Pre-analysis Plan

Estimating the effect of LinkNYC on voter registration and participation

Justin de Benedictis-Kessner and Melissa Sands

"Smart cities" are touted as holding great promise for connecting citizens and government through the use of information and communication technology. Beginning in 2015, New York City began installing LinkNYC kiosks to replace the city's network of payphones. Each kiosk features 55'' high definition displays that show content including advertising and transit information, as well as tablets for accessing city services and maps, free USB charging stations, and phones for making domestic calls. Similar kiosks are being rolled out in London and other U.S. cities.

On September 21st, 2018, LinkNYC launched a voter registration drive to encourage voter registration on kiosk tablets, which connect users to an electronic voter registration page. The campaign would run until the New York voter registration deadline on October 12th. In an era of low voter turnout, this initiative holds promise for increasing voter participation. While these devices are designed to bring internet access and city services to users' fingertips, their ability to promote electoral participation is an open question; their usage for this purpose – and whether messages encouraging participation actually work – has not been studied. Registration is not only a necessary prerequisite to voting, the act of doing so drives subsequent electoral participation (Nickerson 2015). Recent debates over voter registration laws, such as ID requirements, emphasize that removing barriers to participation, especially for historically disenfranchised populations, is an imperative. Thus an important question is whether increasing access to, and the visibility of, voter registration resources improves electoral participation.

The present study is a randomized implementation of registration encouragement using LinkNYC large screens to display either (1) a localized neighborhood-identity based participation appeal; (2) an urgency appeal; or (3) no appeal (a control group). Outcomes measured include individual counts of voter registration site usage, with date-times, duration of use, and location of specific kiosks where each use occurred, as well as voter registration and turnout data.

Census tracts (N = 838 containing kiosks) are block randomized by population density, age, racial breakdown, and income (from the U.S. Census) and number of kiosks per tract, as well as by impressions, a proxy for the number of people exposed to each of the kiosks in a tract. From 9/21/18 through 10/12/18 (registration deadline), all the kiosks in a given census tract will display their assigned message, or lack of a message, at regular intervals, alongside other usual (non-experimental) content. At the end of that period, the implementing organization (CityBridge LLC, which operates the kiosks) will provide timestamped counts of registrations on each kiosk during the study period. Although we will not have information about individual registrants, we will also use the NYC voter file to obtain characteristics of people who registered

during the study period under each of the three conditions based on the census tract in which they live and the kiosks nearest their home address.

Hypotheses

Outcomes are measured from two main sources: (1) counts of individuals who use LinkNYC kiosks to register at the kiosk level, provided by CityBridge, which can be aggregated up to the tract level, and (2) voter information from the NYC voter file. The former includes date and time information, while the latter includes dates, allowing us to determine which registrations occurred during the message campaign and to study heterogeneous effects by proximity to the registration deadline. Moreover, we will obtain an updated NYC voter file after Election Day to determine whether those who registered after exposure to the interventions actually turned out to vote. Though we will be unable to identity which individuals registered on a LinkNYC kiosk, since treatment is randomized at the tract level we can treat all registrants in a tract as having been treated to some degree. The following hypotheses will be tested with data from both sources unless otherwise specified.

H1: Encouragement effect: Registration will be higher in tracts where the messages are shown compared to those in which they were not shown.

H2: Neighborhood-specific vs. countdown message: We will test the efficacy of two different messages. The first appeals to neighborhood identity (e.g., “Greenpoint added XXX new voters in 2017. Tap ‘register now’ on the tablet to join in”) and the second shows a dynamic countdown to the registration deadline (“06 days 04 hours 22 seconds until the voter registration deadline for the November general elections. Tap ‘register now’ on the tablet to join in”).

Heterogeneous effects: Subject to sample size constraints, we plan to explore the possibility of heterogeneous treatment effects by attributes of individuals’ context, as well as by individual characteristics. Are the effect of the messages moderated by local racial or socioeconomic context, or by individual characteristics? Since the literature is largely silent on these effects with respect to voter registration, this exercise is largely exploratory. Specific hypotheses are described below, and are two-sided unless otherwise specified.

H3a: Neighborhood identity: Appeals to neighborhood identity will be more or less effective in various neighborhoods.

H3b: Urgency appeal: The countdown message will be more effective as the registration deadline approaches.

H3c: Local context: Encouragement will vary in effectiveness by local neighborhood characteristics, such as racial and ethnic composition, measures of diversity or heterogeneity (e.g., Herfindahl index), and socioeconomic context (e.g., income, poverty, inequality, and education).

H3d: Individual attributes: Encouragement will vary in effectiveness by imputed race /ethnicity and age (from the voter file), and by imputed SES (based on usage of Aunt Bertha app, which connects users with social services).

H4: Subsequent voting: How likely to vote are people who registered during the messaging campaign in treated tracts vs. those who registered in control tracts? Are people who were exposed to the neighborhood identity message vs. the countdown message more or less likely to vote?

Analysis

The main outcome variable is the tract level registration rate, calculated as the number of new registrants divided by the number of people eligible to register. The numerator comes from either the number of people who register at LinkNYC kiosks, or from the number of people who register via any method during the study period. Subject to data availability we also calculate a registration rate out of the typical number of sessions on each kiosk per day.

H1: Encouragement effect: As a first cut we will conduct t-tests of the difference in means of the tract-level registration rate in treatment tracts (pooled) vs. control tracts. Since we blocked on demographics that correlate with voter registration rates prior to randomization, controlling for these variables shouldn't be necessary. However, we may want to control for pre-treatment registration rates, which we were unable to block on, as well as other tract-level variables including poverty status, educational attainment, age distribution, residential mobility, homeownership, and employment, in order to increase precision. To do so we will regress our outcome variables on treatment indicators and these controls, and conduct two-sided t tests on the estimated coefficients on treatments.

In order to account for upward trending registration rates (regardless of the registration campaign) as the election approaches, we will calculate the difference-in-differences from the NYC voter file: (registration rate in treatment tracts at t0 - registration rate in control tracts at t0) - (registration rate in treatment tracts at t1 - registration rate in control tracts at t1), where t0 includes the 22 days prior to 9/21 and t1 includes 9/21 – 10/12 (same number of days in t0 and t1).

H2: Neighborhood-specific vs. countdown message: We will regress tract level registration rates on indicators for which type of message was received as well as the controls described above.

H3: Heterogeneous effects: We will test whether the different appeals are more effective and where they were most effective using interaction terms in the regression of registration rates on treatment indicators.

H3a: To test whether appeals to neighborhood identity are more or less effective in various neighborhoods, we will include interactions between pre-intervention registration rates and the treatment indicator in our regression of the tract-level registration rate.

H3b: To test whether the countdown message is more effective as the registration deadline approaches, we will include an interaction between days until the registration deadline (e.g., 1-7 or 7-14) and the treatment indicator in the regression of tract-level registration rate. We will show that this relationship is robust to different specifications such as alternative bins of days to the deadline and a continuous measure (though we expect the latter to be underpowered).

H3c: Tracts can be characterized in terms of racial / ethnic makeup and socioeconomic status. These variables will be obtained or constructed from data from the U.S. Census and the American Community Survey. We will test for whether differences between tracts that received encouragement (pooled) versus those that did not, and whether differences between tracts that received the urgency vs. neighborhood identity messages, are moderated by these variables. To do so, we include interaction terms in the regression of registration rates on treatment indicators and context variables.

H3d: Race and ethnicity of individuals in the voter file will be imputed using methods detailed in Imai and Khanna (2016) using Census surname lists and Census block demographics. We will test for whether differences between individuals who received encouragement (pooled) versus those that did not, and whether differences between individuals that received the urgency vs. neighborhood identity messages, are moderated by these variables. We calculate treatment effects among racial and ethnic subgroups using the methods described in H1 above.

Age will be calculated using birth dates in the voter file. We will ask whether the efficacy of the message campaigns is associated with age. We will also report registration statistics among individuals who use “Aunt Bertha” apps, if given access to this data.

H4: Subsequent voting: After the November election, we will obtain an NYC voter file containing indicators for whether each registered individual in fact voted on Election Day. We use this to study whether those who registered during the campaign in treated areas are more or less likely to vote than those who registered during the campaign in control areas. We also compare tracts that received the urgency message to those that received the neighborhood identity message. We regress an indicator for whether the individual voted on treatment indicators and controls for individual characteristics (race, age, gender) and contextual variables, clustering standard errors at the tract level.

Other content: Finally, we plan to account for variation in *other content*, commercial and otherwise, displayed on LinkNYC kiosks during the campaign. We were unable to obtain this information prior to randomization and, at the time, of this writing, do not yet have data on the

extent to which other content varies by kiosk. Thus, we intend to strategize around how to deal with this once we have a sense of whether, and how much, this poses a threat to internal validity. Strategies may include controlling for other content in the analysis and/or dropping observations and/or conducting sensitivity analyses to determine the size of the threat.

Dealing with spillover:

One concern with our study design is the potential for spillover, which we expect to attenuate treatment effects. Some amount spillover from one treatment group to the next is unavoidable; as individuals move about the city, whether to commute to work, run errands, etc., they are likely exposed to multiple treatment conditions. This issue is further complicated by the fact that we do not have access to any information about kiosk registrants; we observe only counts linked to kiosks but do not know where individuals live in relation to the kiosk at which they registered. Eligible voters who reside anywhere in the state of New York can legally register on any kiosk.

In general, we assume that individuals who register do so in the tracts in which they live. We assume that exposure to kiosk content is greatest near one's place of residence, and thus registration is most likely to occur there. This assumption likely generates measurement error, but allows us to report associations between contextual variables (e.g., racial / ethnic composition) and treatment effects.

We have several strategies for diagnosing and dealing with the inevitable spillover post hoc. First, we will conduct sensitivity analyses by imputing potential outcomes under the "worst case" scenario of full spillover, as well as under partial spillover, and quantify the bias in on our average treatment effects resulting from this. Second, we obtain movement data on where people tend to live and work, combined with voter file data, to attempt to identify where spillover is most likely to occur. Finally, we subset to places where we think spillover is least likely (e.g., sparse tracts in outer boroughs) and calculate treatment effects within this subset.