theory, the systematic and treatment-specific errors could run in opposite directions and cancel one another out. In that case, however, $y^\Delta \neq 0$.
3. So long as the validation measures are imperfect and $0 < |\gamma^\nu|$ or $0 < |\delta^\nu - 1|$, the estimates from equation 11 will tend to underestimate measurement error. The confidence interval on θ^Δ also increases with any noise in the validated measure, η .
4. Nonetheless, to the extent that the validation measures are credible, if we validate a random subset of the study sample we can adjust the distribution of y^* (conditional on T or other covariates) or estimate the “true” treatment effect θ using $\hat{\theta} - \theta^\Delta$.

The special case of binary outcomes

We can further refine this empirical strategy based on the fact that our analysis in this paper will be confined to binary outcomes. Taking into account the binary nature of our dependent variable allows us to derive simple characterizations of under- and over-reporting rates ($P(y^s = 0|y^v = 1)$ and $P(y^s = 1|y^v = 0)$, respectively), which we can estimate. For example, suppose we specify a model in which under- and over-reporting rates differ by whether the validated measure is a 0 or 1:

$$y_i^s = \tilde{\beta}_0 + \tilde{\beta}_1 T_i + \tilde{\beta}_2 y_i^v + \tilde{\beta}_3 (y_i^v \times T_i) + \tilde{\mu}_i. \quad (12)$$

Then we can interpret:

- $\tilde{\beta}_0$ as the share of untreated subjects who do not do outcome y according to the validation measure, but report doing it in the survey (over-report).
- $1 - \tilde{\beta}_0 - \tilde{\beta}_2$ as the share of untreated subjects who engage in outcome y according to the validation measure, but do not report it in the survey (under-report).
- $\tilde{\beta}_1$ as the increase in the over-reporting rate due to treatment.
- $\tilde{\beta}_1 + \tilde{\beta}_3$ as the decrease in the under-reporting rate due to treatment.

See section C.2 for the full derivation of 12. While this more complex specification yields several new estimands of interest, these additional model parameters come at a high cost of statistical power. In our case, with 240 observations in total, each parameter is estimated off of roughly 30 observations, putting us on a steep part of the power curve. While we report results of this approach, we choose to focus on the initial, more simplified, approach because of both these power concerns, and because we are most interested in correcting for the average bias in treatment effects using survey data, which we get from Equation 11. Nonetheless, the more flexible approach is an option for instances where the returns to validation are especially high, or where the cost is low, such that the analysis is statistically powered. It is also possible to formulate a version of this approach for non-binary outcome variables, e.g. by specifying one or more threshold values above or below which the accuracy of reporting is of special interest.

4 Validation methodology

We selected six variables for validation, all with recall periods of two weeks. We chose outcomes with varying degrees of salience (or memorability) and potential social stigma and experimenter bias. The variables were:

1. *Stealing*. The survey asked how many times in the last two weeks the respondent stole someone’s belongings or deceived or conned someone of money or goods.¹⁶ Based on our fieldwork, we hypothesized that stealing would be the most salient and least socially desirable of all six measures.
2. *Gambling*. The survey asked how many times in the last two weeks the respondent gambled or bet on sports. Beforehand, we hypothesized gambling had a lower level of salience and sensitivity than stealing, but was still somewhat stigmatized.
3. *Marijuana use*. The survey asked how many times in the last two weeks the respondent smoked marijuana. Marijuana use is not socially acceptable across Liberian society overall, but is fairly prevalent in our target demographic. We initially hypothesized underreporting could arise not so much from social stigma, but from the discouragement of drug use in the therapy treatment.
4. *Homelessness*. The survey asked how many times in the last two weeks the respondent had to sleep outside, on the street, or in a market stall because he had no other place to sleep or stay. This is a salient variable where we hypothesized respondents might have under-reported from embarrassment or over-reported in order to appear more needy (and eligible for more programs).
5. *Phone charging*. In the expenditure section of the survey, the survey asked how many times in the last two weeks the respondent charged his phone for money. This corresponds to taking one’s phone to a kiosk with electricity where one pays a small fee to recharge the battery, a common and routine expense for many Liberians, without stigma and possibly not very memorable. 38% of our sample had a mobile phone at the endline, and 38% reported charging a phone in the last two weeks.
6. *Video Club Attendance*. In the expenditure section of the survey, the survey asked how many times in the last two weeks the respondent went to a video club. These clubs are private businesses where one can go to watch a movie, television show, or

¹⁶The survey also measured more serious forms of theft, such as armed robbery, but our qualitative validation focussed on non-violent theft.

football match for a small fee. This is a popular and socially acceptable pastime, as most Liberians do not have electricity or home entertainment. Salience was unclear, but likely greater than phone charging.

These behaviors also exhibited diversity in program emphasis. Some, like stealing and marijuana use, were highly emphasized in the cognitive behavioral therapy, while others like video club and phone charging were not.

4.1 Validator staff

Eight local staff performed validations over the two years of data collection. We selected validators from the study’s qualitative research staff. These people typically began as survey enumerators, but displayed such skill and rapport with the subjects that we hired and trained them to conduct a separate qualitative research component: longitudinal, formal, open-ended interviews with a different subsample of subjects. All conducted the qualitative validation when they were not working on the formal open-ended interviews.¹⁷ Each validator received at least 10 days of training on the methods, including both classroom learning and extensive field training.¹⁸ Like any qualitative study, we believe staff recruitment and training to have been among the most important tasks and also the largest start-up cost of this method.

4.2 Approach

For each respondent, validators tried to determine whether the respondent had engaged in any of the measured behaviors, even once, in the two weeks preceding the respondent’s survey date, as the survey asked about behaviors occurring during the two weeks prior to the survey. We found it optimal for validators to visit each respondent four times, on four separate days, with each visit or “hangout session” lasting approximately three hours. The validator aimed to begin hanging out the day after subjects completed their quantitative surveys and to conduct all four visits in the 10 days following the respondent’s endline survey date.

Validators deliberately avoided the feeling of a formal interview and would typically accompany respondents as they went about their business.¹⁹ Validators sometimes took

¹⁷All but one were men, and all had a high school education. Two of the men completed roughly half the validations with the remainder doing roughly 10 to 20% each. To find these validators, we trained roughly two to three times the number of people needed from the pool of research staff, selecting only those with the most natural questioning and rapport-building skills for the validation exercise.

¹⁸Details of validator selection and training, team structure, tools and forms are in Appendix B.

¹⁹On the first visit validators would obtain verbal consent. We designed the consent script to be informal, and explained that the goal of hanging out with the respondent was to talk about some of the same things they discussed in the survey. In addition to this verbal consent, the formal consent form that preceded the recent survey said that qualitative staff may come and visit them again to gather more information.

notes during visits, but only in isolated areas out of sight from the respondent.²⁰ The idea follows from basic principles of ethnography, which seeks to study subjects in their natural settings, similar to those the researcher hopes to generalize about (Wilson, 1977). The intent is to reduce the sense of being in an experimental situation, which ethnographers perceive as creating bias.

The main approach was to engage in casual conversation on a wide range of topics, including the six target topics/measures. The target topics were raised mainly through indirect questions while informally chatting. For example, validators typically started conversations with discussions of family. This was both customary among peers in Liberia and a sign of respect and interest in respondents' lives. It was also a stepping stone for discussing the target behaviors—either because the validator can discuss an issue in their family (someone engaging in one of the activities) or how the respondent's family feels about their current lifestyle and circumstances.

In general, validators found it helpful to tell respondents stories or scenarios about another person or themselves, related to the target measures, then steer the conversation to get information about how respondents had behaved in similar situations, eventually discussing the past two weeks. Validators were careful to present these behaviors and incidents in a non-stigmatized light, for instance by discussing a friend who stole in order to get enough to eat, or how they themselves had periods of homelessness or used drugs and alcohol. Validators found that these personal stories (all of which were truthful) and genuineness were essential to building rapport and trust.

Validators might hold these conversations once or twice over the three hours, spending perhaps twenty or thirty minutes in conversation each time, to avoid unnaturally long or awkward conversations. The validator spent the remainder of the three hours in the general vicinity, observing respondents engaging in their daily activities. This could involve taking a rest in the shade or in a tea shop (as is common) or engaging others in conversation. Validators would also try to talk casually with the respondent's friends, relatives, or neighbors to learn about him (although we considered information from these second-hand sources as insufficient to support a conclusion about the respondents' behaviors, but merely as supporting information).

We found that building a rapport with participants in a short space of time was crucial. To develop trusting and open relationships, validators used techniques, including becoming close to respected local community and street leaders, eating meals together, sharing personal information about themselves, assisting subjects with daily activities, and mirroring

²⁰e.g. in a toilet stall or teashop. If validators were unable to find a secluded area in which to take notes, they sometimes recorded information in their cell phones, pretending to send a text message.

participants’ appearances and vernacular, as appropriate. In addition, validators tried to maintain neutrality and openness while discussing potentially sensitive topics. For instance, conveying—through stories or otherwise—that illicit behaviors were not perceived negatively, allowed respondents to feel comfortable sharing their involvement in such activities. Validators did not lie to or deceive respondents, however.

Overall, this approach—trust-building, spending time together over the course of several days, assuming the role of an “insider,” attempting to obtain admission or discussion of the behavior, clandestine but fairly immediate note-taking, and (as discussed below) close examination of the evidence for each respondent with the investigators—was designed to counter the observer bias and selective recall that concern participant observation.²¹ Developing a rapport with respondents, spending time to develop a relationship, and obtaining insider status are considered central to obtaining more honest and valid responses (Baruch, 1981; Bryman, 2003; Fox, 2004). We are not aware of any study, however, that has quantitatively tested this proposition.

4.3 Validation sampling and non-response

In each endline survey round we randomly selected study respondents to be validated, stratified by treatment group.²² Table 2 describes the samples selected for validation in each survey round over the course of the study. In total, we randomly selected 7.3% of all surveys, 297 in total, for validation.

We found 240 (81%) of the 297.²³ This attrition is an identification concern, but there is little evidence of biased attrition. Excess validation attrition (those who were surveyed but not validated) was not robustly associated with baseline characteristics (see Appendix A).

Statistical power. In order to minimize the confidence intervals surrounding any treatment-measurement error correlation, we chose the sample size that maximized the number of in-

²¹For general discussions of validity in qualitative methods, see Wilson (1977); LeCompte and Goetz (1982); Power (1989).

²²For each pair of survey rounds, study participants were randomly divided into blocks (e.g. 1, 2, 3, 4), and block 1 study participants were surveyed before block 2, and block 2 before block 3, etc. Within each block we randomly selected validation subjects using a computer-generated uniform random variable. The selection was performed without replacement in a given pair of survey rounds (e.g. the short-term endline surveys in a given phase), but sampling was performed with replacement across survey rounds. Twenty subjects were validated in more than one round.

²³We could not find 15 for even the endline survey. We could not validate a further 42 because they were difficult to find even immediately after the survey or (more commonly) because they lived a long distance away. In general, we surveyed respondents who had moved far out of Monrovia, but we were unlikely to validate them because of the time and expense and opportunity cost.

Table 2: Validation sample, totals and attrition

Surveys			Validation		Reason for no validation data				
Phase	Round	Target #	Selected	Validated	Unfound at endline	Unfound for validation	% validated (all)	% validated (treatment)	% validated (control)
1	3-week	100	0						
	5-month	100	24	18	2	4	75%	75%	75%
	7-month	100	24	12	1	11	50%	50%	50%
	12-month	100	10	6	3	1	60%	63%	50%
	13-month	100	10	8	2	0	80%	86%	67%
2	3-week	398	26	24	0	2	92%	94%	89%
	5-week	398	27	17	0	10	63%	68%	40%
	12-month	398	28	25	2	1	89%	86%	100%
	13-month	398	44	38	1	5	86%	85%	91%
3	3-week	501	0						
	5-week	501	0						
	12-month	501	35	31	2	2	89%	89%	88%
	13-month	501	69	61	5	3	88%	88%	88%
All		4096	297	240	18	39	81%	81%	80%

Notes: The proportion selected in each round was principally a function of logistical feasibility (e.g. number of available staff), and in some none were selected. As procedures became more familiar and staff more experienced, more could be done over time. The percentage validated in the treatment group includes any treatment (cash, CBT, or both).

interviews we felt qualified validators could manage logistically.²⁴ Post hoc calculations of statistical power confirm the estimates we made at the design stage. With a sample of 240, we can detect general over- or under-reporting greater than 17% of the survey mean (14% of the “true” validated mean).²⁵ Because each treatment arm is a subsample, however, we cannot precisely measure the effect of treatment on misreporting—it is difficult to detect effects greater than 33% of the survey mean (28% of the validated mean). Thus we are principally interested in the sign and magnitude of the treatment effect on misreporting by treatment group.

4.4 Coding validated data

Validators were unaware of the respondents’ survey responses or treatment status, and formed their own opinions (based on the evidence collected) about whether respondents engaged in the six activities during the time period captured by the quantitative survey. Every coding recommendation was then discussed with and vetted by one of the authors.

A core part of the validator training included logical reasoning, supporting reasoning with evidence, and writing this down in a clear and structured manner. After each visit, validators made written notes about the relevant data collected, including evidence to support their conclusions, on a standardized form. At the conclusion of the four visits, the validator coded six indicators, one for each behavior, where “1” meant he had relatively direct evidence the respondent engaged in the behavior during the recall period, and “0” otherwise.²⁶

²⁴In general, the validation sample was a balanced subsample of the full sample (see Appendix A for sampling and balance details). Power calculations, based on roughly the first 60 validator interviews, indicated that there was a modest degree of underreporting of all behaviors, sensitive and non-sensitive, but that the correlation between treatment status and measurement error was uncertain—across outcomes it varied in sign and magnitude, but was about zero on average. Thus the chief advantage of maximizing the sample conditional on time available was to shrink the confidence interval to build confidence in our method and the main outcomes of interest. Further validation was mainly limited by the number of validators we felt could be trained and supervised.

²⁵We calculated this minimum detectible effect (MDE) using a two-sided hypothesis test with 80% power at a 0.05 significance level, using baseline and block controls when calculating the R-squared statistic. We calculated an MDE for both the 0–2 expenditures index and the 0–4 sensitive behaviors index. The expenditures index had a mean of .82 in the survey and an MDE of .13 for general over- and under-reporting and .29 for a treatment effect on misreporting. The sensitive behaviors index had a mean of 1.12 in the survey and an MDE of .2 for general over- and under-reporting and .36 for any treatment effect on misreporting. We estimate that doubling the sample size would have increased power by about a third.

²⁶Over the course of the exercise, different measures offered different experiences and lessons. Because of its relative frequency and visibility, we suspect marijuana use was the easiest to directly observe. But validators found other behaviors straightforward to discuss in conversation. In the survey and (especially) the validation, phone battery charging led to the most confusion—in particular, did simply charging one’s phone count, or did only paying to charge one’s phone count? Paid charging was the focus of the survey question (it appeared in an expenditure survey module), but we were concerned that the validators would use a more expansive definition. We attempted to mitigate such differences through trainings and regular discussions on the coding.

Validators recorded an average of 1.35 “major” pieces of evidence per respondent per behavior to support their coding decision sheets. This was typically the most persuasive piece or pieces of evidence rather than all evidence collected.²⁷ Table 3 reports evidentiary methods by behavior. In general, the validators used some form of direct or indirect questioning to elicit a direct admission of the behavior or persuasive statements that respondents did not engage in the behavior. The validators only witnessed or found direct evidence of the behavior in a fifth of cases, or had third party verification in about 6% of cases. In any event, witnessing or third party verification were not sufficient evidence for a final coding. For instance, witnessing had to be followed by questions confirming that the respondent also engaged in the behavior in the two weeks prior to the survey. This accounts for most of the cases where there was more than one piece of evidence highlighted.

In general, the patterns of evidence are fairly commonsensical. Witnessing is limited to observable behaviors such as marijuana, gambling, homelessness, and phone charging. Stories and scenarios where the respondent is invited to comment or discuss are especially common for the most sensitive subject, stealing. Indirect questioning is most common for everyday topics such as homelessness (e.g. “Is this your house?”) and phone charging (e.g. “I need to charge my phone. Where do you usually charge yours?”).

4.5 Limitations of the approach

While we think, based on our experiences, that this validation exercise gave enough time to gather detailed, accurate information and fostered trust and frankness, there are nonetheless limitations to this approach.

1. POTENTIAL DISRUPTION. The presence, and interactions and conversations with the validators may be intrusive and might disrupt respondents’ daily activities, thereby altering the findings. To mitigate this risk, validators wore clothes that would blend in with their respondent’s environment, and typically accompanied and assisted respondents in their activities as appropriate (e.g. helping a scrap metal collector scavenge).

Homelessness also proved somewhat challenging to measure and validate, as we discovered its definition is subjective. Circumstances arose that were somewhat ambiguous, such as having no home of one’s own but regularly sleeping on a friend’s floor or in an acquaintance’s market stall. To account for the potential variability in perceptions of homelessness, validators were instructed to include as much information as possible about respondents’ living situations in their summary reports. The authors then worked with validators to code a somewhat broad definition of homelessness that included any ambiguous circumstances. Prior to analysis, it was not clear whether survey respondents applied the same definition, and hence we err on the side of finding underreporting in the survey.

²⁷We do not have complete paper records of all evidence collected, and so the 1.35 pieces of evidence is probably an understatement of the full amount of evidence.

Table 3: Evidentiary methods reported by validators, by behavior

Main evidence techniques	Potentially sensitive behaviors				Expenditures	
	Steal	Marijuana	Gamble	Homeless	Video	Phone
	(1)	(2)	(3)	(4)	(5)	(6)
Avg. pieces of evidence	1.1	1.3	1.1	1.7	1.0	1.2
Obs. (All)	240	240	239	240	239	240
Direct question	36%	35%	38%	5%	32%	1%
Indirect question	28%	46%	42%	62%	59%	92%
Story / scenario	36%	6%	13%	12%	2%	1%
Witnessed / found evidence	3%	31%	9%	62%	5%	18%
Third party account	3%	6%	4%	21%	0%	0%
Other / unclear	3%	9%	6%	13%	6%	5%
Obs. (Coded “did not engage” in behavior)	191	118	170	190	93	125
Direct question	38%	44%	39%	5%	34%	0%
Indirect question	26%	46%	44%	60%	58%	98%
Story / scenario	37%	7%	15%	12%	3%	2%
Witnessed / found evidence	2%	3%	1%	65%	2%	1%
Third party account	3%	10%	4%	24%	0%	1%
Other / unclear	2%	1%	1%	14%	4%	0%
Obs. (Coded “did engage” in behavior)	49	122	69	50	146	115
Direct question	29%	25%	36%	4%	30%	2%
Indirect question	33%	46%	38%	70%	60%	86%
Story / scenario	33%	5%	9%	10%	1%	0%
Witnessed / found evidence	10%	59%	28%	52%	7%	37%
Third party account	4%	2%	4%	8%	0%	0%
Other / unclear	8%	17%	17%	6%	8%	10%

Notes: Direct questions imply the validator asked the respondent directly about his engagement in the activity. Indirect questions imply the validator brought up the subject in general conversation (Where do you live?). Stories and scenarios are a form of indirect questioning where the respondent is invited to comment. Witnessing or found evidence implies the validator saw the respondent engaging in the activity in question or found physical evidence that the respondent recently engaged in the activity. Third party accounts imply the validator asked the family and friends of the respondent whether or not he engaged in the activity. Other or unclear methods include a handful of cases of unprompted information from the respondent, and also cases where the behavior could be inferred from other knowledge. Mainly it implies that coding was inconclusive or incomplete but is likely a form of questioning.

2. DIFFERENCES IN RECALL PERIODS. The validation occurred after the time period about which the survey questions had asked, and validators or respondents could have made errors about the relevant window of time (e.g. homelessness could have been observed the week after the survey, and inferred to the time of the survey incorrectly). This is most likely a source of random measurement error.
3. INCONSISTENT QUESTIONS. The survey and validation questions might have been interpreted differently, making it difficult to compare results. However, we used close consultations and reviews of the data, and focus groups with survey and validation staff, to maximize consistency.
4. REVERSE HAWTHORNE EFFECT. Training validators to look for certain behaviors could lead them to overreport those behaviors (akin to the problem of “when you have a hammer everything looks like a nail”). This reverse Hawthorne effect would probably be more of a risk if the validation method relied on passive observation. Rather, validation involved active discussion and (usually) a direct admission of the behavior. Also, one of the authors reviewed and discussed the evidence for every subject with the validator.
5. INCREASING SOCIAL DESIRABILITY BIAS. In principle the participant observation method, by building rapport, could lead to a different source of measurement error by (for example) increasing social desirability bias. Our strong sense is that the opposite is true, that trust and rapport reduced the bias, but this is a subjective interpretation and not independently verifiable.
6. CONSISTENCY BIAS. In principle, respondents could recall their survey response and try to remain consistent despite trust-building. This could motivate randomizing the order of validation and survey in the future.
7. NON-BLINDED VALIDATORS. The researcher is not immune from bias in qualitative research (LeCompte and Goetz, 1982; LeCompte, 1987). We are especially concerned with any bias correlated with treatment. While validators weren’t given the subject’s treatment status, it’s possible and even likely that this could come up during the extended conversations. Thus there is a danger that the validators’ biases will be correlated with treatment. The trust-building and preference for direct admission of the behavior was intended to mitigate this risk, but it still remains.

Most importantly, it seems unlikely that validators would commit most of these errors differentially across study arms. Misreporting correlated with treatment is still a risk under

the consistency bias and non-blinded limitations, but the in-depth focus on a handful of questions, time invested, and trust-building is designed to counteract these biases as much as possible.

Finally, like any qualitative work, this is not an off-the-shelf tool. To select and refine the variables, recruit and train validators, and monitor quality of the data requires that researchers have some familiarity with the context and population and at least basic experience in qualitative data collection.

4.6 Replicability of the approach

There are three reasons to think that this method could be replicated in other developing country field experiments and observational analysis using surveys. First, the expertise needed to implement the method effectively exists in most countries. Indeed, it should be considerably simpler to implement outside of Liberia. After fourteen years of civil war, and with one of the lowest human development indices in the world, Liberia has very low local research capacity, even compared to other poor and post-conflict states.

Second, most social scientists are nearly as well prepared to design and implement the approach as they are a new survey instrument or measure. Like any measure or method, it takes local knowledge, care, and extensive pretesting to develop a credible approach, and can benefit from someone with expertise in the subject area. In our case, one of the field research managers had some background in qualitative work and quality assurance, which we believe improved the quality of training and selection of the validator staff.

Third, the cost of the data collection is not necessarily large relative to many field experiments or large-scale panel surveys. In this instance, the fixed cost of startup was primarily in the recruitment and training of the small number of validators—approximately 2 to 3 weeks of work. We estimate the marginal cost of validation was roughly \$80 per respondent, mainly in wages and transport. By comparison, the marginal cost of surveying a respondent was roughly \$70.²⁸

While this method is considerably more expensive than survey experiments, it is more in line with the depth and cost of commonplace efforts to improve consumption measurement through the use of diaries or physical measurement.²⁹ For crucial measures in large program evaluations, or for statistics informing major policies, the cost is small relative to the inter-

²⁸Both figures were driven by the fact that it typically took one to two days of searching to find each respondent for surveying, plus the time to survey itself. Both surveying and validating in Liberia were expensive by the standards of household surveys, largely because of the cost of operating in a fragile, post-conflict state and the great difficulties in tracking such an unstable population.

²⁹In one extreme example, in the India NSS consumption survey, enumerators physically measure the volume of all food consumption (NSSO Expert Group, 2003).

vention, larger study, or larger purpose. For instance, as a proportion of total expenditures on the study, this validation exercise cost under 3% of all research-related costs, and less than 1–2% of program plus research costs.

5 Results

For each of the six behaviors, we construct indicators for that behavior using survey data and the qualitative validation technique, pooling responses from all endline surveys. We also construct additive indices of the four potentially sensitive behaviors and of the two expenditures. Table 4 reports means in the full sample and each treatment arm, as well as the percentage of times the two measures are in agreement. Table 5 reports estimates of y^Δ , the difference between the survey and validation measures. Table 6 displays estimates of the correlation between treatment and measurement error. Finally, Table 7 reports treatment effects using the original survey data ($\hat{\theta}$); estimates of bias from measurement error (θ^Δ , the correlation between treatment and the survey-validated difference); and adjusted treatment effects that correct for this observed measurement error ($\hat{\theta} - \theta^\Delta$).

5.1 Misreporting

Rates of behavior Overall these are relatively common behaviors within our study sample. According to the survey data reported in Table 4, in the two weeks prior to the survey, 22% stole, 48% used marijuana, 18% gambled, 23% were homeless for at least a night, 42% attended a video club, and 39% paid to charge a mobile phone.

Correspondence in the survey and validator data In general, the survey and validated data are identical about 80% of the time for sensitive measures and about 70% of the time for expenditures (Table 4). Correspondence is lowest for video club expenditures (62% overall), perhaps because attendance is intermittent and has low salience.

On average, the unadjusted validation means were higher than the survey means, suggesting slight underreporting on the survey. The average person reported 1.21 sensitive behaviors and 1.09 expenditures in validation, and 1.12 sensitive behaviors and 0.82 expenditures in the survey.

Note that this is an average, however, and the qualitative validation finds cases of over- and under-reporting in the survey relative to the validation. There are 328 instances where the measures are not equal: 208 cases of survey under-reporting (the survey indicator is a zero and validation indicator is a one), and 120 cases of over-reporting (see Appendix

Table 4: Comparison of survey and qualitative validation means at endline

	Potentially sensitive behaviors					Expenditures			All
	All (0-4)	Steal	Marijuana	Gamble	Homeless	All (0-2)	Video	Phone	(0-6)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
A. FULL SAMPLE									
Survey mean	1.12 (1.14)	0.22 (0.42)	0.48 (0.50)	0.18 (0.39)	0.23 (0.42)	0.82 (0.73)	0.42 (0.50)	0.39 (0.49)	1.93 (1.31)
Validation mean	1.21 (1.18)	0.20 (0.40)	0.51 (0.50)	0.29 (0.45)	0.21 (0.41)	1.09 (0.74)	0.61 (0.49)	0.48 (0.50)	2.30 (1.21)
% in agreement		79%	85%	72%	82%		62%	82%	
B. CONTROL GROUP									
Survey mean	1.25 (1.31)	0.27 (0.45)	0.48 (0.50)	0.23 (0.43)	0.27 (0.45)	0.68 (0.70)	0.37 (0.49)	0.32 (0.47)	1.93 (1.44)
Validation mean	1.30 (1.23)	0.23 (0.42)	0.49 (0.50)	0.34 (0.48)	0.23 (0.42)	1.18 (0.70)	0.65 (0.48)	0.54 (0.50)	2.48 (1.21)
% in agreement		80%	88%	72%	77%		47%	75%	
C. THERAPY ONLY									
Survey mean	1.06 (1.11)	0.19 (0.39)	0.48 (0.50)	0.17 (0.38)	0.22 (0.42)	0.81 (0.75)	0.41 (0.50)	0.41 (0.50)	1.87 (1.35)
Validation mean	1.09 (1.14)	0.17 (0.38)	0.48 (0.50)	0.24 (0.43)	0.20 (0.41)	0.98 (0.76)	0.54 (0.50)	0.44 (0.50)	2.07 (1.24)
% in agreement		80%	89%	74%	80%		72%	81%	
D. CASH ONLY									
Survey mean	1.03 (1.16)	0.21 (0.41)	0.49 (0.50)	0.13 (0.34)	0.21 (0.41)	0.77 (0.71)	0.37 (0.49)	0.40 (0.49)	1.81 (1.35)
Validation mean	1.32 (1.26)	0.23 (0.42)	0.53 (0.50)	0.33 (0.47)	0.24 (0.43)	1.00 (0.81)	0.55 (0.50)	0.45 (0.50)	2.32 (1.33)
% in agreement		76%	82%	74%	90%		56%	85%	
E. THERAPY + CASH									
Survey mean	1.13 (0.98)	0.22 (0.42)	0.48 (0.50)	0.21 (0.41)	0.22 (0.42)	0.98 (0.73)	0.54 (0.50)	0.44 (0.50)	2.11 (1.11)
Validation mean	1.11 (1.11)	0.19 (0.40)	0.52 (0.50)	0.24 (0.43)	0.16 (0.37)	1.17 (0.68)	0.70 (0.46)	0.48 (0.50)	2.29 (1.05)
% in agreement		81%	83%	68%	81%		71%	87%	
Observations	239	238	238	238	239	239	238	239	239

Notes: The table reports the means (standard deviations) of the survey and the qualitatively validated measures for the full sample and by treatment arm. “% in agreement” is the percentage of respondents for whom the survey indicator equals the qualitatively validated indicator.

Table 5: Survey over-reporting, estimated by the mean difference between survey and validation measures (y^Δ)

	Potentially sensitive behaviors					Expenditures		
	All (0-4) (1)	Steal (2)	Marijuana (3)	Gamble (4)	Homeless (5)	All (0-2) (6)	Video (7)	Phone (8)
Full sample	-0.10 <i>0.17</i>	0.02 <i>0.57</i>	-0.03 <i>0.24</i>	-0.11 <0.01	0.02 <i>0.45</i>	-0.27 <0.01	-0.19 <0.01	-0.08 <0.01
Control group	-0.07 <i>0.64</i>	0.03 <i>0.57</i>	-0.02 <i>0.71</i>	-0.12 <i>0.09</i>	0.03 <i>0.60</i>	-0.50 <0.01	-0.29 <0.01	-0.22 <0.01
Therapy only	-0.04 <i>0.80</i>	0.02 <i>0.77</i>	0.00 <i>1.00</i>	-0.07 <i>0.29</i>	0.02 <i>0.77</i>	-0.17 <i>0.08</i>	-0.13 <i>0.07</i>	-0.04 <i>0.53</i>
Cash only	-0.29 0.04	-0.02 <i>0.80</i>	-0.05 <i>0.37</i>	-0.20 <0.01	-0.03 <i>0.42</i>	-0.23 0.03	-0.18 0.03	-0.05 <i>0.32</i>
Therapy + cash	0.02 <i>0.91</i>	0.03 <i>0.57</i>	-0.05 <i>0.37</i>	-0.03 <i>0.66</i>	0.06 <i>0.25</i>	-0.19 0.01	-0.16 0.02	-0.03 <i>0.48</i>
Observations	239	238	238	238	239	239	238	239

Notes: Columns 1 to 8 report the simple mean differences in the survey and validation measures for the full sample and for each treatment arm, along with p values for a t test of whether the mean is different from zero. We bold p values ≤ 0.05 .

C.2). Most of these differences arose from direct or indirect questioning of the participant. Some form of external evidence, such as direct observation or third party confirmation, additionally supported 21% of the underreports found through validation (especially marijuana use, homelessness, gambling and phone charging) and 14% of over-reports (mainly homelessness).³⁰

Underreporting of sensitive behaviors, particularly gambling Table 5 reports our proxy of survey over-reporting: the simple survey-validation differences, with p values from a t test of the difference from zero. Negative values indicate survey under-reporting, assuming the validator measure is more accurate of course. As noted above, we have the statistical power to detect differences greater than about 17% of the survey mean.

Overall, gambling seems to be slightly underreported in every treatment arm, and highly underreported by men in the control and cash only groups. For instance, 33% of the cash only group admitted to gambling during validation, compared to 13% during the survey. Some of this underreporting could be due to ambiguous behaviors being coded as gambling

³⁰Direct observation was more likely to support findings of underreporting than of over-reporting, but this is mechanical since it's not possible to observe a non-behavior. This does not apply to homelessness, however, since the absence of homelessness is having a home, and this is observable.

in validation interviews but not in the survey. But the fact that underreporting is smaller in the therapy arms suggests that the underreporting is not simply an artifact of different definitions, but rather something else, such as a strategic response to treatment status.

If we look at stealing, marijuana use, and homelessness, however, none of the survey-validation differences are statistically significant. There is possibly some slight underreporting of drug use and slight over-reporting of stealing, but the magnitudes are generally small in the sense that they are less than 10% of the survey means reported in Table 4. The sample size is small, however, and hence many of these differences are not precisely estimated.

Underreporting of expenditures We see much stronger evidence of underreporting of expenditures in the survey. The difference for the combined expenditures is -0.27 in the full sample (Table 5, Column 6). This difference is large—about a third of the survey mean reported in Table 4. Expenditure underreporting is largest for the video club measure, but both expenditures appear to be underreported. Interestingly, the mean differences appear to be smaller and less statistically significant if the men received one of the treatments. We return to these differences across treatment arms below.

5.2 Correlation between treatment and measurement error

In order to identify if measurement error is correlated with treatment, we estimate equations 11 and 12 in Table 6. For sensitive behaviors, almost none of the coefficients on treatment indicators or interactions are statistically significant. For the index of four sensitive measures (Panel (a), Column 5), the coefficient on treatment, β_1 , is actually greater than zero for therapy plus cash, implying that the impacts of therapy plus cash are, if anything, larger than the survey data imply. That said, the confidence intervals are relatively large, so we cannot rule out overstatement of treatment effects entirely. Nonetheless, there is almost no evidence of the bias we feared.

The results for our two expenditure measures suggest that all treatment arms are associated with a roughly 0.3 increase in the total number of instances (out of 2) in our proxy for measurement error (Panel (a), Column 8). There is underreporting across all arms, but it is greatest in the control group. As we see below, this implies that any expenditure gains we observe from the interventions may be the result of misreporting.

Before looking at these adjusted treatment effects, we consider the results of the more flexible regression in Panel (b). It does not materially change our conclusions. The effect on sensitive behaviors, in particular, is fairly homogeneous. Treated men who we think did not engage in the sensitive behaviors tend to report them ($\tilde{\beta}_1^{Both} > 0$) more than untreated men. Treated men who did engage in the sensitive behaviors tend to under-report to a lesser extent

than in the control group.³¹ Put differently, treated men reported a slightly higher rate of sensitive behaviors regardless of their measurement in the validation exercise. Of course, this more flexible test has much lower statistical power given our number of observations, and so we interpret it with caution.

5.3 Adjusting treatment effects

Treatment effects using survey data. Blattman et al. (2015) report full treatment effect estimates, short-term and long-term, based on the survey data. These results indicate that cash (alone or in combination with therapy) led to an increase in consumption in the month after the grant, including a fall in homelessness, in part because the men spent the grant directly, but also because they invested in petty business and increased their earnings. After a year, however, these earnings and consumption gains had disappeared, likely because adverse economic shocks eliminated the men’s additional cash, savings and investments.

Therapy, meanwhile, led to self-reported falls in anti-social behaviors ranging from 30% to 50%, especially in interpersonal aggression, drug dealing, and theft. After a month, the falls were similar in both the therapy only and therapy plus cash groups. After a year, however, the fall was only sustained in the therapy plus cash group. The paper hypothesizes that therapy plus cash had a more sustained effect on anti-social behaviors because the cash grant positively reinforced the behavior change and enabled the men to practice their new skills and carry on with their new identity for longer. This large, sustained fall in self-reported anti-social behaviors in the therapy plus cash group is the central finding of the study.

We see the same patterns reflected in our pooled (over time) treatment effects on the sensitive and non-sensitive summary indexes, still focused on the survey data, which are displayed in Table 7. Cash weakly increased expenditures (our nonsensitive index) but had little effect on our sensitive behaviors index. The increased expenditures are driven mainly by short term impacts. Meanwhile, therapy decreased sensitive behaviors such as stealing and gambling. With therapy plus cash, the effects are largest and more statistically significant, in large part because they are sustained in the long run.

Adjusted treatment effects. Table 7 also reports the effect of each treatment on survey over-reporting, θ^Δ . These estimates effectively take the simple survey-validation differences in Panel (a) of Table 6 and estimate the difference across treatment arms. We present

³¹Also, note that, on average, $\tilde{\beta}_0 > 0$, $\tilde{\beta}_2 < 1$, and $\tilde{\beta}_0 + \tilde{\beta}_2 < 1$ for sensitive measures (Column 5). This is consistent with what we observe in Table 4: slight survey underreporting of sensitive behaviors, and 20–30% non-correspondence between survey and validated measures.

Table 6: Estimates of the correlation between treatment and measurement error

(a) Constrained, with round-block fixed effects (Equation 11)

Covariate	Dependent variable (N=239)							
	$y^s - y^v$, Sensitive behaviors					$y^s - y^v$, Expenditures		
	Stealing	Mari- juana	Gambling	Homeless	All (0-4)	Video Club	Phone Charging	All (0-2)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
β_o (Constant)	-0.029 [.087]	0.062 [.061]	-0.109 [.093]	0.093 [.076]	0.015 [.177]	-0.326 [.118]***	-0.194 [.095]**	-0.517 [.158]***
β_1								
Therapy	-0.019 [.084]	0.015 [.057]	0.025 [.097]	-0.025 [.091]	-0.004 [.199]	0.170 [.102]*	0.174 [.085]**	0.339 [.132]**
Cash	-0.038 [.088]	-0.042 [.067]	-0.085 [.090]	-0.077 [.079]	-0.237 [.195]	0.109 [.111]	0.165 [.078]**	0.269 [.134]**
Both	-0.006 [.080]	-0.024 [.062]	0.077 [.095]	0.031 [.089]	0.079 [.183]	0.127 [.103]	0.181 [.075]**	0.304 [.115]***

(b) Unconstrained, with round-block fixed effects (Equation 12)

Covariate	Dependent variable (N=239)							
	y^s , Sensitive behaviors					y^s , Expenditures		
	Stealing	Mari- juana	Gambling	Homeless	All (0-4)	Video Club	Phone Charging	All (0-2)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\tilde{\beta}_0$ (Constant)	0.301 [.140]**	0.098 [.092]	0.231 [.118]*	0.283 [.129]**	0.976 [.287]***	0.367 [.130]***	0.047 [.093]	0.048 [.208]
$\tilde{\beta}_1$								
Therapy	-0.022 [.070]	0.010 [.077]	-0.009 [.074]	-0.036 [.079]	0.154 [.228]	-0.190 [.124]	0.100 [.072]	-0.011 [.207]
Cash	0.003 [.068]	0.025 [.081]	-0.079 [.064]	-0.138 [.070]**	-0.069 [.220]	-0.072 [.139]	0.040 [.068]	0.089 [.219]
Both	-0.013 [.064]	0.025 [.081]	0.064 [.076]	-0.025 [.083]	0.271 [.241]	-0.113 [.138]	0.041 [.064]	-0.150 [.209]
$\tilde{\beta}_2(y^v)$	0.496 [.158]***	0.735 [.096]***	0.315 [.108]***	0.405 [.158]**	0.677 [.108]***	0.038 [.123]	0.504 [.096]***	0.328 [.129]**
$\tilde{\beta}_3$								
Therapy $\times y^v$	-0.166 [.234]	-0.014 [.125]	-0.131 [.176]	-0.020 [.220]	-0.210 [.147]	0.434 [.169]**	0.079 [.143]	0.222 [.169]
Cash $\times y^v$	-0.232 [.208]	-0.114 [.134]	-0.063 [.148]	0.286 [.202]	-0.134 [.144]	0.133 [.174]	0.196 [.137]	0.064 [.169]
Both $\times y^v$	-0.064 [.214]	-0.085 [.129]	-0.241 [.173]	0.066 [.232]	-0.230 [.137]*	0.386 [.171]**	0.234 [.130]*	0.379 [.168]**

Notes: The table reports the degree and direction of bias in our treatment effects. In panel (a), we assume that our measurement error does not vary by whether or not the individual engages in the behavior, which allows for a simple way to use β_2 to adjust our ITT estimates. In panel (b), we relax this assumption and let the measurement error vary by behavior and treatment arm at the cost of reduced statistical power.

Table 7: Estimates of treatment effects ($\hat{\theta}$) and treatment effect bias (θ^Δ) by outcome and treatment

(a) Round and block fixed effects (N=3765/239)							
Treatment arm	Outcome index	ATE using survey data ($\hat{\theta}$)		ATE on measurement error(θ^Δ)		Adjusted ATE ($\hat{\theta} - \theta^\Delta$)	
		β	S.E.	β	S.E.	β	S.E.
		(1)	(2)	(3)	(4)	(5)	(6)
Cash only	Sensitive	-0.057	[.095]	-0.178	[.190]	0.121	[.195]
	Nonsensitive	0.080	[.052]	0.285	[.130]**	-0.205	[.143]
Therapy only	Sensitive	-0.186	[.092]**	0.004	[.199]	-0.190	[.198]
	Nonsensitive	0.000	[.050]	0.335	[.134]**	-0.334	[.154]**
Therapy and cash	Sensitive	-0.398	[.090]***	0.118	[.182]	-0.516	[.196]***
	Nonsensitive	0.076	[.050]	0.314	[.116]***	-0.239	[.134]*

(b) Baseline controls and survey round/block fixed effects (N=3765/239)							
Treatment arm	Outcome index	ATE using survey data ($\hat{\theta}$)		ATE on measurement error(θ^Δ)		Adjusted ATE ($\hat{\theta} - \theta^\Delta$)	
		β	S.E.	β	S.E.	β	S.E.
		(1)	(2)	(3)	(4)	(5)	(6)
Cash only	Sensitive	-0.078	[.071]	-0.313	[.177]*	0.234	[.268]
	Nonsensitive	0.081	[.045]*	0.147	[.136]	-0.066	[.195]
Therapy only	Sensitive	-0.193	[.071]***	-0.215	[.180]	0.022	[.264]
	Nonsensitive	-0.009	[.046]	0.268	[.144]*	-0.277	[.213]
Therapy and cash	Sensitive	-0.402	[.069]***	0.052	[.183]	-0.454	[.270]*
	Nonsensitive	0.073	[.045]	0.267	[.122]**	-0.195	[.197]

Notes: The survey-based average treatment effect (ATE) estimates, $\hat{\theta}$, pool all survey rounds and regress each outcome on treatment indicators. Standard errors are robust and clustered by individual. Estimates of the bias from treatment, θ^Δ , come from a regression of the difference in the survey and validation measures on an indicator for treatment arms. Standard errors are robust and clustered by block. The difference, $\hat{\theta} - \theta^\Delta$, is an estimate of the true treatment effect after adjusting for observed bias. It is calculated as the linear difference of the estimates and the standard error is calculated via bootstrapping (we performed 1000 draws from the sample, with replacement; we calculated $\hat{\theta}$, θ^Δ and $\hat{\theta} - \theta^\Delta$ for each draw; and we generated the standard error on $\hat{\theta} - \theta^\Delta$ using the distribution from these draws).

*** p<0.01, ** p<0.05, * p<0.1

two cases: adjusting for survey round and randomization block fixed effects (Panel a), and adjusting for baseline covariates as well (Panel b). We use these to calculate an adjusted treatment effect, $\hat{\theta} - \theta^\Delta$, for our sensitive and nonsensitive indexes.³²

First, the results imply that the true treatment effect of therapy plus cash on sensitive behaviors is no smaller than what we estimate with self-reported survey data. Indeed, the fall in these sensitive behaviors may even be greater than the survey reports suggest (Columns 5 and 6). This holds true for each of the individual sensitive behaviors, which are shown in Appendix C.3. Despite the large standard errors introduced by the small validation sample, the adjusted treatment effect on an index of all sensitive behaviors is larger in absolute value and significant at the 1% level (when we use round and block fixed effects, in Panel a). When baseline controls are added in Panel b, the coefficient is similar but the standard errors have increased and it is statistically significant at the 10% level only. This is partly because our validation sample is small, and the number of baseline controls is large, reducing the degrees of freedom.³³

In contrast, adjusting expenditures changes the sign of the treatment effect we estimated using survey data, and hence affects our conclusions about the intervention. Based on the survey data, we estimated that the cash grant led to a short term increase in expenditures. But the slight underreporting of expenditures, especially the excess under-reporting by the control group, may have exaggerated the effects of cash on expenditures and incomes (judging by these two expenditures at least). The adjusted (pooled) treatment effect on expenditures is negative for all treatments and both specifications, generally with nontrivial magnitudes but only statistically significant in a couple of instances (Columns 5 and 6).³⁴

6 Discussion and conclusions

Perhaps the most important lesson from this exercise is that structured, in-depth, and representative qualitative work revealed patterns of measurement error that were quite different from our priors, despite extensive experience with the study group. There is little data on measurement error, however, and so (like many) our priors were unavoidably rooted more in

³²Recall that $\theta^\Delta \rightarrow E(\hat{\theta} - \theta)$ as the validation measure approaches the “truth”. Appendix C.3 contains the results for each component of these indexes.

³³Meanwhile, the underreporting of gambling (displayed in Appendix C.3) does not have a statistically significant association with treatment. However, those who received cash alone underreported gambling to the surveyors more often than control group members, and so the measurement error in gambling is probably a combination of a general desirability bias as well as one correlated with treatments. A larger sample size would be needed to separate these more precisely.

³⁴We see a similar pattern with another expenditure-related item, homelessness, in Appendix C.3—the survey-reported decline in homelessness tends to lose a lot of its significance with adjustment.

what seemed like common sense and in common causal identification concerns rather than an informed understanding.

We worried, for instance, that high-risk young men might have special reasons to conceal their behavior—such as suspicion of outsiders, or a desire to receive program benefits in the future. Given that we were focused on measuring the treatment effects of a therapy program that discouraged various anti-social and unhealthy behaviors, we were also concerned that the treated would underreport such behaviors out of experimenter demand or social desirability bias induced by the therapy. Our multi-method approach revealed that the nature of measurement error was quite different.

For this specific field experiment, two findings stand out. First, the qualitative validation suggests that the underreporting in sensitive behaviors was modest, not statistically significant, concentrated in the control and cash only groups, and limited to one behavior in particular (gambling). Meanwhile, expenditures seemed to be broadly underreported in the survey, most of all in the control group.

Based on qualitative interviews, our impression is that these “sensitive” behaviors, while not acceptable within Liberian society as a whole, are not so stigmatized that most men in our sample feel ashamed to report them, perhaps because these men belong to a counterculture in which these activities are common. Moreover, the risk of punishment was minuscule.³⁵ Hence underreporting tended to be modest overall.

An exception was gambling. Gambling, unlike the other sensitive behaviors, is not a defining characteristic of the counterculture to which study participants belong. Furthermore, after receiving a cash handout, it’s possible that men were reluctant to admit they’d gambled some of it away. The same was true of the control group, to a lesser degree, who may have hoped for cash in future.³⁶ Alternatively, the therapy treatment could have increased the familiarity, trust or reciprocity between the subjects and implementers, and so men who received therapy were less likely to underreport.

The second major finding is that the expenditure-related activities were systematically underreported across all arms, and especially in the control group. The effect of treatment on measurement error is large and statistically significant in all arms. This finding is extremely important given that expenditure and consumption surveys are the principal means

³⁵The Liberian police are largely incapable of investigating and prosecuting all but the most grave crimes. Thus, these behaviors are not endangering to most of our sample and their peers, and they discuss them freely.

³⁶There was no future program (this was communicated repeatedly), and the original field experiment actually used countervailing criteria for recruiting subjects, but these features of our program were unusual compared to standard NGO practice.

of measuring material well-being and poverty in most developing countries. We see two main possible explanations:

1. *Strategic behavior.* Since there was underreporting across all treatment arms, every study participant may have had an incentive to exaggerate their neediness in the hopes of future programs. This echoes our gambling result.

Why more so in the control group? It's possible that the fewer treatments a man received, the more strategically he behaved on the survey, trying to appear poorer to encourage eligibility for future treatment. Those who received therapy, for example, might be interested in the cash. Phone charging and going to a video club are considered discretionary spending, and if a respondent wanted to signal destitution, he might underreport spending on these items.

We view this explanation as plausible, although there are caveats. First, the control group did not over-report homelessness to the same degree, which is an obvious indicator of need (although perhaps observable enough that it was perceived as harder to falsify on a survey). Second, while drug use is technically an expenditure, this was not underreported to signal poverty. (One reason may be that the drug users were mainly heavy marijuana users, indeed so heavy that this was somewhat obvious, thus potentially making it less prone to falsification.) Third, in principle those who received one of the earlier treatments also had incentives to behave strategically in the hopes of future programs. Treated men almost universally lobbied for additional assistance.

2. *Salience and recall bias.* Expenditures could be more subject to recall error, because they are less regular and possibly less salient than drug use or crime. There is ample evidence that consumption and expenditure data are underreported, and that underreporting increases with the period of recall, the lower the reported consumption per standardized unit of time, and the less salient the purchase (Deaton and Grosh, 1997; Gibson, 2006; Beegle et al., 2012). Neither video clubs nor mobile phone charging were particularly salient. People may also make cognitive errors when aggregating over a construct such as “the last two weeks.” Finally, the expenditures survey module was long and much more subject to fatigue, compounding underreporting.

Recall bias is plausible, but we are also looking for explanations that would correlate with treatment. There are a few possibilities. Treatment could have increased attention and mindfulness. The therapy was explicitly designed to reduce impulsive behavior and to increase planning. There is some evidence that impulsivity in fact decreased (Blattman et al., 2015). The cash transfer could have had a similar effect for different

reasons. Studies have also shown that recall bias in consumption data increases with poverty (Beegle et al., 2012). This is consistent with evidence that cognition is taxed by poverty and scarcity (Mani et al., 2013). Presumably hunger would affect survey fatigue and mindfulness. The cash grant (and short run decrease in poverty) could have had a similar effect on the margin. Finally, receiving either treatment could have produced enough reciprocity that the treated group exercised more care in recalling less salient data. We regard these explanations with caution, but cannot reject them.

Both explanations are plausible but come with caveats, and so we refrain from a firm conclusion about the sources of measurement error. Given the importance of expenditure surveys in research it bears replication and further research.

In retrospect, we also see that, had the measurement error run in the opposite direction, it would have been difficult to distinguish the large effect of therapy and cash on crime from systematic measurement error given our sample size. The effect of treatment on our proxy for survey over-reporting would have been underpowered. We estimate that doubling the size of the validation sample would have increased power by about a third. The marginal cost of validation per respondent was roughly equal to that of running a survey. Thus we estimate that we could have doubled the number of validations by either increasing the evaluation budget by 3%, or reducing the total sample size by 3%. Given how much the credibility of these types of studies rests on self-reported data, this strikes us as a reasonable investment.

Overall, these results reinforce a fundamental principle of survey methodology: the importance of validating measurements with multiple instruments. To some extent our findings are unexpected (even puzzling), and the explanations are somewhat speculative. With more validation studies, of all kinds, we may start to see systematic patterns. We regard our multi-method approach as largely complementary to list (item count), random response, and endorsement experiments. It is useful to have more methods available.

Like other methods, ours requires a priori assumptions—in this case, that in-depth observation is less prone to bias and does not introduce major new biases. Our method is also more costly to implement, though not necessarily relative to the average cost of large surveys or modest impact evaluations. The stakes are high enough in many panel studies, field experiments, and other impact evaluations that validating a handful of key outcomes seems important for the individual project. Qualitative validation performs well enough, and yields sufficiently important results, that our approach deserves more systematic use and examination, ideally alongside these other methods.

References

- Asher, H. B. (1974). Some consequences of measurement error in survey data. *American Journal of Political Science* 18(2), 469–485.
- Baruch, G. (1981). Moral tales: parents’ stories of encounters with the health professions. *Sociology of Health & Illness* 3(3), 275–295.
- Beegle, K., J. De Weerd, J. Friedman, and J. Gibson (2012). Methods of household consumption measurement through surveys: Experimental results from tanzania. *Journal of Development Economics* 98(1), 3–18.
- Blair, G. and K. Imai (2012). Statistical analysis of list experiments. *Political Analysis* 20(1), 47–77.
- Blattman, C., J. Jamison, and M. Sheridan (2015). Reducing crime and violence: Experimental evidence on adult noncognitive investments in Liberia. *Working paper*.
- Bound, J., C. Brown, and N. Mathiowetz (2001). Measurement error in survey data. In *Handbooks in Economics*, Volume 2, pp. 3705–3843.
- Bryman, A. (2003). *Quantity and quality in social research*. New York: Routledge.
- Bullock, W., K. Imai, and J. N. Shapiro (2011). Statistical analysis of endorsement experiments: Measuring support for militant groups in pakistan. *Political Analysis* 19(4), 363–384.
- de Mel, S., D. J. McKenzie, and C. Woodruff (2009). Measuring microenterprise profits: Must we ask how the sausage is made? *Journal of Development Economics* 88(1), 19–31.
- Deaton, A. and M. Grosh (1997). Consumption. *Designing Household Survey Questionnaires for Developing Countries: Lessons from Ten Years of LSMS Experience*.
- Deming, D. J. (2011). Better schools, less crime? *The Quarterly Journal of Economics* 126(4), 2063–2115.
- Fox, R. C. (2004). Observations and reflections of a perpetual fieldworker. *The ANNALS of the American Academy of Political and Social Science* 595(1), 309–326.
- Freedman, D. A. (1991). Statistical models and shoe leather. *Sociological methodology* 21(2), 291–313.

- Gibson, J. (2006). Statistical tools and estimation methods for poverty measures based on cross-sectional household surveys. *Handbook on Poverty Statistics*.
- Hausman, J. (2001, October). Mismeasured variables in econometric analysis: Problems from the right and problems from the left. *The Journal of Economic Perspectives* 15(4), 57–67.
- Jamison, J. C., D. Karlan, and P. Raffler (2013). Mixed-method evaluation of a passive mHealth sexual information texting service in uganda. *Information Technologies & International Development* 9(3).
- Karlan, D. and J. Zinman (2008). Lying about borrowing. *Journal of the European Economic Association* 6(2-3), 510–521.
- Karlan, D. S. and J. Zinman (2012). List randomization for sensitive behavior: An application for measuring use of loan proceeds. *Journal of Development Economics* 98(1), 71–75.
- LeCompte, M. D. (1987). Bias in the biography: Bias and subjectivity in ethnographic research. *Anthropology & Education Quarterly* 18(1), 43–52.
- LeCompte, M. D. and J. P. Goetz (1982). Problems of reliability and validity in ethnographic research. *Review of educational research* 52(1), 31–60.
- Mani, A., S. Mullainathan, E. Shafir, and J. Zhao (2013). Poverty impedes cognitive function. *Science* 341(6149), 976–980.
- Martinelli, C. and S. W. Parker (2009). Deception and misreporting in a social program. *Journal of the European Economic Association* 7(4), 886–908.
- McKenzie, D. (2012). Beyond baseline and follow-up: The case for more T in experiments. *Journal of Development Economics* 99(2), 210–221.
- NSSO Expert Group (2003). Suitability of different reference periods for measuring household consumption. results in pilot survey. *Economic and Political Weekly* 38(4), 25–31.
- Paluck, E. L. and D. P. Green (2009). Deference, dissent, and dispute resolution: An experimental intervention using mass media to change norms and behavior in rwanda. *American Political Science Review* 103(4), 622–644.
- Power, R. (1989). Participant observation and its place in the study of illicit drug abuse. *British Journal of Addiction* 84(1), 43–52.

- Raghavarao, D. and W. T. Federer (1979). Block total response as an alternative to the randomized response method in surveys. *Journal of the Royal Statistical Society. Series B (Methodological)*, 40–45.
- Ram, P. K., A. K. Halder, S. P. Granger, T. Jones, P. Hall, D. Hitchcock, R. Wright, B. Nygren, M. S. Islam, and J. W. Molyneaux (2010). Is structured observation a valid technique to measure handwashing behavior? Use of acceleration sensors embedded in soap to assess reactivity to structured observation. *The American Journal of Tropical Medicine and Hygiene* 83(5), 1070–1076.
- Scacco, A. (2010). Who riots? Explaining individual participation in ethnic violence. *Dissertation, New York University*.
- Spitzer, R. L., K. Kroenke, and J. B. W. Williams (1999). Validation and utility of a self-report version of PRIME-MD: the PHQ primary care study. *Journal of the American Medical Association* 282(18), 1737–1744.
- Wilson, S. (1977). The use of ethnographic techniques in educational research. *Review of Educational Research* 47(1), 245–265.

Appendix for online publication

A Sample details

A.1 Recruitment and timing

The implementers of the therapy, NEPI, approached roughly 1500 men, and 999 agreed to speak and to enter the sample (we do not have data on those who refused). To avoid recruiting groups of friends and colleagues (i.e. to minimize correlated outcomes and spillovers) NEPI approached just one out of every seven to ten high-risk men they identified. Of those assigned to therapy, nearly all attended at least a day, and two thirds completed it. The higher risk men were the most likely to finish.

Table A.1 describes the structure of the sample and treatment assignment. For logistical reasons we recruited, treated, and studied the men in three phases. The first phase was a small, successful pilot. Based on its results, we decided to scale the program to a further 900 with only minor changes to the interventions and study protocols. To accommodate the large number of men we scaled up in two phases. An advantage of this phased implementation is that we calculate impacts over cohorts with different spatial and time shocks.

A.2 Randomization balance and survey attrition

In general, all randomizations display balance. Table A.2 reports an ordinary least squares (OLS) regression of an indicator for random assignment on baseline variables and block fixed effects (fixed effect coefficients not shown). At the base of the table we report the p-value from an F-test of joint significance of all baseline covariates.³⁷ Columns 1 and 2 look at assignment to cash treatment, Columns 3 and 4 assignment to the CBT treatment, and Columns 5 and 6 assignment to validation.

The baseline covariates are not jointly significant in predicting any treatment (Columns 1 to 4). A small number of baseline variables are individually statistically significant, but no more than would be expected at random. In terms of selection into validation, selection is not very systematic. The covariates (including fixed effects) explain just 14% of the variation, and the baseline covariates are not jointly significant (p-value of 0.1129). 5 of the 28 displayed baseline covariates (18%) have a p-value less than 0.1. Those selected into validation are slightly less likely to be married, support slightly more women, are slightly more likely to have been an ex-combatant, and are slightly less likely to drink alcohol. There

³⁷Randomization block fixed effects are omitted from the table and the F-test.

Table A.1: Study sample and treatment assignment by block and phase

Phase	Start	Block	N	% recruits assigned to:			
	MM/YY			Therapy	Cash	Both	None
1	12/10	Red Light	100	28%	24%	25%	23%
2	06/11	Red Light	219	27%	25%	24%	24%
	06/11	Central Monrovia	179	32%	19%	32%	17%
3	03/12	Clara Town	175	29%	27%	23%	21%
	02/12	Logan Town	86	27%	29%	20%	24%
	02/12	New Kru Town	240	26%	27%	24%	23%

Notes: By design, percentages assigned to each treatment can vary from 25%. Treatments were allocated by public draw of colored chips from a bag, without replacement, in groups (blocks) of about 50 people. There were more chips in the bag than people by design (to avoid the last person having a predetermined probability of assignment).

are few substantively or statistically significant differences in the sensitive behaviors we are interested in measuring and validating. Overall, selection imbalance seems relatively modest.

In addition, attrition appears to be relatively unsystematic. Table A.3 reports OLS regression of an indicator for being unfound on baseline variables and block fixed effects (fixed effect coefficients not shown). At the base of the table we report the p-value from an F-test of joint significance of all baseline covariates. Columns 1 and 2 look at attrition in the endline survey (n=4096), and Columns 3 and 4 selection into the validation sample (n=297), where each endline survey is a separate observation.

For the endline survey, attrition was just 8% on average across all endline survey rounds, and this attrition appears to be unsystematic. Of the 144 men we could not survey, 2 had died, 4 were incarcerated, 2 refused to answer the survey, and the remainder could simply not be found despite repeated attempts over several weeks. As a result, there are 315 surveys that were not completed (of 4,096).³⁸ From Columns 1 and 2, only 4 of the 30 (13%) covariates have a p-value below .1 and jointly the covariates have a p-value of 0.53.

For the validation, 3 of the 30 (10%) covariates have a p-value less than .1 and jointly they have a p-value of .38. Attritors from validation are less educated and less disabled, and are less likely to have received the cash treatment. There are some covariates that show a substantive difference (e.g. attritors are 9 percentage points less likely to be hard drug

³⁸Of the men who are missing at least 1 survey (155), 9 died (7% of missing surveys), 19 were in prison (11% of missing surveys), 6 refused (3%).

Table A.2: Randomization balance to treatments and to selection into validation

Baseline variable	Dependent variable: Assigned to					
	Cash		Therapy		Validation sample	
	Coeff.	SE	Coeff.	SE	Coeff.	SE
	(1)	(2)	(3)	(4)	(5)	(6)
Age	0.001	[.004]	0.000	[.004]	0.001	[.001]
Married/living with partner	-0.024	[.050]	0.019	[.049]	-0.033	[.017]**
# of women supported	0.009	[.027]	-0.040	[.027]	0.033	[.010]***
# children under 15	-0.007	[.006]	-0.001	[.006]	0.000	[.002]
Muslim	-0.014	[.056]	0.074	[.055]	0.016	[.020]
Years of schooling	0.002	[.006]	-0.001	[.006]	0.001	[.002]
Literacy score (0-2)	-0.008	[.024]	0.020	[.024]	0.004	[.008]
Math score (0-5)	-0.009	[.012]	-0.021	[.012]*	-0.003	[.004]
Health index (0-6)	-0.006	[.012]	0.020	[.012]	-0.002	[.004]
Disabled	-0.072	[.064]	0.043	[.063]	0.008	[.019]
Monthly cash earnings (USD)	0.000	[0000]*	0.000	[0000]	0.000	[0000]
Durable assets index, z-score	0.035	[.019]*	-0.013	[.019]	-0.002	[.006]
Savings stock (100s of USD)	-0.027	[.027]	-0.015	[.026]	-0.008	[.008]
Able to get a loan of \$300	0.009	[.055]	0.020	[.054]	0.052	[.021]**
Average weekly work hours in:						
Potentially illicit activities	0.000	[.001]	0.000	[.001]	0.000	[0000]
Agricultural Labor	-0.001	[.004]	0.006	[.004]	0.000	[.001]
Low-skill wage labor	0.001	[.001]	0.000	[.001]	0.000	[0000]
Low-skill business	0.001	[.001]*	0.002	[.001]**	0.000	[0000]
High-skill work	0.004	[.002]**	0.002	[.002]	0.000	[.001]
Ex-combatant	0.037	[.038]	0.010	[.038]	0.026	[.012]**
Currently sleeping on the street	-0.025	[.044]	-0.033	[.043]	-0.008	[.014]
Times went hungry last week	0.013	[.012]	0.004	[.012]	-0.001	[.004]
Sells drugs	-0.002	[.050]	0.017	[.049]	-0.012	[.017]
Drinks alcohol	0.035	[.039]	0.055	[.038]	-0.025	[.014]*
Uses marijuana daily	0.014	[.039]	0.038	[.038]	0.019	[.013]
Uses hard drugs daily	0.071	[.052]	-0.117	[.051]**	-0.017	[.017]
Stole in past two weeks	0.002	[.038]	0.030	[.038]	-0.009	[.013]
Joint significance (p-value)	0.771		0.320		0.112	
R-squared	0.510		0.549		0.136	
Observations	999		999		2994	

Notes: Columns 1 to 4 report the coefficient and standard error on each variable from an OLS regression of a treatment indicator on all variables listed in the table plus block fixed effects. Columns 5 and 6 do the same where the dependent variable is selection into validation. For assignment to treatment the sample is the 999 men in the study. For the validation exercise it is the 4096 endline surveys run at various points in time, although because some rounds have none selected for validation we exclude these rounds and have 2994 remaining. Each endline survey is a separate observation, and robust standard errors are clustered at the individual level.

*** p<0.01, ** p<0.05, * p<0.1

Table A.3: Attrition from the endline survey and from validation

Baseline variable	Dependent variable			
	Unfound at endline		Unable to validate	
	Coeff. (1)	SE (2)	Coeff. (3)	SE (4)
Assigned to cash only	-0.013	[.018]	-0.044	[.061]
Assigned to therapy only	-0.008	[.017]	0.060	[.069]
Assigned to therapy plus cash	-0.014	[.018]	-0.056	[.065]
Age	0.000	[.001]	0.000	[.005]
Married/living with partner	-0.007	[.021]	0.040	[.071]
# of women supported	-0.002	[.009]	0.023	[.052]
# children under 15	-0.003	[.002]*	-0.010	[.007]
Muslim	0.084	[.029]***	-0.047	[.070]
Years of schooling	0.002	[.002]	-0.018	[.009]*
Literacy score (0-2)	0.004	[.009]	0.035	[.038]
Math score (0-5)	0.002	[.004]	-0.013	[.018]
Health index (0-6)	-0.002	[.005]	-0.015	[.018]
Disabled	0.004	[.021]	-0.136	[.077]*
Monthly cash earnings (USD)	0.000	[0000]	0.000	[0000]
Durable assets index, z-score	0.014	[.008]*	0.044	[.029]
Savings stock (100s of USD)	0.011	[.012]	0.008	[.041]
Able to get a loan of \$300	-0.008	[.022]	-0.049	[.077]
Average weekly work hours in:				
Potentially illicit activities	0.000	[0000]	0.000	[.001]
Agricultural Labor	0.000	[.001]	-0.002	[.002]
Low-skill wage labor	0.000	[0000]	-0.002	[.001]*
Low-skill business	-0.001	[0000]**	0.000	[.001]
High-skill work	0.000	[.001]	0.002	[.003]
Ex-combatant	0.002	[.014]	-0.048	[.057]
Currently sleeping on the street	0.004	[.017]	0.102	[.069]
Times went hungry last week	-0.002	[.005]	-0.024	[.018]
Sells drugs	-0.013	[.016]	0.062	[.076]
Drinks alcohol	-0.011	[.015]	-0.036	[.064]
Uses marijuana daily	-0.001	[.016]	0.008	[.059]
Uses hard drugs daily	-0.017	[.020]	-0.090	[.077]
Stole in past two weeks	-0.005	[.015]	0.039	[.057]
Joint significance (p-value)	0.531		0.381	
R-squared	0.167		0.318	
Observations	4096		297	

Notes: Columns 1 to 4 report the coefficient and standard error on each covariate from an OLS regression of an attrition indicator on all covariates listed in the table plus block fixed effects. Each endline survey is a separate observation, and robust standard errors are clustered at the individual level.

*** p<0.01, ** p<0.05, * p<0.1

users, and 10.2 percentage points more likely to be sleeping on the street) but these are not statistically significant.

B Further details of validation methodology

The goal of this section is to expand on certain activities and aspects of the approach in order to facilitate replication or adaption of the method.

B.1 Validator roles and responsibilities

The validation team was comprised of three main roles (in order of seniority): the Project Coordinator, Team Leader, and Validator.

- The primary responsibilities of the Project Coordinator were to: (i) build an effective team, including hiring and training, (ii) identify target respondents, (iii) supervise the team leader and validators, and (iv) monitor field progress.
- The Team Leader was expected to perform all duties, functions, and responsibilities of a qualitative researcher in addition to: (i) assisting with trainings, (ii) overseeing field logistics, such as assigning validators to respondents, (iii) mentoring validators, as well as (iv) monitoring and reporting on the team’s progress.
- Validators were principally responsible for the following duties: (i) locating respondents, (ii) collecting and accurately recording data, and (iii) routinely communicating progress to the Team Leader.

Validators underwent an intensive 10-day training, involving both classroom learning and extensive field training, before being selected.³⁹ The aim of the training was to develop and refine trainees’ skills in acquiring informed consent, building rapport with respondents, collecting and recording data, and analytical reasoning. Trainings were held for eight hours each day and, over the course of 10 days, transitioned gradually from exclusive classroom learning to field trainings with short debriefing sessions. Field trainings provided trainees with opportunities to practice the skills and techniques they had learned.

B.2 Logistical organization

Consent Process Validators obtained informed consent from each respondent prior to commencing the first “hangout” session. Validators were provided with a consent script

³⁹We trained more qualitative researchers than were needed for the exercise. Those who exhibited superior performance during the trainings were selected as validators.

which they recited to respondents by memory to maintain informality. During the consenting process, study participants were advised of the research team’s interest in learning about their lives. However, respondents were not informed of the precise data being collected. This exercise was discussed in the formal consent process to the endline survey, however.

Strategies for assigning respondents Respondents were assigned based on their proximity to the validators’ other assigned respondents or their workload distribution. Respondents whom we had difficulty tracking or getting to answer the survey were often assigned to one of the more experienced validators.

Coordination and communication The validators aimed to begin “hangout” sessions the day after subjects completed their quantitative surveys, and to conduct all four visits within 10 days of respondents having been surveyed. To facilitate coordination and communication, validators were assigned to a survey team. The survey team would alert the validator when the respondent was being surveyed. A validator would go to the location where the respondent was being surveyed, the enumerator would introduce the respondent to the validator if possible, and the validator would set up a time to meet with the respondent the next day. This coordination and communication were essential and typically saved hours or days of tracking.

Workflow management The Team Leader had several strategies for developing work plans. For instance:

- Validators would try to validate two people per day, though this was not always possible because of the time it took to locate respondents.
- Their work hours were flexible, and might include early mornings or evening visits (though late evening visits were discouraged due to safety concerns).
- If a validator was assigned three respondents, he would rotate them, such that all four visits with the three respondents were completed over the course of six days.
- If validators knew a particular respondent was difficult to locate, they would often begin their day early trying to find him. If after a few hours they were not able to locate him, they would have a “backup” respondent for the morning slot in mind (that is, one they knew would be easy to locate). In this way, validators tried to ensure too much time would not elapse between hangouts with “difficult-to-locate” respondents.

- Validators also needed time before and after their “hangout” sessions to locate respondents and record collected data. Thus, if a validator were assigned three respondents, she typically needed eight days to complete the entire exercise (one day for locating respondents, six days for visiting with respondents, and one day for transferring her field notes into the appropriate template).

Strategies Validators were provided with cell phones (and cell phone credit) to contact respondents, each other, and the survey team. Petty cash for transportation to and from hangout sessions was also given to validators on a weekly basis.

Since this population was largely transient, the use of tracking sheets (that specified the whereabouts and contact information of respondents, their relatives and friends) proved crucial. This form was used to locate subjects and was continuously updated as new information became available. Validators also carried a small photograph of their assigned respondent in order to identify him.⁴⁰

Recording Data Finally, the Validators and Project Coordinator kept track of the hangout sessions completed each day and pertinent information about the validation sample (see B.1 for a sample form used) This information was obtained from the Team Leader during debriefing sessions at the end of each day, notes were added, and these were the basis of the review process with the authors for case-by-case final coding.

Training activities

The Liberia training manual is available from the authors on request.

B.3 Evidence collection techniques

Below we outline the kinds of evidence gathering techniques used by the validators and illustrative examples—taken from actual validation data—of each technique.

1. Direct Questioning. The validator directly asked the respondent about his engagement in the activity.
 - Stealing: Having established a strong rapport, and observing that the respondent did not get along well with others in the community, the validator asked whether the respondent ever stole from the community members. The subject replied that

⁴⁰We obtained respondent consent for tracking photos. This photo was not shown to anyone outside of the study team and was destroyed after meeting with the respondent.

Figure B.1: Sample data recording sheet

SDV TEAM MEMBER	TIM WATSON
ACTIVITY	STEALING
How did you bring up the subject?	3
1- Ask him a direct question	
2- Ask him an indirect question	
3- Gave a story/scenario	
4- You saw him do it	
5- You saw proof that he does it	
6- Other	
How did you bring up the subject? [DESCRIBE]	I told the respondent a story about a person stole some money from a wealthy man's car to pay his rent.
What do you know about the respondent and this activity?	1
1- He never does it	
2- He can do it sometimes	
3- He can do it one one time	
4- He can do it daily	
5- Other	
What do you know about the respondent and this activity? [DESCRIBE]	The respondent told me that while some people may resort to stealing when they are faced with problems, he does not steal. He told me that he solves his problems in ways other than stealing. In fact, he told me he's never stolen before, indicating he did not steal during the survey recall period (Dec. 25, 2013 - Jan. 8, 2014).
Has the respondent done it in the 2 weeks before the survey?	0
1- Yes, 0-No	
What are the reasons you have to believe this is true? ONLY ADD NEW INFORMATION, DON'T REPEAT INFORMATION	I believe the respondent did not steal during the survey window (Dec. 25, 2013 - Jan. 8, 2014) because he stated that he has never stolen before. Furthermore, I did not witness the respondent stealing at any time.

he steals when he has the opportunity and explained that he stole two cartons of frozen chicken from a warehouse just last week (i.e. during the survey window).

- Phone Charging: After commenting on how nice the respondent's phone was, the validator asked how often the respondent pays to charge it.

2. Indirect Questioning. The respondent brings up a topic or question closely related to engagement in the activity.

- Homelessness: The validator met the participant in front of a house and asked who lived there. The respondent replied that it was his friend's home, though he had been living there for the last three months. This information was then corroborated in casual conversations with neighbors.
- Video Club: After discussing their enjoyment of football (soccer), the validator asked the participant whether he preferred to watch games from home or the video club. The subject noted that he loves watching games from the video club and had recently watched the Under 20 World Cup at a video club the week prior.

3. Story/Scenario. The validator discusses a personal story or hypothetical scenario involving another person engaging in the activity. The respondent is invited to comment.

- Homelessness: To introduce the topic of homelessness, the validator told the story of a childhood friend who slept on the streets, as opposed to his parents' home, because of the freedom it allowed him. The validator then asked for the subjects' thoughts on the matter, which segued into a conversation about the respondents' own living situation.
- Video Club: The validator noted that if he lived in the respondents' community, he would take advantage of the close proximity to the video club and frequent it to watch football (soccer) games. This scenario led to a conversation about the subject going to the video club the previous week (i.e. during the survey window) with his friends to watch a game.

4. Witnessing or finding evidence. The validator sees the respondent engaging in the activity in question or finds physical evidence that the respondent has recently engaged in the activity. Note that this was almost solely used as supporting evidence.

- Marijuana: While the respondent openly smoked marijuana, a validator asked him to share his experiences smoking – probing to understand when the participant began smoking and how frequently he engaged in the activity. In turn, the

validator offered stories of his struggles with drinking. The respondent then explained that he began smoking in 2002 and had smoked daily for the last several months. In another case, a validator noted that the respondent walked behind a building and when he returned, his eyes were red, he appeared high, and he smelled heavily of marijuana.

- Phone charging: On the first day of a hangout, a validator walked with a respondent to a phone charging booth where the respondent needed to pay to charge his phone.

5. Third party account. The validator asks the family and friends of the respondent whether or not he engages in the activity. Note that this was almost solely used as supporting evidence.

- Homelessness: The respondent took the validator to his home, where the validator met the respondent's family. The family members confirmed the respondent lived in the family home.
- Phone Charging: The respondent told the validator that he did not have a phone. The validator asked the respondents' sister if this was true, and she confirmed that the respondent had lost his phone two months prior.

6. Other or unclear method. This includes unprompted information from the respondent, and also inferring from other knowledge such as when the validator learns about the respondent's life and is able to infer that the respondent did or did not engage in the activity.

- Homelessness: While hanging out with the validator, the respondent announced that he wanted to show the validator the home that he had been living in for the past year.
- Video Club: The validator learned from the respondent that he was in jail throughout the two weeks preceding the survey. Because of this, the respondent could not have attended the video club.
- In a few cases, the validator's notes are unclear about how (s)he prompted the respondent to discuss the activity. In these cases, we give the evidence a separate code, although these are likely cases of questioning.

C Other analysis

C.1 Varying measurement error by treatment status

Another approach to estimating the measurement error is to let it vary by treatment status.

Consider the probability of self-reported behavior y^s conditional on validated behavior ($y^v = 1$) and non-validated behavior ($y^v = 0$). Then,

$$\begin{aligned} E(y^s) &= E(y^s|y^v = 1)P(y^v = 1) + E(y^s|y^v = 0)(1 - P(y^v = 1)) \\ &= P(y^s = 1|y^v = 1)E(y^v) + P(y^s = 1|y^v = 0)(1 - E(y^v)) \\ &= E(y^v)(P(y^s = 1|y^v = 1) - P(y^s = 1|y^v = 0)) + P(y^s = 1|y^v = 0) \end{aligned}$$

If we condition on treatment assignment, we get

$$E(y^s | T) = E(y^v | T)(P(y^s = 1|y^v = 1, T) - P(y^s = 1|y^v = 0, T)) + P(y^s = 1|y^v = 0, T)$$

If we assume that $p(y^s = 1|y^v, T) = p(y^s = 1|y^v)$ (i.e. that survey measure is uncorrelated with treatment once you control for validated measure) and let $p_1 = p(y^s = 1|y^v = 1)$ and $p_0 = p(y^s = 1|y^v = 0)$, we get

$$E(y^s|T) = (p_1 - p_0)E(y^v|T) + p_0$$

This statistic helps show how treatment (T) affects the joint distribution of the validated measurement (“true” behavior, y^v) and self-reporting (y^s) in the survey. It corresponds to the regression

$$y_i^s = \beta_0 + \beta_1 y_i^v + \beta_2 T_i + \beta_3 (T_i \times y_i^v) + \tilde{\mu}_i.$$

If we let $\beta_1 = 1$ and $\beta_3 = 0$, we get

$$y_i^s - y_i^v = \beta_0 + \beta_2 T_i + \tilde{\mu}_i.$$

which corresponds to Equation 11.

C.2 Patterns of misreporting

Table C.1 reports the number of cases where the survey and validation measures do not agree, divided into cases of survey over- and under-reporting relative to the validation measure.

Table C.1: Survey under-reporting and over-reporting by evidence type: With and without any external evidence (direct observation or third party confirmation)

	Sensitive behaviors					Expenditures			
	All (1)	Steal (2)	Marijuana (3)	Gamble (4)	Homeless (5)	All (6)	Video (7)	Phone (8)	All (9)
A. UNDERREPORTS									
Any external evidence	38	2	11	15	10	11	4	7	49
No external evidence	71	21	10	31	9	88	64	24	159
Total count	109	23	21	46	19	99	68	31	208
% with any external evidence	35%	9%	52%	33%	53%	11%	6%	23%	24%
B. OVER-REPORTS									
Any external evidence	23	3	1	2	17	1	1	0	24
No external evidence	63	24	13	19	7	33	22	11	96
Total count	86	27	14	21	24	34	23	11	120
% with any external evidence	27%	11%	7%	10%	71%	3%	4%	0%	20%

C.3 Robustness

Below we report three versions of Table 7 for all index components. First, we run the regressions while controlling only for round and block fixed effects.

Second, we rerun the regressions without controlling for block fixed effects and round fixed effects.

Third we report the same regressions, controlling for baseline covariates, randomization block fixed effects, and survey round fixed effects.

Table C.2: Estimates of treatment effects ($\hat{\theta}$) and treatment effect bias (θ^Δ) by outcome and treatment, round and block fixed effects

	Potentially sensitive behaviors					Expenditures		
	All (0-4) (1)	Steal (2)	Marijuana (3)	Gamble (4)	Homeless (5)	All (0-2) (6)	Video (7)	Phone (8)
A. CASH ONLY								
ATE using survey data ($\hat{\theta}$)	-0.057 [.095]	-0.031 [.029]	0.012 [.040]	0.025 [.029]	-0.062 [.031]**	0.080 [.052]	0.027 [.034]	0.053 [.037]
ATE on measurement error (θ^Δ)	-0.178 [.190]	-0.018 [.087]	-0.046 [.067]	-0.047 [.088]	-0.063 [.076]	0.285 [.130]**	0.106 [.109]	0.185 [.079]**
Adjusted ATE ($\hat{\theta} - \theta^\Delta$)	0.121 [.195]	-0.014 [.087]	0.058 [.074]	0.072 [.093]	0.001 [.076]	-0.205 [.143]	-0.079 [.112]	-0.132 [.082]
B. THERAPY ONLY								
ATE using survey data ($\hat{\theta}$)	-0.186 [.092]**	-0.045 [.028]	-0.022 [.040]	-0.087 [.026]***	-0.031 [.031]	0.000 [.050]	0.001 [.032]	-0.001 [.035]
ATE on measurement error (θ^Δ)	0.004 [.199]	-0.003 [.082]	0.009 [.056]	0.044 [.101]	-0.038 [.088]	0.335 [.134]**	0.169 [.105]	0.171 [.087]*
Adjusted ATE ($\hat{\theta} - \theta^\Delta$)	-0.190 [.198]	-0.041 [.082]	-0.031 [.065]	-0.131 [.106]	0.007 [.086]	-0.334 [.154]**	-0.167 [.118]	-0.172 [.091]*
C. THERAPY + CASH								
ATE using survey data ($\hat{\theta}$)	-0.398 [.090]***	-0.105 [.027]***	-0.069 [.040]*	-0.099 [.026]***	-0.125 [.029]***	0.076 [.050]	0.012 [.033]	0.064 [.035]*
ATE on measurement error (θ^Δ)	0.118 [.182]	0.016 [.076]	-0.021 [.062]	0.099 [.093]	0.031 [.086]	0.314 [.116]***	0.137 [.101]	0.183 [.076]**
Adjusted ATE ($\hat{\theta} - \theta^\Delta$)	-0.516 [.196]***	-0.122 [.077]	-0.048 [.074]	-0.198 [.096]**	-0.156 [.083]*	-0.239 [.134]*	-0.125 [.112]	-0.120 [.078]
Observations, survey/validation	3765 / 239	3764 / 238	3762 / 238	3763 / 238	3765 / 239	3763 / 239	3761 / 238	3759 / 239

Notes: The ATE estimates based on survey data, $\hat{\theta}$, pool all survey rounds and regress each outcome on treatment indicators, controlling for survey round and randomization block fixed effects. Standard errors are robust and clustered by individual. Estimates of the bias from treatment, θ^Δ , come from a regression of the difference in the survey and validation measures on an indicator for treatment arms, controlling for survey round and randomization block fixed effects. Standard errors are robust and clustered by block. The difference, $\hat{\theta} - \theta^\Delta$, is an estimate of the true treatment effect after adjusting for observed bias. It is calculated as the linear difference of the estimates and the standard error is calculated via bootstrapping (we performed 1000 draws from the sample, with replacement; we calculated $\hat{\theta}$, θ^Δ , and $\hat{\theta} - \theta^\Delta$ for each draw; and we generated the standard error on $\hat{\theta} - \theta^\Delta$ using the distribution from these draws). *** p<0.01, ** p<0.05, * p<0.1

Table C.3: Estimates of treatment effects ($\hat{\theta}$) and treatment effect bias (θ^Δ) by outcome and treatment, no block controls

	Potentially sensitive behaviors				Expenditures			
	All (0-4) (1)	Steal (2)	Marijuana (3)	Gamble (4)	Homeless (5)	All (0-2) (6)	Video (7)	Phone (8)
A. CASH ONLY								
ATE using survey data ($\hat{\theta}$)	-0.045 [.099]	-0.028 [.030]	0.015 [.042]	0.027 [.030]	-0.059 [.032]*	0.077 [.052]	0.026 [.034]	0.050 [.037]
ATE on measurement error (θ^Δ)	-0.224 [.195]	-0.050 [.089]	-0.033 [.066]	-0.080 [.089]	-0.066 [.078]	0.274 [.138]**	0.111 [.115]	0.168 [.078]**
Adjusted ATE ($\hat{\theta} - \theta^\Delta$)	0.179 [.199]	0.022 [.085]	0.048 [.070]	0.107 [.091]	0.007 [.072]	-0.198 [.145]	-0.084 [.111]	-0.118 [.077]
B. THERAPY ONLY								
ATE using survey data ($\hat{\theta}$)	-0.180 [.095]*	-0.044 [.029]	-0.021 [.041]	-0.085 [.026]***	-0.030 [.032]	-0.001 [.050]	0.001 [.033]	-0.003 [.035]
ATE on measurement error (θ^Δ)	0.030 [.203]	-0.015 [.089]	0.017 [.059]	0.043 [.095]	-0.015 [.090]	0.333 [.133]**	0.159 [.105]	0.180 [.085]**
Adjusted ATE ($\hat{\theta} - \theta^\Delta$)	-0.209 [.196]	-0.028 [.083]	-0.037 [.066]	-0.127 [.099]	-0.016 [.086]	-0.334 [.145]**	-0.157 [.110]	-0.182 [.086]**
C. THERAPY + CASH								
ATE using survey data ($\hat{\theta}$)	-0.397 [.093]***	-0.105 [.027]***	-0.068 [.041]*	-0.097 [.027]***	-0.126 [.030]***	0.078 [.051]	0.013 [.033]	0.065 [.036]*
ATE on measurement error (θ^Δ)	0.083 [.187]	-0.002 [.083]	-0.031 [.064]	0.085 [.094]	0.030 [.089]	0.310 [.116]***	0.129 [.103]	0.185 [.076]**
Adjusted ATE ($\hat{\theta} - \theta^\Delta$)	-0.480 [.195]**	-0.103 [.079]	-0.037 [.073]	-0.182 [.096]*	-0.157 [.083]*	-0.232 [.123]*	-0.116 [.109]	-0.120 [.074]
Observations, survey/validation	3765 / 239	3764 / 238	3762 / 238	3763 / 238	3765 / 239	3763 / 239	3761 / 238	3759 / 239

Notes: The ATE estimates based on survey data, $\hat{\theta}$, pool all survey rounds and regress each outcome on treatment indicators, without controlling for baseline characteristics, survey round, or randomization block fixed effects. Standard errors are robust and clustered by individual. Estimates of the bias from treatment, θ^Δ , come from a regression of the difference in the survey and validation measures on an indicator for treatment arms, without controlling for baseline characteristics, survey round, or randomization block fixed effects. Standard errors are robust and clustered by block. The difference, $\hat{\theta} - \theta^\Delta$, is an estimate of the true treatment effect after adjusting for observed bias. It is calculated as the linear difference of the estimates and the standard error is calculated via bootstrapping (we performed 1000 draws from the sample, with replacement; we calculated $\hat{\theta}$, θ^Δ , and $\hat{\theta} - \theta^\Delta$ for each draw; and we generated the standard error on $\hat{\theta} - \theta^\Delta$ using the distribution from these draws). *** p<0.01, ** p<0.05, * p<0.1

Table C.4: Estimates of treatment effects ($\hat{\theta}$) and treatment effect bias (θ^Δ) by outcome and treatment, block and baseline controls

	Potentially sensitive behaviors					Expenditures		
	All (0-4) (1)	Steal (2)	Marijuana (3)	Gamble (4)	Homeless (5)	All (0-2) (6)	Video (7)	Phone (8)
A. CASH ONLY								
ATE using survey data ($\hat{\theta}$)	-0.078 [.071]	-0.033 [.024]	-0.007 [.029]	0.020 [.026]	-0.056 [.026]**	0.081 [.045]*	0.021 [.032]	0.061 [.032]*
ATE on measurement error (θ^Δ)	-0.313 [.177]*	-0.116 [.086]	-0.102 [.063]	-0.063 [.080]	-0.025 [.087]	0.147 [.136]	0.002 [.102]	0.150 [.082]*
Adjusted ATE ($\hat{\theta} - \theta^\Delta$)	0.234 [.268]	0.083 [.126]	0.095 [.100]	0.083 [.134]	-0.031 [.114]	-0.066 [.195]	0.019 [.153]	-0.090 [.120]
B. THERAPY ONLY								
ATE using survey data ($\hat{\theta}$)	-0.193 [.071]***	-0.040 [.024]*	-0.046 [.029]	-0.085 [.024]***	-0.023 [.026]	-0.009 [.046]	-0.005 [.032]	-0.004 [.031]
ATE on measurement error (θ^Δ)	-0.215 [.180]	-0.065 [.084]	-0.056 [.064]	-0.013 [.091]	-0.070 [.090]	0.268 [.144]*	0.109 [.107]	0.170 [.083]**
Adjusted ATE ($\hat{\theta} - \theta^\Delta$)	0.022 [.264]	0.025 [.119]	0.010 [.104]	-0.072 [.139]	0.047 [.121]	-0.277 [.213]	-0.114 [.163]	-0.174 [.116]
C. THERAPY + CASH								
ATE using survey data ($\hat{\theta}$)	-0.402 [.069]***	-0.105 [.023]***	-0.089 [.028]***	-0.097 [.024]***	-0.113 [.026]***	0.073 [.045]	0.003 [.032]	0.069 [.031]**
ATE on measurement error (θ^Δ)	0.052 [.183]	-0.030 [.083]	-0.042 [.061]	0.056 [.097]	0.088 [.084]	0.267 [.122]**	0.059 [.102]	0.224 [.074]***
Adjusted ATE ($\hat{\theta} - \theta^\Delta$)	-0.454 [.270]*	-0.075 [.117]	-0.047 [.095]	-0.153 [.140]	-0.201 [.121]*	-0.195 [.197]	-0.056 [.157]	-0.155 [.110]
Observations, survey/validation	3765 / 239	3764 / 238	3762 / 238	3763 / 238	3765 / 239	3763 / 239	3761 / 238	3759 / 239

Notes: The ATE estimates based on survey data, $\hat{\theta}$, pool all survey rounds and regress each outcome on treatment indicators and block fixed effects. Standard errors are robust and clustered by individual. Estimates of the bias from treatment, θ^Δ , come from a regression of the difference in the survey and validation measures on an indicator for treatment arms, controlling for block fixed effects and each endline round. Standard errors are robust and clustered by block. The difference, $\theta - \theta^\Delta$, is an estimate of the true treatment effect after adjusting for observed bias. It is calculated as the linear difference of the estimates and the standard error is calculated via bootstrapping (we performed 1000 draws from the sample, with replacement; we calculated $\hat{\theta}$, θ^Δ , and $\theta - \theta^\Delta$ for each draw; and we generated the standard error on $\theta - \theta^\Delta$ using the distribution from these draws). *** p<0.01, ** p<0.05, * p<0.1