# Analysis Plan

How does Status Seeking impact Conspicuous Consumption Choice?

Field Study, Northern Vietnam, 2026

May 11, 2026

## Contents

# 1 Data Preparation and Variable Definitions

The analysis sample consists of all respondents who complete both the chocolate choice game and the public good game. Respondents who drop out before finishing the questionnaire are excluded; the number and stage of attrition will be reported. For the main specifications we retain every observation regardless of performance on attention checks, but we flag those who fail the attribute recall check and later run a robustness specification that excludes them, to verify that inattentiveness does not drive the results. If any covariate has more than five per cent missing values, we replace the missing cells using multiple imputation with chained equations (ten imputed datasets); otherwise observations with missing covariates are listwise deleted.

The two co-primary outcomes are binary indicators derived from the chocolate choice game. Let organic_choice $= 1$ if the respondent selected the organic-certified chocolate (Option B) and zero otherwise (Options A or C). Let luxury_choice $= 1$ if the respondent selected the imported luxury chocolate (Option A) and zero otherwise (Options B or C). These variables are mutually exclusive by construction. For descriptive completeness we also record the three-category variable choice_cat $\in \{1, 2, 3\}$ (luxury, organic, conventional), which will be used in a supplementary multinomial logit model.

The status-seeking orientations that form our core explanatory variables are measured through multi-item Likert scales administered after the chocolate choice. All scales are constructed as the respondent's mean score across the constituent items, provided at least three-quarters of the items are non-missing; Cronbach's $\alpha$ is reported for each scale, and any scale with $\alpha < 0.65$ will be inspected item-by-item with transparent documentation of any subsequent adjustments.

The green signalling motivation scale (green_signal) averages items 1–4 of Section 1.1, capturing the degree to which the respondent views green consumption as a vehicle for self-presentation. Perceived social rewards for green consumption (social_reward) averages items 5–8, measuring the belief that green choices confer social approval. Identity alignment via consumption (identity_align) averages items 9–12, reflecting the extent to which product choices are seen as expressions of personal identity. Status and social comparison orientation (status_compare) combines items 13–15 from Section 2.1, tapping the general tendency to seek status through consumption and to compare oneself to others. Additional control scales include environmental self-identity (env_selfid), practical and economic considerations (practical_econ), and single-item measures of social media visibility and prior green consumption.

Because the attitudinal scores are self-reported and measured after the incentivised choice, they may be subject to both common-method bias and consistency-driven rationalisation. To mitigate the endogeneity concern we pre-specify three community-level items as candidate instrumental variables: the belief that people in the community pay attention to what others buy (item 26), the importance of reputation in the community (item 27), and the perception that environmentally friendly behaviours are noticed (item 28). These variables are intended to capture the local social observability norms

that shape status-signalling motivation while being plausibly exogenous to an individual's inherent taste for chocolate type. We will use them in a two-stage least squares framework as a robustness check, but we recognise that their validity as instruments is debatable; we therefore treat the IV analysis as exploratory and report all specification tests transparently.

The experimental manipulation is captured by two treatment indicators: treat_luxury = 1 for assignment to the luxury salience arm and treat_organic = 1 for assignment to the organic salience arm; the control arm (no attribute script, no visibility framing) is the reference category. In the luxury and organic arms respondents receive a brief spoken description of the relevant product attribute, an explicit mention that their choice will be visible to others in the room, and labelled display cards next to the chocolates. The control arm receives only the basic choice invitation, with none of these amplifying elements. Thus the treatment bundle comprises both attribute salience and social-visibility salience, and the estimated average treatment effects should be interpreted as the combined effect of making the focal attribute cognitively accessible while also signalling that the choice is socially observed.

The vector of control variables $\mathbf{X}_i$ includes all socio-demographics: age (linear and quadratic), gender, education, occupation, marital status, household size, number of children, monthly household income, home ownership, length of residence, urban/rural self-perception, community involvement, smartphone ownership, and social media hours. We also include the amounts allocated to the environmental project and to the general charity fund in the public good game, because they capture baseline pro-environmental and prosocial preferences net of the status-signalling dynamics elicited in the chocolate game. Finally, the credibility and effectiveness ratings of the environmental programme (items 23–24) enter as controls for trust in the public-good task. All models include enumerator-day fixed effects to absorb unobserved session-level influences.

# 2 Descriptive Checks and Baseline Balance

Before testing hypotheses we characterise the sample and verify that randomisation produced comparable groups. A balance table reports means and standard deviations of all outcome variables, attitudinal scales, socio-demographic characteristics, and behavioural controls, separately for each treatment arm and for the pooled sample. For each variable we test the difference between each active arm and the control arm using a two-sided $t$-test (or a $\chi^2$ test for categorical variables). Variables that exhibit a significant imbalance at the five per cent level are noted; in a subsequent robustness check we add their interaction with the treatment indicators to assess whether the baseline imbalance alters the estimated treatment effects.

We also report the proportion of respondents who pass the attribute recall check, the mean perceived social visibility by arm, and test for differences across arms using ANOVA. These checks serve both as a manipulation validation (does the visibility framing raise perceived visibility?) and as a attention filter for the planned sensitivity analysis.

# 3 Primary Analysis: Status-Signalling, Salience, and Product Choice

Our central question is whether individual status-signalling orientation, measured through attitudinal scales, predicts a revealed-preference choice between conspicuous consumption and conspicuous conservation, and whether this relationship is amplified when the relevant product attribute is made salient. We estimate linear probability models (LPM) for the two binary outcomes because the coefficients of the interaction terms are directly interpretable as changes in the probability of choice, which is natural for communicating effect sizes in a field-experiment context. Logistic regressions are reported alongside to confirm that the functional form does not drive the conclusions.

The base specification regresses each binary outcome on the treatment dummies and the full set of controls, without interactions:

$$Y_i = \alpha + \beta_1 \text{treat\_luxury}_i + \beta_2 \text{treat\_organic}_i + \gamma \mathbf{X}_i + \varepsilon_i, \tag{1}$$

where $Y_i \in \{\text{organic\_choice}, \text{luxury\_choice}\}$. This yields the average treatment effect of the combined salience-and-visibility manipulation on the probability of selecting each chocolate type, net of the rich set of socio-demographic and behavioural controls.

The theoretically central specification adds the continuous attitudinal score and its interactions with treatment:

$$\begin{aligned} Y_i = \alpha &+ \beta_1 \text{treat\_luxury}_i + \beta_2 \text{treat\_organic}_i + \beta_3 \text{score}_i \\ &+ \beta_4(\text{score}_i \times \text{treat\_luxury}_i) + \beta_5(\text{score}_i \times \text{treat\_organic}_i) \\ &+ \gamma \mathbf{X}_i + \varepsilon_i, \end{aligned} \tag{2}$$

where $\text{score}_i$ is one of the attitudinal scales green_signal, social_reward, identity_align, or status_compare.

The hypotheses map onto specific coefficients in this model. First, the *conspicuous conservation hypothesis* states that green-signalling motivation should predict organic choice, and that this relationship is stronger when the organic attribute is made salient. Operationally, we expect $\beta_5 > 0$ in the organic choice model when the score is green_signal, social_reward, or identity_align. Conversely, the *conspicuous consumption hypothesis* expects that status orientation predicts luxury choice, and that the luxury salience amplifies this link: $\beta_4 > 0$ in the luxury choice model with score = status_compare. Finally, we test a *cross-motivation suppression* pattern: if the treatments activate competing status motives, the organic salience condition may weaken the link between status orientation and organic choice ($\beta_4 < 0$ for organic choice with status score), while the luxury salience may weaken the green-signalling to luxury-choice link. Testing these predictions simultaneously across the two outcomes allows us to assess whether the salience manipulation shifts the entire pattern of status-motivated choice or merely intensifies an existing preference.

We will report interaction coefficients with heteroskedasticity-robust standard errors, and supplement them with marginal effects plots showing the predicted probability of choice at each level of the attitudinal score, separately for the three treatment arms, with 95% confidence intervals. This graphical presentation makes it easy to see whether the slopes differ in the predicted direction.

We test four scores across two outcomes, yielding eight primary hypothesis tests. Given the strong theoretical priors and the fact that each test corresponds to a distinct substantive hypothesis, we present results both without and with a multiple-comparison correction. The transparency correction uses the Benjamini–Hochberg procedure to control the false discovery rate across the family of eight tests; sharpened $q$-values are reported alongside unadjusted $p$-values.

# 4 Robustness and Sensitivity Checks

The robustness of the primary results is assessed along several dimensions. First, we re-estimate all interaction models using logistic regression to verify that the LPM does not produce artefacts due to its linear functional form. Second, we rerun the main specifications after excluding respondents who failed the attribute recall check; this isolates the effect on participants who demonstrably attended to the product attributes. Third, we examine sensitivity to covariate imbalance: for any variable that shows a significant difference between arms at baseline, we include its interaction with treatment in the model and test whether the interaction coefficients of interest change meaningfully. Fourth, although treatment is assigned at the individual level, respondents are seated in small groups of two or three during the session. To account for any within-session correlation, we cluster standard errors at the enumerator-day level, which yields approximately 30–40 clusters. Results are robust if the statistical significance of the key coefficients is unchanged relative to the heteroskedasticity-robust standard errors.

# 5 Exploratory Instrumental-Variables Analysis

The attitudinal scores are measured after the chocolate choice and may therefore reflect ex-post rationalisation or consistency motives, which would bias the ordinary least squares estimates. We explore an instrumental-variables strategy that uses community-level observability norms as instruments. The three instruments are IV_comm_obs (item 26), IV_reputation (item 27), and IV_notice_green (item 28). The identifying assumption is that these community-level norms influence chocolate choice only through their effect on the individual's status-signalling motivation, conditional on the included controls. While this exclusion restriction is strong and cannot be fully tested, the overidentification test based on Hansen's $J$ statistic provides a statistical plausibility check.

The first stage regresses the endogenous attitudinal score on the three instruments,

the treatment dummies, and all controls:

$$\begin{aligned} S_i = \delta_0 &+ \delta_1 \text{IV\_comm\_obs}_i + \delta_2 \text{IV\_reputation}_i + \delta_3 \text{IV\_notice\_green}_i \\ &+ \delta_4 \text{treat\_luxury}_i + \delta_5 \text{treat\_organic}_i + \gamma \mathbf{X}_i + u_i. \end{aligned} \tag{3}$$

We report the first-stage $F$-statistic for the excluded instruments, the partial $R^2$, and the Angrist-Pischke multivariate $F$-test where relevant. A first-stage $F$ below 10 would signal weak instruments, and we would interpret the subsequent IV estimates with appropriate caution.

The second stage replaces the observed score with its predicted value $\widehat{S}_i$ from the first stage, and includes the same interaction terms:

$$\begin{aligned} Y_i = \alpha &+ \beta_1 \text{treat\_luxury}_i + \beta_2 \text{treat\_organic}_i + \beta_3 \widehat{S}_i \\ &+ \beta_4 (\widehat{S}_i \times \text{treat\_luxury}_i) + \beta_5 (\widehat{S}_i \times \text{treat\_organic}_i) \\ &+ \gamma \mathbf{X}_i + \varepsilon_i. \end{aligned} \tag{4}$$

Standard errors are bootstrapped to account for the generated regressor. The IV results are not intended to provide the primary basis for causal inference; rather, they serve as a robustness check on the direction and significance of the interaction coefficients when the potential endogeneity of the attitudinal score is explicitly addressed. Given the likely efficiency loss, we expect wider confidence intervals, and we will interpret the comparison with the OLS estimates mainly in terms of sign consistency.

# 6 Mediation by Perceived Social Visibility

Because the treatment bundle includes an explicit social-visibility framing, we expect that part of the treatment effect operates through increasing the perceived observability of the choice. The post-choice manipulation check (Check 1) asks respondents to rate how visible their choice is to others on a five-point scale. We first estimate a simple OLS regression of this visibility score on the treatment dummies and controls, to test whether the luxury and organic arms indeed raise perceived visibility relative to the control arm.

Building on this, we test whether perceived visibility moderates the effect of status-signalling attitudes on choice. In the primary interaction model (Equation 2) we replace the treatment interaction with a three-way interaction among score, treatment, and visibility, or, more straightforwardly, we estimate treatment-by-score interactions separately for respondents with high and low perceived visibility (split at the median). The expectation, grounded in the status-signalling literature, is that attitude–behaviour consistency is strongest when individuals believe their choices are socially observed.

For a formal causal mediation decomposition we apply the algorithm of Imai, Keele, and Tingley (2010), which allows us to estimate the Average Causal Mediation Effect (ACME) of the treatment on choice probability that operates through perceived visibility, separate from the direct effect. We compute bootstrapped confidence intervals

for the ACME. Because visibility is measured post-treatment, the sequential ignorability assumption is required; we discuss its plausibility and report sensitivity analyses.

We also explore, purely descriptively, whether the treatment shifted the attitudinal scores themselves. A finding that, for example, the organic salience condition raises green-signalling self-ratings would be consistent with a priming or consistency effect, but we will not attach causal interpretations to such post-treatment attitude differences.

# 7 Secondary Analyses

Several additional analyses complement the primary results. First, we estimate a multinomial logit model for the three-category choice variable choice_cat to examine the simultaneous substitution pattern between luxury, organic, and conventional options, rather than analysing each binary outcome in isolation. This allows us to see, for instance, whether the organic salience draws respondents primarily from the conventional or from the luxury option.

Second, we examine whether the attitudinal scores that predict chocolate choice also predict contributions in the public good game. A positive correlation between green-signalling motivation and the share of the endowment directed to the environmental project (as opposed to the charity fund) would suggest that the status-signalling orientation captured by our scales generalises beyond the chocolate choice context. Conversely, a divergence between the two behavioural measures would underscore the distinctiveness of the real-product revealed-preference design.

Third, we conduct exploratory subgroup analyses to assess whether the pattern of results differs by gender, urban versus rural residence, age (median split), and income level. These tests are underpowered and are presented as suggestive only.

Finally, we document attrition and non-response at every stage, and compare completers with dropouts on available baseline characteristics to evaluate the threat of selection bias.

# 8 Pilot Study (Strasbourg, June 2026)

The pilot is designed to validate the product stimuli and the questionnaire instruments before the main fieldwork, not to test hypotheses. We will check that the three chocolates are perceived as intended—specifically, that the Lindt bar is seen as significantly more of a luxury/high-status product, the organic-certified bar as significantly more environmentally friendly, and the conventional bar as neutral on both dimensions. Mean ratings will be compared with paired $t$-tests. We will compute Cronbach's $\alpha$ for every multi-item scale and refine any item that degrades reliability. We will also examine the distribution of the visibility manipulation check and the attribute recall check to ensure that the experimental procedure generates sufficient variance and attention. Any modifications to the protocol or questionnaire resulting from the pilot will be fully documented.

# 9 Reporting Standards

All estimations will be performed in Stata (version 18) and R (version 4.4). The paper will report point estimates, robust standard errors, 95% confidence intervals, and exact $p$-values. Regression tables will include both the LPM and logit specifications side by side. A CONSORT-style flow diagram will summarise participant inclusion and exclusion. Results will be shown with and without the Benjamini–Hochberg correction, so that readers can assess the influence of multiple comparisons. Any deviation from this pre-analysis plan will be clearly marked in the final manuscript as exploratory.

*This analysis plan is submitted to the AEA RCT Registry. It may be updated if unforeseen data characteristics require adjustments; all versions will be time-stamped and documented.*