San Diego Pothole Project:
Pre-Analysis Plan

David Schönholzer, Benjamin Vatter, Yifei Wang

January 2026

1 Overview and research questions

This study evaluates whether (i) simple, low-cost information about the City of San Diego's "Get-It-Done" (GID) service request system and (ii) exposure to recent pothole repairs affect administrative service request behavior, beliefs, and attitudes about City services and pothole repairs. The study combines a street-segment-level randomized controlled trial of door hangers (DH) with quasi-experimental variation in pothole fixes (PF), along with a separate door-to-door survey (D2DS) used to measure beliefs and attidues, and administrative outcomes for subsequent service uptake.

Key comparisons. The randomized variation is in DH saturation at the street-segment level (Full-DH, Partial-DH, Control-DH). Pothole fixes (PF) are *not* randomized, though we are familiar with the allocation mechanism and will atempt to exploit it for identification.

Main hypotheses.

1. DH may increase GID awareness, improve and increase future service requests (pothole and/or non-pothole) by reducing informational/transaction costs.
2. Households on segments with timely pothole fixes (PF; main specification: within 2 days) have more favorable views of government than those experiencing long lags in pothole repairs.
3. Among segments with recent pothole service requests, the effect of DH may differ depending on whether those requests are fixed quickly (PF).
4. A small monetary reward conditional on a successfully repaired pothole report increases the probability that a D2DS respondent submits a GID pothole service request, and the submission probability is (weakly) monotone in the reward amount.

2 Pilot

We piloted DH and D2DS logistics in one neighborhood, Linda Vista. Linda Vista is excluded from the study sample and all analyses.

3 Sampling frame

3.1 Universe and segment/address construction

We start from all street segments in the City of San Diego. Street segments are a natural unit because they are the unit used by the City for road repair planning.

1. Assign residential addresses to segments if the address lies within 40 meters of the segment.
2. Require each segment to have at least two addresses within 40 meters.
3. Restrict to addresses in primarily single-family residential zoning categories (currently `RS-1-2` through `RS-1-14`, plus `RM-1-1`), and further restrict to addressable records without explicit unit numbers, to ensure physical accessibility for DH and D2DS (apartment buildings may not be accessible).

3.2 Bundling

We group eligible segments into compact, walkable *bundles* of approximately 80 addresses each. We drop bundles whose total segment length is above 3km, making them very hard complete in the allocated interviewer walk time. A bundle typically consists of 1–19 segments (median 8 segments). Let j index bundles, i index segments, and t index Saturdays (field weeks).

4 Field design, assignment, and sampling schedule

4.1 Door hangers (DH): segment-level randomized saturation

When a bundle is selected for DH distribution, each segment in that bundle is randomized into one of three DH saturation conditions:

- **Full-DH** with probability 1/4: all addresses on the segment receive a door hanger;
- **Partial-DH** with probability 1/4: exactly half of addresses on the segment receive a door hanger (simple random sample without replacement);
- **Control-DH** with probability 1/2: no addresses on the segment receive a door hanger.

For most administrative analyses, we define $DH = 1$ for Full-DH and Partial-DH segments and $DH = 0$ for Control-DH segments. We additionally pre-specify saturation/spillover comparisons using separate Full and Partial indicators (Section 9.5).

4.2 Fieldwork schedule

There are six interviewers. Fieldwork occurs every Saturday starting January 10, 2026, through August 8, 2026 (i.e., every Saturday from January 10 onward during this period, except May 23, 2026 and July 4, 2026), yielding 29 completed field weeks.

4.3 Week 1 (January 10, 2026): DH only

On the first field day, we sample 24 bundles (4 per interviewer) for DH distribution only. Of these 24 bundles, 4 are sampled from bundles with at least one pothole patched in the preceding week (conditional), and 20 are sampled pothole-unconditionally from all available bundles. D2DS does not occur in Week 1. This ensures that, starting in Week 2, D2DS can be conducted in areas where DH may have been delivered in the prior week, and provides a sufficient pool of previously-sampled DH bundles (including the 4 conditional bundles) for D2DS reuse.

4.4 Weeks 2+ (January 17, 2026 onward): DH and D2DS

Each Saturday $t \geq 2$:

- **DH bundles:** sample one DH bundle per interviewer (6 total), from bundles not previously sampled for DH, with composition:
  - four out of six DH bundles have at least one pothole patched on one of their segments during the seven days prior to the Saturday field day ("this week"). This bundle selection criterion ensures sufficient variation in exposure to recent pothole patches and (non-random) time-to-fix, while recognizing that selected bundles will contain a mix of segments: some with pothole patches, some with no pothole activity, and some with pothole reports not yet serviced. In case there are less than 4 bundles with pothole patches from this week, we sample unconditionally to fill in the remaining bundles.
  - the remaining two out of six DH bundles are sampled pothole-unconditionally. We do this so as to not skew the sample entirely towards pothole-ridden bundles.
- **D2DS bundles:** sample one D2DS bundle per interviewer (6 total), from bundles not previously sampled for D2DS, with composition:
  - four of the six D2DS bundles correspond to the four DH bundles with at least one pothole patched drawn last week;
  - the remaining two D2DS bundles are drawn from bundles that have not received any DH but have had potholes patched in the *preceding* week. We focus D2DS on pothole-ridden bundles so as to maximize the sample of pothole-affected households in the survey.
  - in either case, if there are not enough pothole-conditional bundles, we again sample pothole-unconditionally to fill in the remaining bundles.

4.5 D2DS pothole-report incentive (RI)

A subset of the D2DS instrument tests whether a small monetary reward for *successfully* reporting a pothole through the City's GID platform increases the probability that a household engages with the system.

Mechanism. Each D2DS respondent who completes the consent module is offered a printed business card containing: (i) a study-issued unique ID, (ii) a QR code linking to a study-run landing page (`sdpotholes.org`), and (iii) a printed reward amount. To claim the reward, the respondent submits a pothole service request through the City's GID platform, then enters their City Service Request ID and email at the landing page. The study team matches each submission against the City's internal repair records (`notification_activities.csv`, `POTHOLE PATCHED (EA)`; see Section 6.2). If the submission is matched and the underlying pothole is patched within 60 days of submission, a gift card of the printed amount is emailed to the respondent. Submissions outside the 60-day window or unmatched to a repair record are ineligible for payment. This conditional-on-repair structure was a precondition of the City's approval and serves as a fraud-prevention measure.

Randomization. Each card's reward amount is pre-randomized as follows: with probability 1/2, the card is a \$0 *placebo*; with probability 1/2, the card displays a positive reward drawn uniformly from $\{\$5, \$10, \$15, \$20\}$ (so the marginal probability of each positive amount is 1/8). Cards with a printed amount of \$0 are visually indistinguishable from positive cards, advertise the GID submission process, and route to the same landing page, but do not pay out. Cards are pre-printed in fully shuffled order; interviewers hand out cards in order to consenting respondents, so the realized randomization is at the respondent level conditional on consent. Let $RI_h \in \{0, 5, 10, 15, 20\}$ denote the reward amount printed on respondent h's card, and let $POS_h \equiv \mathbf{1}\{RI_h > 0\}$.

Unit of randomization and identification. RI is randomized within D2DS bundles, conditional on a respondent reaching the incentive module. Identification of the RI effect therefore comes from within-bundle, across-respondent variation in RI_h, holding fixed the DH and bundle-week environment.

5 Non-randomized pothole fixes (PF)

Pothole fixes are not randomized. We exploit quasi-idiosyncratic variation in time-to-fix that arises from SDTD's FIFO logic and routing constraints, conditional on pavement condition index (PCI) categories and other pre-specified controls.

Exploratory note on instruments (not pre-registered as main analysis). We may explore instrumental-variable approaches for PF timing (e.g., leave-segment-out measures of crew workload or nearby fix intensity). These IV approaches are not part of the main pre-registered estimands. An IV design is hard to anticipate because of the nature of sampling, which is a function of the frequency and distribution of potholes.

6 Data sources and units of observation

6.1 Administrative data (GID)

We use City administrative data from the GID platform. Because service requests are difficult to assign to individual addresses, the primary administrative unit of observation is the street segment-week. We may attempt such a linking, but it would likely be subject to considerable noise.

6.2 Pothole repair data for sampling

For bundle sampling, we use pothole repair data from the City's internal notification activities system (`notification_activities.csv`), filtered to records with activity code "POTHOLE PATCHED (EA)". These are repair work orders initiated by the City, with dates corresponding to when repair work was notified or initiated (`NOTIFICATION_DATE`), rather than citizen service request dates. We will also use the publicly available service request data.

Bundles are selected based on whether at least one segment within the bundle has at least one pothole patched during the relevant time window. Importantly, bundles selected through this criterion contain a heterogeneous mix of segments: some segments have pothole patches, some have no pothole-related activity, and some may have pothole reports (either citizen-initiated or City-identified) that have not yet been serviced by the field date. This selection strategy ensures exposure to recent pothole repair activity while maintaining representativeness across segments with varying pothole conditions.

6.3 GIS data and segment covariates

Segment length and PCI are obtained from City street segment data. We also use GIS-based road classification (e.g., arterial/collector vs. local) and, where available, associated speed-class proxies. [check if this kind of data exists].

6.4 Survey data (D2DS)

D2DS outcomes are measured at the respondent (address) level. In regressions, D2DS outcomes will be analyzed with fixed effects at the bundle-by-week level and controls at the segment level. In case we are able to obtain address-level outcomes at a later stage, we will be able to run address-specific regressions exploiting the partial DH treatment.

7 Outcomes

7.1 Primary administrative outcomes (segment-week)

For each segment i and week t, we construct:

- pothole service requests (count and indicator),
- non-pothole service requests (count and indicator),

- any service requests (count and indicator).

We also construct cumulative outcomes for each segment over horizons of 1 week, 1 month, and 3 months following the segment's DH week.

7.2 Primary D2DS outcomes

We pre-specify the following City Service Trust/Quality Index: the average of standardized survey items $G1$ (overall satisfaction), $G2$ (tax dollars used efficiently), and $G3$ (tax dollars used fairly across neighborhoods), coded so higher values indicate more positive perceptions.

All remaining D2DS items (pothole salience, willingness-to-pay, response-time beliefs, neighbor beliefs, etc) will be taken as analyzed as primary outcomes as well.

7.3 Primary D2DS outcome: GID submission (RI sample)

For each D2DS respondent h who is offered an RI card (Section 4.5), define $Submit_h$ as an indicator equal to 1 if the respondent's unique card ID appears at `sdpotholes.org` paired with a valid GID Service Request ID within 60 days of card issuance, and 0 otherwise. As secondary outcomes we construct: (i) $SubmitMatched_h$, equal to 1 only if the SR ID is matched to the City's request data; (ii) $Repaired_h$, equal to 1 only if the matched request is patched within 60 days; (iii) $DaysToSubmit_h$, the number of days from card issuance to first submission (right-censored at 60).

8 Key explanatory variables

8.1 Door-hanger event time

Let t_i^* denote the Saturday in which segment i's bundle is sampled for DH (the segment's "DH week"). For segments in DH-sampled bundles (including Control-DH segments), define event time $s = t - t_i^*$ and an indicator DH_{is} for being assigned DH in that DH week.

The event-time window is $s \in \{-4, -3, -2, 0, 1, \ldots, 12\}$, omitting $s = -1$ as the reference period.

8.2 Pothole service request (PSR)

Define PSR_{it} as an indicator equal to 1 if segment i has at least one pothole service request created during the seven days preceding Saturday t (the "request week" window), and 0 otherwise.

8.3 Pothole fix within x days

For segments with at least one pothole service request in the relevant request window, define $PF_i(x)$ as an indicator equal to 1 if at least one such request is fixed within x days of request creation, and 0 otherwise. A request is considered "fixed" if it can be matched to a pothole patch record in the City's internal notification activities data (`POTHOLE PATCHED (EA)`). The number of days to fix is calculated from the public service request creation date to the corresponding repair completion date. In the main static specification below, we set $x = 2$. As sensitivity analyses, we will re-estimate the same specification for each $x \in \{1, \ldots, 7\}$. In cross-sectional analyses involving pothole fixes, PF_i (without an argument) denotes whether at least one such request has been fixed by the time outcomes are measured.

9 Estimation

9.1 Main administrative event study: DH dynamics

Our baseline specification is an event-study style regression of an administrative outcome:

$$Y_{it} = \alpha_i + \gamma_{j(i),t} + \sum_{s \in \{-4,-3,-2,0,1,\ldots,12\}} \beta_s \, DH_{is} + \varepsilon_{it}, \tag{1}$$

where α_i are segment fixed effects and $\gamma_{j(i),t}$ are bundle-by-week fixed effects. Coefficients β_s trace the dynamic effect of DH relative to the omitted reference period $s = -1$. [Important to check: how many bundles have only one segment? These will be dropped in this specification! Also, can we have a histogram of the number of segments in our bundles? And maybe also the treatment/control split of segments within bundles: i.e. a histogram of the share of treated segments for each bundles. Alternatively, maybe we should use PCI-bin-by-week, e.g. good-by-week-2, or fair-by-week-3.]

9.2 DH event-study robustness with external controls

As a robustness check, we will also estimate a stacked event study in which each DH-treated cohort is compared only to control segments in the same PCI category drawn from bundles that are never sampled for DH during the study period and are outside the treated segment's bundle. This specification is intended to reduce the scope for within-bundle spillovers.

We do not pre-specify a separate PF event study. PF enters the pre-specified analysis only through the static cross-sectional specifications below.

9.3 Static cross-sectional specification (admin and D2DS)

For administrative outcomes aggregated over horizons (1 week, 1 month, 3 months) and for D2DS outcomes, we estimate:

$$Y_i = \delta_{j(i)} + \beta_1 DH_i + \beta_2 PF_i(2) + \beta_3 DH_i \times PF_i(2) + X_i' \mu + \varepsilon_i, \tag{2}$$

where $\delta_{j(i)}$ are bundle fixed effects and X_i is a vector of pre-specified controls. For administrative outcomes, Y_i is the cumulative count over a post-DH horizon (1 week, 1 month, or 3 months) measured from the segment's DH week. For D2DS outcomes, DH_i indicates whether the segment was assigned DH in the week preceding the D2DS survey. These regressions are estimated on the sample of segments with at least one pothole service request in the relevant request window, so that $PF_i(2)$ is well-defined. As sensitivity analyses, we will re-estimate the same specification for each $x \in \{1, \ldots, 7\}$, replacing $PF_i(2)$ with $PF_i(x)$.

9.4 Heterogeneity

We pre-specify two primary heterogeneity dimensions, both defined at the Census block-group level using the most recent ACS 5-year estimates and matched to each segment via its centroid:

- **Below-median household income:** $LowInc_i \equiv \mathbf{1}\{MedHHInc_{bg(i)} < m_Y\}$, where m_Y is the median of $MedHHInc_{bg(i)}$ taken over the universe of segments eligible for the study (sampling frame, Section 3).
- **Below-median share white:** $LowWhite_i \equiv \mathbf{1}\{ShareWhite_{bg(i)} < m_W\}$, where m_W is the median of $ShareWhite_{bg(i)}$ over the same eligible-segment population.

Both cut points are fixed ex ante from the eligibility frame, before unblinding, and are applied identically to the administrative and D2DS samples. For each primary outcome, we re-estimate equation (2) interacting DH_i, $PF_i(2)$, and $DH_i \times PF_i(2)$ with $LowInc_i$, and (in a separate regression) with $LowWhite_i$.

Beyond these two pre-specified cuts, we may explore additional heterogeneity dimensions; any such analyses will be clearly labeled as exploratory.

9.5 Within-segment saturation / spillover analysis (D2DS sample)

To assess within-segment spillovers under partial saturation, we use the D2DS sample and estimate:

$$Y_h = \delta_{j(h)} + \theta_F \mathbf{1}\{FullDH_{ih}\} + \theta_P \mathbf{1}\{PartialDH_{ih}\} + X'_{i(h)}\mu + \varepsilon_h, \tag{3}$$

where h indexes D2DS respondents and the right-hand-side saturation indicators are defined at the segment level (control omitted). We will also report the linear combination $\theta_P - \frac{1}{2}\theta_F$ as a simple diagnostic for departures from proportional scaling with saturation, consistent with within-segment spillovers or nonlinearities.

9.6 Reporting-cost incentive (RI) regressions

On the sample of D2DS respondents offered an RI card, we estimate three pre-specified regressions of $Submit_h$ on reward indicators with bundle-by-week fixed effects $\delta_{j(h),t}$ and the same control vector $X_{i(h)}$ used in the static D2DS specification.

The primary specification is a binary contrast,

$$Submit_h \; = \; \delta_{j(h),t} \; + \; \pi^{POS} \, POS_h \; + \; X'_{i(h)}\mu \; + \; \varepsilon_h, \tag{4}$$

where $POS_h = \mathbf{1}\{RI_h > 0\}$. By construction, POS_h has marginal probability 1/2, so this contrast carries the most power.

10 Controls

We pre-specify two control sets.

Primary (minimal) controls. Segment length and PCI-category fixed effects.

Secondary (augmented) controls (robustness). In addition to the primary controls: seasonally-adjusted historical pothole service requests (last four years), road class / speed-class proxies (e.g., arterial vs. local), and continuous PCI where available. We may also merge Census neighobrhood characteristics to segments as further controls.

No household-level controls pre-specified. All controls are defined at the segment level; the $X_{i(h)}$ vector in the D2DS and RI specifications is the segment-level control vector indexed by the respondent's segment. We do not pre-specify household-level covariates as controls. We may attempt to merge parcel-level characteristics (such as hedonics or prices) as controls.

11 Inference

Event studies. Standard errors are clustered at the segment level to account for serial correlation in the segment-week panel.

Cross-sectional regressions. For our primary segment-level cross-sectional regressions, we report heteroskedasticity-robust standard errors. For D2DS individual-level regressions, we report standard errors clustered at the segment level.

Multiple inference. We do not plan to adjust standard errors or p-values for multiple testing. In settings with many estimates (event-time coefficients; PF(x) sensitivity), we will present full coefficient paths/curves and emphasize pre-trends diagnostics (event study) and sensitivity plots (PF(x)) rather than selecting individual coefficients ex post. Since our set of survey outcomes is relatively small, we prefer to be transparent about the full set of results across all survey responses and interpret inference accordingly.

12 Secondary specifications

1. **Across-segment spillovers within bundles:** re-estimate key specifications replacing bundle-by-week fixed effects with council-district-by-week fixed effects.
2. **D2DS effects on administrative outcomes:** estimate event studies on D2DS event-time indicators and add a D2DS indicator (and interactions) to the static administrative specification.
3. **Heterogeneity:** explore heterogeneous effects by PCI category (good/fair/poor/etc.).

13 Ethics / IRB

This study has been reviewed and determined exempt by the University of Southern California IRB (Study ID UP-25-00768; exemption determination under 45 CFR §46.104(d)(2); approval dated Oct 7, 2025) and by MIT COUHES (E-7213; exempt categories 2 and 3; submission dated Oct 10, 2025).

14 Registration and deviations

This document is intended for AEA RCT Registry pre-registration. Any deviations from the design or analysis plan described here will be documented in an amendment prior to outcome analysis when feasible, and clearly labeled in any dissemination materials.