What Shapes Perceptions of Fraud in the Social Safety Net?
Analysis Plan

1. Overview

In this survey experiment, we use a conjoint experimental design to disentangle the causal factors that shape perceptions of fraud in the social safety net. Approximately 1,200 U.S. adults, recruited via Prolific and representative on race, gender, age, and political affiliation, will evaluate pairs of government benefits recipients (Module A) and pairs of government benefits programs (Module B), each varying randomly on attributes such as beneficiary demographics, program type, benefit generosity, and administrative requirements. Our primary outcome is perceived fraud; secondary outcomes include support for additional verification burdens and support for continued program funding.

2. Research Questions

Primary

RQ 1: How do beneficiary characteristics and program design features shape perceptions of fraud in the social safety net?

RQ 2: How do the effects of beneficiary and program design features on fraud perceptions differ by respondent characteristics?

Secondary

RQ 3: How do beneficiary characteristics and program design features shape support for imposing additional burdens?

RQ 4: How do beneficiary characteristics and program design features shape support for program funding?

3. Hypotheses

We have no a priori hypotheses.

4. Study sample

We will recruit approximately 1,200 U.S. adults (aged 18 or older) through Prolific's representative sample feature, which matches respondents to U.S. Census margins on sex, age, ethnicity, and political affiliation. Respondents must be currently living in the United States.

The primary analytic sample excludes: (a) respondents who fail both attention checks embedded in the survey, (b) respondents with duplicate IP addresses or duplicate Prolific participant IDs, and (c) respondents who are flagged as bots by qualtrics.

5. Study design

The full survey instrument is attached in Appendix A.

The survey is administered via Prolific. Participants view a consent form, confirm U.S. residency, and complete an attention check. Participants then view an example conjoint table to familiarize themselves with the task format; whether they answer correctly on their first attempt is recorded for robustness checks.

This is a choice-based paired conjoint experiment with two modules, counterbalanced across respondents. Each respondent completes 2 tasks per module (4 tasks total). In each task, respondents view two side-by-side profiles and select one profile for each outcome measure. Each profile varies on 16 attributes with two to four levels, randomly assigned with no restrictions on combinations.

Module A displays individual government benefits recipients. Module B displays government benefits programs. Attributes and levels are shown in Table 1.

Table 1: full set of attributes and levels for both modules

	MODULE A - BENEFICIARY		***MODULE B - PROGRAM***
Bundle 1: Beneficiary characteristics			
Race/ethnicity	**[race_black_a]** Black **[race_white_a]** White **[race_hispanic_a]** Hispanic	*Racial composition of the program*	**[race_black_b]** 30% are White, 60% are Black, 10% are Hispanic **[race_white_b]** 70% are White, 20% are Black, 10% are Hispanic **[race_hispanic_b]** 30% are White, 10% are Black, 60% are Hispanic
Gender	**[gender_male_a]** Male **[gender_female_a]** Female	*Gender composition of the program*	**[gender_male_b]** 26% women **[gender_female_b]** 70% women
Age	**[age_young_a]** 25 years old **[age_adult_a]** 40 years old **[age_older_a]** 61 years old	*Primary age group of recipients*	**[age_young_b]** The average beneficiary is in their 20s **[age_adult_b]** The average beneficiary is in their 40s **[age_older_b]** The average beneficiary is in their 60s
Marital status	**[marital_single_a]** Single **[marital_married_a]** Married	*Marital status of recipients*	**[marital_single_b]** 70% are single households **[marital_married_b]** 35% are single households
Parental status	**[children_low_a]** No children **[children_high_a]** Has children	*Parental status*	**[children_low_b]** 29% have children **[children_high_b]** 80% have children
Citizenship status	**[origin_us_a]** Born in the US **[origin_naturalized_a]** Naturalized US citizen **[origin_permanent_a]** Permanent resident	*Citizenship composition*	**[origin_us_b]** 70% are US citizens; the rest naturalized or permanent residents **[origin_perm_nat_b]** 30% are US citizens; the rest naturalized or permanent residents

Disability status	**[disable_yes_a]** Disabled **[disable_none_a]** Not disabled	*Disability among recipients*	**[disable_yes_b]** 69% have a disability **[disable_none_b]** 10% have a disability
Employment status	**[emp_effort_a]** Works full-time **[emp_effort_nocontrol_a]** Unemployed, looking for work **[emp_noeffort_a]** Unemployed**,** not looking for work	*Employment among recipients*	**[emp_high_b]** 77% are employed **[emp_low_b]** 10% are employed
Annual Income	**[income_low]** Less than $10,000 per year **[income_higher]** Less than $58,000 per year	*Income of recipients*	**[income_low]** Less than $10,000 per year **[income_higher]** Less than $58,000 per year
Bundle 2: Program characteristics			
Receives assistance for	**[program_housing]** Housing expenses **[program_health]** Healthcare expenses **[program_food]** Food expenses **[program_gen]** Cash assistance	*The program provides assistance for*	**[program_housing]** Housing expenses **[program_health]** Healthcare expenses **[program_food]** Food expenses **[program_gen]** Cash assistance
Benefit amount	**[benefit_low]** Up to $200 per month **[benefit_high]** Up to $1300 per month	*Benefit amount*	**[benefit_low]** Up to $200 per month per recipient **[benefit_high]** Up to $1300 per month per recipient
Benefit type	**[type_cash_unrestrict]** Unrestricted cash **[type_cash_restricted]** Cash restricted to specific products **[type_in_kind]** In-kind (e.g., subsidies or payments sent directly to service providers)	*Benefit type*	**[type_cash_unrestrict]** Unrestricted cash **[type_cash_restricted]** Cash restricted to specific products **[type_in_kind]** In-kind (e.g., subsidies or payments sent directly to service providers)
Duration	**[duration_short]** Has received assistance for 6 months **[duration_long]** Has received assistance for 5 years	*Duration*	**[duration_short]** Beneficiaries can receive assistance for up to 6 months **[duration_long]** Beneficiaries can receive assistance for up to 5 years
Application requirements	**[burden_low]** No action necessary. Benefits are received automatically for eligible participants **[burden_medium]** Submit an application and provide documentation of income and residency **[burden_high]** Submit an application, provide documentation of income and residency, and attend an in-person interview	*Application requirements*	**[burden_low]** No action necessary. Benefits are received automatically for eligible participants **[burden_medium]** Submit an application and provide documentation of income and residency **[burden_high]** Submit an application, provide documentation of income and residency, and attend an in-person interview
Work requirements	**[work_none]** There are no work requirements to receive benefits **[work_high]** Must work or participate in job training at least 20 hours per week to receive benefits	*Work requirements*	**[work_none]** There are no work requirements to receive benefits **[work_high]** Must work or participate in job training at least 20 hours per week to receive benefits
Recertification	**[recer_none]** No recertification required **[recer_high]** Must recertify eligibility every 6 months	*Recertification*	**[recer_none]** No recertification required **[recer_high]** Must recertify eligibility every 6 months

The order of Module A and Module B is counterbalanced across respondents. The order of attribute bundles (beneficiary characteristics and program characteristics) is randomized across respondents but held constant across tasks within respondent.

After each task, respondents are asked a series of outcome measures (defined in Section 6). The order of outcome measures is randomized across respondents but held constant across tasks within respondent. All outcomes use a forced-choice format: Person [Program] A or Person [Program] B.

After completing both modules, respondents complete a second attention check, then complete the covariates and demographics modules.

Covariates (order randomized across respondents):
1. Burden tolerance (4 items)
2. Social dominance orientation (2 items)
3. Attributions for poverty (1 item)
4. Trust in government (1 item)

Demographics: age, gender, party identification, political ideology, race/ethnicity, current and past government benefits receipt, employment status, education, and income.

6. Outcomes and Data

All outcomes are coded at the profile level. Within each task, the selected profile is coded as 1, the non-selected profile as 0. This yields two observations per task per outcome.

The continued funding outcome is measured in Module B only.

Table 2: outcomes across both modules

Measure	Definition	Modules
Fraud perception (primary)	"Which person is more likely to be cheating the system to receive benefits?" (Module A) "Which program has more recipients who are cheating the system to receive benefits?" (Module B)	A, B
Additional burdens (secondary)	"If the government is going to choose one person to conduct extra verification checks to ensure eligibility, which person should they choose?" (Module A) "The government is considering imposing additional requirements to make it more difficult to access benefits. For which program should the government impose more requirements?" (Module B)	A, B
Continued funding (secondary)	"If the government can only continue funding one program, which program should continue funding?"	B only

7. Covariates

The survey includes the following variables that will be used as moderators and covariates:

- Burden tolerance
    - Continuous index (1-7) that's the average of four items
    - Binary: 1= Above median; 0 = At or below median
- Social dominance orientation
    - Continuous index (1–7) that's the average of two items.
    - Binary: 1= Above median; 0 = At or below median
- Poverty attributions
    - Binary: 1 = dispositional ("laziness and lack of willpower") ; 0 = structural ("unfair society")
- Trust in government
    - Binary: 1 = Just about always; Most of the time ; 0 = Only some of the time; Never
- Race/ethnicity
    - Categorical: White, Black, Hispanic, Other (combining all remaining categories)
- College Education
    - Categorical: No HS; HS graduate; Some College; 2-year degree; 4-year degree; graduate degree
    - Binary: 1 = 4-year degree of graduate degree; 0 = less than 4-year degree.
- Political ideology
    - Categorical: Liberal, Moderate, Conservative, Other (including not sure)
    - Binary: 1 = Conservative ; 0 = Liberal
- Party identification
    - Categorical: Democrat, Republican, Other (combining Independent, Other party, Prefer not to say, I don't know).
- Gender
    - Categorical: Male, Female, Other
    - Binary: 1 = Male ; 0 = Female/Other.
- Income
    - Categorical: 10 levels from from < $10,000 to $100,000+ (in $10,000 increments).
    - Binary: 1 = household income $80,001 and above ; 0 = household income $0–$80,000
- Current or past benefits receipt
    - Binary: 1 = currently receive or received in the past; 0 = does not receive benefits
- Employment status
    - Categorical: Full-time employed; Part-time employed; Not currently employed
    - Binary: 1 = employed full-time ; 0= not employed full time

For all categorical variables, we will collapse categories that comprise <1% of the analytic sample into an "other" category.

8. Analyses

a. Analytic sample

The primary analytic sample includes all respondents who meet the inclusion criteria in Section 4.

b. Missingness

All conjoint items and moderators are required for survey completion. We do not expect missing data conditional on survey completion after applying the exclusions specified in Section 4.

c. Statistical models for each RQ

RQ 1: How do beneficiary characteristics and program design features shape perceptions of fraud in the social safety net?

We will estimate separate models for each module as follows:

Module A (Beneficiary framing)

$$Y_{ijk} = \alpha + \beta_{1-2} race_{ijk} + \beta_3 female_{ijk} + \beta_{4-5} age_{ijk} + \beta_6 marital\ status_{ijk} + \beta_7 parental\ status_{ijk} + \beta_{8-9} citizenship_{ijk} + \beta_{10} disability_{ijk} + \beta_{11-12} employment_{ijk} + \beta_{13} income_{ijk} + \beta_{14-16} program_{ijk} + \beta_{17} amount_{ijk} + \beta_{18-19} type_{ijk} + \beta_{20} duration_{ijk} + \beta_{21-22} enrollment\ req_{ijk} + \beta_{23} work\ req_{ijk} + \beta_{24} recertification_{ijk} + respondent_i + \epsilon_{ijk}$$

where Y_{ijk} is an indicator equal to 1 if respondent i selected profile k in task j as more likely to involve fraud, and 0 otherwise. The model includes indicator variables for all attribute levels shown in Table 1, with reference categories: White, male, age 25, single, no children, born in U.S., not disabled, works full-time, income less than $10,000, housing assistance, $200/month, unrestricted cash, 6 months duration, automatic enrollment, no work requirement, and no recertification. $respondent_i$ is a respondent fixed effect. The coefficients β_1 through β_{24} are interpreted as the average marginal component effect (AMCE) of each attribute level on the probability that a profile is selected as more likely to involve fraud, relative to the reference level for that attribute. We will use standard errors clustered at the respondent level and will reject the null hypothesis of no effect of a given attribute level if $p < 0.05$.

Module B (Program framing)

$$Y_{ijk} = \alpha + \beta_{1-2} race\ comp_{ijk} + \beta_3 female\ comp_{ijk} + \beta_{4-5} age\ comp_{ijk} + \beta_6 marital\ status\ comp_{ijk} + \beta_7 parental\ status\ comp_{ijk} + \beta_8 citizenship\ comp_{ijk} + \beta_9 disability\ comp_{ijk} + \beta_{10} employment\ comp_{ijk} + \beta_{11} income\ comp_{ijk} + \beta_{12-14} program_{ijk} + \beta_{15} amount_{ijk} + \beta_{16-17} type_{ijk} + \beta_{18} duration_{ijk} + \beta_{19-20} enrollment\ req_{ijk} + \beta_{21} work\ req_{ijk} + \beta_{22} recertification_{ijk} + respondent_i + \epsilon_{ijk}$$

where Y_{ijk} is defined as above. The model includes indicator variables for all program composition attributes shown in Table 1, with reference categories: 70% White composition, 26% women, average age 20s, 70% single-parent households, 29% with children, 70% U.S. citizens, 10% disabled, 77% employed,

income less than $10,000, and program characteristics as in Module A. We will use standard errors clustered at the respondent level and will reject the null hypothesis of no effect of a given attribute level if $p < 0.05$.

RQ 2: How do the effects of beneficiary and program design features on fraud perceptions differ by respondent characteristics?

As exploratory analyses, we will examine how the effects of recipient and program characteristics on fraud perceptions, support for additional verification burdens, and support for continued funding differ by respondent characteristics and demographics. Specifically, we will explore heterogeneity by:

a. Political ideology
b. Burden tolerance
c. Social dominance orientation
d. Trust in government
e. Poverty attribution beliefs
f. Race/ethnicity
g. Gender
h. Income
i. Education
j. Current or past government benefits receipt

These analyses will follow the same interaction model structure as the primary analyses, estimated separately for each moderator and each module.

RQ 3: How do beneficiary characteristics and program design features shape support for imposing additional burdens?

We estimate separate models for Module A and Module B using the same specifications as RQ 1, with Y_{ijk} redefined as an indicator equal to 1 if respondent i selected profile k in task j as the profile for which the government should impose additional verification checks, and 0 otherwise. The coefficients β_1 through β_{24} (Module A) and β_1 through β_{22} (Module B) are interpreted as the AMCE of each attribute level on the probability that a profile is selected for additional verification, relative to the reference level for that attribute.

RQ 4: How do beneficiary characteristics and program design features shape support for program funding?

This outcome is measured in Module B only. We estimate the same specification as RQ1 Module B, with Y_{ijk} redefined as an indicator equal to 1 if respondent *i* selected program *k* in task *j* as the program that should continue to receive funding, and 0 otherwise. The coefficients β_1 through β_{22} are interpreted as the AMCE of each attribute level in the probability that a program is selected for continued funding, relative to the reference level for that attribute.

Robustness checks

Robustness checks will exclude: (a) respondents who failed the comprehension check on the example conjoint table on their first attempt, (b) speeders (completing the survey in less than 45% of the median completion time), and (c) respondents who selected the same profile (A or B) on every conjoint task across both modules.

9. Exploratory Analyses (if applicable)

To explore differences across modules on the program characteristics attributes, we will pool observations from both modules and estimate a linear probability model regressing the forced-choice fraud indicator on program characteristic attribute levels, a Module B indicator, and their interactions, with respondent fixed effects and standard errors clustered at the respondent level. The interaction terms capture whether the AMCEs of program characteristics differ between the beneficiary framing (Module A) and the program framing (Module B). We will reject the null hypothesis of no differential effect if $p < 0.05$.

For beneficiary characteristics, which are operationalized differently across modules — as individual traits in Module A and as program composition in Module B — direct statistical comparison is not appropriate. We will instead compare AMCEs across modules descriptively, plotting estimates side-by-side and assessing whether patterns differ substantively.

10. Power calculations

We performed power calculations using the tool developed by Lukac & Stefanelli (2020). Calculations are based on the per-module design: 1,200 respondents, 2 tasks per respondent, and 2 profiles per task, estimated separately for each module. We assume $\alpha = 0.05$ and target 80% power.

Attribute levels	Respondents	Tasks	Effect Size (MDE)	Power
2	1,200	2	0.04	85%
3	1,200	2	0.05	89%
4	1,200	2	0.05	81%

The study is powered to detect AMCEs as small as 4 percentage points for attributes with two levels, 5 percentage points for attributes with 3 levels, and 5 percentage points for attributes with four levels. All attributes in the conjoint have between two and four levels (see Table 1), so the study is adequately powered to detect effects of this magnitude across all primary analyses.

11. IRB approval

Harvard University Institutional Review Board, protocol #IRB22-1397