

Learning without Teachers? Evidence from a Randomized Experiment of a Mobile Phone-Based Adult Education Program in Los Angeles

Christopher Ksoll, Jenny C. Aker, Danielle Miller, Karla Perez, Susan L. Smalley*

July 2014

Abstract: Over 755 million adults worldwide are unable to read and write in any language. Yet the widespread introduction of information and communication technology offers new opportunities to provide standardized distance education to underserved illiterate populations in both developed and developing countries. Using data from a randomized experiment of an innovative mobile phone-based adult education program (*Cell-Ed*) in Los Angeles, we find that the *Cell-Ed* program significantly increased students' basic and broad reading scores, equivalent to a 2-4 year increase in reading levels over a four-month period. The program also increased participants' self-esteem by 7 percent as compared with the comparison group. These results are robust to correcting for non-random attrition using a variety of non-parametric methods, including using the phase-in design to tighten the Lee bounds. Our results suggest that there is great scope for using information technology as a means of improving educational skills for illiterate adults.

JEL codes: D1, I2, O1, O3

* Christopher Ksoll, School of International Development and Global Studies, University of Ottawa, 120 University, Ottawa, ON, Canada; christopher.ksoll@uottawa.ca. Jenny C. Aker, Department of Economics and The Fletcher School, Tufts University, 160 Packard Avenue, Medford, MA 02155; Jenny.Aker@tufts.edu. Danielle Miller, Dornsife College of Letters, Arts and Sciences, University of Southern California and Cell-Ed. Karla C. Perez-Mendoza, Graduate School of Education and Information Studies (GSE&IS), University of California Los Angeles; kcperez77@ucla.edu. Susan L. Smalley, Department of Psychiatry and Behavioral Science, UCLA and Cell-Ed; sue@suesmalley.com. We thank Michael Clemens, Jose Galdo, Louis-Philippe Morin, Concepcion Valadez, Bruce Wydick, as well as seminar participants at the University of San Francisco and Tufts University for helpful comments. We thank Ariadna Cornei, Stephanie Del Mel and Wei Yang for excellent research assistance. Ksoll acknowledges funding from a Canadian Social Sciences and Humanities Research Council (SSHRC) Insight Grant.

Education is positively correlated with many dimensions of well-being, both at the individual and aggregate levels. Yet educational achievements are still remarkably low in many parts of the world: It is estimated that there are over 755 million adults worldwide who are unable to read and write in any language (UNESCO 2012). This is not merely an issue that is confined to developing countries: In the US, approximately 14 percent of adults are considered to be functionally illiterate (National Center for Education Statistics 2003).¹

Adult education programs have the potential to bridge this gap, but they are often characterized by low enrollment and high drop-out rates, potentially because the opportunity costs associated with such programs are too high for adults with families and busy work schedules. Moreover, these opportunity costs might not yield sustainable returns, as evidence suggests that adult learners' skills often depreciate after completing the program. The failure for literacy gains to persist may be due to the irrelevancy of such skills in daily life or limited opportunities to practice such skills in the learner's native language (Aker, Ksoll and Lybbert 2012).

More than four billion individuals have access to mobile phones, with developing countries representing some of the fastest-growing markets (ITU 2013). The widespread growth of mobile phone coverage has the potential to facilitate skills acquisition of illiterate adults, as well as increase the scope and scale of such programs. While a number of technology-based adult education programs have been developed, most require Internet access or smart phones, which are often not easily available to poorer populations in both developed and developing countries. Furthermore, educational programs delivered

¹ A person is considered illiterate if they are unable to read and understand basic prose in any language) A person is considered functionally literate if they can “engage in all those activities in which literacy is required for effective functioning of his group and community and also for enabling him to continue to use reading, writing and calculation for his own and the community's development.” (UNESCO 2005)

via information technology are often complements to teachers, rather than substitutes. As such, they are heavily dependent upon teacher availability and quality, a major constraint in many contexts.

We report the results of a randomized adult education program in Los Angeles, where a mobile phone-based adult education program in Spanish (*Cell-Ed*) was offered to illiterate adults. Unlike many other technology-enhanced education programs, the *Cell-Ed* learning curriculum was completely provided via a series of voice and SMS-based operations on the mobile phone, and therefore did not require teacher instruction or in-situ learning. Seventy Spanish-speaking adult students were randomly assigned to the treatment (*Cell-Ed*) or control group, with the control group phased into the program after a three-month period.² We find that the program substantially improved learning outcomes: The increase in *Cell-Ed* participants' reading skills after four months was equivalent to the reading skills children acquire after 2–4 years of schooling. We posit that these increases may be due in part to the flexibility of the curriculum, as learners opted to learn at all times of the day and for short durations, in stark contrast to the fixed schedules of many adult education programs.

A key claim of many adult education programs is that they empower students, primarily by providing them with the knowledge and skills necessary to acquire new labor market opportunities and live more independent lives. We investigate whether the *Cell-Ed* program affected empowerment, as measured by the Rosenberg self-esteem score (RSES) and the General Self-efficacy Scale (GSES). Our results provide more speculative evidence that the *Cell-Ed* program increased participants' self-esteem scores. Using real-time data on learning progress on the platform and weekly measures of self-esteem, we also provide insights into the dynamic relationship between learning progress (or lack thereof) and self-esteem.

²The minimum time spanned by different episodes of access to the platform is three weeks.

A potential threat to the internal validity of our findings is differential attrition between the treatment and control groups. To deal with non-random attrition, we implement two non-parametric approaches (Manski 1990, Lee 2009). Our first approach employs Lee's methodology, exploiting the randomized phase-in design of the program to further tighten the traditional Lee bounds by identifying the control group respondents and non-respondents after they have been phased into the program.³ As a second approach, we combine the monotone treatment response (MTR) and monotone treatment selection (MTS) assumptions to bound the treatment effect (Manski 1997, Manski and Pepper 2000). Using these corrections, we are able to conclude that the *Cell-Ed* program is highly effective in improving reading levels, even when imposing minimal assumptions. Our paper thus proposes an alternative set of assumptions for bounding the treatment effect and suggests that phase-in designs may confer important identification benefits.

Our results contribute to the growing debate on the effectiveness of ICT-assisted learning in other contexts. While Barrera-Osario and Linden (2009) and Fairlie and Robinson (2013) find that computers have either no or mixed effects on learning outcomes, Banerjee et al (2007) found that computers increased students' math scores and were equally effective for all students. Barrow, Markman and Rouse (2009) find that students randomly assigned to a computer-assisted program obtained significantly higher math scores, primarily due to more individualized instruction. Most relevant to our study, Aker, Ksoll, and Lybbert (2012) find that learning how to use mobile phones in adult education classes in Niger increased test scores by .19-.25 s.d. Our paper provides additional evidence that a simple, self-directed educational device can lead to substantial improvements in learning, on the order of magnitude of 1.3-3.2. s.d.

³DiNardo et al (2006) show an alternative approach to tightening Lee bounds. This entails creating two groups with different survey attrition rates by randomizing the levels of effort exerted in recontacting survey respondents.

Our paper also speaks to the potential relationship between education and empowerment. The UNFPA, for example, states that "Education is one of the most important means of empowering women with the knowledge, skills and self-confidence necessary to participate fully in the development process", a sentiment that has often been repeated by donors, policymakers and practitioners alike. Yet what is surprising about these claims is the relative paucity of evidence to support them. There is significant literature measuring the impact of education programs on educational and labor markets outcomes, as well as health and fertility (see, for example, Case 2005). A separate strand of literature assesses the relationship between empowerment and economic development (see, for example, Duflo 2012, Doepke and Tertilt 2014) or the impact of different interventions on intra-household decision-making and empowerment (Ashraf 2009, Ashraf et al 2010). Yet there are few studies that explicitly link education and empowerment. Our paper contributes to this literature by measuring the impact of an educational program on self-esteem, as well as the dynamic relationship between learning and self-esteem over time.

The remainder of the paper is organized as follows. Section I provides background on the setting and the research design. Section II describes the different datasets and Section III outlines the estimation strategy. Section IV discusses the results, whereas Section V looks at the potential learning mechanisms. Section VI concludes.

I. Research Setting and Design

While the United States is one of the wealthiest countries in the world, it is estimated that 1 in 7 U.S. adults are functionally illiterate, defined as being unable to read and understand basic prose (National Center for Education Statistics 2003). The Hispanic population accounts for a disproportionate share of the functionally illiterate, with approximately 44 percent of Hispanics having

literacy levels considered to be at or “below basic” level. Hispanic populations also represent 63% of the U.S’ “least literate” adults, defined as those who fail the simplest screener questions.⁴

Greater Los Angeles has the highest rate of so-called “undereducated adults” in any major U.S. metropolitan area (Literacy Network of Greater LA and United Way 2004).⁵ Similar to national figures, a disproportionate number of the low-literate and illiterate populations are immigrants, primarily Spanish-speaking immigrants who are unable to read or write in Spanish or English.

In an effort to address this issue, the U.S. Adult Education and Family Literacy Act provides approximately US\$ 600 million per year in funding for states to implement adult education programs. A majority of these adult education programs are taught using teacher-based personal instruction, often in “adult schools” or community centers. Despite these efforts, adult education programs reach fewer than 16 percent of the low-literacy population in Los Angeles, and drop-out rates are well over 50 percent (Literacy Network of Greater LA and United Way 2004).

A. Cell-Ed Intervention

The widespread growth of mobile phone coverage has the potential to facilitate skills acquisition of illiterate adults, as well as increase the scope and scale of such programs. While a number of technology-based adult education programs have been developed, most require Internet access or smart

⁴ The National Assessment of Adult Literacy is a nationwide representative survey on literacy in the US. Survey respondents representing approximately 4 million people (around 1.5 percent) were unable to communicate in English or Spanish, so that the assessments could not be administered. The assessment first asks respondents 7 very simple literacy questions. If respondents fail this initial screening (about 3% representing about seven million people did), they participate in a survey for the “least literate adults”. 63% of the respondents in this category are Hispanics. Of the three percent, only 57% could read letters, 46% could read words.

⁵ Literacy Network of Greater LA and United Way (2004) define literacy and numeracy levels slightly differently from the 2003 National Assessment of Adult Literacy (National Center for Education Statistics 2003). In their definition, individuals with a literacy “Level 1” are unable to engage in basic daily tasks, such as find an intersection on a map. Approximately 32 percent of the population in Greater LA (or 2.3 million people) fall into this category, as compared with 20 percent at the national level.

phones, which are often not available among poorer populations in Los Angeles. *Cell-Ed* is a platform that provides basic educational instruction via simple mobile phones. The platform uses voice (audio) and SMS messages to deliver 437 adult education lessons (called “micro-modules”) to learners. Each micro-module consists of three components: 1) audio instruction: an audio lesson on a particular concept (vowels, consonants, words), and varying from 1-3 minutes in length, is introduced when the learner calls a designated number; 2) written instruction: a SMS message reinforcing the voice lesson is sent to the participant; and 3) interactive quiz: a SMS question is sent to the participant asking them about the lesson that they recently learned, and the participant must text a response. A correct response to the question triggers the beginning of the next micro-module, whereas an incorrect response leads to a repetition of the same micro-module until the user succeeds. To activate the program and each micro-module, participants call the *Cell-Ed* phone number from their own mobile phone.⁶ Students could access the program 24 hours a day, seven days a week, allowing them to learn when, where and how they wished.

The *Cell-Ed* curriculum for this study was based upon a traditional Spanish adult education program (LEAMOS!) developed by the Centro Latino in Los Angeles. The traditional LEAMOS! curriculum is comprised of 43 lessons and teaches students simple Spanish letter and word recognition, as well as reading and writing skills. The curriculum is typically taught by an in-classroom teacher and takes approximately 150 hours to complete over a four-month period, or 9 hours per week. The 43 LEAMOS! lessons were adapted into 437 micro-modules for the *Cell-Ed* platform, with each micro-module including recorded audio instructions, SMS messages and queries for interactive testing. In

⁶ If Cell-Ed participants did not own a mobile phone, they were provided with a simple phone as part of the program. All participants (regardless of mobile phone ownership) were also provided with an unlimited voice and SMS plan if they did not already have one.

addition to the micro-modules, pre-recorded audio messages were sent to each learner to offer encouragement at various points in the process.

B. Experimental Design

Prior to the introduction of the program, *Cell-Ed* partnered with two schools and five community resource centers with large Hispanic populations in the Los Angeles area. A variety of recruitment methods were used to recruit potential participants, including informational flyers, door-to-door visits by parent volunteers and presentations at community meetings, school fairs and school events. Using these recruitment methods, we identified 250 individuals and conducted an initial screening process via phone. The screening process collected information on individuals' socio-demographic characteristics, reading ability, mobile phone access and eyesight. Participants were excluded from further participation if they were over the age of 80 years, could read sentences, or required eyeglasses and did not own them, thereby reducing the sample size to 124 participants.⁷

The initial 124 participants were scheduled for an additional in-person baseline evaluation, of which only 89 attended (71 percent). Those who accepted to participate in the baseline evaluation were provided an ID number and asked to complete a battery of assessments, including the Woodcock-Muñoz Achievement Battery.⁸ Participants who scored sufficiently low on the reading test (with a score lower than “basic” in two of three reading sub-tests) and did not suffer from neurological disorders were considered eligible for participation in the study. These exclusions further reduced the eligible sample to

⁷ This initial screening process implies that our reduced sample was less literate than the initial sample of 250 individuals, but had access to eyeglasses, if necessary.

⁸ For each survey round, participants were provided with compensation in the form of a US\$50 gift card to a local supermarket chain. In addition, there was a US\$100 incentive for completing the program in 4.5 months. Thus, a treatment participant that completed all survey rounds and finished the program in 4.5 months would receive US\$ 200, whereas a control participant who completed all survey rounds and finish the program in 4.5 months (after being phased in) would receive US\$ 250.

70 participants. After finishing the screening process, participants were informed of their treatment status via a previously sealed letter (linked to their initial ID) that indicated whether they had been assigned to the treatment or control condition. This letter was read to them.⁹ The ID numbers had been divided into strata and randomly assigned to either the treatment (the *Cell-Ed* program) or the control group (to be phased into the *Cell-Ed* program approximately 3-4 months later) with equal probability.¹⁰ The random assignment took place before the evaluation, and enumerators were not informed of participants' treatment or control status.

The treatment thereby followed a randomized phase-in design. After the initial baseline evaluation, those in the treatment group were asked to return for a second evaluation approximately 4.5 months later, in order to insure that they had sufficient time to complete the *Cell-Ed* program.¹¹ The control group was eligible to start the *Cell-Ed* program approximately three months after the baseline, after having completed a midline survey. A timeline of the implementation and data collection activities is provided in Figure 1.

Overall, the screening process outlined above suggests that our sample is less educated than the broader population of Spanish-speaking immigrants in Los Angeles, as they did not perform well on initial reading tests. In light of the low levels of literacy among our sample, participants reported that

⁹As the control group was informed that they would receive the treatment at a later time, this could have changed their behavior between the first and second survey rounds (the John Henry effect). If the control group exerted less effort in learning than they would have otherwise, our results could potentially provide an upper bound for the treatment effect. If, however, this letter encouraged the control group participants to exert greater effort in learning during the first round, then our results would provide a lower bound on the treatment effect. As reading skills are difficult to acquire without some type of external support, we do not think that the John Henry effect is a primary concern in this context.

¹⁰In order to ensure that balance between treatment and controls were achieved for earlier and later potential participants, we stratified the random assignment by pre-screening ID, in the following strata: 1-22,23-46,47-70,70-100,100+. Within strata the ex ante probability of selection into treatment was 50 percent. We control for strata fixed effects in all OLS as the ex post probability differed due to ineligible applicants.

¹¹Early piloting of the program suggested that most participants could complete the program in a 4.5 month time period. This time period was also similar to the time frame for most Leamos! learners.

they relied on families and friends to complete tasks that required reading and writing or simply avoided such tasks altogether. Most participants were employed in the informal economy, and stated that they were employed in jobs in which literacy was not essential.

II. Data

This paper relies upon four primary datasets. First, we administered comprehensive reading tests to measure the impact of the *Cell-Ed* on learning outcomes. Second, we collected data on student and household characteristics, as well as some qualitative data on their experiences as in Los Angeles. Third, we conducted short weekly phone calls to both treatment and control participants. Fourth, we collected real-time *Cell-Ed* usage data, providing user-specific information on how much time was spent on the program, as well as performance statistics (i.e., whether the student failed or passed a specific module).

A. Test Score Data

Reading tests were administered to all participants during the baseline, providing a baseline sample of over 89 students (of which 70 were retained for the survey, as outlined above). We administered a second round of reading tests for both groups approximately 3-4.5 months after the baseline, with a third round of follow-up tests for the control group only approximately 4.5 months after they started the *Cell-Ed* program. The second set of test scores allows us to measure the impact of the *Cell-Ed* program on learning, as the control group had not yet started the program. The third set of test scores for the control group allows us to assess whether the impacts were similar in the control group as they were in the initial treatment group.¹²

¹²For treatment individuals who completed the program in less than 4.5 months, the first follow-up survey was done as soon as possible after completion of the program. The main purpose of the control group was to ascertain how outcomes change over time in the absence of treatment. Typically, follow-up surveys would be implemented at the same time for both the

We administered the Spanish language equivalent of the widely used Woodcock-Johnson battery of literacy tests, the *Woodcock-Muñoz III Language Survey (WMLS-III)*. The typical survey includes a screener (three tests) and a more-comprehensive seven-test battery, although our survey did not include the former.¹³ We administered the cognitive battery of tests, which provides a composite verbal IQ score, as well as the reading achievement battery of tests, which were used to calculate two composite scores: the *basic reading score*, which is comprised of two sub-tests (letter-word identification and word attack); and the *broad reading score*, which is comprised of three sub-tests (letter-word identification, reading fluency and passage comprehension). The basic reading score covers literacy skills as narrowly defined, namely reading decoding and phonetic coding (identifying letters and building them up to words) (McGrew et al. 2007). The broad reading score can be interpreted as a measure of more advanced reading skills, as it also contains dimensions such as fluency and comprehension. Each composite score is calculated as an “age equivalent”, which indicates the typical age of persons in the population who obtain a given score. For example, if a student’s performance on the test of reading comprehension is equal to an age equivalent of 8.5 years, this means that his or her obtained *raw score* is equivalent to the predicted average for 8-year, 6-month old children in the norm group. We use the age equivalent throughout our paper since it has an intuitive interpretation.¹⁴ As Spanish is a phonetic language, where each letter maps into one sound, becoming literate may be simpler than in English, where the mapping between letters and sounds is not unique (Abadzi, 2013).

treatment and control groups. In our case, the change over the course of 3 months in reading ability in the absence of treatment is zero, implying that this would be the case 4.5 months later as well, unless the comparison group learned how to read otherwise. Testing the treatment group when they finished the program (with a maximum date of 4.5 months) alters the interpretation of the estimate: It is not the impact on learning after 4.5 months, but the impact of *Cell-Ed* at the individually defined end of the program, censored after 4.5 months. Some learners continued after the 4.5 months, and this learning is not captured in our tests.

¹³ The typical seven tests involve picture vocabulary, verbal analogies, letter-word identification, dictation, understanding directions, story recall and passage comprehension. Testing times range from approx. 25 minutes (screener) to 55 minutes (complete battery).

¹⁴ We also provide results with z-scores based on the norming population of the tests.

B. Socio-Demographic Characteristics and Measures of Empowerment

The second primary dataset includes information on student and household characteristics. Among all eligible participants, we conducted a baseline survey, as well as a follow-up survey 3-4.5 months later. For the control group only, we conducted a third survey 4.5 months after their midterm survey. Each survey collected detailed information on students' demographics (age, gender, birthplace), education, employment, household size, mobile phone ownership and usage. We use these data primarily for balance checks and additional controls.

While collecting data on socio-economic characteristics, we also included separate survey questions on two measures of psychological dimensions of wellbeing: self-esteem and self-efficacy, as measured by the Rosenberg self-esteem scale and the general self-efficacy score. The Rosenberg self-esteem scale (RSES) is a series of statements, designed to capture different aspects of self-esteem (Rosenberg 1965). Five of the statements are positively worded, while the other five statements are negatively-worded. Each answer is assigned a point value, with higher scores reflecting higher self-esteem (the maximum is 30). The General Self-Efficacy Scale (GSES) (Schwarzer and Jerusalem 1995) is also a 10-item psychometric scale that is designed to assess whether the respondent believes he or she is capable of performing new or difficult tasks and to deal with adversity in life. The scale ranges in value from 12-60, with higher scores reflecting higher perceived self-efficacy. We use these results to measure the impact of the *Cell-Ed* program on participants' perceptions of empowerment.

C. Weekly phone calls

In addition to the in-person surveys, we contacted all participants via weekly phone calls. Treatment participants were asked whether anyone assisted them with *Cell-Ed* that week, whether they had technical difficulties, what they liked and disliked about the curriculum and whether there were any impediments to their studies. They were also asked about mobile phone usage and self-esteem. As the

phone call might be intervention in itself, phone calls were also made to the control group. Control participants were asked the same questions on mobile phone usage and self-esteem, but not about the *Cell-Ed* program.

D. Cell-Ed Usage Data

The final dataset is comprised of records from the *Cell-Ed* platform, which logs every interaction between the student and the platform. This log includes the date and time the participant called in, the mini-module they accessed, the timing and content of SMS messages received and sent and their performance on the test question for that micro-module. Using these data, we calculate a number of statistics, including the number of days between different modules; how the student performed on each module (including the number of attempts to obtain a correct response); how quickly each student completed the program; and whether the student stopped calling (and hence dropped out of) the *Cell-Ed* program. We use these data to provide insights into the mechanisms through which the *Cell-Ed* program affected learning outcomes.

E. Pre-Program Balance

Table 1 provides an overview of student characteristics by treatment status. The evidence suggests that the randomization was successful in creating comparable groups along observable dimensions. Differences in students' characteristics between the treatment and control group before the program started are small and not statistically significant (Table 1, Panel A). Participants were 47 years old, and a majority of participants were women. While 51% percent of participants had some schooling, the average duration of schooling was quite low: 1.22 years. 51 percent of households in the sample owned a mobile phone, broadly in line with other studies on mobile phone ownership among Hispanic

populations with lower levels of education.¹⁵ Among all of these variables, only the difference in mobile phone ownership is statistically significant at the 10 percent level, with a higher mean in the treatment group as compared with the control group. As mobile phone ownership could be positively correlated with the treatment and our outcomes of interest, we control for baseline levels of mobile phone ownership as a robustness check.

Panel B shows the means of key outcome variables (broad and basic reading age equivalent, the Rosenberg score and the self-efficacy score) by treatment status. During the baseline, participants had mean basic and broad reading scores of 6, suggesting that their reading levels were equivalent to a six year old. There was no statistically significant difference between the two groups. While the results of the Generalized Self-Efficacy score were similar across treatment and control groups, participants in the treatment group had slightly higher baseline levels of self-esteem, with a statistically significant difference at the 10 percent level. As this measure is potentially correlated with treatment and the outcomes of interest, we control for baseline levels of self-esteem in the robustness checks as well.

A key concern with the balance tests is that the absence of any statistically significant differences may be due to our limited number of observations, hence increasing the size of our standard errors. While this might be a potential concern, we note that most of the differences in means (with the exception of years of schooling) represent less than 20 percent of the mean of the control group, suggesting that the magnitude of the differences is not large between the two groups for most of the baseline characteristics.

¹⁵While data on mobile phone ownership among Spanish-speaking immigrants in Los Angeles is not readily available, a survey by the Pew Trust estimated that mobile phone ownership among Hispanic households with lower income and educational levels ranged from 56-77%. <http://www.pewhispanic.org/2013/03/07/closing-the-digital-divide-latinos-and-technology-adoption/>.

F. Non-Response and Drop-outs

Table 2A shows the rates of survey non-response by treatment status and across all survey rounds. While definitions of survey non-response (attrition) and drop-out can differ, in our case, dropping out of the program at some point was correlated with eventual survey non-response, and so we treat them as interchangeable (and herein refer to them as non-respondents). Drop-outs are comprised of: 1) treated participants who could not make it past the first ten lessons;¹⁶ 2) treated participants who made it past the first ten lessons but refused a follow-up interview; or 3) control participants who refused a follow-up interview, either because they refused a follow-up interview or dropped out after the first ten lessons (after being phased into the program).

Overall, survey non-response was 17 percent during the first round of the program. Yet there are significant differences by treatment status: Twenty-five percent of the treatment group did not participate in the follow-up survey as compared with 8.8 percent in the control group, with a statistically significant difference between the two. The rate of drop-out was higher during the second period (once the control group started the program), where average non-response was 32 percent. While these rates of attrition are high, we note that they are lower than what is typically observed in most adult education programs. The Literacy Network of Greater LA and the United Way (2004) estimate that 50 percent adult education students drop out within the first three weeks of such programs, with similar rates reported by Romain and Armstrong (1987).¹⁷

Table 2B presents more detailed information on learning and attrition. All students spent some time on the *Cell-Ed* platform, although there are marked differences in the duration of time spent by

¹⁶Follow-up data were not collected from this group of non-respondents, and hence that data cannot be included in the analysis.

¹⁷ Literacy agencies covering almost 600,000 learners responded to their survey, of which only 55 percent tracked information on learner retention.

respondents (ie, those who completed the surveys) and those who eventually dropped out (non-respondents). Focusing on the treatment group (Panel A), respondents interacted with the platform over the course of 123 days (four months), as compared with 92 days (three months) for those in the treatment group who eventually dropped out (non-respondents). Respondents in the treatment group spent approximately 50 hours (3144 minutes) on the platform, with a minimum of 8 hours and a maximum of 148 hours. By contrast, treatment non-respondents spent only 8 hours on the platform, with a minimum of 2 hours and a maximum of 30 hours. Treatment respondents completed 402 modules (out of 437), about 90 percent of the curriculum, whereas treatment non-respondents completed only 31 micro-modules (7 percent of the curriculum). Once the control group phased into the *Cell-Ed* program, the difference in usage patterns between respondents and non-respondents follows a similar pattern (Panel B). Unsurprisingly, these results suggest that there are substantial differences in the *Cell-Ed* usage patterns among respondents and non-respondents, the latter of whom are primarily comprised of people who dropped out of the program.

The difference in means for the baseline characteristics of respondents and non-respondents are presented in the first three columns of Table A1, with the estimates from a logistic regression of non-response in Table A2.¹⁸ Overall, non-respondents were largely similar along observable dimensions as compared with respondents, although part of this could be explained by our limited power to detect a statistically significant difference.¹⁹ Non-respondents were more likely to have attended formal school and scored slightly higher on the baseline self-efficacy score (Table A1). These correlations are confirmed when looking at the determinants of drop-out: Attriters were more likely to have attended

¹⁸Columns 3-5 of Table A1 show the differences in baseline characteristics between treatment and control respondents during the initial phase of the program, before the control group was phased into the program.

¹⁹ Looking at a broader range of covariates, initial treatment non-respondents differed at the 10% level from respondents on 8 variables out of 68. There were no statistically significant differences at the 5% level.

school and have a higher self-efficacy score than non-attriters (Table A2). In addition to these covariates, Table A2 also suggests that older participants are more likely to attrit and males were less likely to drop out. These comparisons confirm that non-differential attrition is a threat to the validity of our results, a topic that we address in more detail below.²⁰

III. Estimation Strategy

All of the participants who were initially assigned to the *Cell-Ed* platform were provided access to the platform. The minimum time between the first and last contact with the platform among all treatment observations (non-respondents included) was 42 days.²¹ As a result, there was no imperfect compliance: All of those *assigned* to the initial treatment participated in the platform (even if they dropped out at a later time), and none of the initial control group participants accessed the *Cell-Ed* platform during the initial phase. As a result, our treatment assignment variable is equivalent to our treatment participation variable, and our intention to treat (ITT) effect is equivalent to the average treatment effect on the treated (ATT).

Let $test_i$ be the composite basic or broad reading test score attained by student i after the program. $Cell-Ed_i$ is an indicator variable for whether individual i is assigned to the Cell-Ed intervention ($Cell-Ed=1$) or the control ($Cell-Ed=0$) in the first period. θ_R are fixed effects that indicate the randomization strata, which were based on the ID number. \mathbf{X}'_i is a vector of student-level baseline covariates, primarily gender, age and a proxy for IQ. Ignoring attrition, we first estimate the following specification:

²⁰An additional potential threat to the validity of our findings is spillovers. We cannot rule out the possibility that some treatment and control group participants knew each other before the program. In this case, however, our treatment effect estimates would represent a lower bound on the treatment effect in the absence of spillovers.

²¹ When the control group was phased in to receive treatment, the minimum number of days between first and last interaction was 21 days.

$$(1) \quad test_i = \beta_0 + \beta_1 CellEd_i + \mathbf{X}'_{i0} + \theta_R + \varepsilon_i$$

The coefficient of interest is β_1 , which captures impact of the *Cell-Ed* program on learning outcomes. We also modify this specification to include the baseline outcome variable (a value-added specification), covariates that were statistically significant at baseline (mobile phone ownership and self-esteem), as well as a difference-in-differences (DD) estimation strategy.

While randomization would normally imply that β_1 has a causal interpretation, this is only the case if we assume that non-respondents would have progressed in their learning as much as those who continued with the program.²² As noted in the methodological appendix, an unbiased estimate of the ATT can be estimated under the assumption that non-response is random – or observations are missing at random - which is not a reasonable assumption in our context. Nevertheless, we first proceed under this assumption before using various techniques to deal with non-random attrition.

IV. Results

Figure 2A depicts the mean reading scores for treatment and control respondents for the first round (prior to the phase in of the control group) before and after the program. Before the program, both groups had basic reading scores corresponding to an age equivalent of 6.5 years, without a statistically significant difference between the two (Table A3). This suggests that participants had a first to second grade reading level before the program. After the program, *Cell-Ed* participants had basic reading skills of 11.5 years, suggesting that they moved from a first grade to a sixth grade reading level, with a statistically significant difference at the 1 percent level. For the more comprehensive measure of

²²An alternative assumption would be if attrition were equally high in treatment and control groups and we assume that treatment did not affect dropout for any survey participant. Then the estimated parameter would have the interpretation of the average treatment effect on the treated. In our survey there is a statistical difference in attrition, so this alternative identification assumption is untenable.

reading scores, broad reading scores, *Cell-Ed* participants moved from a first-grade to a fourth-grade broad reading level, whereas the comparison group showed no improvements in reading during this time. Thus, the *Cell-Ed* participants had broad reading scores that were 2-3 years higher than those in the control group, with a statistically significant difference at the 1 percent level. These results are similar for the control group once they phased into the program (Table A3), suggesting that the treatment effect is stable over time.²³ The distributional effects of the program are shown in Figures 2B-2C, suggesting that the treatment group's reading score distribution shifted sharply to the right.

A. Impact of Cell-Ed Program on Learning Outcomes

Table 3 presents the results of equation (1) for survey respondents. Controlling only for the *Cell-Ed* program and the randomization strata, the *Cell-Ed* program increased students' basic reading scores by 5.65 years, with a statistically significant effect at the 1 percent level (Column 1). These effects are robust to the inclusion of baseline covariates, namely age, gender and IQ (Column 2), a value-added specification (Column 3), a difference-in-differences specification (Column 4) and controlling for individual fixed effects (Column 5). The effects are also robust to the inclusion of the covariates that were different at baseline, as well as using normalized test scores as the dependent variable (Table A4). Overall, these results suggest that participating in the *Cell-Ed* program increased students' basic reading scores by 4.76 - 5.84 years within a four-month period, with a statistically significant effect at the 1 percent level.

²³As there are only small baseline differences between the initial treatment and control group, and no significant changes in the control group's learning outcomes between the baseline and follow-up survey, any difference in impacts between these two groups would point to technical difficulties with the platform, spillover effects or learning on the part of the implementers. Given that the program is highly standardized through the platform, the absence of a difference is not too surprising.

The estimates are smaller in magnitude for the broad reading scores, which capture a broader set of reading ability including fluency and comprehension: the *Cell-Ed* program increased broad reading scores by 2.58 years when controlling only for the treatment variable and strata fixed effects, with a statistically significant effect at the 1 percent level (Column 6). These effects are similar when controlling for baseline covariates (Column 7), a value-added specification (Column 8), a DD specification (Column 9) and controlling for individual fixed effects (Column 10). The results are also robust to the inclusion of the baseline mobile ownership and self-esteem scores and using normalized test scores as the dependent variable (Table A4). The smaller impacts of the program on broad reading scores relative to basic reading scores is not surprising, as the *Cell-Ed* program focused on basic reading skills, rather than fluency and reading comprehension.²⁴

Overall, the results in Table 3 suggest that *Cell-Ed* significantly increased participants' reading outcomes over the four-month period, with an increase that is comparable to a child's increase in reading levels over 2.5 to 5 years. Comparing our effects with those of other rigorous evaluations of adult education outcomes is difficult, as such evaluations are relatively rare and often use different measures of learning. Nevertheless, a large scale evaluation of an adult education program in India finds that the program increased learning outcomes by .06 – 0.15 s.d. (Banerji et al 2013), whereas Aker, Ksoll and Lybbert (2012) find that a mobile phone-enhanced adult education program increased writing and math test scores by .20-.25 s.d. as compared with a traditional adult education program. When using normalized test scores (Table A4), our effects suggest a treatment effect of 1.3-3.9 s.d. for broad and

²⁴Table A3 presents evidence that the initial control group benefits as much from the treatment when they are phased-in. Column 4 tests for difference in outcomes between treated treatment group respondents and treated control group phase-in respondents and cannot reject that they are equal.

basic reading scores, respectively, which is considerably larger than the typical impacts of adult education programs.

B. Dealing with Attrition

The previous results ignore the differential attrition between the treatment and control groups, and hence provide biased estimates of the impact of the *Cell-Ed* program. We first address attrition by bounding the treatment effects under the assumption of monotonicity and implementing Lee bounds, and then use the phase-in design of the program to further tighten these bounds.

Standard Lee Bounds

The initial treatment group in our context was surveyed over two rounds and is divided into two groups: respondents (with a sample size (N) of 27) and non-respondents (N=9). The control group was surveyed over three rounds, and hence can be divided into three groups: The “Always non-respondents” (those who dropped out of the survey after the baseline, N=3); the phase-in non-respondents (those who participated in baseline and second round, but dropped out by the third round, N= 10); and the phase-in respondents (or those control group participants within who participated in all three survey rounds, with N=21).²⁵ Figure 3A shows the phase-in design of our specific program, indicating the proportion of baseline respondents who participated in each round of the survey.

Lee bounds assume monotonicity, in other words, that the likelihood of non-response is monotonically related to receiving the treatment (Lee 2009). This is equivalent to stating that receiving the treatment either makes *all* observations (weakly) *more* likely to respond to the survey, or *all*

²⁵ There is no ranking of outcomes associated with this graph, as observations in the phase-in non-respondents could well have expected outcomes that are higher than the lower bound of the phase-in respondents.

observations (weakly) *less* likely to respond.²⁶ This assumption rules out the possibility that the treatment may affect different sub-groups differently, such as increasing the likelihood of response for one specific sub-group while decreasing the likelihood for another. This also implies that the non-respondents in the group with lower non-response would not have responded if their treatment status were changed. When attrition is higher in the treatment group, as is the case in our experimental set-up, the monotonicity assumption implies that the control observations who did not respond in the second survey would also not have responded had they immediately received treatment (Figure 3B).^{27,28}

Under these assumptions, we first estimate the upper and lower bounds by assuming the “best” and “worst” case scenarios.²⁹ In our case, the treatment group has an attrition rate of 25 percent and the control group has an attrition rate of 8.8 percent, so the difference in drop-out is 16.2 percentage points. To estimate the *lower bound*, we trim the lowest 17.7 percent of the *observations* from the control group.³⁰ The lower bound is thus equivalent to assuming that the outcomes of additional treatment non-respondents correspond to those control group observations who make up the bottom end of the control group distribution: We compare the observed observations in the treatment group (who make up 75% of all the treatment group observations) to the 82.3% “best” observed observations in the control group (who make up 75% of all control group observations) (Figure 3C). As is usual, the confidence intervals around this estimated *lower bound* are presented. We present bootstrapped confidence intervals around all bounds. To create the *upper bound*, we trim the highest 17.7 percent of observations from the control

²⁶ We write “weakly” because it is possible that the treatment leaves some observations’ likelihood unchanged.

²⁷ To be exact, the term “always non-respondent” thus already embodies an assumption we are making.

²⁸ In the appendix, we discuss a weaker assumption than this.

²⁹ This assumes that the additional percentiles of non-respondents in the treatment group correspond to, respectively, the best and worst respective percentiles of observed outcomes in the control group (Lee 2009).

³⁰ Note that we calculate the proportion of observed control group observations that is trimmed as the difference in attrition divided by the number of control group observations that are observed is $17.7 = 16.2 / (100 - 8.8)$. Because of indivisibilities, the exact trimming proportion is 19.35 in the empirical application.

group (equivalent to 16.2 % unconditionally). The upper bound is equivalent to assuming that the additional non-respondents are the top learners in the control group, and is constructed by comparing the observed treatment group outcomes to the 82.3% worst control group observations (who make up 75% of all control group observations) (Figure 3D).

Adjusted Lee Bounds

The phase-in design of the program allows us to observe which control group participants later dropped out of the program once they received the treatment (Figure 3A). Using this information, we further tighten the traditional Lee bounds by assuming that the treatment non-respondents would not have participated in the third round in the survey had they been assigned to the initial control group (Figure 3E), either because they are *always non-respondents* or because they would also not have responded after receiving treatment (Figure 3E). This rules out a case in which the treatment affects the probability of non-response in opposite directions for those in the treatment group and those in the control group when they become treated.^{31,32}

Using this additional assumption, we can then construct the lower bound on the treatment effect by discarding the worst cases among the phase-in non-respondents and comparing the mean outcomes of the treatment group with the mean outcomes of the newly-constructed control group (Figure 3F). The upper bound similarly discards the best cases among the phase-in non-respondents (Figure 3G). Compared with standard Lee bounds, these bounds are tighter, as we are trimming the same proportion

³¹As such, it is really implied by the monotonicity assumption, but we have to also assume that the monotonicity assumption holds across time. It can't be that treatment makes a control person more likely to respond just because he or she waited for three months for the program, rather than less as for anyone else.

³²This is equivalent to assuming that the control group phase-in non-respondents are ordered in their likelihood of non-response between the "always non-respondents" (who respond in one survey) and the "phase-in respondents" (who respond in three surveys), *even if they had been in the treatment group*.

of respondents but in a smaller group. The methodological appendix contains the formal proof that the bounds from this methodology, herein called the “adjusted Lee bounds”, are tighter than the standard Lee bounds.^{33,34}

A phase-in design also has an additional advantage relative to standard experimental designs in that it most likely reduces the attrition in the control group if they are aware that they will receive the treatment later. As many social experiments are not blinded (i.e. participants *do* know if they are receiving treatment or not), researchers often worry that attrition in the control group tends to be higher than in treatment groups, as treatment respondents may (incorrectly) associate the survey with the benefits of the program. By providing an incentive for control participants to remain in contact with the program, attrition may be reduced.

Bounding under the Monotone Treatment Selection (MTS) and Monotone Treatment Response (MTR) Assumptions

Using the phase-in design, we can also implement two alternative assumptions on the potential outcomes of non-respondents: The monotone treatment selection (MTS) assumption and monotone treatment response (MTR) assumption. The *monotone treatment selection* assumption (Manski and Pepper 2000) implies that program drop-outs have lower outcomes on average than those who take up the treatment. More specifically, it assumes that those who select into the treatment have higher

³³Lee (2009) shows that conditioning on any discrete variable will weakly tighten the bounds. In our case, we prove it strongly tightens the bounds. To provide an intuition why this leads to a tighter upper bound (Figure 3G), note that the mean grade of the top six students in a sub-group (control group phase-in non-respondents) is always lower than that of the top six in the larger group (all control group respondents), unless the top six students in the sub-group are also the top six students in the larger group. We can reject the latter exception with our data. Since we are subtracting the non-trimmed (who now have a higher mean) from the treatment group, we get a lower upper bound.

³⁴When there is treatment effect heterogeneity, the bounded estimate measures the treatment effect on the treatment respondents. In our context, this is highly correlated with continuing the *Cell-Ed* program until the survey (and perhaps completing it) and being surveyed.

expected outcomes under both treatment and non-treatment than those who do not. This allows us to bound the outcomes (from above) of the control group drop-outs during the second round.³⁵ The *monotone treatment response* assumption (Manski, 1997), on the other hand, specifies that treatment cannot make anyone worse off. This can be implemented by bounding the missing treatment observations. Given the very high correlation between the baseline and follow-up outcomes in the control group (with a correlation coefficient of 0.9 and a small difference in means), this is a plausible strategy. We thus construct plausible lower bounds for the average treatment effect by implementing MTS combined with MTR, discussing the assumptions and methodology more thoroughly in the methodological appendix.³⁶

C. Bounding Results

This section first provides the bounding results using the traditional and “adjusted” Lee bounds, and then provides lower bounds on the treatment effect using the MTS and MTR assumptions.

Table 4 shows both the traditional and adjusted Lee bounds for basic and broad reading scores. The relevant bounds on the treatment effects are in bold, namely, the lower bound of the 95% confidence interval (around the lower bound), as well as the upper bound of the 95% confidence interval (around the upper bound).³⁷ In all cases, the treatment effects are bounded away from zero: The point estimates of the lower bound for basic and broad reading scores are 4.43 and 2.08 years, respectively, with a statistically significant effect at the 1 percent level (Panel A, Columns 1 and 4). Focusing on the

³⁵Specifically, we bound the second round outcome of the control group dropouts by including the larger of two values: the minimum of those who eventually took up treatment and the observations own baseline value. This is an even more conservative approach.

³⁶The panel structure of the data allows for many different reasonable assumptions to be implemented. For example, we can implement MTR based on the changes due to the treatment; as well as bounding changes between two surveys in the control group dropout by changes of the observed changes in the control group among respondents who do not drop out.

³⁷The lower bound of the 95% confidence interval of the lower bound and the upper bound of the 95% confidence interval of the upper bound are conservative estimates of the 95% confidence interval around the estimate (Lee, 2009). As noted above, we bootstrap the confidence intervals.

conservative lower bound of the 95% confidence interval, the treatment effect is bounded below by 1.58 and 0.72 years for basic and broad reading scores, respectively (Panel A, Columns 2 and 5).

As predicted, the adjusted Lee bounds are tighter (Panel B). The point estimates for the lower bound are 4.54 and 2.04 for basic and broad reading scores, respectively, with a statistically significant effect at the 1 percent level (Panel B, Columns 1 and 4). The lower bounds of the 95% confidence interval are 1.73 and 0.84 years (Panel B, Columns 2 and 5). Thus, making weak assumptions on the non-respondent outcomes, both the Lee and adjusted Lee bounds provide strong evidence of the positive and statistically significant impacts of the *Cell-Ed* program on students' learning.

Table 5 presents the lower bounds on the treatment effect using the MTS and MTR assumptions. The point estimates on the lower bound of the treatment effect are 3.31 and 1.54 years for basic and broad reading scores respectively, with a statistically significant effect at the 1 percent level (Columns 1 and 4). The lower bounds of the 95% confidence intervals confirm the patterns observed in Table 4, suggesting that *Cell-Ed* increased basic and broad reading scores by 1.07-.5 years, respectively. While considerably smaller in magnitude, these bounded estimates still suggest that the *Cell-Ed* program resulted in an 8-16% increase in baseline basic and broad reading scores over a four-month period.³⁸

³⁸The difference in the magnitude of the lower bounds in Tables 4 and 5 can be explained by two factors: the strength of the assumptions made on the drop-outs and the different estimations of the treatment effects. With regards to the latter, we note that the treatment group splits into two groups: 1) the treatment respondents, namely, the group who spent an average of 50 hours and completed at least 184 mini-modules. For this group, the *Cell-Ed* program “worked”; and 2) the group that spent on average 8 hours on the platform and completed a maximum of 104 mini-modules. For this group, the *Cell-Ed* “worked less”. The Lee bounds in Table 4 bound the treatment effect for the treatment respondents, and suggest that the point estimate for basic reading scores was 4.54, with a lower bound of 1.73. The MTS and MTR assumptions of Table 5 bound the average treatment effect for the whole population, and find a point estimate of 3.31 and a lower bound of 1.07 for basic reading scores.

D. Impacts of the Cell-Ed Program on Empowerment

The previous results have shown that the *Cell-Ed* program increased students' basic and broad reading scores in a relatively short period of time. Yet beyond learning, a key claim of many adult education programs is one of empowerment: educated adults are able to make better decisions, read street signs and prescriptions and search for jobs that require literacy skills. This section investigates whether the increased learning due to *Cell-Ed* led to improvements in self-esteem and self-efficacy, using the RSES and the GSES.

Table 6A reports the results of the same specifications as those in Table 3, using measures of self-esteem and self-efficacy as the dependent variables. Overall, the *Cell-Ed* program increased students' self-esteem score by 2.49 points, with a statistically significant effect at the 5 percent level (Table 6A, Column 1). These results are robust to the inclusion of additional covariates (Column 2), as well as a value-added specification (Column 3). However, the result is smaller in magnitude and no longer statistically significant once the DD (Column 4) and fixed effects specifications (Column 5) are used.

While overall self-efficacy scores are higher in the treatment group, there is no statistically significant impact of the *Cell-Ed* program on self-efficacy for the first three specifications (Columns 6-8). This is perhaps unsurprising, as self-efficacy scores were relatively higher in the control group at baseline. Using DD and fixed effects specifications (Columns 9 and 10), the *Cell-Ed* program is associated with an increase of 3.75-4.75 points in students' self-efficacy scores, with a statistically significant effect at the 10 percent and 1 percent levels, respectively.^{39,40}

³⁹Table A5 shows that these results are robust to including baseline mobile ownership and self-esteem, as well as using normalized empowerment scores as the dependent variable.

As was the case with the test score results, the results in Table 6A are likely to be biased due to non-random attrition. Table 6B corrects for non-random attrition using the Lee and adjusted Lee bounds. Unlike the reading results, the results are not robust to using the Lee bounds: the lower Lee bounds for both empowerment measures are not statistically significant from zero. Using the adjusted Lee bounds, the lower bound for the basic reading score is statistically significant at the 10% level, while the confidence interval around the lower bound for broad reading scores includes zero.⁴¹ Bounding the results using the MTR and MTS assumptions (Table 6C), the estimate of the lower bound of the *Cell-Ed* treatment effect on the self-esteem score is 2.3 and statistically significant at the 5% level, whereas the effect on the self-efficacy is not statistically significant from zero. Overall, these results suggest that the *Cell-Ed* program is associated with an increase participants' self-esteem, but we are unable to conclude that there is a statistically significant effect on self-efficacy.⁴²

While Table 6 shows the impact of the *Cell-Ed* program on self-esteem and self-efficacy at the end of the program, we are also interested in the relationship between the *Cell-Ed* program, learning and self-esteem over time. Table 7 investigates this relationship using data from the weekly self-esteem measurements and the *Cell-Ed* platform. Overall, self-esteem is negatively associated with lack of learning progress during the previous week (Column 1). In other words, as a student's proportion of incorrect responses during the previous week increases, the student's self-esteem decreases. When disaggregating these results by gender (Columns 2-4), the relationship is driven entirely by men. This

⁴⁰Table A6 uses an empowerment index comprised of the normalized RSES and the normalized GSES as the dependent variable. The results are broadly consistent with those in Table 6.

⁴¹The estimated lower bound for the index of empowerment is significantly different from zero at the 5% level of significance when implementing the adjusted Lee bounds, but is insignificant for the estimated Lee bounds.

⁴²The estimated lower bound for the index of empowerment is significantly different from zero at the 5% level of significance when implementing the MTS and MTR assumptions.

suggests that while overall self-esteem is higher for *Cell-Ed* participants by the end of the program, perceptions of self-esteem may change over time, particularly when experiencing learning failures.

While the results in Table 7 provide important insights into the dynamic relationship between learning and self-esteem, Table A8 shows how learning is associated with the likelihood of dropping out of the program. *Cell-Ed* participants who had a higher proportion of errors over the course of their interactions with the platform were more likely to stop interacting with the platform and drop out of the subsequent survey round (Column 1).⁴³ Columns 2 and 3 focus on the sub-sample of women: Female participants with a proportion of errors exceeding a threshold of 40% were more likely to drop out of the program. The effects seem to be particularly strong among participants with a high proportion of errors, as suggested by the quadratic specification (Columns 4-6).⁴⁴

These results, taken together, provide some additional insights into the relationship between learning, self-esteem and non-response in the context of our non-parametric bounds. First, the heterogeneous effects in Table 7 suggest that conditioning on gender would further tighten the bounds for the empowerment regressions in Tables 6B and 6C. Unfortunately, our sample size does not allow us to do this. Second, these results raise some questions about the validity of our bounding assumptions for the empowerment outcomes. For example, if a larger proportion of learning errors leads to lower self-esteem, and a learner is more likely to drop out after these errors have lowered his or her self-esteem below their baseline level, then the MTR assumption would not be valid. We do not believe that this is the case for men, as few men drop out and there seems to be no pattern linking men's errors to attrition (although we acknowledge the low power to detect such a pattern). Yet a dynamic in which

⁴³The monetary incentive provided for completing the program may play a role.

⁴⁴We note that the sample for men is too small for robust results, in particular, because only 12.5 percent of men (2 out of 16) drop out, as compared to almost one third for women.

women whose self-esteem was lowered by their lack of progress in the program and eventually drop out is consistent with the above findings. Table A9 investigates this further. We find that female *Cell-Ed* participants who eventually drop out do not have lower levels of self-esteem immediately prior to dropping out (as compared with their own baseline levels of self-esteem). This suggests that the threat to the validity of the MTR assumption is not a primary concern for our self-esteem bounding results.⁴⁵

V. Mechanisms

Why might the *Cell-Ed* program improve students' learning outcomes, especially in the absence of a teacher? A key hypothesis of the program was that a mobile phone-based program might reduce the opportunity costs of investing in adult education for busy adults by allowing them to learn when, where and how they wished. We are able to test this hypothesis by using the *Cell-Ed* platform data, which provides information on when *Cell-Ed* students interacted with the platform.

Figure 4 shows the learning patterns of *Cell-Ed* participants. Overall, participants learned at all times during the day and night, at times much later than a standard adult education class would operate, suggesting that the mobile-based course was more appropriate for adults' work schedules (Figure 4A). Figure 4B suggests that *Cell-Ed* participants spent less time on the platform during weekends as compared with weekdays, perhaps due to their work schedules. Finally, Figure 4C shows the distribution of interactions with the platform in terms of the amount of time spent: most interactions are quite short, approximately 10 minutes, suggesting that learners use the platform for short learning episodes. This is in contrast to most adult education programs, which hold classes for several hours on a

⁴⁵Regarding the reading outcomes, we posit that participants' reading skills could not actually decrease in response to an adult education treatment.

weekly basis. On average, though, time spent on learning is similar to standard literacy programs: Cell-Ed learners were spending an average of 2.95 hours per week (3.75 for respondents) on the platform. The target for in-situ learning is about 4-6 hours per week. Accounting for in-situ absences, the 50 percent dropout rate and travel time suggests that *Cell-Ed* learners spent slightly more time (on average) on literacy acquisition than students in a standard adult education program.⁴⁶

While we do not know whether this flexibility is indeed the key to the large improvements in reading outcomes, the usage patterns highlight that people did make use of it.

VI. Conclusion

Information technology, and in particular mobile phones, enables individuals to access information at any time and location. Yet despite the potential, there is limited evidence that simple mobile phone devices can be used to teach basic educational skills, such as literacy and numeracy. Using a randomized control trials of a mobile phone-based adult education program (*Cell-Ed*), we find that the program significantly increased adult students' reading levels within a short period of time. These results are robust to correcting for significant non-random attrition using a variety of non-parametric methods. We also find that the *Cell-Ed* program is correlated with an increase in students' self-esteem by the end of the program, and that there is a dynamic relationship between learning progress and self-esteem.

On a methodological side, our paper contributes to the literature on the use of non-parametric approaches to bounding treatment effects. The phase-in design of the program allows us to observe

⁴⁶Figure 5 shows learning patterns by baseline employment status. Employed participants have somewhat different patterns of accessing the platform, learning for longer periods on weekends and in the evenings as compared to their unemployed counterparts.

which control group participants later dropped out of the program once they received the treatment. It probably also contributed to reducing attrition in the control group. Exploiting the randomized phase-in design, we are able to tighten the Lee bounds. We are also able to show that this experimental set-up provides a number of identification possibilities that are absent with simple random assignment. An additional benefit of phase-in designs is that they address concerns by non-governmental partner organizations of excluding observations from treatment, which however, comes at the expense of longer-term treatment effects.

Admittedly, our experimental set-up has several limitations. First, we are unable to compare learning via the *Cell-Ed* platform with learning in a traditional adult education program, or the interaction between the two. As a result, we are unable to conclude whether such programs are complements or substitutes for teachers and in-classroom learning. Second, our small sample size greatly limits the external validity of our results. Nevertheless, our results show that a distance learning program via a simple mobile phone significantly improved adults' learning outcomes in this context, and suggests that there is great scale and scope for using these technologies in education programs in both developed and developing countries.

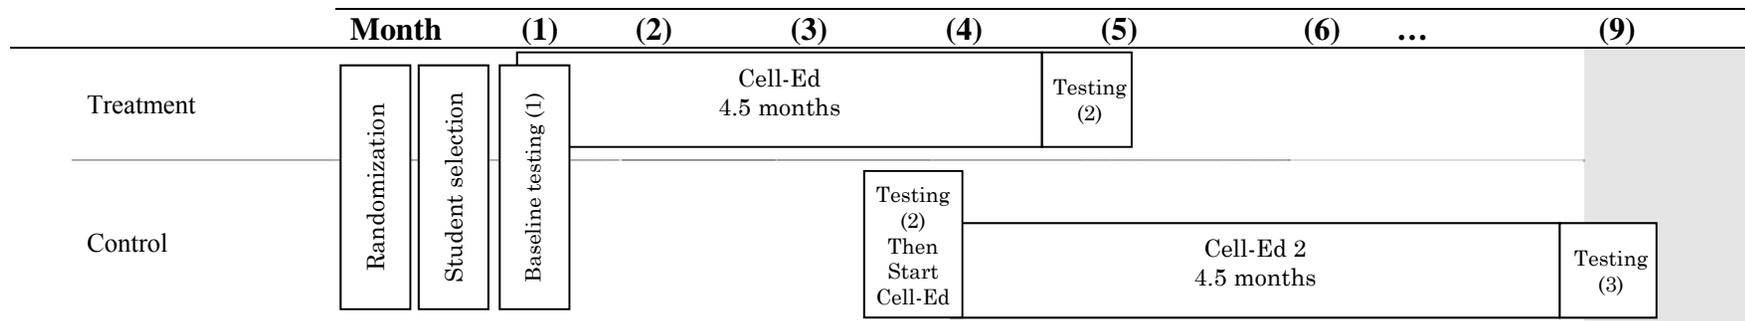
References

- Abadzi, Helen.** 1994. "What We Know About Acquisition of Adult Literacy: Is There Hope?," In *World Bank discussion papers*, ix, 93 p. Washington, D.C.: The World Bank.
- Abadzi, Helen.** 2013. Literacy for All in 100 Days? A research-based strategy for fast progress in low-income countries," *GPE Working Paper Series on Learning No. 7*
- Aker, Jenny C., Christopher Ksoll and Travis J. Lybbert.** October 2012. "Can Mobile Phones Improve Learning? Evidence from a Field Experiment in Niger." *American Economic Journal: Applied Economics*. Vol 4(4): 94-120.
- Andrabi, Tahir, Jishnu Das, Asim Ijaz Khwaja, and Tristan Zajonc.** 2011. "Do Value-Added Estimates Add Value? Accounting for Learning Dynamics." *American Economic Journal: Applied Economics*, 3(3): 29–54.
- Angrist, Joshua D., and Jorn - Steffen Pischke.** 2009. *Mostly Harmless Econometrics: An Empiricist's Companion* (Princeton, New Jersey: Princeton University Press, 2009).
- Ashraf, Nava, Dean Karlan, and Wesley Yin.** March 2010. "Female Empowerment: Further Evidence from a Commitment Savings Product in the Philippines." *World Development* 38, no. 3 (March 2010): 333–344. □
- Ashraf, Nava.** September 2009. "Spousal Control and Intra-Household Decision-Making: An Experimental Study in the Philippines." *American Economic Review* 99, no. 4 (September 2009): 1245–1277.
- Banerjee, Abhijit, Shawn Cole, Esther Duflo and Leigh Linden.** 2007. "Remedying Education: Evidence from Two Randomized Experiments in India." *The Quarterly Journal of Economics*, 122(3), pp. 1235-64.
- Rukmini Banerji, James Berry and Marc Shotland.** 2013. "The Impact of Mother Literacy and Participation Programs on Child Learning: A Randomized Evaluation in India."
- Barrow, Lisa, Lisa Markman and Cecilia Elena Rouse.** 2009. "Technology's Edge: The Educational Benefits of Computer-Aided Instruction." *American Economic Journal: Economic Policy*, 1(1), pp. 52-74.
- Blunch, Niels-Hugo and Claus C. Pörtner.** 2011. "Literacy, Skills and Welfare: Effects of Participation in Adult Literacy Programs." *Economic Development and Cultural Change*. Vol. 60, No. 1 (October 2011): 17-66.
- Carron, G.** 1990. "The Functioning and Effects of the Kenya Literacy Program." *African Studies Review*, pp. 97-120.

- Case, Anne.** 2005. "The Primacy of Education," In *Understanding Poverty*, ed. Abhijit Vinayak Banerjee, Roland Benabou and Dilip Mookherjee. Oxford ; New York: Oxford University Press.
- De Grip, A. and J. Van Loo.** 2002. "The Economics of Skills Obsolescence: A Review," In *Research in Labor Economics*, ed. J. Van Loo and K. Mayhew, 1-25.
- DiNardo, J., J. McCrary, and L. Sanbonmatsu.** 2006. "Constructive Proposals for Dealing with Attrition: An Empirical Example." Working paper, University of Michigan.
- Duflo, Esther.** 2012. "Women Empowerment and Economic Development." *Journal of Economic Literature*. 50(4). 1051-1079.
- Doepke, Mathias and Michele Tertilt.** 2014. "Does Female Empowerment Promote Economic Development?" *NBER Working Paper 19888*, NBER, Inc.
- Heckman, James.** 1979. "Sample Selection Bias as a Specification Error," *Econometrica*, January 1979, 47 (1), 153–162.
- Horowitz, Joel L. and Charles F. Manski.** 2000. "Nonparametric Analysis of Randomized Experiments With Missing Covariate and Outcome Data," *Journal of the American Statistical Association*, 95, 77–84.
- Imbens, Guido and Charles F. Manski.** 2004. "Confidence Intervals for Partially Identified Parameters," *Econometrica*, November, 72 (6), 1845–1857.
- Kim, Y., R. Telang, W. Vogt and R. Krishnan.** 2009. "An Empirical Analysis of Mobile Voice Service and SMS: A Structural Model." *Management Science*, 56(2), pp. 234-52.
- Lee, David S.** 2009. "Training, Wages, and Sample Selection: Estimating Sharp Bounds on Treatment Effects" *The Review of Economic Studies*, 6, 1072-1102.
- Linden, Leigh.** 2008. "Complement or substitute? The effect of technology on student achievement in India." JPAL Working Paper.
- Literacy Network of Greater LA and United Way.** 2004. *Literacy @ work. Skills Today, Jobs Tomorrow*. Los Angeles: Literacy Network of Greater LA and United Way.
- Manski, Charles.** 1990. "Nonparametric Bounds on Treatment Effects," *American Economic Review Papers and Proceedings*, 1990, 80, 319–323.
- Manski, Charles.** 1997. "Monotone Treatment Response," *Econometrica*, 65, 1311–1334.
- Manski, Charles, John Pepper.** 2000. "Monotone Instrumental Variables: With an Application to the Returns to Schooling" *Econometrica*, 68, 997-1010.
- McGrew, Kevin, Frederick Schrank and Richard Woodcock.** 2007. *Technical Manual. Woodcock-Johnson III Normative Update*. Rolling Meadows, IL: Riverside Publishing.

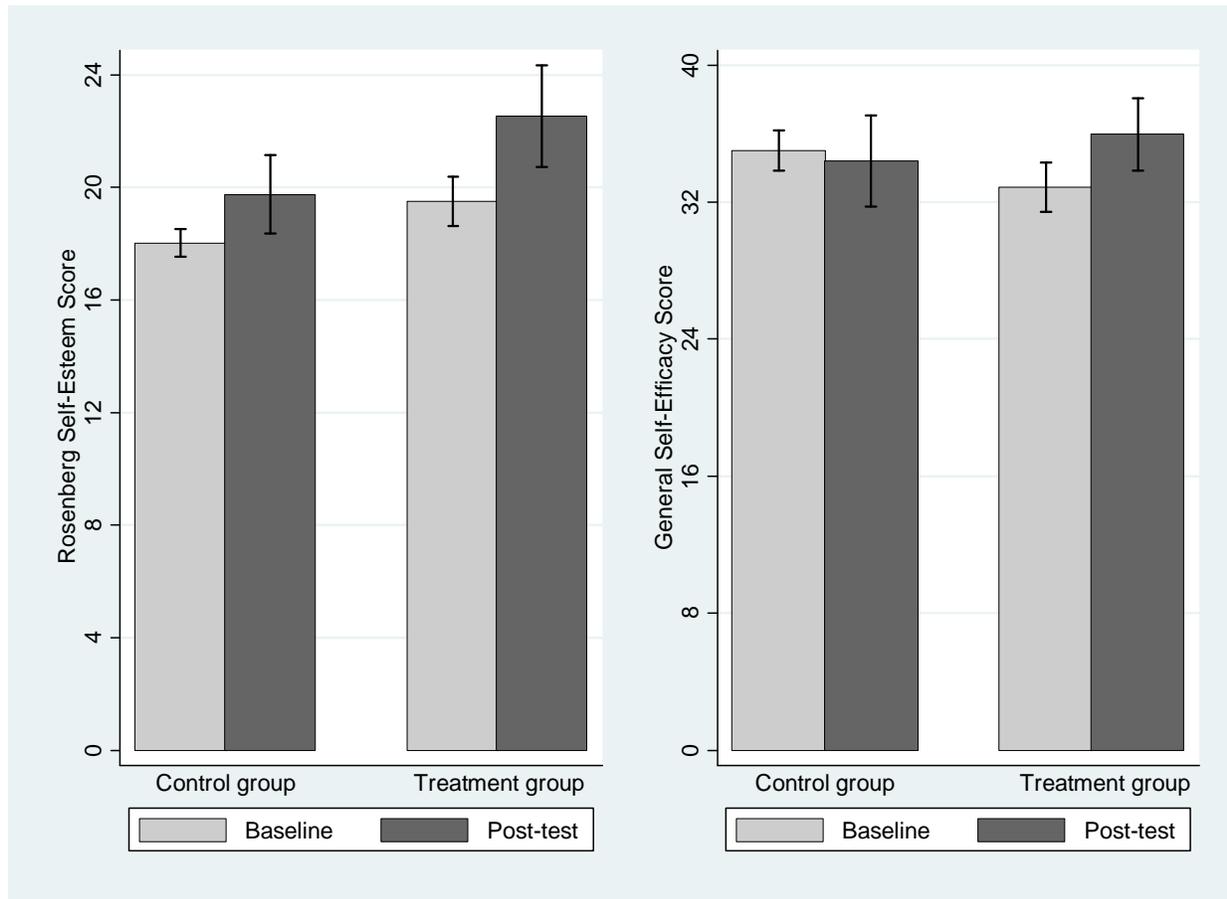
- National Centre for Education Statistics. 2003.** "The 2003 National Assessment of Adult Literacy", Washington DC, USA: National Centre for Education Statistics, U.S. Department of Education.
- Ortega, Daniel and Francisco Rodríguez. 2008.** "Freed from Illiteracy? A Closer Look at Venezuela's Mision Robinson Literacy Campaign." *Economic Development and Cultural Change*, 57, pp. 1-30.
- Osorio, Felipe, and Leigh L. Linden. 2009.** "The use and misuse of computers in education: evidence from a randomized experiment in Colombia." *The World Bank Policy Research Working Paper Series*.
- Oxenham, John, Abdoul Hamid Diallo, Anne Ruhweza Katahoire, Anne Petkova-Mwangi and Oumar Sall. 2002.** *Skills and Literacy Training for Better Livelihoods: A Review of Approaches and Experiences*. Washington D.C.: World Bank.
- Romain, R. and L. Armstrong. 1987.** *Review of World Bank Operations in Nonformal Education and Training*. World Bank, Education and Training Dept., Policy Division.
- UNESCO. 2005.** *Education for All: Global Monitoring Report. Literacy for Life*. Paris: UNESCO.
- UNESCO. 2008.** *International Literacy Statistics: A Review of Concepts, Methodology and Current Data*. Montreal: UNESCO Institute for Statistics.
- UNESCO. 2012.** *Education for All: Global Monitoring Report. Youth and Skills: Putting Education to Work*. Paris: UNESCO.

Figure 1. Timeline of Data Collection and Adult Education Activities



Notes: This figure represents the timeline for the *Cell-Ed* program. Testing (1), Testing (2) and Testing (3) refer to the first, second and third round of testing, respectively.

Figure 2A: Basic and Broad Reading Scores (by Treatment Status and Survey Round)



Notes: Graph pictures the means scores on the Woodcock-Munoz III reading assessment, for basic (left panel) and broad (right panel) reading scores, by treatment status. Baseline survey measurement means are in lighter color, Round 2 measurements in darker color. Error bars depicted are the 95% confidence intervals around the mean.

Figure 2B: Distribution of Basic Reading Score by round and treatment status

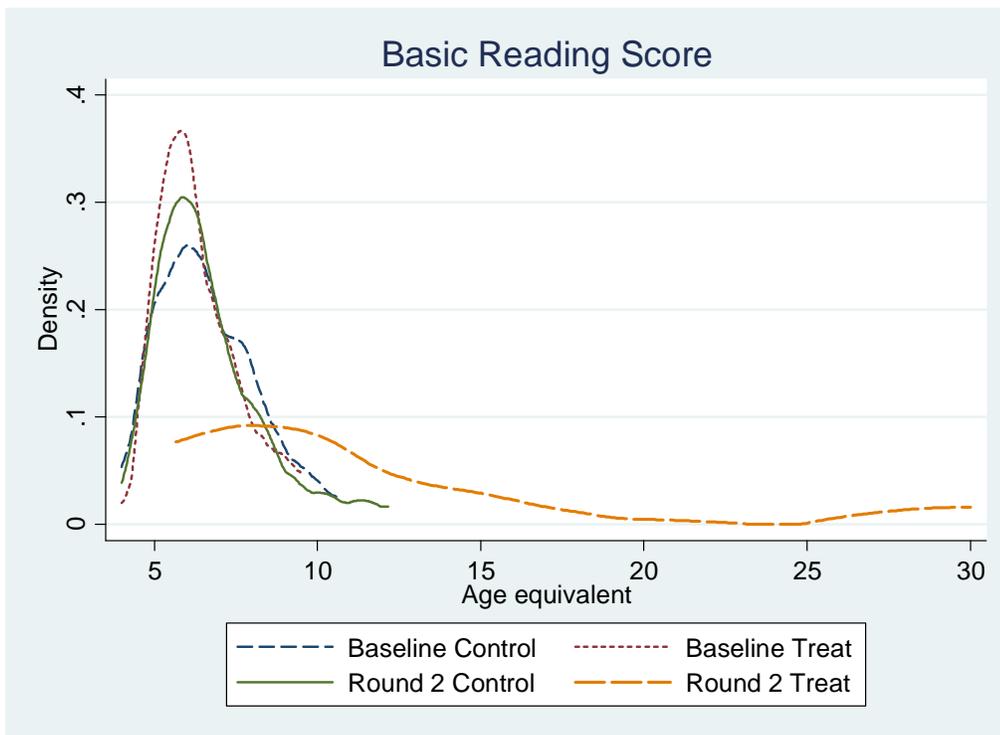


Figure 2C: Distribution of Broad Reading Score by round and treatment status

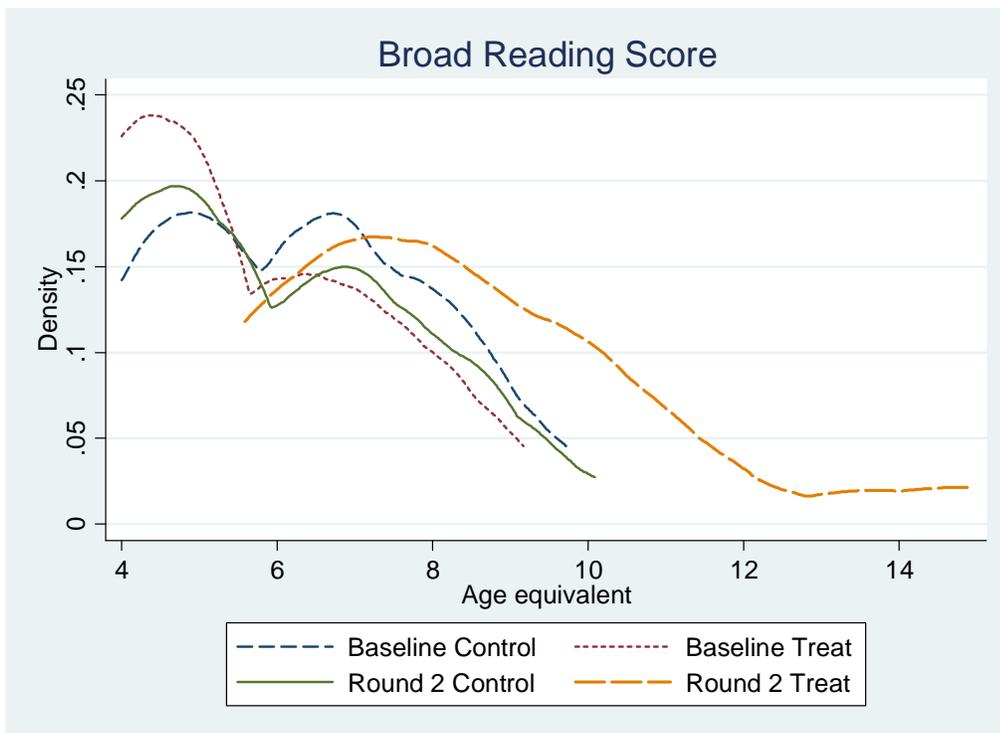


Figure 3A. Phase in Design and Respondents

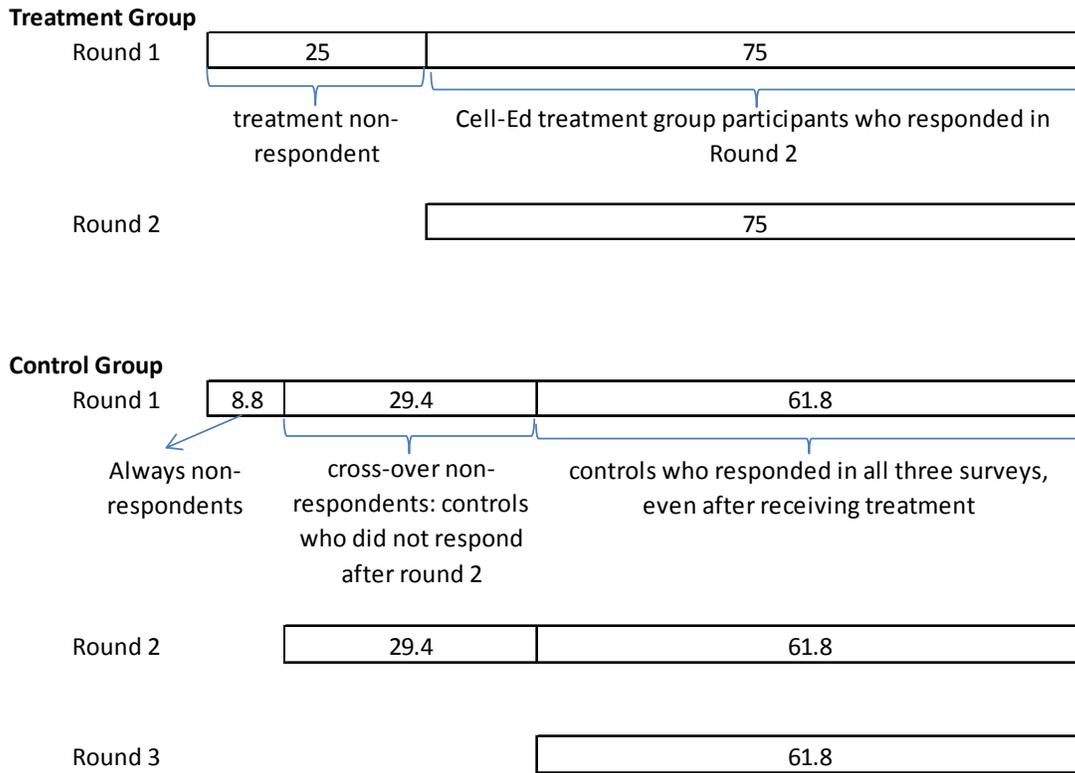


Figure 3B. Traditional Lee Bounds and attrition

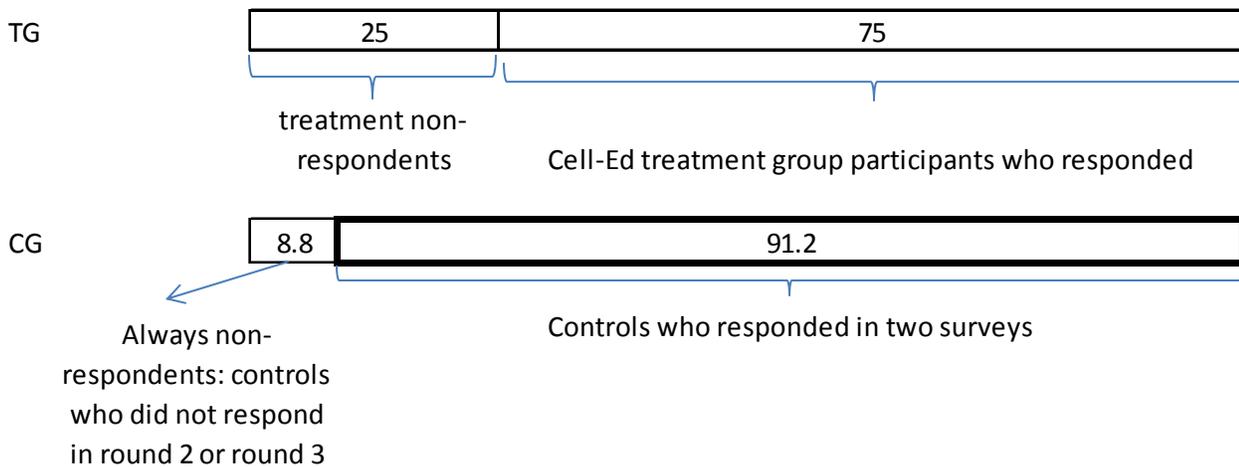
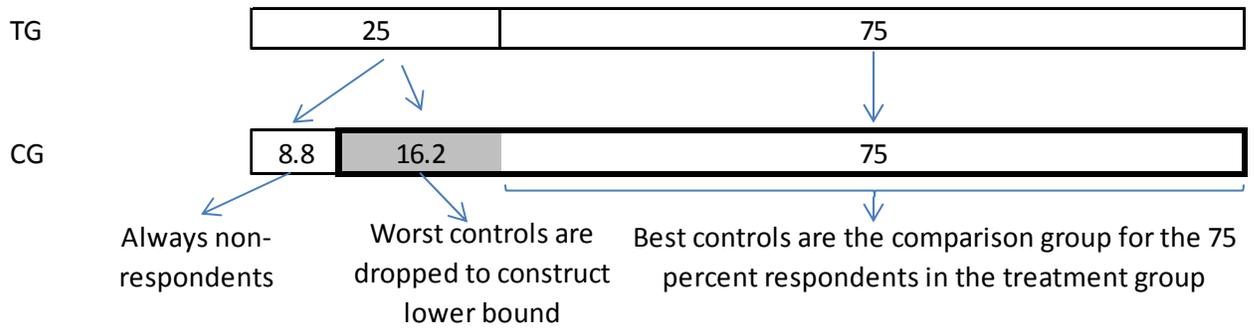
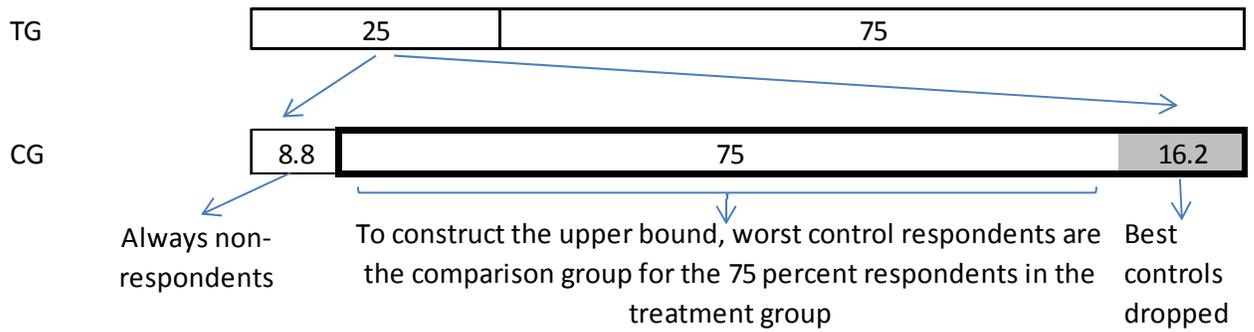


Figure 3C. Traditional Lee Lower Bounds Computation



Notes: Thick black outline indicates control respondents (some of whom are trimmed). The shaded area indicates observations that are trimmed.

Figure 3D. Traditional Lee Upper Bounds Computation



Notes: Thick black outline indicates control respondents, some of whom are trimmed. The shaded area indicates observations that are trimmed.

Figure 3E. Adjusted Lee Bounds: respondents and non-respondents

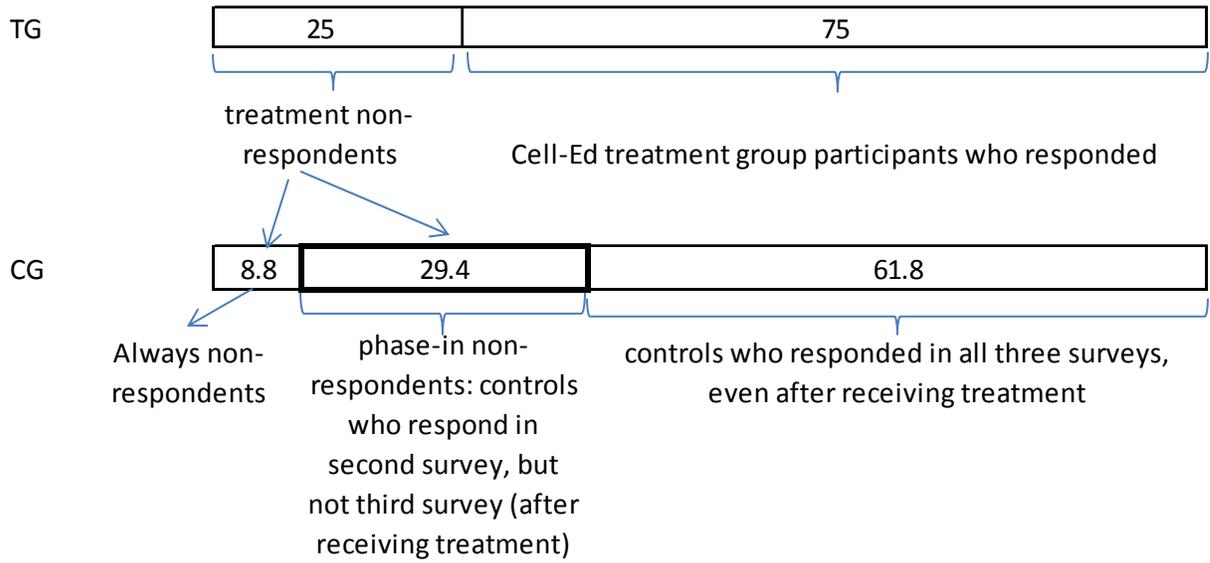
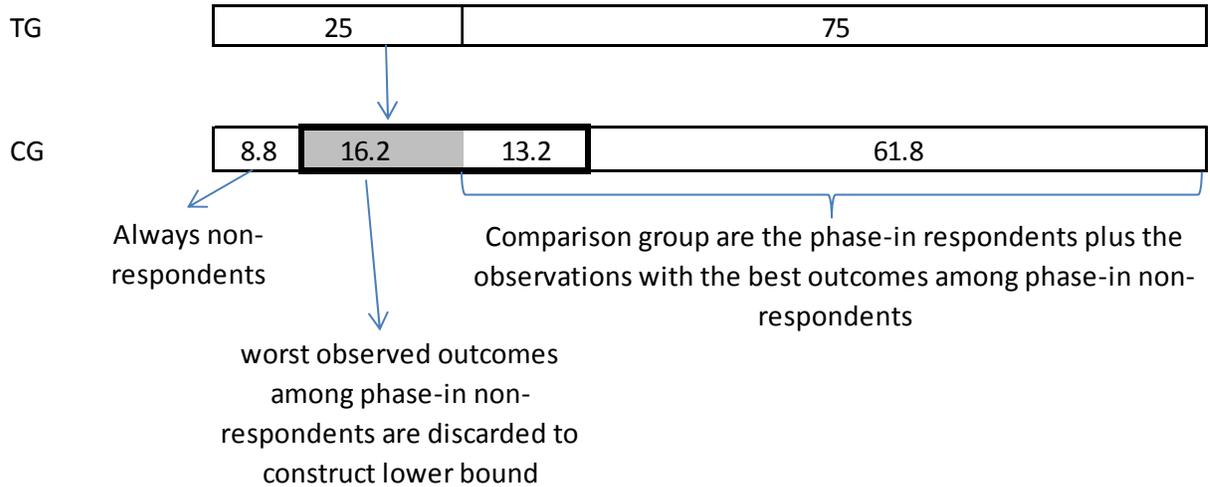
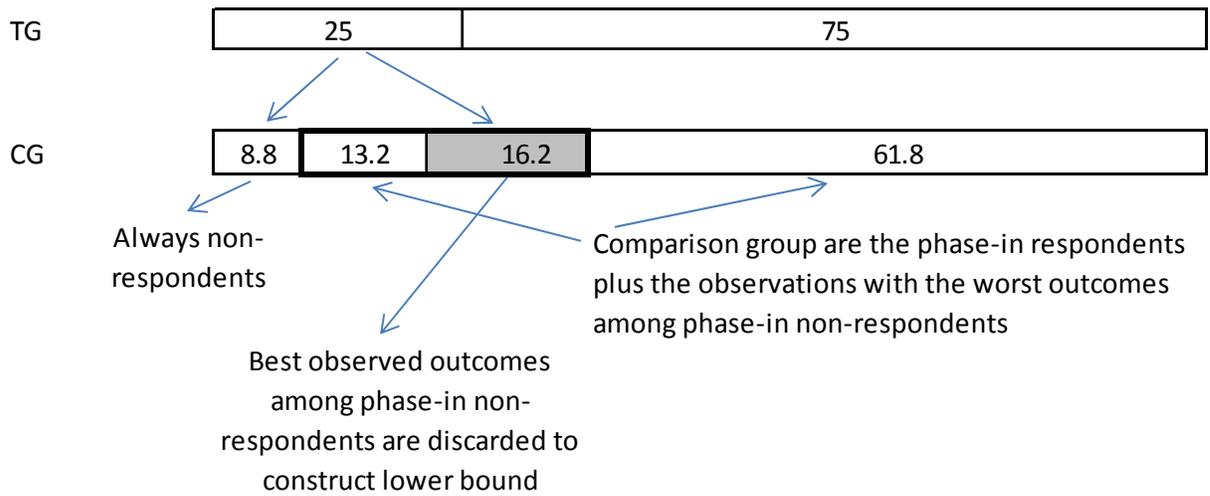


Figure 3F. Lee Lower Bound using the Phase-in Non-Respondents



Notes: Thick black outline indicates control observations who did not respond after receiving treatment (phase-in non-respondents). This is the group of observations from which some trimmed. The shaded area indicates those observations which are trimmed.

Figure 3G. Lee Upper Bound using the Phase-in Non-Respondents



Notes: Thick black outline indicates control observations who did not respond after receiving treatment (phase-in non-respondents). This is the group of observations from which some trimmed. The shaded area indicates those observations which are trimmed.

Figure 4A: When Learners Learn: Call Start and End Times

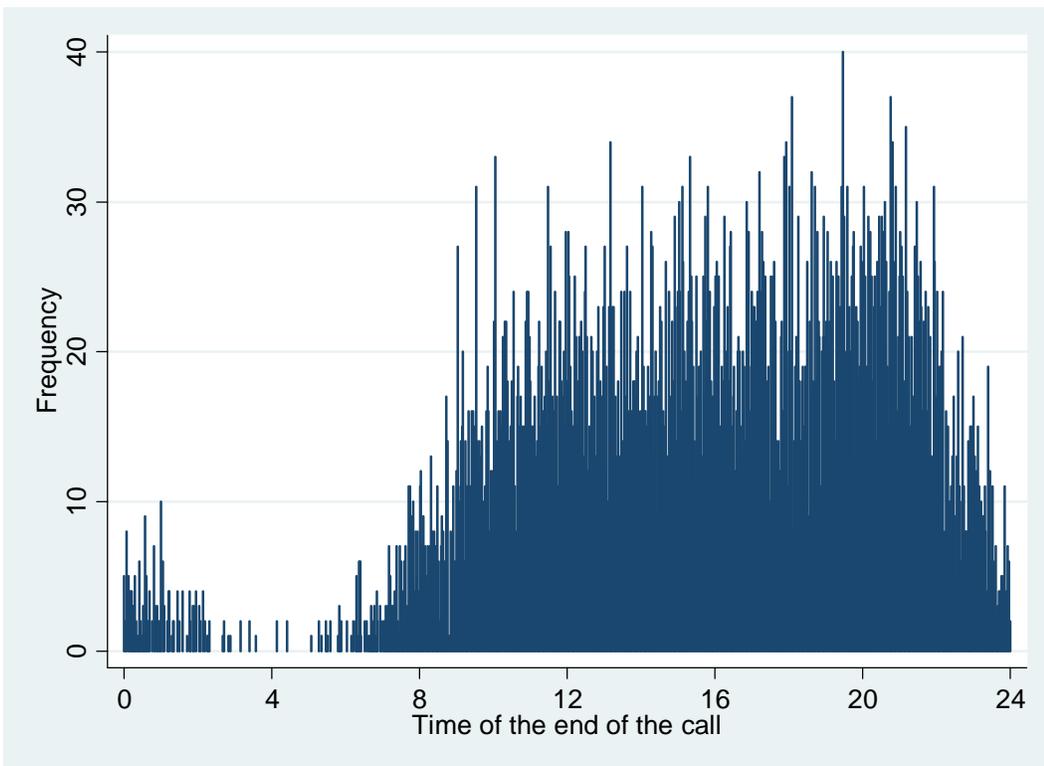
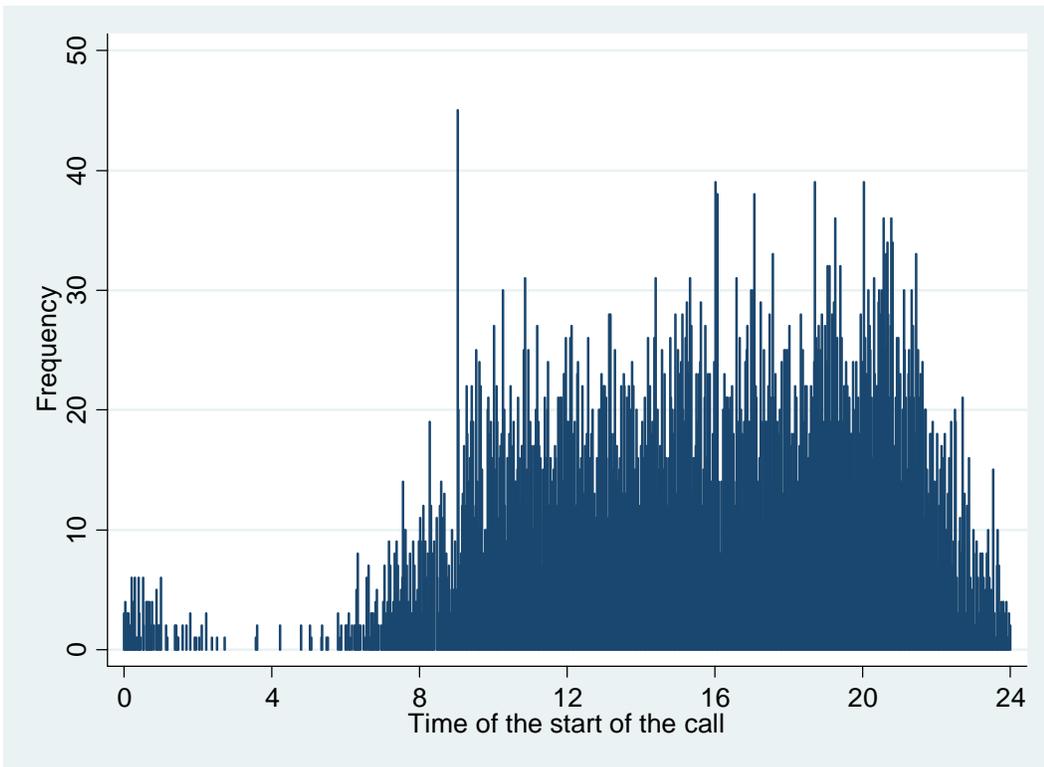


Figure 4B: When Learners Learn: Day of the Week

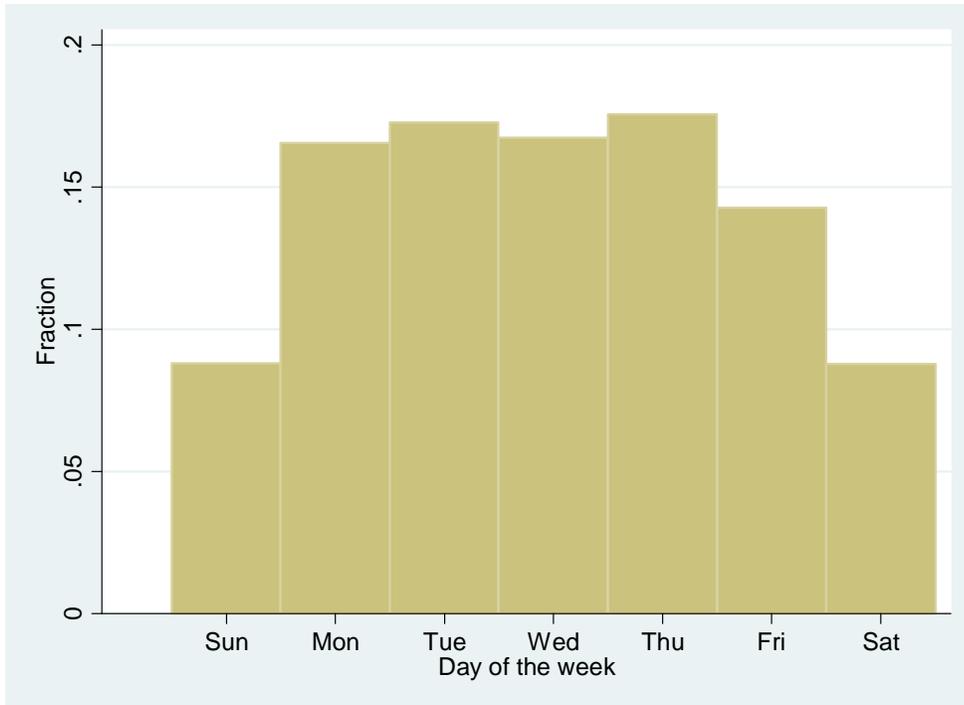


Figure 4C: Duration of Calls in Hours

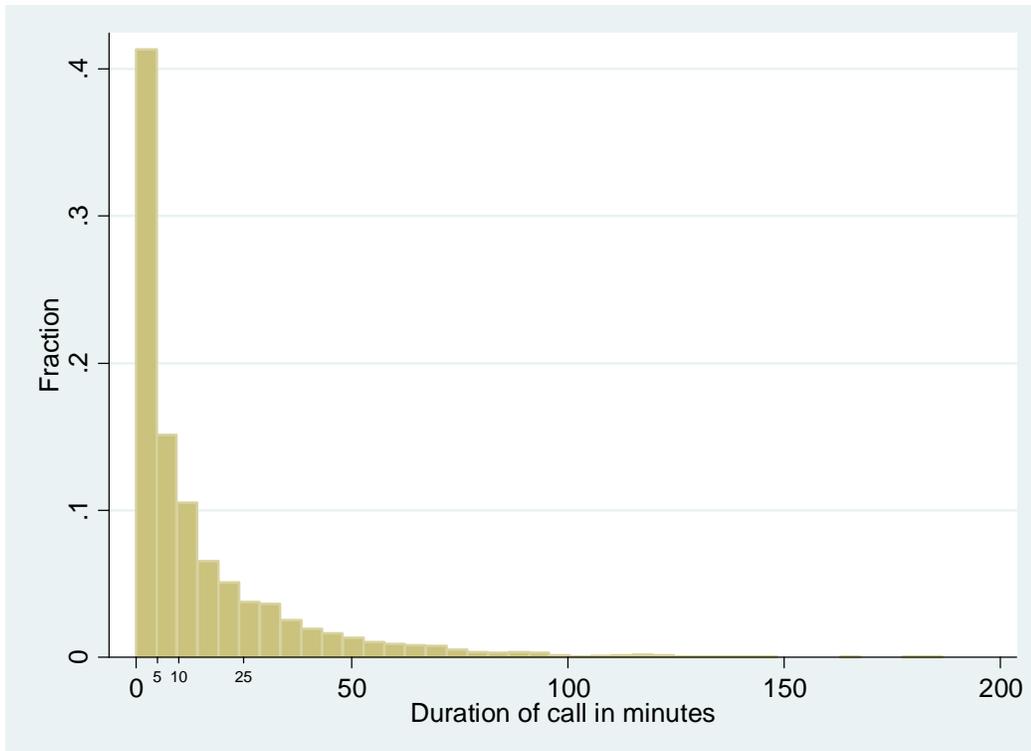


Figure 5A: Time of calls, by employment status

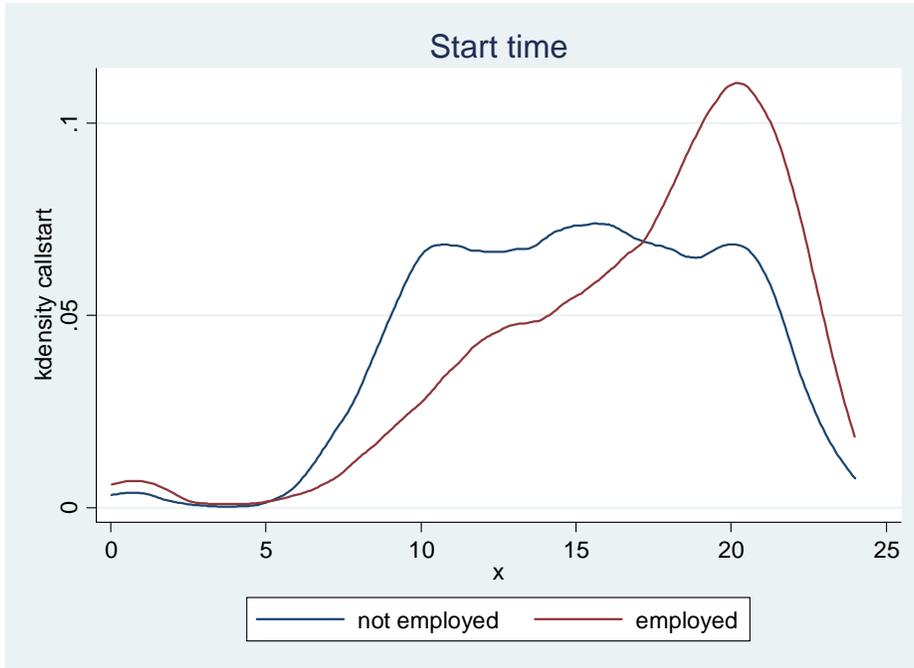


Figure 5B: Day of week for calls, by employment status

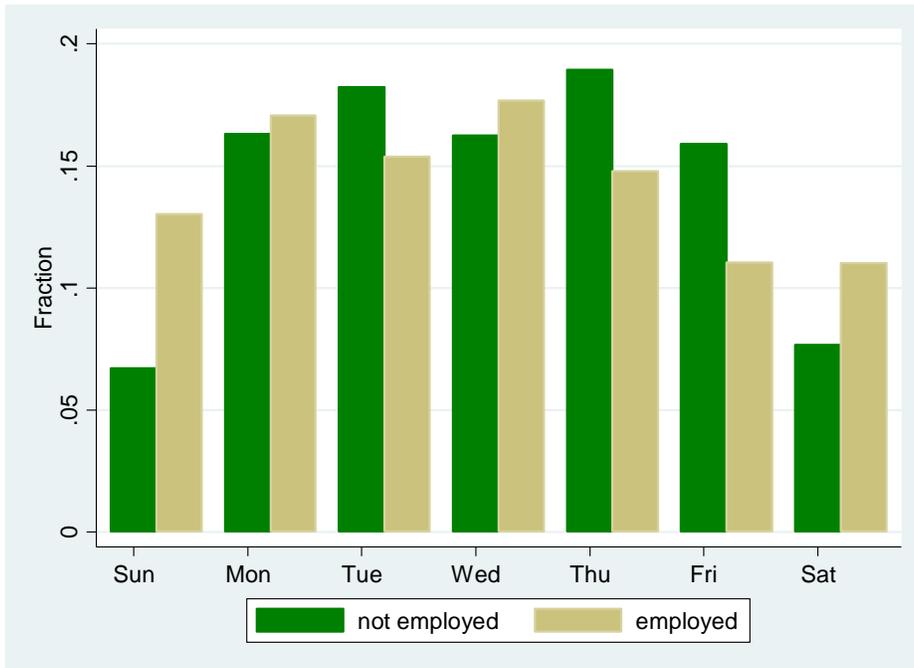


Figure 5C: Call duration in minutes, by employment status

