

Beyond Demographic Profiling: Cross-Economy Variation in Payment Responses to Unsolicited Money Requests

Phan Duong

Thanh Dong University

Le Huyen

Thanh Dong University

Le Minh

Thanh Dong University

Vu Duyen

`duyenvtm@thanhdong.edu.vn`

Thanh Dong University

Research Article

Keywords: consumer vulnerability, demographic profiling, consumer protection, unsolicited money requests, common-composition standardization, cross-economy comparison

Posted Date: October 20th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11233480/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Abstract

Consumer-protection strategies often use demographic characteristics to identify groups that may face elevated marketplace risks, yet consumer-vulnerability research cautions against treating demographic identity as vulnerability itself. This study examines whether differences in observed demographic composition account for cross-economy variation in reported payment responses following unsolicited money requests. Using Global Findex Database 2025 public-use microdata, we analyze 13,766 respondents aged 18 years or older across 75 economies who reported receiving a survey-defined unsolicited request for money. Economy-specific reported-payment probabilities are estimated using released survey weights and then standardized to a common distribution of sex, age, education, workforce status, and urbanicity. The equal-economy 90th-to-10th percentile spread was 16.75 percentage points before standardization and 15.86 points afterward. The resulting attenuation was 0.89 percentage points, with a paired-bootstrap 95% confidence interval of -0.57 to 1.45. The remaining dispersion was robust across alternative weighting schemes, age specifications, covariate extensions, nonlinear models, leave-one-economy-out analyses, and a restriction to economies with more payment events. These findings establish a boundary of demographic explanation rather than showing that demographics are irrelevant. Routinely observed demographic composition provides limited descriptive leverage for accounting for cross-market variation in this post-contact payment response, motivating consumer-protection research that examines solicitation characteristics, decision environments, verification opportunities, and other contextual mechanisms directly.

1. Introduction

Unsolicited requests for money pose a consumer-protection challenge that extends beyond whether consumers are contacted. Once such a request reaches a consumer, subsequent responses may include ignoring the solicitation, engaging with it, seeking verification, or transferring money. These stages need not share the same determinants. Empirical research on financial fraud similarly distinguishes exposure to a solicitation from subsequent engagement and monetary outcomes, showing that predictors can differ across stages and scam types (DeLiema et al. 2023; Xu et al. 2022; Yu and Fan 2024). For consumer policy, this distinction matters because interventions that reduce exposure need not be the same as those that help consumers evaluate or interrupt a payment decision after contact has occurred.

Identifying which consumers may require greater protection presents a second challenge. Demographic characteristics such as age, sex, education, employment status, and income are observable and therefore attractive as practical indicators of consumer risk. Yet consumer-vulnerability research cautions against equating such characteristics with vulnerability itself. Baker et al. (2005) conceptualize consumer vulnerability as arising from interactions between individual characteristics, states, and external conditions, while Hill and Sharma (2020) emphasize limitations in consumers' access to and control over resources in particular marketplace contexts. More recent consumer-policy scholarship similarly stresses that vulnerability may be situational and shaped by the environments in which consumers make decisions rather than being a stable property of predefined demographic groups

(Dodsworth et al. 2023). Research on financial vulnerability and digital payments also shows that relationships observed for particular demographic groups or technologies in one setting need not reproduce straightforwardly in another (Seldal and Nyhus 2022).

This distinction becomes especially important in cross-market comparisons. Suppose consumers who have encountered the same survey-defined class of money request differ substantially across economies in age, education, workforce participation, urban residence, or other observable characteristics. Raw differences in subsequent payment behavior may then partly reflect differences in the composition of exposed populations rather than differences that remain among observationally comparable consumers. Conversely, if substantial cross-economy variation remains after exposed populations are evaluated on a common demographic composition, routinely observed demographic characteristics have limited descriptive leverage for accounting for the comparative pattern. This does not imply that demographic characteristics are irrelevant to consumer vulnerability. Rather, it places a boundary on how far demographic composition alone can account for variation across consumer-market contexts.

Existing fraud and consumer-risk research has typically approached this issue from one of two directions. One line of work estimates whether individual characteristics are associated with exposure, engagement, or victimization; another documents differences in fraud experiences across populations and countries. DeLiema et al. (2023), for example, show that correlates differ across stages and scam types, while Houtti et al. (2024) document substantial international variation in scam exposure, losses, contact vectors, and reporting. Raval (2021) likewise shows that demographic patterns vary across consumer-protection fraud cases. These approaches are valuable, but an adjusted demographic coefficient answers a different question from the one relevant here: it does not directly show how much of the overall cross-economy distribution would remain if every economy were evaluated on the same observed consumer composition. That distributional question is central to assessing the descriptive portability of demographic profiling across markets.

The Global Findex Database 2025 offers an unusual opportunity to examine this question. In the relevant economies, mobile-phone owners were asked whether, during the previous 12 months, they had received a call, SMS, or text message from someone they did not know asking them to send money. Respondents reporting such contact were then asked whether they sent the requested money. The survey therefore permits the analysis to condition explicitly on reported exposure rather than treating unexposed respondents as consumers who successfully resisted a solicitation. At the same time, the item has important interpretive boundaries. It does not independently establish the sender's fraudulent intent, adjudicated fraud victimization, the amount transferred, subsequent recovery, or the complete solicitation narrative. We therefore analyze reported payment following a survey-defined unsolicited money request, not verified fraud loss.

We analyze 13,766 respondents aged 18 years or older across 75 economies who met the survey-routing, outcome, and complete-case eligibility criteria. For each economy, we first estimate the survey-

weighted probability of sending the requested money among exposed respondents. We then use common-composition standardization to evaluate every economy on the same distribution of sex, age, education, workforce status, and urbanicity. Rather than treating the coefficients on these characteristics as the primary result, we compare the cross-economy distribution of payment probabilities before and after demographic composition is harmonized. The research question is: To what extent does reported payment following unsolicited money requests vary across economies among exposed adult consumers, and how much of that variation remains after the exposed populations are standardized to a common observed demographic composition?

The study makes three contributions to consumer-policy research. First, it provides a boundary test of demographic approaches to consumer vulnerability by evaluating whether five widely available demographic markers are jointly sufficient to account for a substantial portion of cross-market variation in a defined post-contact financial behavior. Second, it shifts the comparative estimand from individual regression coefficients to the distribution of economy-specific probabilities before and after common-composition standardization, making the empirical question directly relevant to the cross-market portability of demographic profiling. Third, it clarifies the interpretation of residual cross-economy heterogeneity. Variation that remains after demographic harmonization does not identify culture, regulation, institutional quality, scam sophistication, or another contextual mechanism as its cause. Instead, it identifies the limits of routinely observed demographic composition and motivates richer measurement of solicitation characteristics, verification opportunities, consumer resources, and payment environments.

2. Literature and Conceptual Framework

2.1 Consumer Vulnerability Beyond Demographic Profiling

Consumer vulnerability provides the central conceptual lens for the present study. Early consumer research cautioned against treating vulnerability as synonymous with membership in a predefined demographic group. Baker et al. (2005) conceptualize consumer vulnerability as arising from the interaction of individual states and characteristics with external conditions, while Hill and Sharma (2020) emphasize consumers' access to and control over resources required to function effectively in the marketplace. Under this view, characteristics such as age, sex, education, employment status, or place of residence may correlate with vulnerability in particular settings, but they do not constitute vulnerability itself.

This distinction has become increasingly important in contemporary consumer policy. Dodsworth et al. (2023) argue that digital environments can create or reinforce vulnerability through market architecture, regulation, and enforcement rather than merely reveal pre-existing vulnerable groups. Their account shifts attention from identifying 'vulnerable people' toward examining how interactions between consumers and their environments generate vulnerability. This contextual view is consistent with

broader arguments that vulnerability can vary across situations and need not be a permanent attribute of particular consumers.

Empirical research on financial vulnerability similarly shows why observable demographic characteristics should be treated cautiously. Seldal and Nyhus (2022), for example, examine financial vulnerability in relation to financial literacy and digital payment technologies while controlling for demographic characteristics rather than defining vulnerability through those characteristics themselves. Raval (2021), using victims identified from 23 consumer-protection enforcement cases, shows that demographic patterns vary across fraud contexts rather than forming one uniform consumer profile. Recent JCP research on scam susceptibility similarly moves beyond demographics toward psychological processes and potentially modifiable decision conditions (Waschka and Cunha 2026).

These perspectives motivate a deliberately limited role for demographics in the present study. Sex, age, education, workforce status, and urbanicity are treated as observable composition markers, not as indicators in a latent measurement model of consumer vulnerability and not as an exhaustive profile of vulnerable consumers. The empirical question is correspondingly narrower: whether differences in the observed composition of exposed consumers have sufficient descriptive leverage to account for cross-economy differences in a subsequent payment response. A finding that demographic standardization leaves substantial comparative dispersion would identify a boundary of demographic explanation rather than establish that demographics are irrelevant to vulnerability.

2.2 From Solicitation Exposure to Payment Response

Consumer vulnerability also needs to be located within the stage at which a consumer decision is observed. Contact with a potentially deceptive solicitation, engagement with it, transfer of money, realized financial loss, reporting, and recovery are analytically distinct events. Research on consumer fraud increasingly recognizes this sequencing. DeLiema et al. (2023) distinguish responses to fraud solicitations from subsequent victimization and show that their correlates vary across scam types. Xu et al. (2022) likewise separate fraud exposure from victimhood, while Yu and Fan (2024) use a two-stage approach to distinguish exposure from subsequent consumer financial fraud outcomes. These studies imply that determinants of being contacted need not be the same as determinants of behavior once contact has occurred.

The distinction is important for consumer protection because interventions may operate at different margins. Measures that block suspicious communications or reduce their reach operate primarily on exposure. By contrast, warnings, verification prompts, trusted-contact mechanisms, transaction frictions, cooling-off periods, and related safeguards can intervene after contact but before a payment decision. Robb and Wendel (2023), using experiments with more than 4,000 U.S. adults, find that inoculation-based training improves consumers' ability to distinguish digital imposter scams from legitimate communications, although the effect varies by communication mode and diminishes over time. Waschka and Cunha (2026) similarly show experimentally that introducing a period of incubation can

reduce engagement with fraudulent offers, suggesting that post-contact decision environments themselves can be modified.

The present study focuses on this post-contact analytical margin, but with a narrower observational outcome. Respondents who did not report receiving the survey-defined solicitation are not coded as consumers who resisted it; they are outside the analytical risk set. Among respondents who did report receiving the request, the outcome records whether they subsequently reported sending the requested money. Conditioning on exposure therefore avoids combining two different processes: variation in receiving or recognizing a solicitation and variation in payment behavior after reported exposure.

At the same time, this staged logic should not be used to overstate what the Global Findex measure captures. The survey does not independently establish that every unknown sender acted fraudulently, that every transfer generated a realized financial loss, or that a transfer was subsequently adjudicated as fraud. The empirical outcome is therefore described throughout this study as reported payment following a survey-defined unsolicited money request. The fraud and scam literature motivates the staged analytical design, but the observed outcome is not equated with independently verified fraud victimization.

2.3 Population Composition and Cross-Market Portability

Comparative consumer research faces an additional problem when the populations being compared differ in observable composition. International evidence indicates that fraud-related experiences vary substantially across settings. Houtti et al. (2024), for example, document considerable differences across 12 countries in scam exposure, financial loss, scam types, communication vectors, and reporting. However, observed cross-market differences can reflect several components simultaneously. One component is the composition of the consumers observed in each setting; another is variation that remains even when populations are made comparable on observed characteristics.

Consider two economies in which consumers reporting exposure to an unsolicited money request differ substantially in age, educational attainment, workforce participation, or urban residence. Their raw payment probabilities may differ partly because different kinds of consumers populate the two exposed samples. Conversely, two economies with similar observed demographic profiles may retain substantially different payment probabilities. Raw comparisons alone cannot distinguish these possibilities.

This distinction is especially important when demographic information is used to inform consumer-risk profiling. Evidence that a characteristic is associated with an outcome within a regression model does not establish that differences in that characteristic account for a meaningful portion of the distribution of outcomes across markets. Raval's (2021) findings across multiple consumer-protection cases illustrate that demographic associations can vary across fraud settings rather than forming a uniform consumer profile. More generally, the contextual approach to vulnerability advanced by Dodsworth et al.

(2023) cautions against assuming that characteristics of individuals provide a complete account of risks generated in different market environments.

The present study therefore focuses on what we call the cross-market portability of demographic profiling. We use portability descriptively rather than as a previously established theoretical construct. It refers here to the extent to which a common set of routinely observable demographic characteristics has sufficient explanatory leverage to account for comparative differences when applied across multiple market settings. Greater compression of the cross-economy distribution after composition standardization would indicate greater descriptive leverage from population composition. Limited compression would indicate that the same demographic profile provides only limited leverage for accounting for cross-market variation.

This question differs from conventional covariate adjustment. Regression coefficients summarize conditional associations under a specified statistical model. The present estimand instead asks what happens to the distribution of economy-specific payment probabilities when every economy is evaluated against the same observed composition. We therefore compare the raw cross-economy distribution with the distribution obtained after harmonizing sex, age, education, workforce status, and urbanicity. The primary quantity is not an individual demographic coefficient, nor a ranking of particular economies, but the extent to which comparative dispersion changes after common-composition standardization.

2.4 Policy Interpretation and Inference Boundary

The consumer-policy importance of the analysis lies not only in whether demographic standardization changes comparative dispersion, but also in what can legitimately be inferred from the variation that remains. Contemporary consumer-vulnerability scholarship suggests that marketplace vulnerability can reflect interactions among consumer resources, information environments, choice architectures, market practices, and regulatory safeguards. Dodsworth et al. (2023), for example, emphasize that vulnerability in digital markets may be produced or reinforced by systemic and architectural conditions rather than originating solely in personal characteristics. Intervention research likewise demonstrates that features of the decision environment can alter scam-related responses: inoculation training improves detection of imposter scams (Robb and Wendel 2023), while an incubation period can reduce susceptibility to fraudulent offers (Waschka and Cunha 2026).

These literatures provide plausible classes of explanation for why demographic composition may fail to account fully for cross-market variation. Consumers may encounter different solicitation narratives, urgency cues, impersonated identities, requested amounts, persuasion strategies, opportunities for verification, payment infrastructures, or protective interventions. Consumers may also differ in scam-specific knowledge, trust, cognition, emotional state, previous experience, social support, or access to advice. However, the Global Findex variables used in this study do not measure these mechanisms directly.

Consumer vulnerability therefore functions here as a conceptual framework for delimiting the interpretation of demographic standardization, not as a fully observed causal mechanism. If harmonizing demographic composition substantially compresses the cross-economy distribution, observable composition has meaningful descriptive leverage for the comparative pattern. If substantial dispersion persists, the result indicates that these demographic markers provide an incomplete account and that richer consumer, solicitation, transaction, and contextual measures are required.

Crucially, the latter finding cannot identify the mechanism responsible for the remaining variation. Persistent standardized dispersion is not evidence that culture causes payment behavior, that one regulatory system is superior to another, that institutional quality explains the pattern, or that scammers are more sophisticated in particular economies. Nor can it establish the effectiveness of warnings, verification systems, payment delays, or other interventions that are not directly observed. The study therefore uses residual comparative dispersion as a diagnostic boundary: it identifies what routine demographic composition does not account for while leaving the causes of that remaining variation to research designs capable of measuring or identifying them.

3. Data and Methods

3.1 Data Source and Adult Analytical Population

This cross-sectional secondary analysis uses the public-use world microdata from the Global Findex Database 2025 (reference WLD_2024_FINDEX_v02_M), collected in 2024 by the World Bank in collaboration with Gallup (Development Research Group, Finance and Private Sector Development Unit 2025; Klapper et al. 2025). The underlying survey covers the civilian, non-institutionalized population aged 15 years or older within participating economies. The public world file analyzed here contains 144,090 respondents and 140 unique economy identifiers. The unsolicited-money-request item used in this study was administered in 75 economies; consequently, all estimands refer to respondents in these 75 economies rather than to the Global Findex population as a whole, and no claim of worldwide prevalence is made.

Questionnaire routing defines the initial analytical risk set. Mobile-phone owners were asked whether, during the preceding 12 months, they had received a phone call, SMS, or text message on their mobile phone from someone they did not know asking them to send money (con21). Respondents answering yes were then asked whether they sent the money to that person (con22). Although the dataset metadata identify these as scam-related items, the literal questions do not independently establish fraudulent intent, adjudicated fraud, realized financial loss, or subsequent recovery. We therefore describe con21 as reported exposure to the survey-defined unsolicited money request and con22 as the subsequent reported payment response.

The full file contains 144,090 observations. Of these, 76,898 respondents are located in the 75 economies in which the item was administered; 64,663 are mobile-phone owners to whom con21 was

administered; 64,509 provide a valid yes/no response to con21; 14,359 report receiving the request; and 14,347 provide a valid yes/no response to con22. Requiring complete information on sex, age, education, workforce status, urbanicity, the released respondent weight, and economy identifier yields 14,242 observations. For the present secondary analysis, we impose an additional adult-only eligibility criterion of age 18 years or older. This excludes 476 respondents aged 15–17 and yields the primary analytical sample of 13,766 adults across 75 economies, of whom 1,188 report sending the requested money. Study size is determined by all respondents meeting these routing, complete-case, and age eligibility criteria; no prospective sample-size calculation was performed.

The analysis is therefore exposure conditioned. Respondents who did not report receiving the request are outside the risk set and are not coded as consumers who declined or resisted the request. The estimands describe reported payment behavior among exposed adult mobile-phone owners in the 75 observed economies and do not estimate susceptibility in the general adult population or correct for nonrandom selection into exposure.

3.2 Outcome and Composition Variables

The binary outcome equals one when an exposed adult respondent reported sending the requested money and zero when the respondent reported not sending it. We refer to this outcome as reported payment following exposure. It should not be interpreted as independently verified fraud victimization or confirmed financial loss because the public-use survey does not establish fraudulent intent, the amount transferred, adjudication, or recovery.

Five observed characteristics define the primary composition-adjustment set: sex, age, education, workforce status, and urbanicity. These variables were selected for cross-economy availability, comparability, and interpretability rather than because of favorable adjusted associations with the outcome. They are treated as population-composition markers, not as a measurement model of consumer vulnerability. Age is represented using a restricted cubic spline with fixed boundary knots at 18 and 63 years and internal knots at 27 and 39 years. Retaining these fixed knots after restricting the analytical population to adults preserves the previously defined functional form rather than selecting a new age specification after inspection of the adult-only results. Education is entered categorically.

Within-economy household-income quintile is excluded from the primary specification because it is unavailable for all observations from Lesotho and would therefore alter the comparative economy universe; it is included in a 74-economy sensitivity analysis. A secondary digital-characteristics extension additionally includes phone type (con9), ability to read a text message (con14), use of a PIN or password on the mobile phone (con18), and recent internet use (internet_use). These variables are treated as descriptive extensions rather than pre-exposure protective factors because their temporal ordering relative to the preceding 12-month solicitation and payment events cannot be established.

3.3 Comparative Estimands and Equal-Economy Reference Distribution

For each economy c , the raw estimand, $\theta_{\text{raw},c}$, is the Findex-weighted probability of reported payment among respondents in the exposure-conditioned adult analytical sample. Cross-economy dispersion is summarized by the equal-economy 90th-to-10th percentile spread, $H_{\text{raw}} = P90(\theta_{\text{raw},c}) - P10(\theta_{\text{raw},c})$. The percentile spread is preferred to the full observed range because estimates from the most extreme economies are particularly sensitive to sparse payment counts.

The standardization target is an analyst-defined equal-economy reference distribution. Within each economy, the relative released Findex respondent weights are retained; across economies, weights are normalized so that every economy contributes the same total mass. The resulting analyst-constructed weight is not an official World Bank weight. It operationalizes the comparative estimand by preserving within-economy weighting while preventing economies with larger exposed samples from dominating the common reference distribution. The reference distribution is a synthetic comparative population and should not be interpreted as the population distribution of the world or of the 75 economies combined.

Equal-economy weighting is therefore part of the substantive estimand rather than a correction to the official World Bank weights. It treats each economy as one comparative consumer-market setting rather than weighting the cross-market analysis by the number of exposed respondents observed in the public-use file. Alternative weighting schemes are examined as sensitivity analyses because they correspond to different pooled comparative populations.

3.4 Common-Composition Standardization

The primary standardization model is an equal-economy-weighted linear probability model containing economy fixed effects and the five primary composition variables. The model serves primarily as a prediction engine for marginal standardization; individual regression coefficients are not the primary estimands.

For each target economy, the economy indicator is assigned to that economy for every respondent in the common adult reference distribution while the observed composition variables are retained. Predicted probabilities are then averaged over the common reference distribution. The resulting quantity, $\theta_{\text{std},c}$, represents the model-based probability of reported payment for economy c if every economy were evaluated on the same observed adult composition. The standardized cross-economy spread is $H_{\text{std}} = P90(\theta_{\text{std},c}) - P10(\theta_{\text{std},c})$, and $\Delta H = H_{\text{raw}} - H_{\text{std}}$ denotes descriptive attenuation following composition harmonization.

Standardization is descriptive rather than causal. The primary model allows economy-specific intercepts but imposes common covariate slopes across economies. Standardized probabilities are therefore model-based comparative quantities under this working specification rather than nonparametrically identified causal economy effects. Standardization does not eliminate selection into solicitation exposure, equalize the solicitations encountered across economies, or identify contextual mechanisms responsible for residual differences.

A linear probability model is used in the primary specification because the estimands are defined directly on the probability scale. Its suitability is assessed by the proportion of individual fitted values outside the unit interval and by the range of standardized economy probabilities. Fixed-effects logit and probit models provide nonlinear sensitivity analyses.

3.5 Uncertainty and Secondary Demographic Diagnostics

Uncertainty in Delta H is evaluated using 999 paired respondent-level bootstrap replicates. Respondents are resampled with replacement within economy while each economy's observed adult sample size is preserved, and the released respondent weights are retained in re-estimation. The same replicate is used to calculate raw and standardized dispersion so that uncertainty in their difference is evaluated jointly.

The public-use Global Findex file provides final respondent weights but does not contain the original primary sampling-unit and stratum identifiers required to reconstruct full multistage survey variance. The bootstrap interval should therefore be interpreted as a within-economy respondent-resampling assessment conditional on the released weights and observed public-use data structure, not as a fully design-based survey interval.

Associations between the five demographic families and reported payment are reported only as secondary pooled diagnostics. HC3 covariance estimates are used and multiplicity is controlled with the Holm procedure across sex, age, education, workforce status, and urbanicity. These diagnostics are not interpreted as causal effects, do not carry the primary comparative claim, and are not tests of universal demographic irrelevance.

3.6 Sensitivity Analyses and Sparse-Event Diagnostics

Sensitivity analyses examine whether the standardized cross-economy spread is materially altered by reasonable analytical alternatives. The weighting checks use (a) the released Findex weights without equal-economy normalization, (b) equal respondent weights within equally weighted economies, and (c) released weights trimmed at the within-economy 99th percentile before equal-economy normalization. Age is alternatively represented linearly and using categories of 18–24, 25–34, 35–44, 45–54, 55–64, and 65 years or older. Additional specifications introduce household-income quintile, the four digital characteristics described above, and fixed-effects logit and probit response models. Influence is assessed by omitting each economy in turn.

Sparse-event robustness is assessed by restricting the analysis to economies containing at least 10 reported payment events. This restriction retains 42 economies and 9,330 respondents in the adult analytical population. These sensitivity analyses are stress tests of the descriptive pattern; they are not identification strategies for the causes of residual heterogeneity.

3.7 Analysis-Plan Transparency

Descriptive feasibility auditing preceded the original adjusted analysis and revealed questionnaire routing, sample sizes, missingness, raw outcome frequencies, weight distributions, and sparse-event

structure. The original multivariable specification and robustness sequence were fixed before inspection of the primary composition-standardized result, although the study was not preregistered before access to the data.

During subsequent manuscript revision, the analytical population was conservatively restricted to respondents aged 18 years or older. This age restriction was introduced after the original analysis and is therefore not described as prespecified. Once the adult-only boundary was adopted, the primary model and the previously defined robustness sequence were rerun without result-driven modification of the underlying estimands, covariate set, spline knots, or weighting logic.

3.8 Ethical and Research-Integrity Considerations

This study is a secondary analysis of de-identified public-use survey microdata from the Global Findex Database 2025. The authors did not recruit, contact, or intervene with respondents and did not collect identifiable individual-level information. The source survey covers respondents aged 15 years or older; however, the present secondary analysis is restricted to adults aged 18 years or older, and respondents aged 15–17 are excluded from all primary and reported robustness analyses in the revised study.

No additional ethics approval was obtained for the present secondary analysis.

During preparation of the manuscript, OpenAI ChatGPT was used to assist with literature organization, language and structural refinement, and drafting and debugging reproducible analysis code. The tool did not generate, modify, or replace the underlying research data. The authors independently checked analytical decisions, code, statistical outputs, cited sources, and the final manuscript and take responsibility for the accuracy, originality, and integrity of the submitted work.

Table 1
Characteristics of the Exposure-Conditioned Adult Analytical Sample

Characteristic	Unweighted n (%) / Mean (SD)	Equal-economy reference % / Mean (SD)
Age, years	37.39 (14.39)	36.39 (14.01)
Sex		
Female	6,881 (50.0%)	44.8%
Male	6,885 (50.0%)	55.2%
Education		
Primary or less	3,902 (28.3%)	38.2%
Secondary	8,012 (58.2%)	48.8%
Tertiary or more	1,852 (13.5%)	13.0%
Workforce status		
In workforce	10,096 (73.3%)	71.1%
Out of workforce	3,670 (26.7%)	28.9%
Residence		
Rural	7,673 (55.7%)	57.1%
Urban	6,093 (44.3%)	42.9%
Respondents	13,766	
Economies	75	
<i>Note. The equal-economy reference column applies the within-economy Finindex respondent weights and assigns each economy equal total mass. It describes the analyst-defined standardization population and is not an estimate of the population composition of the world or of the 75 economies combined.</i>		

4. Results

4.1 Analytical Sample and Outcome Structure

The primary analytical sample comprised 13,766 respondents aged 18 years or older who reported receiving the survey-defined unsolicited money request across 75 economies and met the complete-case eligibility criteria. Of these respondents, 1,188 reported sending the requested money. The adult-only restriction excluded 476 respondents aged 15–17 from the otherwise eligible complete-case sample while retaining all 75 economies.

Analytical sample sizes varied across economies. The median adult analytical sample was 156 respondents per economy, with a 10th-to-90th percentile range of approximately 64 to 323 respondents. The median number of reported payment events was 12 per economy. Thirty-three of the 75 economies contained fewer than 10 such events, and one economy contained none. These features reinforce the rationale for summarizing the cross-economy distribution using the 90th-to-10th percentile spread rather than the full observed range and motivate the sparse-event sensitivity analysis reported below. Table 1 describes the composition of the adult analytical sample and the equal-economy reference distribution used for standardization.

4.2 Raw and Common-Composition-Standardized Cross-Economy Dispersion

Reported payment probabilities varied substantially across the 75 economies before composition standardization. The equal-economy median raw probability was 8.13%. The 10th and 90th percentiles of the Findex-weighted economy-specific probabilities were 2.39% and 19.14%, respectively, yielding an equal-economy 90th-to-10th percentile spread of $H_{\text{raw}} = 16.75$ percentage points.

After every economy was evaluated on the same observed distribution of sex, age, education, workforce status, and urbanicity, the corresponding 10th and 90th percentiles were 2.79% and 18.66%, and the standardized median was 7.84%. The standardized spread was therefore $H_{\text{std}} = 15.86$ percentage points. Common-composition standardization reduced the spread by $\Delta H = 0.89$ percentage points. The paired respondent-level bootstrap 95% interval for this attenuation ranged from -0.57 to 1.45 percentage points. Because the interval includes zero, the data do not establish a precisely positive reduction in cross-economy dispersion following demographic composition standardization. The more stable empirical feature is the magnitude of the dispersion that remains: the cross-economy distribution continues to span approximately 16 percentage points after all economies are evaluated against the same observed demographic composition.

The primary linear probability model also satisfied the prespecified prediction-range diagnostic. Approximately 2.09% of individual fitted probabilities fell outside the unit interval, below the 5% diagnostic threshold, while the standardized economy probabilities remained within the required technical range.

Secondary pooled demographic diagnostics did not provide multiplicity-adjusted evidence of a common association for any of the five demographic families. The adjusted female-male difference was -0.56 percentage points (95% HC3 CI -2.14 to 1.02); secondary versus primary-or-less education was -1.69 points (-3.70 to 0.33); tertiary versus primary-or-less education was -1.13 points (-3.73 to 1.46); workforce participation was $+1.35$ points (-0.57 to 3.27); and urban versus rural residence was -0.74 points (-2.54 to 1.05). The joint age-spline test and the other demographic-family tests did not survive Holm adjustment. These secondary results are pooled descriptive associations and do not constitute tests of universal demographic irrelevance.

Table 2
Cross-Economy Dispersion in Reported Payment Following Unsolicited Money Requests

Estimand / specification	N	Economies	P10	P90	P90-P10 spread
Panel A. Primary comparative estimands					
Raw reported-payment probability	13,766	75	2.39%	19.14%	16.75 pp
Common-composition standardized probability	13,766	75	2.79%	18.66%	15.86 pp
Descriptive attenuation, Delta H	-	75	-	-	0.89 pp
95% paired-bootstrap interval for Delta H	-	-	-	-	-0.57 to 1.45 pp
Panel B. Robustness of standardized spread					
Primary equal-economy specification	13,766	75	-	-	15.86 pp
Original Findex respondent weights	13,766	75	-	-	16.05 pp
Equal economy, unweighted within economy	13,766	75	-	-	15.23 pp
99th-percentile within-economy weight trim	13,766	75	-	-	15.85 pp
Linear age specification	13,766	75	-	-	15.85 pp
Categorical age specification	13,766	75	-	-	15.89 pp
Household-income extension	13,697	74	-	-	15.82 pp
Digital-characteristics extension	13,660	75	-	-	15.38 pp
Fixed-effects logit	13,627	74 identified	-	-	15.56 pp
Fixed-effects probit	13,627	74 identified	-	-	15.84 pp
Leave-one-economy-out	-	74 each	-	-	15.44–15.97 pp
<i>Note. Delta H is the difference between the raw and common-composition-standardized P90-P10 spreads. It is a descriptive attenuation measure, not a variance decomposition, causal explained share, or estimate of the proportion of consumer vulnerability attributable to demographics. All specifications use respondents aged 18 years or older. The digital extension additionally includes phone type, ability to read a text message, use of a PIN or password on the mobile phone, and recent internet use.</i>					

Estimand / specification	N	Economies	P10	P90	P90-P10 spread
Panel A. Primary comparative estimands					
Economies with ≥ 10 reported payment events	9,330	42	-	-	13.93 pp
<i>Note. Delta H is the difference between the raw and common-composition-standardized P90-P10 spreads. It is a descriptive attenuation measure, not a variance decomposition, causal explained share, or estimate of the proportion of consumer vulnerability attributable to demographics. All specifications use respondents aged 18 years or older. The digital extension additionally includes phone type, ability to read a text message, use of a PIN or password on the mobile phone, and recent internet use.</i>					

Note

Each line represents one of the 75 economies in the analytical universe. Raw probabilities retain the released within-economy Findex respondent weights. Standardized probabilities evaluate every economy on the same equal-economy adult reference distribution of sex, age, education, workforce status, and urbanicity. Thick vertical segments mark P10-P90 and diamonds mark medians. The figure describes comparative dispersion and should not be interpreted as a ranking of causal economy effects.

4.3 Robustness of the Remaining Cross-Economy Dispersion

The magnitude of the standardized cross-economy spread was stable across alternative weighting schemes, age specifications, covariate extensions, nonlinear response models, and influence checks (Table 2). Changing the weighting scheme produced only modest differences. The standardized spread was 16.05 percentage points when the released Findex respondent weights were used without equal-economy normalization, 15.23 points when respondents were equally weighted within equally weighted economies, and 15.85 points after the released weights were trimmed at the within-economy 99th percentile. Alternative representations of age yielded similarly close estimates: 15.85 percentage points under a linear age specification and 15.89 points when age was represented categorically.

Expanding the adjustment set also left substantial dispersion. Adding within-economy household-income quintile in the 74-economy sample produced a standardized spread of 15.82 percentage points. Adding phone type, ability to read text messages, mobile-phone PIN or password use, and recent internet use produced a spread of 15.38 points across all 75 economies. These extensions therefore did not materially compress the comparative distribution.

Results were also similar under nonlinear probability models. Economy fixed-effects logit and probit models, estimated across the 74 economies in which the models were identified, yielded standardized spreads of 15.56 and 15.84 percentage points, respectively. Moreover, no single economy accounted for

the primary result: across 75 leave-one-economy-out re-estimations, the standardized spread ranged from 15.44 to 15.97 percentage points.

Finally, restricting the analysis to economies with at least 10 observed reported-payment events retained 42 economies and 9,330 respondents. Even under this substantially more restrictive sample, the standardized cross-economy spread remained 13.93 percentage points. The reduction relative to the full adult sample indicates that sparse cells contribute to the width of the observed distribution, but they do not account for the persistence of substantial cross-economy dispersion after common-composition standardization.

4.4 Empirical Summary and Inference Boundary

Taken together, the analyses reveal a consistent descriptive pattern. Reported payment following the survey-defined unsolicited money request varies substantially across the observed economies, and harmonizing the exposed adult populations on sex, age, education, workforce status, and urbanicity produces little compression of that distribution. The primary standardized spread remains 15.86 percentage points, and estimates remain broadly similar across weighting, age functional form, covariate extensions, nonlinear models, leave-one-economy-out analyses, and a restriction to economies with more reported payment events.

This persistence should be interpreted as a limit of demographic explanation, not as evidence for a particular alternative mechanism. The analyses do not identify whether remaining cross-economy differences arise from solicitation characteristics, consumer resources, verification opportunities, payment environments, survey or reporting processes, institutional safeguards, or other contextual conditions. Nor do they establish causal economy effects. The empirical result is narrower: routinely observed demographic composition does not account for most of the width of the observed cross-economy distribution under the present standardization design.

That boundary motivates the consumer-policy questions considered in the Discussion: what demographic profiling can and cannot accomplish as a basis for identifying consumer vulnerability across markets, and whether protection following solicitation exposure may require greater attention to decision environments and safeguards that do not depend on a stable demographic risk profile.

5. Discussion

5.1 Principal Finding: A Boundary of Demographic Explanation

The principal finding of this study is persistence rather than attenuation. Among 13,766 respondents aged 18 years or older who reported receiving the survey-defined unsolicited money request across 75 economies, the equal-economy 90th-to-10th percentile spread in reported payment probabilities was 16.75 percentage points before composition standardization and remained 15.86 percentage points

after every economy was evaluated on the same observed distribution of sex, age, education, workforce status, and urbanicity. The estimated attenuation was only 0.89 percentage points, and its paired-bootstrap confidence interval included zero.

This result is best interpreted as a boundary of demographic explanation. Differences in the demographic composition of exposed populations could, in principle, generate substantial raw cross-economy variation if some economies contain larger shares of people with characteristics associated with payment behavior. The common-composition exercise asks directly how much of that comparative distribution remains once those compositional differences are removed statistically. In the present data, harmonizing five routinely observed characteristics produces little compression of the distribution.

This does not imply that age, sex, education, workforce status, or urbanicity are irrelevant to consumer vulnerability. The secondary pooled analyses provide no multiplicity-adjusted evidence of a common association for the five demographic families under the specified model, but such pooled associations are neither universal tests of demographic relevance nor the primary estimand of the study. Demographic relationships may differ by solicitation type, decision stage, population, or setting. The narrower conclusion is that the observed distribution of these five characteristics has limited descriptive leverage for accounting for the width of the cross-economy distribution in this particular post-contact payment outcome.

The distinction between demographic association and demographic composition is therefore central. A variable can be associated with an outcome within individuals without differences in that variable accounting for much of the distributional variation between economies. Conversely, weak pooled coefficients do not by themselves establish that demographic composition is unimportant. The present contribution comes from comparing the raw and standardized economy distributions directly rather than inferring the importance of composition from a set of individual coefficients.

5.2 Consumer Vulnerability Beyond Demographic Profiles

The findings are consistent with a contextual interpretation of consumer vulnerability. Consumer-vulnerability theory has long cautioned against equating vulnerability with membership in predefined demographic groups. Baker et al. (2005) emphasize interactions between individual states and external conditions, while Hill and Sharma (2020) conceptualize vulnerability in terms of consumers' access to and control over resources needed to function effectively in marketplace situations. From this perspective, demographic characteristics may mark differences in resources or exposure to particular conditions without constituting vulnerability itself.

This distinction is particularly relevant in digital environments. Dodsworth et al. (2023) argue that consumer vulnerability can arise from the architecture and conditions of digital markets rather than simply from stable characteristics of particular groups. The consumer-policy implication is not that individual characteristics cease to matter, but that vulnerability can emerge from interactions between people and the environments in which decisions are made.

The present findings are consistent with this contextual view, but they do not constitute a direct test of the mechanisms proposed by consumer-vulnerability theory. Sex, age, education, workforce status, and urbanicity are used here as observable composition markers. The Global Findex data do not measure many resources and conditions that may be closer to a post-contact payment decision, including scam-specific knowledge, cognitive or emotional pressure, trust, access to independent advice, capacity to verify the sender, or characteristics of the payment environment. The standardized distribution therefore cannot reveal which omitted dimension of vulnerability accounts for the remaining variation.

This qualification is important because the empirical pattern should not be transformed into the opposite demographic stereotype. Finding that demographic harmonization explains little of cross-economy dispersion does not imply that all consumers face equal risks or that demographic targeting is always ineffective. Rather, it indicates that a stable cross-market account of vulnerability based primarily on the five observed demographic characteristics is incomplete.

Related JCP research reinforces the value of looking beyond static demographic profiles. Robb and Wendel (2023) show that consumers' ability to distinguish digital imposter scams from legitimate communications can be changed through inoculation-based training and that effectiveness varies with communication mode and time since training. Waschka and Cunha (2026) similarly show that modifying the decision process through an incubation period can reduce engagement with fraudulent offers, while individual differences condition some intervention effects. These studies demonstrate mechanisms and interventions that the present observational data cannot test, but they illustrate why consumer vulnerability need not be understood solely through demographic identity.

5.3 Implications for Consumer Protection

The findings have implications for how consumer protection is targeted after suspicious or unsolicited requests reach consumers. A demographic targeting strategy can be attractive because demographic information is relatively easy to obtain and can help identify populations that may merit additional attention in particular contexts. The present results do not provide a basis for abandoning such targeting. They do, however, caution against assuming that one demographic risk profile will provide sufficient or portable guidance across different market settings.

A first implication is therefore complementarity rather than substitution. Demographic targeting may be combined with safeguards that respond to characteristics of the transaction or decision situation itself. Examples include warnings triggered by unusual payment requests, recipient or account verification, opportunities to consult a trusted contact, enhanced confirmation for transfers to unfamiliar recipients, and proportionate delays or frictions in circumstances associated with elevated risk. The present data do not test any of these interventions, so their effectiveness cannot be inferred from the cross-economy results.

A second implication concerns the timing of consumer protection. Because this study conditions on respondents who already reported receiving an unsolicited request, it does not evaluate measures

intended to prevent the request from reaching them in the first place. Filtering, communication blocking, account disruption, and other exposure-reduction strategies operate at an earlier margin. The outcome studied here is more closely aligned with interventions that operate after contact but before or around the payment decision. Separating these margins avoids treating reduced exposure and improved post-contact decision support as interchangeable policy objectives.

Experimental research provides examples of such post-contact interventions. Robb and Wendel (2023) find that targeted inoculation can improve fraud detection, although effects depend partly on communication medium and decay over time. Waschka and Cunha (2026) report that an incubation period can reduce consumers' engagement with fraudulent offers and discuss policy-relevant applications such as delays and disclosure timing. These findings cannot be generalized mechanically to the heterogeneous requests represented in Global Findex, but they illustrate a broader policy principle: protection can modify the decision environment rather than depending exclusively on accurately identifying who is vulnerable in advance.

A third implication concerns cross-market portability. The modest compression produced by demographic standardization suggests caution in transferring a demographic targeting rule developed in one setting directly to another. This does not demonstrate that any particular economy requires a different policy. It indicates instead that the same observable composition leaves substantial variation in reported payment behavior across the settings studied. Consumer-protection systems may therefore benefit from combining population targeting with locally validated evidence about solicitation channels, payment practices, verification opportunities, communication environments, and intervention effectiveness.

Importantly, none of these implications requires attributing responsibility to consumers who send money. The survey does not observe the pressure, urgency, impersonation, trust-building, information asymmetry, or manipulation preceding the reported decision. The policy relevance of the findings lies in identifying opportunities for external safeguards and better decision environments, not in treating payment as evidence of gullibility or poor judgment.

5.4 Limitations and Inference Boundary

Several limitations define the scope of the conclusions. First, both the reported exposure and the subsequent payment response are self-reported. The questionnaire does not independently verify the identity or intent of the person requesting money. Sending the requested money therefore cannot be equated with independently verified fraud victimization or confirmed financial loss. The data do not report the amount transferred, whether the transfer was recovered, or whether the event was subsequently adjudicated as fraudulent. This is why the study consistently refers to a reported payment response rather than verified fraud loss.

Second, exposure is nonrandom. Respondents may differ in their probability of receiving a request, noticing it, recalling it, or reporting it. Conditioning on reported exposure is necessary for the post-

contact estimand examined here, but it does not correct for selection into exposure. Consequently, the findings do not estimate the probability that members of particular demographic groups are targeted and should not be interpreted as evidence about susceptibility to solicitation exposure.

Third, the questionnaire does not observe the content and circumstances of the solicitation in sufficient detail. It does not identify the type of request, impersonated actor, requested amount, urgency cues, persuasion strategies, legitimacy signals, trust formation, emotional pressure, or opportunities for verification. These omitted features may be important to payment decisions but cannot be separated from other sources of cross-economy variation in the present data.

Fourth, the five standardization variables are deliberately limited. They were selected because they are available and interpretable across all 75 economies, not because they constitute a complete measurement model of consumer vulnerability. Household income is available only for a reduced economy universe and is therefore examined as a sensitivity extension; the digital variables are contemporaneous measures whose temporal ordering relative to the previous 12 months of solicitation exposure cannot be established. The persistence of standardized dispersion therefore applies to the observed adjustment set, not to all possible individual-level explanations.

Fifth, the relevant questionnaire module was administered in 75 economies rather than throughout the full Global Findex universe. The findings are therefore comparative estimates for these observed economies and should not be interpreted as global prevalence estimates.

Sixth, survey administration and cross-language measurement may contribute to observed differences. Global Findex relies on standardized survey procedures, but the public-use analysis does not itself establish full measurement equivalence of the solicitation and payment items across languages, interview modes, and economies. To the extent that recall, interpretation, disclosure, or response processes vary across survey settings, part of the remaining comparative dispersion may reflect measurement processes rather than differences in underlying behavior. The present analysis cannot quantify that component separately.

Seventh, the public-use dataset provides final respondent weights but not the original primary sampling-unit and stratum identifiers needed to reconstruct the full multistage sampling variance. The reported bootstrap intervals therefore represent within-economy respondent resampling conditional on the released weights rather than fully reconstructed design-based uncertainty.

Eighth, common-composition standardization remains model based. The primary specification uses economy-specific intercepts and common covariate slopes, although alternative age specifications, weighting schemes, and nonlinear models produce similar standardized spreads. Robustness across specifications reduces concern about dependence on one working model but does not make the standardized probabilities nonparametric or causal.

Finally, and most importantly for interpretation, the residual economy variation is not an estimate of an economy effect. The study is cross-sectional and observational. The standardized differences cannot establish causal effects of culture, regulation, institutional quality, economic development, financial infrastructure, scam prevalence, enforcement, or consumer-protection policy. None of those explanations is identified merely because demographic composition fails to account for the remaining dispersion.

5.5 Future Research

The remaining comparative dispersion defines a research agenda rather than a mechanism. A first priority is to collect more detailed information on the solicitation itself. Cross-market surveys should distinguish communication channel, solicitation type, impersonated actor, requested amount, urgency, persuasion strategy, legitimacy cues, and payment method. Such measures would allow future research to determine whether apparently cross-economy differences partly reflect differences in the kinds of requests consumers encounter rather than differences in consumer characteristics.

A second priority is to measure theoretically closer dimensions of consumer vulnerability. These include scam-specific knowledge, financial and digital competence, trust, cognitive load, emotional pressure, previous experiences with fraud, access to independent advice, social support, verification capacity, and resources available at the moment of decision. This would permit direct tests of multidimensional or contextual vulnerability rather than relying on demographic characteristics as observable proxies. Evidence from financial-vulnerability research already demonstrates that relationships between digital payment use, demographic characteristics, literacy, and vulnerability can differ across settings, underscoring the need for comparative evidence beyond one national context (Seldal and Nyhus 2022).

A third priority is to distinguish exposure, engagement, payment, realized loss, reporting, and recovery within the same longitudinal or staged design. The present study observes exposure and a subsequent reported payment response but does not observe the full process. Repeated or prospective measurement would help determine whether the sources of cross-market variation differ by stage and whether the same protective factors operate before and after contact.

A fourth priority is to move from comparative diagnosis to intervention identification. Randomized or strong quasi-experimental designs could test warnings, verification prompts, trusted-contact systems, cooling-off periods, transaction delays, recipient verification, and other decision safeguards. The effectiveness of such interventions should be evaluated not only on average but also across demographic groups, communication channels, payment contexts, and market settings. Existing experiments on inoculation and incubation demonstrate that post-contact consumer responses can be modified, but they also indicate that effects may vary by context and intervention design (Robb and Wendel 2023; Waschka and Cunha 2026).

Finally, future comparative work should test contextual explanations prospectively and selectively rather than by searching across large numbers of country indicators after observing the outcome distribution.

Measures of regulation, payment-system design, enforcement, institutional protection, or digital-market conditions may eventually explain part of the remaining dispersion, but such hypotheses should be grounded in explicit mechanisms and designs capable of distinguishing contextual effects from composition, measurement, and selection.

The present study therefore stops at a deliberately limited conclusion: substantial cross-economy variation in reported payment remains after routinely observed demographic composition is harmonized. Determining why it remains, and which consumer-protection measures reduce harmful payment responses, requires richer measurement and designs capable of identifying those mechanisms.

6. Conclusion

This study examined whether differences in routinely observed demographic composition account for cross-economy variation in reported payment responses following unsolicited money requests. Among 13,766 exposed adults across 75 economies, the equal-economy 90th-to-10th percentile spread in reported-payment probabilities was 16.75 percentage points before common-composition standardization and remained 15.86 percentage points after every economy was evaluated on the same observed distribution of sex, age, education, workforce status, and urbanicity. The estimated attenuation was 0.89 percentage points, and its paired-bootstrap 95% confidence interval ranged from – 0.57 to 1.45. The persistence of substantial standardized dispersion was also robust across alternative weighting schemes, functional forms, covariate extensions, nonlinear response models, influence checks, and a restriction to economies with more observed payment events.

The contribution of this result is a boundary rather than an explanation of the residual. The findings do not show that demographic characteristics are irrelevant to consumer vulnerability, nor do they identify culture, institutions, regulation, economic development, or any other economy-level condition as the cause of the remaining variation. Instead, they show that harmonizing five widely observed demographic characteristics does little to compress the comparative distribution for this specific post-contact payment outcome. For consumer-policy research, this limits how far a common demographic risk profile can be expected to account for cross-market variation and strengthens the case for directly studying the solicitation, transaction, verification, and decision environments in which consumer vulnerability arises. The next empirical task is not to label the remaining variation, but to identify the mechanisms and consumer-protection safeguards that actually change post-contact payment decisions.

Declarations

Author Contribution

Phan Huy Duong: Writing – review and editing of the initial manuscript. Le Hong Huyen: Methodology. Le Quang Minh: Formal analysis. Vu Thi Mai Duyen: Writing – review and editing; manuscript revision and

finalization. Vu Thi Mai Duyen serves as the corresponding author. All authors critically reviewed the manuscript, approved the final version, and agree to be accountable for the work.

References

1. Baker, S. M., Gentry, J. W., & Rittenburg, T. L. (2005). Building understanding of the domain of consumer vulnerability. *Journal of Macromarketing*, 25(2), 128–139. <https://doi.org/10.1177/0276146705280622>
2. DeLiema, M., Li, Y., & Mottola, G. (2023). Correlates of responding to and becoming victimized by fraud: Examining risk factors by scam type. *International Journal of Consumer Studies*, 47(3), 1042–1059. <https://doi.org/10.1111/ijcs.12886>
3. Development Research Group, Finance and Private Sector Development Unit (2025). The Global Findex Database 2025: Connectivity and Financial Inclusion in the Digital Economy [Data set]. World Bank. <https://doi.org/10.48529/BK9N-8R43>.
4. Dodsworth, T., Riefa, C., Saintier, S., & Twigg-Flesner, C. (2023). Editorial: Digital vulnerability. *Journal of Consumer Policy*, 46, 413–417. <https://doi.org/10.1007/s10603-023-09555-6>
5. Hill, R. P., & Sharma, E. (2020). Consumer vulnerability. *Journal of Consumer Psychology*, 30(3), 551–570. <https://doi.org/10.1002/jcpy.1161>
6. Houtti, M., Roy, A., Gangula, V. N. R., & Walker, A. M. (2024). A survey of scam exposure, victimization, types, vectors, and reporting in 12 countries. *Journal of Online Trust and Safety*, 2(4). <https://doi.org/10.54501/jots.v2i4.204>
7. Klapper, L., Singer, D., Starita, L., & Norris, A. (2025). The Global Findex Database 2025: Connectivity and Financial Inclusion in the Digital Economy. World Bank. <https://doi.org/10.1596/978-1-4648-2204-9>
8. Raval, D. (2021). Who is victimized by fraud? Evidence from consumer protection cases. *Journal of Consumer Policy*, 44(1), 43–72. <https://doi.org/10.1007/s10603-020-09466-w>
9. Robb, C. A., & Wendel, S. (2023). Who can you trust? Assessing vulnerability to digital imposter scams. *Journal of Consumer Policy*, 46(1), 27–51. <https://doi.org/10.1007/s10603-022-09531-6>
10. Seldal, M. M. N., & Nyhus, E. K. (2022). Financial vulnerability, financial literacy, and the use of digital payment technologies. *Journal of Consumer Policy*, 45, 281–306. <https://doi.org/10.1007/s10603-022-09512-9>
11. Waschka, L., & Cunha, M. Jr. (2026). Reducing consumer fraud and scam victimization through incubation. *Journal of Consumer Policy*, 49. <https://doi.org/10.1007/s10603-026-09619-3>. Article 15.
12. Xu, L., Wang, J., Xu, D., & Xu, L. (2022). Integrating individual factors to construct recognition models of consumer fraud victimization. *International Journal of Environmental Research and Public Health*, 19(1), 461. <https://doi.org/10.3390/ijerph19010461>
13. Yu, Z., & Fan, J. X. (2024). Migrant status and consumer financial fraud in China: A two-stage approach using a representative household survey. *SAGE Open*, 14(2), 1–16.

Figures




Figure 1

Raw and Common-Composition-Standardized Economy Probabilities of Reported Payment

Note. Each line represents one of the 75 economies in the analytical universe. Raw probabilities retain the released within-economy Findex respondent weights. Standardized probabilities evaluate every economy on the same equal-economy adult reference distribution of sex, age, education, workforce status, and urbanicity. Thick vertical segments mark P10-P90 and diamonds mark medians. The figure describes comparative dispersion and should not be interpreted as a ranking of causal economy effects.