Pre-Analysis Plan

Does Cost Information Shift Carsharing Adoption?

A Randomized Field Experiment in Bergen, Norway

1. Overview and Research Question

This pre-analysis plan (PAP) describes the design, outcomes, hypotheses, and statistical methods for a randomized field experiment. The experiment aims at testing whether an online cost-comparison calculator, comparing the total cost of car ownership with the costs and benefits of carsharing, can impact carsharing adoption behavior and preferences among households in Bergen, Norway.

Summary

This project examines whether providing tailored information about own economic costs of car ownership and the potential costs and practical benefits of carsharing can influence carsharing adoption behaviour and preferences among car-owning households in Bergen, Norway. Approximately one third of the city's adult population is invited to participate in a short survey, where the identified target group is randomly assigned to either a treatment or control group.

The intervention consists of an online calculator that compares the total cost of owning a private car with the cost and convenience of carsharing. Participants in the treatment group are invited to use the calculator, while the control group instead answers unrelated questions designed to equalize time and effort. Finally, both groups are offered the opportunity to request a discounted coupon to try carsharing.

We implement a randomized encouragement design in which car owners are randomly exposed to the calculator. Using requesting carsharing coupons as the primary outcome, random assignment identifies the effect of the offer, and a two-step instrumental variable analysis isolates the effect among individuals whose exposure to calculator content was changed by the offer.

In addition, we study how the information provided by the calculator affects preferences about carsharing and car ownership using a within-subject survey design. Finally, we examine longer-term behavioural responses by tracking whether individuals redeem the coupons and try carsharing, and whether the intervention leads to changes in car ownership status over time.

Central research question: *To what extent does tailored cost and convenience information comparing car ownership with carsharing affect (i) coupon-requesting behavior, (ii) coupon redemption, (iii) carsharing usage, (iv) stated preferences about carsharing and car ownership, and (v) actual car ownership status?*

2. Study Population and Eligibility

2.1 Sampling frame

The sampling frame consists of all adults aged 18–75 in Bergen municipality reachable by email: approximately 75,000 individuals, accounting to about one third of the city's adult population. The sample is provided by the Norwegian tax authority.

2.2 Eligibility criteria

Respondents who complete the baseline survey are included in the analysis sample if they satisfy all three conditions:

1. Do not hold a disabled parking permit (HC card).
2. Either (i) the household owns one or more cars, or (ii) the household plans to buy a car within the next 12 months (including replacing a car or acquiring an additional car).
3. Are not frequent carsharing users at baseline: defined as having used a carsharing service fewer than 3 times in the past 12 months.

2.3 Randomization timing

To prevent treatment-induced selection into the sample, random assignment follows the baseline survey and eligibility determination. The sequence is: (1) email invitation, (2) baseline survey including all eligibility items, (3) eligibility check, (4) random assignment within strata, (5) delivery of treatment or control experience. All baseline measures are therefore pre-treatment.

2.4 Stratification

We stratify on **gender** and **access to carsharing**, defined at the number of carsharing cars available in 2025 in all (142) postcodes in Bergen. Before randomization, postcodes are classified into two strata split at the median: high and low carsharing access (the median is in the low category). The median number of cars available in each postcode is 4. This variable is frozen before any data collection and is not altered during the study.

3. Experimental Design

3.1 Unit of randomization

The unit of randomization is the individual. Within each of the four strata (high/low carsharing access × female/not female), eligible individuals are randomly assigned 55% to treatment and 45% to control using a pre-committed randomization algorithm.

> ▶ *Two adults from the same household may both receive an invitation and be assigned to different arms, since randomization is at the individual level. Their outcomes are likely correlated through shared car ownership and joint household decisions. The clustering decision rule in Section 7.1 accounts for this.*

3.2 Treatment arms

Treatment (calculator offer): After completing the baseline survey, treated households receive a link to an online calculator delivering four types of information:

(a) a generic cost comparison between car ownership and carsharing shown to all visitors;

(b) a personalized total cost report for own vehicle, obtained by entering own license plate number, or alternatively chose from six representative vehicles;

(c) information on the geographic availability of carsharing in Bergen; and

(d) practical information on how carsharing works.

Control: Control households do not receive a calculator link. Their survey is extended with additional non-treatment-related questions (about freight transport and investment in infrastructure) to equalize total response effort across arms.

> ▶ *The treatment is the information. The primary estimand (LATE) captures the effect of this information on those who completed the calculator. Mechanism analyses (Section 11) explore which components matter.*

3.3 Coupon offer

Both arms are offered an identical coupon for a discounted carsharing trial. The coupon appears at the end of the control-group survey and on the final page of the treatment-group calculator. Content and value are as such made to be as similar as possible, across arms.

3.4 Preference and belief elicitation

Both groups are asked preference (willingness to try carsharing, Y_try_cs;) and belief (car ownership necessity, Y_need_car) questions twice. For both arms, T0 is measured at the beginning of the baseline survey, right after the eligibility questions. T1 is measured at the end: for the control group, at the end of the baseline survey after the effort-equalization questions; for the treatment group, at the end of the calculator after viewing the cost report. This symmetric design means both arms provide a T0 and a T1 measurement, enabling a causal difference-in-differences analysis of belief change (see Section 7.3).

3.5 Design summary

Group	Eligibility filter	Baseline survey	Treatment	Coupon offer
Control	No HC permit; owns car or plans to buy; not frequent CS user	Beliefs + preferences (T0)	**None:** Extra questions to equalize effort (freight transport)	Offered in survey after completion, Beliefs + preferences re-measured (T1)
Treatment	Same as control	Beliefs + preferences (T0)	**Calculator and information:** Cost comparison + personalized car cost report + geographic availability + practicalities of carsharing	Offered on final calculator page; Beliefs + preferences re-measured (T1)

4. Outcome Variables

4.1 Primary outcome

Our main outcome variable measures whether respondents request a coupon to try carsharing.

- Y_req = 1 if the participant submits the coupon request;
- Y_req = 0 otherwise, including for non-completers.

Y_req is recorded in the survey program (Walr) for control respondents and from website logs for treatment responders. All userid and discount codes will be saved in a dedicated file for both groups.

Outcome	Variable	Definition	Unit	Window	Missing data rule
Coupon requested	Y_req	=1 if participant submits coupon request form, 0 otherwise. Recorded for all eligible respondents regardless of survey completion or calculator engagement, which means that those that do not complete the survey or calculator gets 0 on this variable.	Binary	Same session or within end of the survey period.	No log event = 0

4.2 Secondary and exploratory outcomes

Outcome	Variable	Definition	Unit	Window	Missingness rule
Willingness to try CS (secondary)	Y_try_cs	0–100 slider: 'How likely are you to try carsharing in the next 12 months?' Asked at T0 (both arms) and T1 (treatment completers only).	0–100	T0: both arms (survey). T1: treatment completers (end of calculator); control completers (end of survey)	Non-response = 50% (default value)
Car ownership necessity (secondary)	Y_need_car	Likert 1–5: 'Owning a car is a necessity for me in the foreseeable future.' Asked at T0 (both arms) and T1 (treatment completers only).	1–5	T0: both arms (survey). T1: treatment completers (end of calculator); control completers (end of survey)	Non-response = 3 (default value)
Coupon redeemed (exploratory)	Y_red12m	=1 if unique coupon code redeemed with any carsharing partner within	Binary	8 months	Admin merge failure = missing.

Outcome	Variable	Definition	Unit	Window	Missingness rule
		31 december 2026. Contingent on data-sharing agreements with partners, very likely.			Robustness: treat unmatched as 0.
Any CS usage (exploratory)	Y_use12m	=1 if any trip/booking in partner admin data within 12 months (including the coupon redemption). Contingent on data availability, unsure.	Binary	1 year	Not reported if data are not collected
Car ownership at 12m (exploratory)	Y_car_own12m	≥1 if individual owns ≥1 car at 12 months; admin registry for those who write their number plate (both treatment and control).	Binary and n.of cars owned	12m	We can only observe outcomes if respondents disclose their car number plate.
Belief error (descriptive)	Y_belief_err	Y_belief0 minus calculator-implied annual cost. Used in mechanism analyses only.	NOK/year	After report	Missing if no report completed.

> ▶ *Y_try_cs and Y_need_car are both secondary belief outcomes (H2), analyzed using ANCOVA, primary). Both are observed symmetrically in both arms at T0 and T1. Y_red12m, Y_use12m, and Y_car_own12m are exploratory and contingent on data availability. Y_belief0 and Y_belief_err are used in mechanism analyses only and carry no confirmatory status.*

4.3 Exposure and compliance variables

The following variables are constructed deterministically. They describe treatment take-up and are the basis for the primary analysis. The primary IV take-up variable is D_report. D_personal is a prespecified robustness variant.

Variable	Name	Definition (deterministic from server log)
Visited landing page	D_click	=1 if calculator landing page URL accessed by this participant ID
Entered vehicle	D_input	=1 if license plate entered OR one of the 6 generic vehicles selected
Viewed cost report	D_report	=1 if participant reaches the final results page (report-rendered server event). PRIMARY IV take-up variable.
Personalized report	D_personal	=1 if report is based on license plate lookup (not a generic vehicle). Secondary IV robustness variant.

4.4 Handling repeated calculator sessions (pre-committed rule)

A session is defined as a continuous sequence of page views in the website calculator. Multiple sessions by the same participant are possible when the user arrives to the final page and starts again.

We will consider sessions where the user has arrived to the final page: D_report=1. If there are multiple finalized sessions within the same user, we will use the following rule.

Primary rule: The calculator result (savings figure) used in mechanism analyses is drawn from the first completed report in the first session. The rationale is that people may want to "play" with the calculator to see how it changes and may want to try different cars if they have several.

Robustness: secondary analyses use the last session and a random one respectively. These are prespecified and do not alter primary conclusions.

5. Estimand Hierarchy

Label	Estimand	Analysis	Interpretation
Primary (H1): Voucher (ITT)	E[(Y_req|Z=1)−(Y_req|Z=0)]	A: ITT OLS/LPM	Causal effect of being offered the calculator on coupon requests. Full eligible sample. Causal interpretation only if the attrition is not different between groups (see section 9 about attrition).
Primary (H1): Voucher (LATE)	E[Y_req (1)- Y_req(0) | compliers]	B: 2SLS — random assignment Z instruments D_report; both arms used	Causal effect of viewing the calculator on coupon requests.
Secondary (H2&H3): beliefs (ANCOVA)	E[Y_try_cs_T1 | completer, T=1] vs. control E[Y_owncar_T1 | completer, T=1] vs. control	C: ANCOVA — regress Y_T1 on treatment indicator and Y_T0 as covariate.	Causal effect on end-of-session preferences and beliefs for the person who completes the calculator.

▶ *Why can ITT on Y_req possibly not be causal? In this design, the coupon is offered at the end of the calculator in the treatment arm, and at the end of the survey in the control arm. Treatment respondents who do not complete the calculator never see the coupon and receive Y_req = 0 by friction rather than by revealed preference. Example: With expected calculator completion of approximately 30%, around 70% of the treatment group is structurally excluded from the possibility of Y_req = 1. Under these conditions, the ITT on Y_req could become a negatively biased estimate of the calculator's effect on coupon interest: it captures both the effect of the calculator content and the effect of the additional friction, inseparably. The LATE — estimated by instrumenting calculator completion with random assignment — isolates the effect of actually viewing the report among those induced to do so, which is the scientifically meaningful quantity. The ITT is still reported as a descriptive statistic alongside the LATE for transparency, but confirmatory inference is based on the LATE.*

6. Pre-Registered Hypotheses

All hypotheses are specified before data collection. We distinguish three levels: confirmatory primary (one hypothesis, formally tested, strong claims if supported); secondary (one hypothesis, causally interpreted with stated identification caveat); and exploratory (no confirmatory claims).

6.1 Primary hypothesis: LATE on coupon requests (H1)

H1 : Viewing the report increases the probability of requesting a carsharing coupon.

Analysis A and B, see below.

6.2 Secondary hypothesis: willingness to try carsharing and belief about car ownership (H2 H3)

H2: Exposure to the calculator increases willingness to try carsharing (Y_try_cs).

H3: Exposure to the calculator decreases the belief that owning a car is a necessity (Y_need_car).

Both outcomes are measured at T0 (start of baseline survey, both arms) and T1 (end of calculator for treatment completers; end of baseline survey for control completers).

Analysis C, see below.

6.3 Exploratory behavioral hypotheses (H4–H6)

H4 (exploratory): Being offered the calculator increases the probability of redeeming a carsharing coupon within 8 months. Contingent on data availability from carsharing partners.

H5 (exploratory): Being offered the calculator increases the probability of any carsharing usage within 12 months. Contingent on data availability.

H6 (exploratory): Being offered the calculator reduces the probability of owning a car at 12 months. The experiment is only powered to detect effects of approximately 4 percentage points or larger.

> ▶ *Whether H3, H4, and H5 can be tested depends on data access from carsharing partners and administrative registers. This will be documented in the deviations log (Section 14) once confirmed.*

6.4 Exploratory heterogeneity hypotheses (H7)

We test whether the treatment effect on Y_req (H1) differs across subgroups defined by baseline characteristics. All moderators are pre-specified and measured before treatment. The study is not powered to detect subgroup differences — these are hypothesis-generating. All tests are twosided.

H7a: The effect is larger for women who are the primary household decision-maker for car matters.

H7b: The effect is larger among households planning to buy a car in the next 12 months.

H7c: The effect is larger among households with more than one car.

H7d: The effect differs among respondents anticipating a major life change (moving, new job, new child) in the next 12 months. (We test the different life events separately)

H7e: The effect is different between owners of diesel or petrol cars and EV owners.

H7f: The effect is larger among respondents who find parking near their home difficult.

H7g: The effect is smaller among respondents uncomfortable using apps, because a practical barrier to acting on the information remains.

H7h: The effect is larger among respondents who already use public transport or active modes as their main commute mode.

H7i: The effect is larger among lower-income households.

H7j: The effect is larger among those who insert their number plate because they get more tailored information.

6.5 Mechanism hypotheses (H8, exploratory)

These analyses are conducted within the treatment group only, using treatment completers. They are descriptive and not causally identified. They explore the pathway from calculator content to belief and preference change.

H8a (mechanism): The magnitude of financial savings shown in the calculator is positively correlated with the shift in Y_try_cs and negatively correlated with Y_need_car from T0 to T1 and positively correlated with Y_req.

H8b (mechanism): The gap between the respondent's baseline cost estimate and the calculator's implied annual cost (Y_belief_err) is positively correlated with the shift in Y_try_cs and negatively correlated with Y_need_car from T0 to T1 and positively correlated with Y_req. This tests whether surprise — not just information — is the active ingredient.

> ▶ *H8a–b are descriptive dose-response analyses within treatment completers. They have no counterfactual and do not identify treatment effects.*

7. Empirical Methods

7.1 Coupon requests (H1)

We estimate the effect of the treatment on the outcome variable "coupon request" (Y_req) to answer H1, in two ways:

- **Approach A:** The ITT captures the effect of **offering** the carsharing calculator and associated communication on coupon request behavior. This estimand reflects the full intervention as implemented, including all behavioral channels such as attention, salience, framing, and navigation frictions. As such, it could be a policy-relevant effect of scaling the intervention in practice.

  However, the ITT can conflate the effect of the calculator content with the effect of the different friction required to reach the coupon, if attrition is significantly different between treatment and control. Under low take-up, the ITT on Y_req could become negative even if the calculator strongly increases coupon interest among those who complete it.
- **Approach B:** The LATE — estimated by instrumenting calculator completion (D_report) with random assignment — isolates the effect of actually viewing the cost report for those induced to do so by the offer. LATE estimated via 2SLS using random assignment as the instrument for D_report.

- **Choice rule, analysis based on 9.2:**
    - **If attrition is different** between control and treatment: The ITT is reported as a descriptive and does not carry causal interpretation. LATE is the primary approach.
    - **If attrition is not different** between control and treatment**:** The ITT is interpreted as causal. Then LATE and ITT answer slightly different research questions and carry the same wight.

A: ITT

Estimand: Intention to treat ITT. Both arms are used. The control group, which cannot access the calculator, provides the counterfactual.

$$Y\_req_i = \alpha + \tau \cdot T_i + \delta\_s(i) + X_i'\beta + \varepsilon_i$$

Where Y_req_i is coupon requested, T_i is the treatment assignment indicator, δ_s_i are strata fixed effects (always included), and X_i is a prespecified baseline covariate vector for precision improvement. The estimator is OLS / linear probability model (LPM). Strata fixed effects control for all time-invariant, unobserved differences across strata.

Standard errors and clustering: The unit of randomization is the individual.

Two adults from the same household may be assigned to different arms. Following Abadie et al. (2023), we cluster at the level where non-independence actually arises — the residential address — but only if household overlap in the achieved sample is non-negligible. We precommit to the following data-driven rule, applied before any outcome data are examined:

1. Compute the mean number of respondents per residential address in the achieved analysis sample.
2. Test whether this mean is significantly greater than 1 using a one-sided t-test ($\alpha = 0.05$).
3. If the test rejects: cluster standard errors at the residential address level (primary specification).
4. If the test does not reject: use heteroskedasticity-robust standard errors (primary specification).
5. In all cases: report a deduplicated sensitivity analysis retaining one randomly selected respondent per residential address.

> ▶ *Rationale (Abadie et al. 2023): if mean household size in the sample is not significantly greater than 1, the design effect from address clustering is trivially close to 1 regardless of the intra-household correlation coefficient, and clustering adds no value. If household overlap is non-negligible, address-level clustering directly targets the source of non-independence from shared car ownership decisions.*

Covariate selection. The primary specification uses Post-Double-Selection (PDS) LASSO (Belloni, Chernozhukov, and Hansen 2014) to select from the prespecified covariate list below, with stratum fixed effects always forced in. An OLS regression with the full covariate set is reported in the appendix (Lin 2013 specification).

Prespecified covariate list (all measured before treatment):

1. Age (continuous and squared)
2. Gender (female indicator)

3. Household car ownership count (0, 1, 2, 3+)
4. Plans to buy a car in the next 12 months (binary)
5. Self-reported carsharing accessibility (5-point scale)
6. Postcode-level carsharing availability (stratification variable, continuous share)
7. Baseline belief: estimated annual car ownership cost (NOK)
8. Baseline preference: likelihood of trying carsharing (0–100%)
9. Baseline preference: car ownership necessity (1–5)
10. Indicator for anticipated major life event in next 12 months
11. Household decision-maker for car (self / partner / joint)

B: LATE, IV (H1, primary)

Estimand: the local average treatment effect (LATE) for compliers: the average effect of viewing a personalized cost report among those whose viewing status was changed by the offer. Both arms are used. The control group, which cannot access the calculator, provides the counterfactual.

Instrument validity: (i) Relevance: assignment substantially increases D_report, tested via first-stage F-statistic (below 10 flagged as weak). (ii) Exogeneity: assignment is random. (iii) Exclusion restriction: assignment affects outcomes only through report viewing; control respondents are not sent to the calculator. (iv) Monotonicity: no defiers assumed; share of never-takers in treatment arm reported as diagnostic.

First stage:

$$D_i = \pi_0 + \pi_1 \cdot Z_i + \delta_{s(i)} + X_i{}'\kappa + u_i$$

Second stage:

$$Y_i = \beta_0 + \beta_1 \cdot \widehat{D}_i + \delta_{s(i)} + X_i{}'\gamma + e_i$$

where $Z_i = T_i$ (assignment), $D_i = D_{report}$ (endogenous take-up), and $\widehat{D}_i$ is the fitted value from the first stage. β_1 is the LATE, interpreted as applying to compliers only. Standard errors follow the precommitted clustering rule.

7.2 Willingness to try carsharing and belief about car ownership (H2 & H3)

Because both arms have T0 and T1 measurements, we can use the control group's change (T1−T0) as the counterfactual for what beliefs would have done absent the calculator. The treatment group's change (T1−T0) reflects the calculator effect plus any pure survey-taking effect. The difference between these two changes isolates the calculator's effect by netting out survey-taking effects that apply to both arms equally.

C: ANCOVA

Estimand: the effect of being offered the calculator on belief change, for the type of person who completes the calculator. Both Y_try_cs and Y_need_car are analyzed using the same model.

Primary model: ANCOVA is the recommended approach for pre-post experimental designs (Frison and Pocock 1992; McKenzie 2012). It regresses the endline score directly on the treatment indicator and the baseline score as a covariate.

$$Y_{i,T1} = \alpha + \tau \cdot T_i + \gamma \cdot Y_{i,T0} + \delta_{s(i)} + X_i'\beta + \varepsilon_i$$

$Y_{i,T1}$ is the endline belief score; T_i equals 1 for the treatment arm and 0 for control; $Y_{i,T0}$ is the baseline score measured at the very start of the survey; $\delta_{s(i)}$ are stratum fixed effects (always included); and X_i are additional baseline covariates for precision.

The coefficient τ is the treatment effect: how much higher (or lower) the endline belief score is in the treatment arm compared to the control arm, after adjusting for where each person started at baseline. By including $Y_{i,T0}$ on the right-hand side, we net out pre-existing differences in beliefs between arms. τ therefore captures only the within-person change attributable to treatment.

D: Different attrition

If attrition is different between control and treatment we will do sample and IPW. The model is estimated on all survey completers in both arms. T1 is only observed for treatment respondents who reached the belief questions in the calculator; those who never clicked the link or dropped out before the questions have no T1. If we find that there are large, systematic differences in completion, we will use inverse probability weights (IPW) to reweight the sample. This corrects for observable differences between completers and non-completers. The unweighted estimate is always reported alongside as a robustness check. Residual bias from unobserved motivation cannot be ruled out and is stated whenever H2 and H3 results are reported.

Also reported: (i) raw mean T1−T0 by arm (ii) ANCOVA with full covariate set, only strata, and PDS Lasso selected covariates; (iii) a balance check comparing T0 scores across arms after reweighting — these should be equal if IPW is working correctly.

8. Multiple Hypothesis Testing

We define outcome families ex ante and apply separate error-rate control procedures within each family. The table below is the authoritative reference for what constitutes a primary conclusion.

Family	Outcomes	Error rate	Procedure	Decision rule
Primary	Y_req (H1) (ITT and LATE)	FWER	N/A — single outcome	$p < 0.05$ on 2SLS LATE coefficient β_1; positive sign required; first-stage $F > 10$.
Secondary beliefs	Y_try_cs (H2) and Y_need_car (H3)	FWER	Romano–Wolf stepdown (preferred); Holm as fallback	Romano–Wolf adjusted $p < 0.05$ on ANCOVA τ for Y_try_cs and Y_need_car (tested jointly). IPW caveat stated throughout. Labeled secondary.
Exploratory behavioral	Y_red12m (H4), Y_use12m (H5), Y_car_own12m (H6)	None	Unadjusted; flagged exploratory	No confirmatory claims. Subject to data availability.

Family	Outcomes	Error rate	Procedure	Decision rule
Heterogeneity	All prespecified interaction terms (H7a–i)	FDR	Benjamini–Hochberg (q = 0.05)	Adjusted p < 0.10; labeled exploratory.

Different families warrant different error-rate concepts and procedures. We do not do multiple hypothesis adjustment on LATE versus ITT. They are using the same variable (not different outcomes), the same treatment, and this is not the same as subgroup analysis. The secondary belief family (H2 and H3, two correlated outcomes) uses Romano–Wolf (2005) stepdown, which controls the family-wise error rate (FWER) while accounting for the positive correlation between Y_try_cs and Y_need_car through resampling. Because these outcomes move together, the effective number of independent tests is less than two, and Romano–Wolf exploits this to be more powerful than Holm while maintaining the same FWER guarantee. The exploratory heterogeneity family (H7a–i, nine tests) uses Benjamini–Hochberg (1995). FWER control would be too conservative for nine exploratory tests: requiring that no false positives occur at all would rarely flag anything interesting. FDR control instead limits the expected proportion of false discoveries among those flagged, which is appropriate for hypothesis-generating subgroup analyses. For all families, unadjusted p-values are always reported alongside adjusted ones.

▶ *The primary family contains one outcome (H1, Y_req), estimated as a LATE via 2SLS and ITT. A confirmatory claim requires $p < 0.05$ on the LATE coefficient β_1, a positive sign, and a first-stage F-statistic above 10. The ITT on Y_req is reported in addition. H2 (Y_try_cs) and H3 (Y_need_car) are secondary: tested jointly under Romano–Wolf, reported alongside H1. A secondary confirmatory claim requires Romano–Wolf adjusted $p < 0.05$ on each outcome.*

9. Missing Data, Attrition, and Compliance

9.1 Administrative vs. survey outcomes

Y_req is recorded for the full eligible sample, and it is coded 0 for those who do not complete.

9.2 Y_req denominator: pre-commitment

We precommit to the following: Y_req = 0 for all eligible respondents who did not submit a coupon request, including survey non-completers. The LATE denominator (first stage) uses all eligible respondents. Y_req = 0 for treatment non-completers is accurate by construction — the coupon page is only reachable by completing the calculator, so non-completers structurally cannot have Y_req = 1.

> ▶ *Differential attrition risk. If the calculator redirect causes higher dropout in the treatment arm, more treatment respondents contribute Y_req = 0 by friction. This attenuates the ITT (already expected to be negative/zero) and compresses the first stage, increasing the LATE MDE. We precommit to reporting survey completion rates by arm and testing for differential completion. If treatment completion is more than 3 percentage points lower than control, Lee (2009) attrition bounds are reported.*

9.3 Survey non-response

We pre-commit to the following protocol for survey-based outcomes:

1. Report response and completion rates separately by treatment arm, and test for differential attrition using a regression of the response indicator on treatment assignment with strata FE.
2. Test covariate balance among survey respondents and report standardized mean differences.
3. Primary analysis: unweighted complete-case for all survey outcomes.
4. Robustness: inverse-probability weighting (IPW) using only pre-treatment predictors (age, gender, postcode access). IPW weights winsorized at the 99th percentile.
5. Sensitivity: if survey attrition differs across arms substantially for main outcome, Lee (2009) bounds are reported.

9.4 Take-up funnel

We report a complete take-up funnel for transparency:

Invited → Started baseline survey → Eligible → Assigned → [Treatment: visited landing page → entered vehicle → viewed cost report → changed assumptions] → Coupon requested → Coupon redeemed

Numbers and percentages are reported at each stage, separately by arm. Differential rates are flagged and investigated.

9.5 Protocol deviations

The following are prespecified as material protocol deviations requiring documentation: (i) randomization code contains an error; (ii) the calculator is inaccessible for more than 48 hours during the treatment window. Any deviation is logged with time and date and disclosed in the final paper.

Ineligible respondents who pass through eligibility screening are excluded from the primary eligible-sample analysis but included in an all-randomized robustness analysis.

10. Power and Realistic Sample Projections

10.1 Realistic response funnel

The table below presents realistic projections for each stage of the response funnel, based on Norwegian survey response experiences and the specifics of this design. These are pre-study estimates used for power planning. Actual numbers will be reported in the paper.

Stage	Assumption	Est. N	Per arm
Invited	Full frame	75,000	—
Complete survey	~15%	~11250	—
Eligible (randomized)	~80% pass eligibility	~9000	T= 4950 C=4050
Treatment: click calculator link (D_click)	~70% of treatment arm	~3465	treatment only assumed 100% for control

Stage	Assumption	Est. N	Per arm
Treatment: complete calculator (D_report = 1)	~70% of clickers; ~50% of treatment arm	~2425	treatment only assumed 100% for control

> ▶ *Basis for assumptions. Complete survey rates have been reported around 15% in previous data collection at TØI both at city and national level. Eligibility criteria at 80% assumes this is the rate of household with a car or planning to buy a car. Probability of continue to the calculator and finalize are set to 70%, but are empirically unknown.*

10.2 Design parameters and MDE

All analyses use $\alpha = 0.05$ (two-sided) and target power = 0.80. Based on the funnel in Section 10.1, the expected eligible sample is approximately 9000 respondents. The LATE MDE scales as MDE_LATE ≈ MDE_ITT / π_1, where π_1 is the first-stage take-up rate (~0.30). Randomization is at the individual level.

10.3 MDE table

Outcome	Baseline rate / SD	MDE	N (ICC = 0)
Coupon requested (Y_req, primary)	p_0 = 30% sd = 0.46	5 pp	N ≈ 9000 eligible
Willingness to try CS (Y_try_cs, secondary)	p_0 = 4.78 (scale 1-7) sd = 1.82	2% or 0.096	N ≈ 2425 eligible

11. Heterogeneity and Mechanism Analyses

For each heterogeneity hypothesis H7a–i, we estimate a single interaction model:

$$Y\_i = \beta_0 + \beta_1 H\_i + \beta_2 T\_i + \beta_3 (H\_i \times T\_i) + \delta\_s(i) + \gamma' X\_i + \varepsilon\_i$$

Where H_i is the heterogeneity and mechanism variables, T_i is the treatment variable and β_3 is the heterogeneity estimate. The outcomes for heterogeneity tests are Y_req, Y_try_cs and Y_need_car. All independent variables are measured at baseline. Multiplicity is controlled within the heterogeneity family using Benjamini–Hochberg (Table in Section 8).

12. Data Pipeline and Replication Package

12.1 Pipeline overview

1. Finalize PAP + complete AEA RCT Registry entry.
2. Archive randomization code, seed, and assignment file before data collection.
3. Launch survey: using a predefined plan that spans over several days
4. Download data from survey program and from website. Merge data.

5. Analyze data and start writing results
6. Administrative merge: discount redemption data several times during the year
7. Administrative merge at 12 months: final discount redemption data + carsharing use + car ownership registry.
8. Estimation + robustness checks + reweighted beliefs + sensitivity analyses.
9. Replication package assembled: raw data + clean data + code + README + codebook.

13. Ethics and Data Protection

13.1 Regulatory framework

This study processes personal data (email addresses, vehicle license plates, etc) and is subject to the Norwegian Personal Data Act and GDPR. The study is registered via Sikt's meldeskjema system and approved by SIKT. NESH guidelines are followed throughout. The project has also been approved by TØI's ethical review board before the data collection started.

13.2 Personal data handling

Personal data will be stored in TØI's secure research zone. This runs on modern hyperconverged infrastructure in our own local datacenter. The storage areas are isolated from our PCs, and we can only work on the information via locked down and encrypted RDP sessions. Login to the zone requires two-factor authentication. There is no open communication to and from the insured the research zone. All communication regarding software updates or communication with external services is done via Whitelisting, or through our structured data reception. All information to be entered and exited from the secured research zone is logged detailed information about what, who and when information is moved. The information is covered by access control and the information is encrypted during transmission. Access is limited to selected project staff.

- Email addresses: used solely to send the survey invitation and (if requested) the coupon. Stored on a safe TØI server with restricted access and logs.
- Vehicle license plates: Stored on a safe TØI server with restricted access and logs
- Server logs: contain participant IDs (pseudonymous tokens), page-visit timestamps, and interaction events. No direct personal identifiers stored.
- Administrative partner data (coupon redemption): received as records matched by anonymous user id and coupon code only.
- All personal data will be anonymized by 2030 as described in the protocol submitted to SIKT and explicitly consented by the respondents.

13.3 Participant information and consent

All participants receive a clear information notice stating the study purpose and funder, the types of data collected and why, the voluntary nature of participation, the right to withdraw, the identity of the data controller, storage duration, and DPO contact. Informed consent is obtained by affirmative action (proceeding past the consent screen).

14. Deviations from This Plan

Any deviation from this PAP is documented in a deviations log (deviations.md in the replication repository) and disclosed in the final paper. Deviations are classified as:

4. Pre-specified deviations: changes made before data collection or before outcome data are reviewed, necessitated by unforeseen design issues. Documented with justification and timestamp before data is examined.
5. Post-collection, pre-unblinding deviations: changes made after data collection but before treatment status is revealed to the analyst. Treated as confirmatory only if approved by a pre-registered amendment.
6. Post-hoc analyses: any analysis conducted after the primary results are known. Labeled exploratory in all outputs; given no confirmatory weight.