Signaling at a Loss: Experimental Evidence on Conspicuous Consumption in Tanzania

Stine Helmke*

November 8, 2026

Click here for the latest version

Abstract

Individuals in low-income settings invest considerable resources in signals like conspicuous attire. But signaling does not always pay off. The same signal may be beneficial in some domains and disadvantageous in others. Using two experiments in Dar es Salaam, Tanzania, I show that senders signal even when it backfires. In the first experiment, I establish that conspicuous attire enhances professional prestige, leaves credit approval unaffected, and reduces friendship formation and informal help. In the second experiment, senders believe wearing conspicuous attire is beneficial in every domain and give up about 60% of daily per-capita income to dress conspicuously even where it backfires. A structural decomposition attributes at least 74% of what they give up to wrong beliefs, rather than intrinsic motives. Senders appear to hold misperceptions about both which traits the signal conveys and how these traits are valued. Reducing the welfare losses from signaling requires correcting both.

*Department of Economics, University of Zurich. Contact: stine.helmke@econ.uzh.ch. The paper was previously circulated as 'Stuck in Signaling: Experimental Evidence on Conspicuous Consumption in Tanzania'. I am especially grateful to my advisors Lorenzo Casaburi, Roberto Weber, Arun Chandrasekhar and Marek Pycia for their ongoing guidance and support. I thank the enumerator team led by Ibrahim Kisinza, Mansura Shemela, and Lenna Mutalemwa for outstanding field work in Dar es Salaam. I further thank Marina Agranov, Ingvild Almås, Sandro Ambuehl, Douglas Bernheim, Ana Brás Monteiro, Katherine Coffman, Esther Duflo, Pascaline Dupas, Marcel Fafchamps, Ernst Fehr, Matteo Filippi, Thomas Graeber, Andreas Hefti, Ronak Jain, Eliana La Ferrara, Melanie Morten, Muriel Niederle, Sara Rabino, Frank Schilbach, Florian Schneider, Alexandra Schubert, Philipp Sternal, Alessandra Voena, Joachim Voth, David Yanagizawa-Drott, Andreas Ziegler, and seminar participants at the University of Zurich, Stanford University, ESA, NCBEE, DENS, and ESRF Dar es Salaam for helpful comments and feedback. This research was funded by the Excellence Foundation Zurich for Economic and Social Research, the URPP Equality of Opportunity initiative, and the Graduate Research Campus of the University of Zurich. All parts of this study received IRB approval from the University of Zurich and were covered by a research permit from COSTECH (Commission for Science and Technology, Tanzania). Both experiments are preregistered in the AEA RCT Registry (AEARCTR-0017572 and AEARCTR-0018398). All errors are my own.

1 Introduction

Individuals invest considerable resources in signaling, trying to shape how others see them (Bloch et al., 2004; Banerjee and Duflo, 2007; Bursztyn et al., 2018). This is particularly true in low-income settings where conspicuous goods alone account for seven to thirteen percent of monthly household consumption (Burger et al., 2015; Jaikumar et al., 2018).[1] In the standard economic view, signaling is a rat race: wasteful for society, but pays off for the individual (Veblen, 1899; Akerlof, 1976; Frank, 1985; Hopkins and Kornienko, 2004). This paper shows that signaling can be wasteful for the individual too.

Knowing when signaling pays off is hard. The same signal may be beneficial in some domains and disadvantageous in others: a woman's engagement ring signals stability, raising her social standing; but it also signals inflexible working hours, decreasing her chances of being hired. Luxury attire signals wealth and success, raising professional esteem; but it also signals social distance, decreasing cooperation. Knowing when signaling pays off requires understanding two things: (i) which traits a signal conveys, (ii) how these traits are valued in different domains.

This paper asks whether individuals signal even where it backfires and why. I provide joint experimental evidence on the differential effects of one signal, conspicuous attire, across four socio-economic domains, and test whether senders anticipate these domain-specific effects. The setting of this paper is Dar es Salaam, Tanzania, where signaling with attire is pervasive. Participants in my sample spend 23% of their monthly income on clothing, while living on incomes of about $7 PPP a day per household member. In the first experiment, I measure the causal effect of wearing conspicuous versus neutral attire on evaluations in four socio-economic domains, by relevant decision-makers. Wearing conspicuous attire pays off in some domains but not in others: it enhances professional prestige, leaves credit approval unchanged, and lowers friendship formation and informal help. In the second experiment, a representative sample of adults from Dar es Salaam predicts how they would be evaluated wearing conspicuous versus neutral attire in each domain. They know where conspicuous attire pays off most and where it pays off least but believe it pays off in every domain. For friendship formation, they believe that wearing conspicuous attire helps when it actually backfires. When asked to make an incentivized choice of attire, individuals are willing to give up about 60% of daily per-capita income to be evaluated in conspicuous attire, even in domains where it backfires. A structural decomposition shows that the bulk of this spending

[1] Conspicuous consumption is understood as the practice of spending money on ostentatious goods in order to signal wealth and enhance socio-economic status (Bagwell and Bernheim, 1996).

reflects wrong beliefs about the payoff from conspicuous attire, not an intrinsic taste for wearing it. Signaling is therefore wasteful at the individual level: senders lose money, investing more in the signal than it benefits them. Finally, I ask where these misperceptions come from. Individuals misjudge both which traits the signal conveys and how these traits are valued. Both errors inflate the role of wealth. Individuals overestimate threefold how much wealth conspicuous attire conveys, and overstate by a factor of 1.4 how much evaluators value wealth relative to traits like trustworthiness and kindness. Reducing the welfare losses from signaling requires correcting both.

The first experiment establishes that conspicuous attire is beneficial in some domains and disadvantageous in others. I recruit 655 participants (*senders*) and create two standardized profiles for each of them. Profiles contain a photograph showing the participant in either neutral or conspicuous attire paired with standardized demographic information, held constant across both profiles. Neutral attire consists of standardized plain, unbranded clothing; conspicuous attire is a standardized bundle of visibly branded clothing and accessories. Profiles are evaluated in four outcome domains: professional prestige, credit approval, willingness to form a friendship with the profiled individual, and willingness to help them. In each domain, profiles are assessed by a domain-specific group of relevant decision-makers (*evaluators*). Each evaluator rates a series of randomly allocated profiles on a seven-point Likert scale. Since the only difference between the two profiles of the same individual is their attire, this design identifies the causal effect of conspicuous versus neutral attire on evaluations in each domain.

Relative to neutral attire, conspicuous attire increases the probability of a positive evaluation of professional prestige by 5.7 percentage points, has no detectable effect on evaluations of credit approval, reduces the probability of a positive evaluation of friendship formation by 6.7 percentage points, and of willingness to help by 12 percentage points. This provides joint evidence that the same visible signal can generate sharply different effects in both sign and magnitude across socio-economic domains. These findings extend the prevailing view that visible signals of status pay off: luxury brand labels on clothing elicit preferential treatment from strangers (Nelissen and Meijers, 2011), high-status appearance protects against social marginalization (Dupas et al., 2024), and obesity improves access to credit (Macchi, 2023). Measuring the effect of the same signal across several domains establishes that status signals can also backfire: the attire that raises professional prestige lowers friendship formation and willingness to help. Trait ratings from the same evaluators point to why: conspicuous attire raises perceived wealth by 11.3 percentage points but lowers perceived trustworthiness by 4.8 and kindness by 8.1 percentage points.

The second experiment shows that senders overestimate the payoff from conspicuous

attire in every domain, but know where it pays off most and where it pays off least. I recruit 502 participants from the original pool of senders. For each of the four outcome domains, I separately elicit incentivized beliefs about the probability that each of the two profiles receives a positive evaluation from evaluators. The difference between the two profiles is the predicted payoff from dressing conspicuously. Senders overestimate this payoff by 27 percentage points for professional prestige, by 29 percentage points for credit approval, by 28 percentage points for friendship formation, and by 14.6 percentage points for informal help. Senders therefore recover the correct ranking of payoffs across domains but are wrong about the levels, shifting the entire payoff distribution uniformly upward, above zero. As a result, they believe it pays off to dress conspicuously in every domain. For friendship, this shift crosses zero: senders perceive a positive payoff where the realized payoff is negative.

In every domain, senders give up more money to be evaluated in conspicuous attire than it earns them. To obtain a measure of actual monetary investment in conspicuous attire, I use a multiple price list to elicit willingness to accept (WTA) to have the conspicuous rather than the neutral attire profile submitted to a payed future evaluation. I benchmark WTA against potential earnings: how much more the conspicuous profile would earn than the neutral profile, based on the evaluations from the first experiment. Senders' WTA exceeds these potential earnings in every domain. For professional prestige, mean WTA of TZS 4,028 ($6 PPP) compares to potential earnings of TZS 537 ($0.81 PPP); for credit approval, mean WTA is TZS 3,245 ($5 PPP) when potential earnings are indistinguishable from zero. For friendship formation, senders exhibit a positive WTA of TZS 2,857 ($4 PPP) when potential earnings are −TZS 619 (-$1 PPP). This means senders are willing to give up 60% of one day's average income per household member for submitting the conspicuous profile, when it lowers expected earnings by a further 13%. For informal help, mean WTA is negative: −TZS 317 (−$0.48 PPP), but still falls short of the potential earnings of −TZS 1,107 (−$2 PPP).

Whether senders actually waste money on signaling depends on why they overspend: if WTA reflects intrinsic taste or identity (Akerlof and Kranton, 2000), the money buys utility; if it reflects wrong beliefs about the payoff from conspicuous attire, it buys nothing and goes to waste. I show that wrong beliefs about payoffs account for the bulk of WTA. I model sender utility as consisting of a belief-driven instrumental component and an intrinsic utility component. This structure is supported by strong correlations between beliefs and WTA. A structural decomposition of WTA calibrates sender utility with elicited beliefs. It shows that the belief-driven utility component accounts for at least 74% of WTA in professional prestige, credit approval, and friendship. In help, beliefs pull towards conspicuous attire, while intrinsic motives partially offset them. Overall, this means that misperceptions largely drive spending on conspicuous attire. To quantify how much money senders lose on signaling, I compare the

belief-implied share of WTA to potential earnings, adjusting for risk aversion: misperceptions are costly. Senders lose TZS 950 ($1.4 PPP) per signaling choice in help and up to TZS 2,900 ($4 PPP) in the other three domains, equivalent to 60% of daily per-capita income. Individuals are known to misperceive the payoffs to consequential economic investments such as schooling (Jensen, 2010; Dizon-Ross, 2019) and migration (McKenzie et al., 2013; Baseler, 2023). I show that this extends to investments in signaling, at measurable welfare costs.

Misperceptions are costly, but what can correct them depends on which layer of signaling they sit in. In a final step, I locate the misperceptions, writing the payoff from signaling as the product of its two layers: (i) which traits the attire conveys (wealth, trustworthiness, kindness), and (ii) how these traits are valued across domains (e.g., professional prestige, friendship). I examine possible misperceptions in each layer in turn.

Complementary experimental data on trait evaluations and beliefs show that misperceptions sit in both layers: senders overestimate how much wealth the attire conveys, and how much evaluators value wealth relative to other traits. I establish this in two steps. First, do senders misperceive which traits conspicuous attire conveys? Using elicitations from the same evaluators and senders as in the main experiments, I measure the causal effects of conspicuous attire on perceived wealth, trustworthiness, kindness, and competence. Senders predict accurately how much trustworthiness, kindness and competence conspicuous attire conveys, but overestimate threefold how much wealth it conveys. A first source of misperceptions thus lies in the traits conveyed by the signal. Second, do senders misperceive how these traits are valued? Since beliefs shift the realized payoff distribution upward almost uniformly, leaving the ranking of domains intact, misperceptions in this layer must sit in the domain-invariant part of the signal. Using data on the perceived valuation of wealth in each domain, I construct the belief a sender would hold if wealth were the only valued trait in each domain. This benchmark tracks senders' beliefs closely, and a back-of-the-envelope calibration suggests that senders overstate by a factor of 1.4 the valuation of wealth relative to other traits. A second source of misperceptions thus lies in the valuation of traits. Senders overweight the role of wealth, the most salient trait of conspicuous attire throughout and underweight less salient, unfavorable traits, such as lower trustworthiness and kindness (Bordalo et al., 2013). Disentangling the two sources of misperceptions matters for policy, as they call for different remedies.

This paper speaks to the literature on costs of signaling. Economists have long viewed signaling as a rat race (Veblen, 1899; Akerlof, 1976; Frank, 1985; Hopkins and Kornienko, 2004). When everyone signals, social ranks stay the same, so for society the resources spent on signals are wasted. For the individual, however, signaling is rational: those who can afford to signal and do not lose standing to those who do. And in low-income settings,

individuals spend considerable resources on conspicuous goods (Bloch et al., 2004; Banerjee and Duflo, 2007; Jaikumar and Sharma, 2021) despite costs such as redistributive claims from kin (Jakiela and Ozier, 2016; Beekman et al., 2015; Squires, 2024; Carranza et al., 2025). I show that signaling can be wasteful for the individual too: in some domains, conspicuous attire lowers the very evaluations it is meant to raise, senders do not anticipate this, and they pay for it nonetheless. These findings are consistent with related work in psychology documenting negative social returns to conspicuous consumption (Srna et al., 2022; Garcia et al., 2019). I go beyond this work by establishing that the same signal generates both positive and negative payoffs across domains, demonstrating that misperceptions of these payoffs lead to measurable welfare losses, and tracing the sources of these misperceptions.

The payoff variation that I document also differs from known sources of heterogeneity in signaling payoffs. Payoffs are known to vary by the sender's type (Bagwell and Riordan, 1991; Feltovich et al., 2002) and by context (Bellezza et al., 2014). In these cases, senders do not need to know the full payoff schedule of the signal to benefit from it: they need only learn the optimal strategy for their own type, or move to a context where their signal works. The variation I document in this paper is different: a sender who wants friendship as well as prestige must know the payoff schedule across domains. Farrell and Gibbons (1989) consider a sender whose message reaches two audiences at once, so that a single message must balance preferences of both. There, a cost to signaling is imposed through a constraint in the signaling environment. I study the unconstrained case, in which each signal can be tailored to a single domain. Any losses I document therefore originate within the sender, and are inherent to signaling itself rather than to the environment.

Reducing the welfare losses from signaling requires correcting both sources of misperceptions, and the two differ in reach. Misperceptions of what a signal conveys can be corrected at the level of the signal, by making common knowledge what others infer from conspicuous attire. Misperceptions of how the conveyed traits are valued attach to the traits themselves and travel with every signal of them; correcting them requires making evaluation criteria public and binding. Absent such policies, individuals in many low-income settings remain stuck in costly signaling.

The rest of this paper is structured as follows. Section 2 describes the setting and the experimental flow. Section 3 presents the *Evaluations Experiment*; profile creation, design, and results. Section 4 presents design and results of the *Beliefs Experiment*. Section 5 structurally decomposes willingness to accept. Section 6 develops the conceptual framework and locates misperceptions of senders, and Section 7 concludes.

2 Setting and Experimental Setup

2.1 Setting

I study the signaling payoffs from conspicuous versus neutral attire in Dar es Salaam, Tanzania. Signaling plays an important role in shaping socio-economic opportunities in this context as formal institutions that would otherwise substitute for visible cues remain sparse. Credit is screened informally: only 31.4% of Dar es Salaam households hold a formal bank account, and just 7% of Tanzanians are aware that credit reference bureaus exist (Financial Sector Deepening Tanzania (FSDT), 2023; National Bureau of Statistics (NBS), Tanzania, 2018). Without credit histories, lenders rely on observable proxies for creditworthiness, including physical appearance (Macchi, 2023). Similarly, the informal sector accounts for 62% in Dar es Salaam (National Bureau of Statistics (NBS), Tanzania, 2021), such that labor-market opportunities are frequently allocated without formal credentials. Finally, social support and insurance in many low-income settings runs through informal networks rather than formal institutions (Townsend, 1994; Udry, 1994).

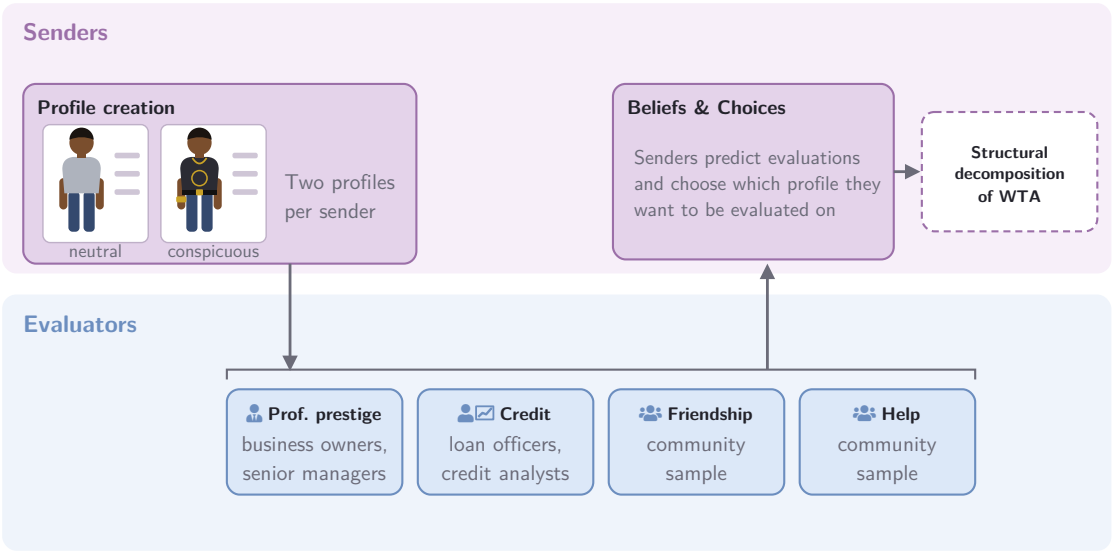
Luxury attire and branding are the signal of choice in many emerging African consumer markets, including among low-income consumers (Boston Consulting Group, 2021). Tanzania specifically is one of the fastest-growing markets for luxury goods consumption globally (Statista, 2024) with Dar es Salaam being the country's commercial hub: its port handles approximately 95% of Tanzania's international trade and serves as the primary gateway for imported goods, including branded and second-hand clothing (Tanzania Ports Authority, 2026).

Spending on informal signals like attire is economically meaningful, as financial strain is widespread in this setting. Two thirds of Tanzanian adults report often or always struggling to keep up with regular expenses (Financial Sector Deepening Tanzania (FSDT), 2023). This is consistent with the study sample in this project which is broadly representative of the adult population of Dar es Salaam (Table A1): households support an average of four people on a monthly income of approximately TZS 564,206 ($853 PPP), equivalent to roughly TZS 142,548 ($215 PPP) per household member per month. Borrowing and debt are widespread; in pilot data, 74% of participants report having taken a loan from a bank, microfinance institution, or other lender at some point in their lives. Nonetheless, households invest considerable resources in attire. In the study sample, individuals spend on average TZS 33,446 ($51 PPP) per month, on clothing and footwear, which accounts for roughly 23% of an individual's per-capita share of household income.

2.2 Experimental Setup

Studying signaling payoffs in this setting requires observing both how conspicuous attire is evaluated and whether the people who wear it, the *senders*, anticipate these evaluations. The experimental strategy of this paper therefore combines two experiments that unfold in three stages. The *evaluation experiment* covers the first two stages. In the first stage I create profiles of Dar es Salaam residents, each photographed once in neutral and once in conspicuous attire. In the second stage, these profiles are shown to *evaluators* who rate them in four outcome domains: business owners and senior managers evaluate professional prestige, loan officers evaluate credit approval, and two community samples evaluate willingness to form a friendship and willingness to offer help, respectively. The *beliefs experiment* constitutes the third stage, in which I return to the profiled individuals from the first stage. For each domain, I elicit their beliefs about how each of their two profiles was rated in evaluations, which profile they prefer to submit to a future paid evaluation, and their willingness to accept (WTA) compensation for submitting their less preferred profile to the same evaluation instead.[2] Subsequently, a structural decomposition estimates which share of this WTA is explained by instrumental motives reflecting sender's beliefs that one profile will be rated more favorably than the other, or by intrinsic motives such as taste or social image concerns. The complete experimental setup is summarized in Figure 1.

Figure 1: Experimental setup



[2]All parts of this study received IRB approval from the University of Zurich and were covered through a research permit from COSTECH (Commission for Science and Technology, Tanzania). Both experiments are preregistered; pre-analysis plans are available upon request.

3 Evaluations Experiment

The main goal of the evaluations experiment is to measure the causal effect of conspicuous relative to neutral attire on evaluations in four domains: professional prestige, credit approval, friendship formation and informal help, and to test whether this effect differs across domains. In addition, I use supplementary evaluations to test what attire conveys about the personality traits of the person wearing it, and whether differential evaluation effects across domains are driven by the different themes of the domains or by the different evaluator groups of the domains.

3.1 Profile Creation

3.1.1 Sampling

I recruit 655 adults aged 18 and above in 13 wards spanning all five districts of Dar es Salaam. In January 2026, enumerators approached participants in the selected wards while filling quotas on gender, district, age, household income, and occupation sector, with targets chosen to match the adult population of Dar es Salaam. The resulting sample is broadly representative along these dimensions: as shown in Table A1, it closely matches population benchmarks for gender, district and age. Senders feature a 7.8 pp larger share of individuals in the lowest income bracket and a 7.4 pp larger share of individuals working in the service industry. Overall, the sample spans the socio-economic range of the city's adult population, including its large lower-income majority.

3.1.2 Profile Content

A profile consists of a photograph of the sender together with four pieces of information about the sender, namely nationality, place of residence, age, and occupation sector. Each sender is photographed twice, once in neutral attire, meaning plain, unbranded clothing together with one plain wristband, and once in conspicuous attire, a bundle of one branded garment and three accessories common in Dar es Salaam. Both attire bundles were provided by the enumerators for purposes of standardization. the average market value is roughly TZS 15,000 ($23 PPP) for a neutral attire bundle and TZS 60,000 ($91 PPP) for a conspicuous

attire bundle.[3] A total of 48 different attire bundles were provided for both neutral and conspicuous attire. 94.8% of senders report prior exposure to similar outfits in Dar es Salaam. Both photographs were taken in the same session with the same background and pose. The accompanying information is held constant across the two profiles such that the two profiles of the same sender consequently differ only in attire. This yields one profile pair per sender and 1,310 profiles in total. Figure A2 shows an example of two profile pairs, one of a male participant and one of a female participants. All profiles were created before the start of the evaluations experiment.

3.2 Design

3.2.1 Evaluators and domains

Evaluations were conducted between February 2026 and March 2026. They were conducted separately for each of the four outcome domains: professional prestige, credit approval, friendship formation and offering of informal help. Domains are chosen to span a range of locally relevant social and economic contexts, based on qualitative evidence from pilots. Senders themselves confirm their relevance: asked how important each domain is to them, large majorities rate each of the four domains as somewhat or very important (Appendix, Figure A1).

Within each domain, profiles are assessed by domain-specific evaluators. The reasoning to recruit domain-specific evaluators is that the people who actually make the relevant decision in real-life differ across domains: a loan is granted by a loan officer, an invitation into professional circles is extended by those already in them, and friendship or spontaneous help can come from anyone living in local communities. Matching evaluators to their domain therefore brings each evaluation as close as possible to the real-life decision it stands for. For professional prestige, evaluators are business owners, senior managers, and other high-income professionals. Starting from an initial set of direct contacts of high-income business owners spread across Dar es Salaam, participants were asked to recommend and connect the study team to others in their professional networks, and the sample was built through these referrals. For credit, evaluators are loan officers and credit analysts. Using a public registry of microfinance and loan companies in Dar es Salaam, the study team connected with close to 50

[3]Conspicuous attire bundles included both real and counterfeit brand garments. The distinction between original and counterfeit clothing does not appear to be a relevant margin for local evaluations: fewer than 0.2% of the open-ended justifications evaluators provide for each evaluation (Section 3.2.2) mention whether clothing is genuine or counterfeit, and these mentions are not concentrated among conspicuous-attire profiles. Moreover, both types of garments are sold ata premium on local markets.

distinct credit institutions and surveyed their loan officers and credit analysts. For friendship and help, where the relevant decision maker is an ordinary community member, evaluators were drawn from representative samples of Dar es Salaam residents, recruited under the same protocol as the senders. Evaluator samples in both domains match population benchmarks for gender, district, age and income bracket. Both feature a lower share of individuals working in the industry sector and a larger share working in transport. Evaluator characteristics are shown in Table A1.[4] Table 1 summarizes questions and evaluator samples for each domain. Verbatim question wordings for each domain are provided in Appendix C.1.

The question scenarios in all four domains are framed as first encounters with a stranger. This is the context in which visible signals carry the most weight and where benchmarking their returns is therefore most informative. First impressions are moreover highly decisive for socio-economic opportunities in this setting: in a city of Dar es Salaam's size ($\approx$ 7 million), loan applications, business gatherings, and everyday encounters routinely take place between people meeting for the first time. Throughout all evaluation surveys I ensure that evaluators did not personally know any of the displayed individuals on the profiles.[5]

Table 1: Outcome domains and evaluator samples

Domain	**Evaluation question**	**Evaluators**	**N**
Professional prestige	"Would you invite this person to an exclusive dinner for influential business people?"	Business owners, senior managers, members of business associations, high-income professionals	225
Credit	"Would you approve this person for a TZS 1 million business loan?"	Loan officers, credit analysts	228
Friendship	"After meeting this person for the first time at a social gathering, would you begin a friendship with them?"	Representative community sample	228
Help	"You are running late and see this person struggling with heavy bags. Would you stop to help them?"	Representative community sample	227

[4] 88.9% of all evaluators report familiarity with either all or most of the outfits shown in profiles.

[5] One loan officer was excluded from the credit evaluation sample after reporting prior acquaintance with a profiled individual. Results are unchanged by this exclusion.

3.2.2 Evaluations Procedure

Each evaluator assesses a series of 16 randomly assigned profiles with respect to the question of their domain, on a seven-point Likert scale.[6]. The unit of randomization is the profile evaluation within evaluator. Profile assignment is stratified by attire and gender; no evaluator sees more than one profile of the same individual, and no evaluator is shown more than one profile featuring the same outfit. Within this assigned set, the order of profiles is randomized. The same pool of profiles is used across all four domains. Since assignment of profiles to evaluators is randomized, exogenous variation in attire within individual allows me to causally attribute any difference in evaluations to attire alone. The randomization algorithm was preregistered and implemented prior to data collection, subject to pre-specified balance and coverage constraints, and does not condition on evaluator characteristics or outcomes. Across evaluators, each profile was assessed between 0 and 4 times per domain, with a median of 3 evaluations per profile.[7]

3.2.3 Outcomes

For every profile, I elicit two outcomes per evaluator. The first one is the evaluator's own rating of the profile. I refer to this as the *first-order evaluation*; it is the preregistered primary outcome and yields a total of 14,528 evaluations. The second outcome is the evaluator's guess of the modal first-order evaluation that the other evaluators give to the same profile in the same domain. I refer to this as the *second-order belief* (SOB) of the evaluator. The distinction matters for incentives. A first-order evaluation is a subjective judgment with no correct answer against which it could be scored, so it cannot be incentivized for accuracy. A second-order belief can be incentivized: evaluators are paid for correct guesses.[8] Comparing the two outcomes therefore allows me to assess whether the presence of incentives changes evaluation behavior. Evaluators complete all first-order evaluations before proceeding to the second-order elicitation.

At the end of the survey, each evaluator also cross-evaluates two randomly selected profiles for any of the other three outcome domains. This serves the purpose of shedding light on whether it is the evaluator sample or the outcome domain itself driving possible differential effects in evaluations. Moreover, to partially illuminate the channels underlying evaluations, two other randomly selected profiles are additionally assessed on four perceived personality

[6]The Likert scale spans seven points: "Extremely unlikely", "Unlikely", "Somewhat unlikely", "Neither likely nor unlikley", "Somewhat likely", "Likely", "Extremely likely"

[7]For each outcome, more than 99.4% of all profiles were evaluated at least once.

[8]Evaluators receive TZS 250 ($0.38 PPP) per correctly guessed modal response.

traits: competence, assertiveness, trustworthiness, and kindness (Fiske, 2018), as well as perceived wealth.[9]

3.3 Estimation

The preferred specification is a linear probability model (LPM) which takes as an outcome a binary indicator equal to one if the evaluation falls in the positive range of the Likert scale.[10] Another specification captures the entire seven point Likert-scale. However, the binary variable is used in the preferred specification for three reasons: 1) the LPM provides for an easier interpretation of the regression coefficients, 2) it is comparable to the beliefs elicited in the second experiment, discussed in Section 4, 3) I observe that the majority of evaluators sticks to the extremes of the Likert scale such that the specification captures most of the variation from the seven-point scale. The main regression equation reads:

$$Y_{ij}^{d} \;=\; \beta^{d}\, CC_{ij} \;+\; \mu_{i}^{d} \;+\; \lambda_{j}^{d} \;+\; \varepsilon_{ij}^{d}, \tag{1}$$

where Y_{ij}^{d} is an indicator equal to one if evaluator j gives a positive evaluation to profiled individual i in domain $d \in \{\text{prestige}, \text{credit}, \text{friendship}, \text{help}\}$; CC_{ij} is an indicator equal to one if evaluator j sees profiled individual i in conspicuous rather than neutral attire; and μ_{i}^{d} and λ_{j}^{d} denote profiled-individual and evaluator fixed effects, respectively. Standard errors are clustered at the profiled individual $\times$ attire level. Equation (1) is estimated separately for each domain d. The main coefficient of interest is β^{d}, which captures the evaluation return to conspicuous relative to neutral attire in domain d. To assess whether the return to conspicuous attire differs across domains, I additionally test the joint null hypothesis $H_0 : \beta^{\text{prestige}} = \beta^{\text{credit}} = \beta^{\text{friendship}} = \beta^{\text{help}}$. I implement this as a joint Wald test of equality of the domain-specific treatment effects.

3.4 Results

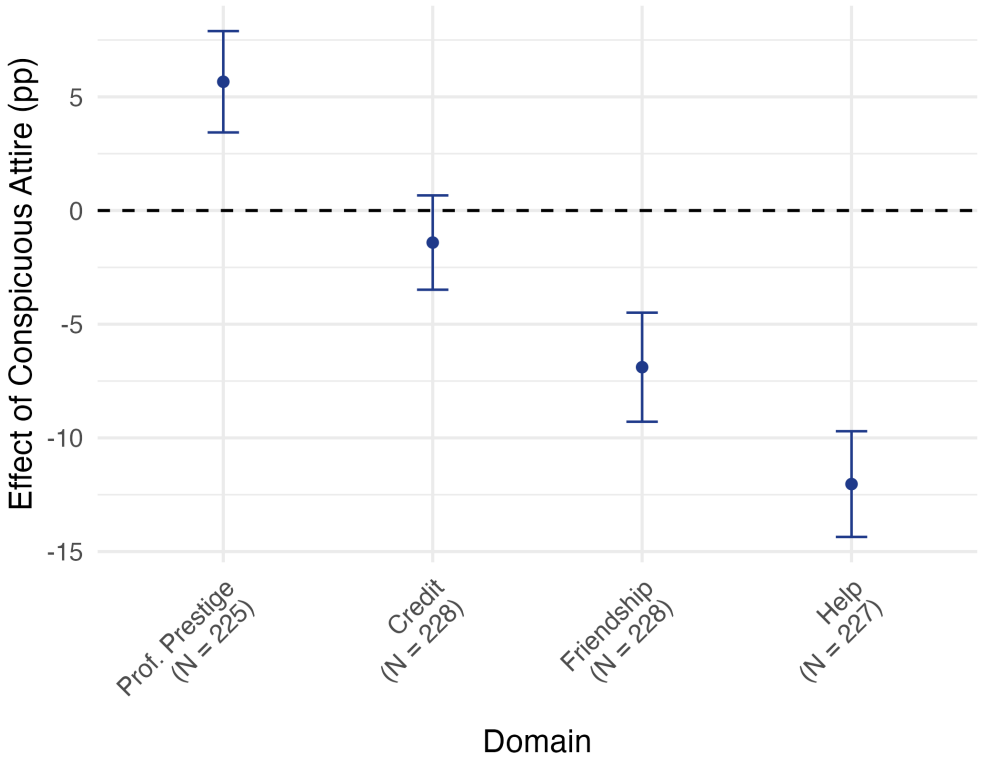
Relative to neutral attire, conspicuous attire increases the likelihood of a positive professional-prestige evaluation by 5.7 percentage points, leaves credit approval unchanged, and reduces the likelihood of a positive evaluation by 6.9 percentage points for friendship formation and by 12.2 percentage points for receiving informal help. Figure 2 reports the estimated

[9] Perceived wealth was not elicited in the main Evaluations experiment but in several preceding pilot studies. I integrate these wealth evaluations into the analysis below.

[10] The positive range includes answers expressing a positive likelihood including "Somewhat likely", "Likely", "Extremely likely".

treatment effects across all four outcome domains. In addition, Table A2 shows that a Wald test strongly rejects equality of effects across domains ($p < 0.001$). The same table shows that the pattern of effects is unchanged when the regression outcome is the full seven-point Likert score rather than the binary positive-evaluation indicator in the LPM.[11]

Figure 2: Effect of conspicuous attire on evaluations, by domain



Notes: Effect of conspicuous vs. neutral attire on probability of giving a positive first-order evaluations across outcome domains. Coefficients are from domain-specific linear probability models (Equation 1), estimated separately for each domain. The dependent variable is a binary indicator equal to one if the evaluation falls in the positive range of the seven-point Likert scale. The treatment indicator equals one for profiles featuring conspicuous attire. All specifications include profiled individual fixed effects and evaluator fixed effects. Standard errors are clustered at the profiled individual × attire level. Horizontal bars indicate 95% confidence intervals.

The central result of this experiment is that the same signal generates both positive and negative payoffs in different domains: conspicuous attire is beneficial for professional prestige, irrelevant for credit approval, and disadvantageous for friendship formation and receiving informal help. This matters for how signaling payoffs are understood. Existing evidence documents positive returns to conspicuous consumption in status, credit, and labor-market settings (Nelissen and Meijers, 2011; Macchi, 2023; Dupas et al., 2024); the pattern here shows

[11]The coefficients for credit approval are marginally significant in this specification.

that these gains are not a property of the signal but of the domain. To my knowledge, this is the first evidence in the economics literature that a single visible signal carries simultaneous positive and negative payoffs across different domains.

Second-order beliefs display an equivalent pattern to first order evaluations (Figure A3 in the Appendix): evaluators' guesses of the modal evaluation track the first-order results domain by domain, including the sign reversals. This has two implications. First, it indicates that the payoff structure is not an artifact of unincentivized responses. Second, it shows the payoff structure is common knowledge within the evaluator population: receivers know how other receivers respond to the same signal in the same domain.

3.4.1 Cross-domain evaluations

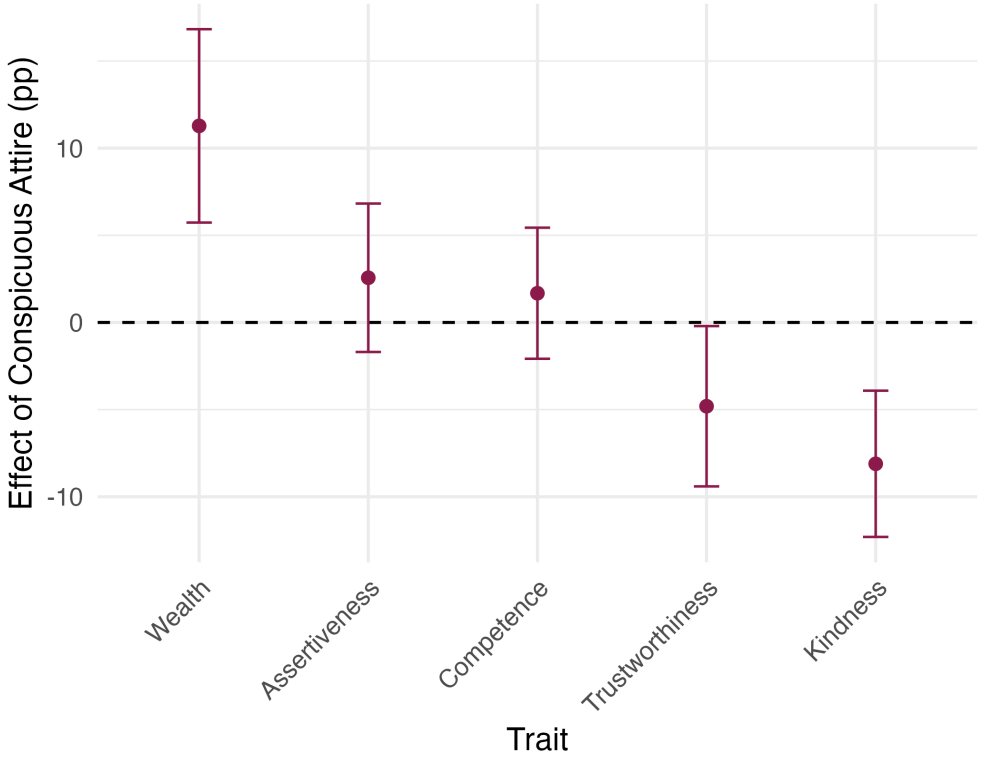
Differential siganling payoffs across domains could in principle be driven by differences in both the theme of the outcome domain (e.g. credit approval vs. friendship formation) and the evaluator composition in each domain (e.g. loan officers vs. community sample). To assess whether the latter could be driving the differential payoffs, I exploit the cross-evaluation design at the end of the survey, where each evaluator assesses two profiles with respect to the three outcome domains outside their own assigned domain. I pool all evaluations of both community samples and investigate them by domain. Figure A4 in the Appendix reports the results. Holding the evaluator pool fixed and varying only the domain reproduces the differential returns to conspicuous attire: conspicuous attire reduces positive evaluations by 11.1 pp for help and 6.2 pp for friendship, has no effect on credit (+0.5 pp), and carries a positive but imprecise effect on professional prestige (+3.3 pp). For prestige and credit, the pooled community estimates rest only on the two cross-evaluations each evaluator completes at the end of the survey, whereas the friendship and help estimates also draw on the community samples' main evaluation task; the former are accordingly less precise, though both main estimates lie well within their confidence intervals. The effects remain statistically distinct across domains: the prestige and credit effects each differ significantly from the help effect. The countervailing returns to conspicuous attire therefore appear to be a property of the domain theme rather than of the evaluator composition.

3.4.2 Traits evaluations

In terms of traits, conspicuous attire enhances the likelihood of being evaluated as wealthy by 11.3 percentage points, diminishes perceptions of trustworthiness by 4.8 pp and kindness by

8.1 pp and has no detectable effect on assertiveness and competence.[12] Figure 3 reports the estimated treatment effects across all traits, pooled across domains. Figure A15 additionally shows that inferences of the four personality traits do not statistically differ across evaluators in different domains. I return to these results in Section 6.

Figure 3: Effect of conspicuous attire on evaluations, by trait



Notes: Effect of conspicuous attire on perceived traits. Coefficients are from trait-specific linear probability models, estimated analogously to Equation 1. The dependent variable is a binary indicator equal to one if the evaluation falls in the positive range of the seven-point Likert scale. The treatment indicator equals one for profiles featuring conspicuous attire. All specifications include profiled-individual fixed effects; evaluator fixed effects are not included because each evaluator assesses only two profiles per attribute. Standard errors are clustered at the evaluator level. The wealth estimate is not from the main Evaluations experiment, in which perceived wealth was not elicited, but from preceding pilot studies in which participants evaluated the same profiles on perceived wealth ($N = 1,135$ evaluations); its specification likewise includes profiled-individual fixed effects, with standard errors clustered at the participant × profiled-individual level. Horizontal bars indicate 95% confidence intervals.

[12]This shows that conspicuous attire actually fits the definition of conspicuousness, namely wealth signaling.

3.4.3 Heterogeneities

To test whether the attire effects differ across evaluators, senders, or evaluator–sender pairs, I interact the attire treatment in Equation 1 with eight evaluator characteristics, three sender characteristics, and a measure of evaluator–sender homophily.[13] Tables A3 and A4 in the Appendix report the results. Each cell reports, for one heterogeneity dimension and one domain, the p-value of a Wald test that the attire effect is the same across all categories of the heterogeneity dimension (i.e., that all interaction coefficients are jointly zero). The p-value in brackets adjusts for multiple hypothesis testing across the entire table.

On the evaluator side, the picture that emerges is one of homogeneity. A handful of unadjusted tests are significant, but they follow no consistent pattern across domains and none survives the multiple-testing correction. No dimension systematically moderates the returns in any domain. The differential returns documented above therefore appear to be a broad feature of how evaluators read attire signals, rather than an artifact of particular evaluator or sender compositions. This is consistent with the cross-domain evaluation results discussed in Section 3.2.1.

On the sender side, a similar picture arises, though female profiles appear to receive somewhat more favorable evaluations of conspicuous attire in prestige, credit, and friendship, and slightly less favorable ones in help, this does not change the substantive patterns in differential returns to conspicuous attire and does not survive the multiple-testing adjustment.

4 Beliefs Experiment

The beliefs experiment investigates whether the individuals depicted on the evaluated profiles (senders), anticipate the countervailing returns to conspicuous attire documented in Section 3. First, I elicit sender's incentivized beliefs about how their own two profiles are evaluated in each of the four domains. Second, I elicit senders' incentivized preferences over which of their two profiles they would prefer to submit to a future round of evaluations. Third, I elicit their willingness to accept for forgoing their preferred profile choice, using a multiple-price-list design. I benchmark all findings against the evaluations results from Section 3.

[13] For homophily, I construct a level indicator for homophily (counts the number of observable characteristics that equal the same characteristic of the evaluator).

4.1 Design

4.1.1 Sampling

Between May and June 2026, four months after profile creation, enumerators attempted to re-contact all 655 initially recruited senders (Section 3.1.1).[14] Of these, 502 (76.6%) were successfully reached and completed the beliefs experiment.[15] Table A1, column three summarizes the characteristics of the senders who completed the beliefs survey. This subsample approximates the original pool of senders closely, with the exception of a lower mass of senders now in the lowest income bracket and a larger mass in the lower middle income segments. To verify that attrition in the sender sample is unrelated to the evaluation returns, I re-estimate Equation 1 using only evaluations of profiles belonging to senders who completed the beliefs survey. This regression reproduces the full-sample effects, as shown in Figure A5.[16]

4.1.2 Elicitation Procedure

At the start of the survey, senders are informed that both of their profiles, the conspicuous and the neutral attire profile, may be evaluated by other people in Dar es Salaam in the future. I explain that an evaluation involves someone (i) seeing the sender's profile, (ii) answering a question about it, and (iii) seeing, in the same survey, the profiles of other participants. Senders then go through the four domains one at a time, in random order. Within each domain, the sequence is always the same. First, senders report their beliefs about how evaluators would assess each of their two profiles, one after the other, in random order. Second, they state which of the two profiles they prefer to have evaluated in that domain. Third, they report their willingness to accept (WTA) for having the other profile evaluated instead. Only then do they move on to the next domain.

4.1.3 Beliefs

For each profile and domain, I elicit senders' beliefs about how their profile will be evaluated. Beliefs are measured as the expected number of evaluators, out of 100, to evaluate the profile

[14] Participants verified their identity by confirming name and year of birth before being shown their profiles. They then additionally certified that these profiles indeed contained their two profile images and their correct demographic information.

[15] The remaining 153 senders (i) could not be reached under any of the contact details provided, (ii) had moved outside the study area, or (iii) were away from Dar es Salaam for extended work stays during the survey period.

[16] Of these, 487 had both profiles evaluated at least once in all four outcome domains; the remaining 15 had one outcome domain missing for one profile in the evaluation study. Since profiles were randomly paired with evaluators in the first experiment, this subsample was constructed randomly and is included in all analyses.

positively. To rule out that the evaluations procedure itself induces misperceptions in senders, I reproduce the exact screen from the evaluations experiment during the beliefs elicitation. Figure 4 shows an example of what senders see during the beliefs elicitation, in the credit domain. The screen consists of three components. At the top of the screen, senders are asked to assume that the research team will show their profile to 100 evaluators from the group relevant for the domain in question: influential business professionals for professional prestige, loan officers and credit professionals for credit approval (as featured in the example), and community members in Dar es Salaam for friendship formation and informal help. In the middle of the screen, senders see exactly the same screen the evaluators would see in the evaluations experiment. It consists of the sender's own profile, the evaluation question, and the seven-point Likert answer scale. At the bottom of the screen, senders are asked to predict how many of the 100 evaluators would give a positive evaluation, on the green marked part of the answering scale.[17] The beliefs elicited here are therefore second order in the sense that they are beliefs over evaluations of others.

Beliefs are incentivized. Before submitting their predictions, senders are told that there is a real chance their profiles will be shown to 100 evaluators, and that they will receive a bonus of TZS 100,000 ($151 PPP) if their prediction is within five evaluators of the realized number of positive evaluations.[18]

4.1.4 Preferences and Willingness to Accept

After senders have reported their beliefs for both profiles in a given domain, I inform them that in case neither of their profiles is selected for evaluation by 100 evaluators, one of their profiles may still be evaluated by a single evaluator from the same group. If this evaluation is carried out, the sender earns TZS 10,000 ($15 PPP) if the evaluation is on the positive side of the Likert scale, and nothing if it not. Senders then go through two choice stages.

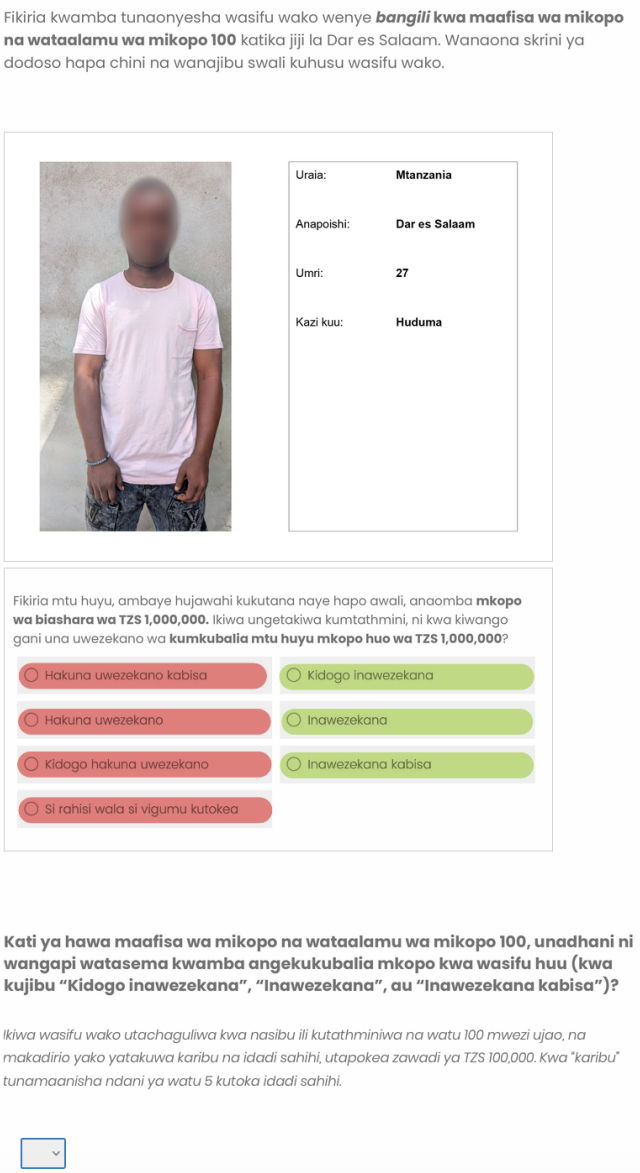
In the preference stage, the goal is to simply elicit which profile senders prefer to submit to the single evaluation. To this end, senders make a binary choice of which of their two profiles they would prefer to submit to the evaluations. The two profiles are listed in random order.

In the willingness to accept (WTA) stage, the goal is to elicit the amount of money senders

[17] In particular, senders are told that this part of the scale includes the following answers: "Somewhat likely", "Likely", "Extremely likely".

[18] To implement this, I randomly draw one domain and randomly select a subset of profiles. These profiles are shown to 100 new evaluators, who are recruited in the same way as the evaluators for the same domain in the evaluations experiment. Senders with a profile in the selected subset whose prediction falls within the correct range of positive evaluations received the bonus.

Figure 4: Example Survey Screen shown to senders during beliefs elicitation



Notes: Belief elicitation screen (Swahili original). The introductory text above the profile reads: "Imagine that we are showing your profile with neutral attire to 100 loan officers and in the city of Dar es Salaam. They see the survey screen below and answer a question about your profile." The screen then shows one of the participant's profiles. Below the profile, participants see the same seven-point Likert scale shown to evaluators in the evaluations experiment, here for the credit domain: "Imagine this person, whom you have never met before, is applying for a business loan of 1 Million Shilling. If you are asked to assess them, how likely are you to approve this person for the loan?" Participants are then asked: "Out of these 100 loan officers and credit professionals, how many do you think will approve you for the loan with this profile (answer: "Somewhat likely", "Likely", or "Extremely likely")?" Participants enter their numerical belief on a scale from 0 to 100 using a drop-down menu. The incentive scheme is stated in italics at the bottom: "If your profile is randomly selected for evaluation by 100 people next month and your guess is close to the correct number, you will receive a prize of TZS 100,000. By close we mean within 5 people of the correct number."

are willing to invest into the conspicuous attire profile over the neutral attire profile. WTA is elicited using an incentive-compatible multiple price list (MPL): in each row, senders choose between submitting the conspicuous profile or the neutral attire profile for evaluation. In particular, they choose between submitting their previously stated more preferred profile and they less preferred profile to which a cash transfer is attached which varies across rows between TZS 3,000 ($5 PPP) and TZS 9,000 ($14 PPP). The midpoint of the interval at which a sender switches, identifies the transfer at which they are indifferent between the two profiles.[19] I code WTA as positive when senders prefer the conspicuous profile and negative when they prefer the neutral profile. Following Jack et al. (2022), the format of the MPL is randomized at the domain level within each sender: both the ordering of cash transfers (ascending or descending) and the placement of the two profiles (left or right column).

Both the preference and the WTA choice are consequential. Senders are informed that, for most senders, the binary preference task determines which profile is submitted to the evaluation. For a randomly selected minority of senders, one randomly drawn row of the MPL is implemented instead (Becker et al., 1964); the profile chosen in this row is then submitted to the evaluation and the sender receives the potentially attached cash transfer. This payment rule addresses a potential strategic concern. In the MPL, the cash transfer is always attached to the profile the sender did *not* choose in the binary preference task. In principle, a sender could therefore misreport their preference in the binary choice, so that the transfer in the MPL is attached to the profile they actually prefer. Because the binary choice is implemented for the majority of senders, such misreporting is costly.

4.2 Estimation

4.2.1 Beliefs

The outcome in the beliefs elicitation measures the sender's subjective probability that an evaluator rates their profile positively. I estimate the following linear probability model:

$$\hat{Y}_{ip}^{d} \;=\; \beta_{B}^{d}\, CC_{p} \;+\; \mu_{i}^{d} \;+\; \varepsilon_{ip}^{d}, \tag{2}$$

where $\hat{Y}_{ip}^{d} \in [0,1]$ is the share of evaluators that sender i expects to evaluate profile p positively in domain $d \in \{\text{prestige}, \text{credit}, \text{friendship}, \text{help}\}$; CC_p is an indicator equal to one if profile p shows the sender in conspicuous rather than neutral attire; and μ_i^d denotes sender fixed effects.

[19] The elicitation placed no restrictions on choice patterns; nevertheless, no sender switches more than once in any price list, consistent with high task comprehension.

Unlike in the evaluations experiment, there are no separate evaluator fixed effects: each sender forms beliefs about their own two profiles, so the profiled individual and the person reporting the belief coincide, and both are absorbed by μ_i^d. Standard errors are clustered at the sender level. Equation (2) is estimated separately for each domain d. The main coefficient of interest is β_B^d, which captures the perceived return to conspicuous relative to neutral attire in domain d.

The coefficients β_B^d and β^d from Equation 1 measure the same quantity: how much more likely a profile is to be rated positively if it features conspicuous rather than neutral attire (in percentage points). The only difference is who averages over evaluators. In the evaluations experiment, the regression averages over the ratings of actual evaluators. In the beliefs experiment, the sender does this averaging themselves, when predicting how many out of 100 evaluators would rate their profile positively. The two coefficients are therefore directly comparable. To measure misperceptions, I compare the two coefficients. For each domain d, I define the misperception as $\beta_B^d - \beta^d$, the difference between the perceived returns and realized returns. I test the null hypothesis of no misperception, $H_0 : \beta_B^d = \beta^d$, separately for each domain and jointly across all four domains.

4.2.2 Benchmarking Willingness to Accept

The WTA data allow me to address the central research question of this paper: do senders pay for conspicuous attire even when it does not pay off? I approach this question in two steps. First, I ask whether senders give up more to submit their conspicuous attire profile than it can be expected to earn them in monetary terms. Second, I ask whether senders also overspend in utility terms, that is, whether overspending persists once their preferences and risk attitudes are accounted for.

To answer the first part of this question, I compare senders' WTA to a simple benchmark: the WTA of a sender who is perfectly informed about how evaluators assess their profiles, who only cares about the evaluation bonus, and who is risk-neutral. For such a sender, the value of submitting a profile equals its expected bonus, that is, TZS 10,000 times the probability that the profile receives a positive evaluation. The benchmark WTA is then the difference in expected bonus between the conspicuous and the neutral profile. I impute the probability of a positive evaluation for each profile from its share of positive evaluations in the evaluations experiment (Section 3.4). If senders' WTA exceeds this benchmark, they pay more for the conspicuous profile than it can be expected to return in monetary terms. I will in the following refer to this benchmark WTA as *potential earnings*.

To answer the second part of the question, I distinguish three factors that can drive such

overspending: (i) misperceived returns: senders overestimate how positively their conspicuous attire profile is evaluated or underestimate how positively their neutral attire profile is evaluated; (ii) non-instrumental preferences: senders value being evaluated in conspicuous attire beyond its monetary returns, for instance due to personal taste or social image concerns; (iii) risk aversion: the cash transfer attached to the less preferred profile is certain, whereas the evaluation bonus is uncertain, so a risk-averse sender values the two options differently even if their expected values are the same. Only (i) implies that senders invest in conspicuous attire even when it does not pay off in utility terms. In Section 5, I struturally decompose WTA accordingly.

4.3 Results

4.3.1 Beliefs

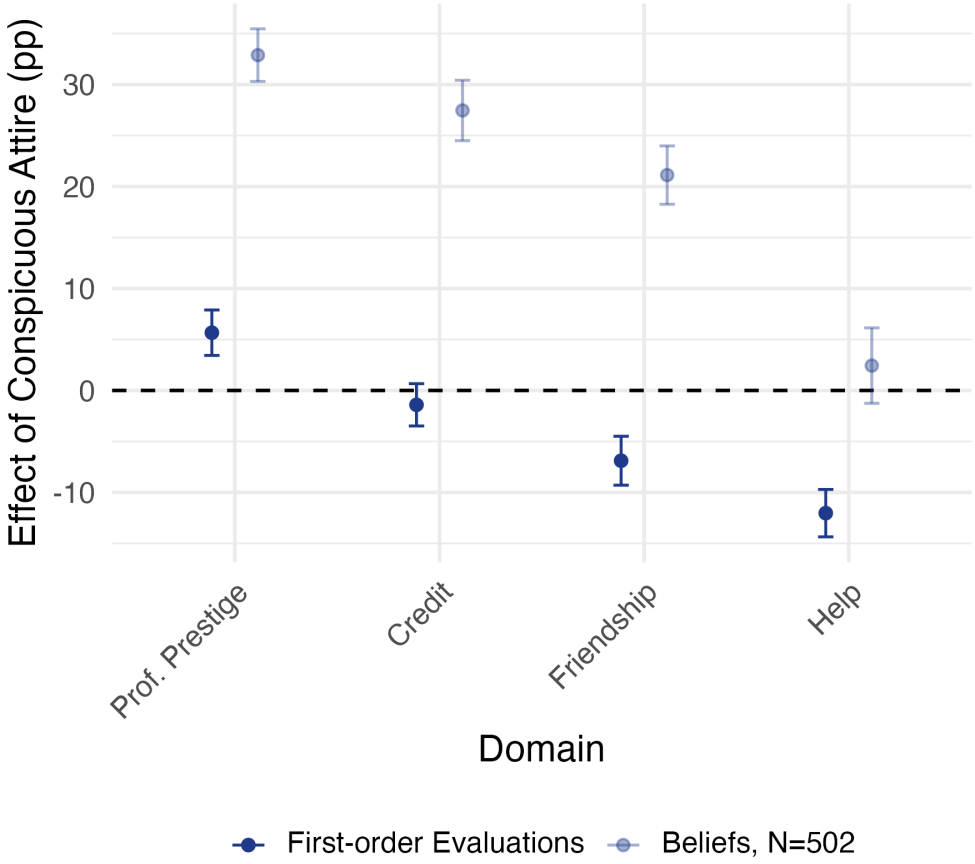
Senders systematically overestimate the returns to conspicuous attire across all four domains. Figure 5 compares the realized and the perceived effect of conspicuous relative to neutral attire on the probability of receiving a positive evaluation, by domain. Realized effects are estimated from the evaluations experiment (Equation (1)) and reproduced from Figure 2; perceived effects are estimated from senders' elicited beliefs (Equation (2)). Specifically, senders overestimate the return to conspicuous attire by 27.2 pp for professional prestige, by 28.8 pp for credit approval, by 28.0 pp for friendship, and by 14.6 pp for help. These misperceptions are statistically significant in all domains (Table **??**). At the same time, senders recover the correct ranking of returns across domains: they correctly perceive that conspicuous attire pays off most in professional prestige and least in help, but are wrong about the level of returns in an almost uniform way. For friendship, this shift reverses sign: perceived returns to conspicuous attire are positive, while realized returns are negative.

4.3.2 Preferences and Willingness to Accept

Consistent with beliefs, a large majority of senders prefer to submit the conspicuous attire profile for an upcoming evaluation in professional prestige (90.0%), credit approval (83.1%), and friendship (78.5%), while preferences are more evenly split for help (48.4%). These results are summarized in Figure 6. Preferences are correlated with beliefs: a participant who believes that their conspicuous-attire profile is 1 pp more likely to receive a positive evaluation than their neutral-attire profile is 0.9 pp more likely to submit it to the future evaluation.

WTA significantly exceeds potential earnings in all domains. Figure 7 compares mean

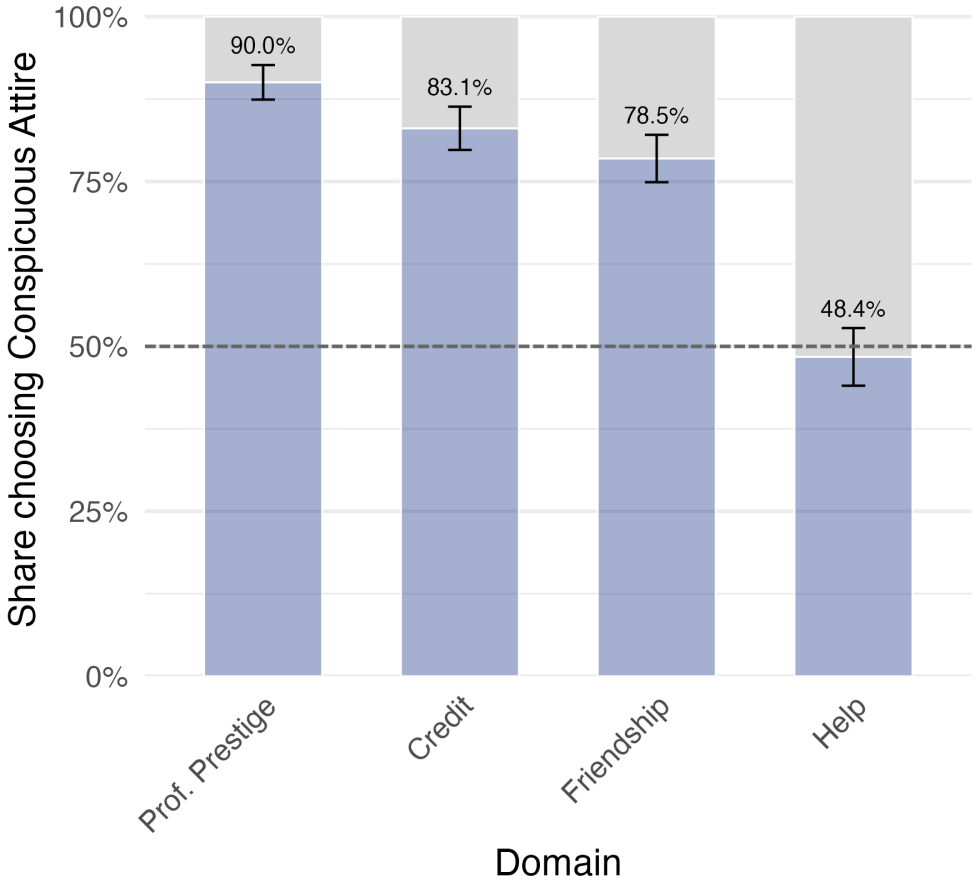
Figure 5: Perceived vs. realized effect of conspicuous attire on evaluations, by domain



Notes: Perceived versus realized evaluation returns. Coefficient estimates show the effect of conspicuous relative to neutral attire on the probability of receiving a positive evaluation, separately for each outcome domain. Dark blue circles show first-order evaluation effects from the evaluations experiment, estimated as the average difference in positive evaluation rates between conspicuous and neutral attire profiles across evaluators. Light blue circles show mean belief returns elicited in the beliefs and WTA experiment ($N = 502$), defined as the difference in elicited probabilities of a positive evaluation for the conspicuous versus neutral profile, averaged across participants. Both measures are expressed as percentage point differences (CC minus Neutral). Error bars show 95% confidence intervals.

WTA for conspicuous over neutral attire to imputed potential earnings as described in Section 4.2.2. For professional prestige, participants are willing to forgo an average TZS 4,028 ($6 PPP) to have their conspicuous profile evaluated, against imputed potential earnings of TZS 537 ($0.81 PPP). For credit, participants are willing to forgo an average TZS 3,245 ($5 PPP) compared to near-zero potential earnings of − TZS 47 (-$0.07 PPP) from conspicuous attire. The starkest case is friendship. Here, senders exhibit a mean WTA of TZS 2,857 ($ 4 PPP) to have their conspicuous rather than neutral attire profile evaluated, against an average potential earnings of −TZS 619 (−$1 PPP). Senders are thus paying for conspicuous

Figure 6: Preference for conspicuous attire, by domain



Notes: Share of senders choosing to submit their conspicuous rather than their neutral profile to a future round of evaluations, by outcome domain ($N = 502$ senders). The darker segment of each bar marks the share choosing the conspicuous profile; vertical bars indicate 95% confidence intervals. The dashed line marks indifference at 50%.

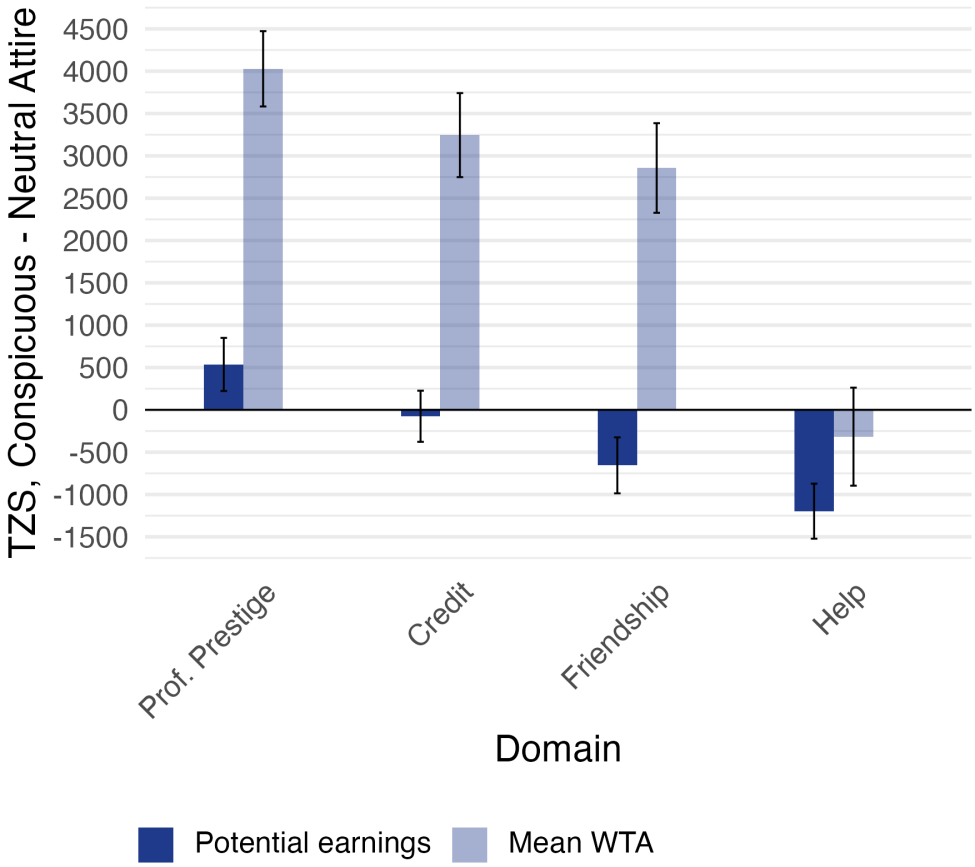
attire even though it actively reduces the probability of friendship formation. To put these magnitudes in perspective, senders are willing to give up about 60% of one day's average income per household member for the conspicuous attire profile, even though submitting it lowers expected earnings by a further 13% of a day's income.[20] For help, mean WTA of −TZS 317 (−$0.48 PPP) comes closest to potential earnings of −TZS 1,107 (−$2 PPP).

4.3.3 Robustness Checks

Misperceived Evaluators. One may worry that senders misperceive not the returns to attire, but the composition of evaluators. In the beliefs elicitation, senders are told which group

[20] Average monthly income per household member in the sender sample is TZS 142,548 (TZS 4,750 per day), computed from average reported household income and reported household size.

Figure 7: Willingness to accept for conspicuous attire versus neutral attire profile, by domain.



Notes: Dark blue bars show imputed potential earnings — the expected monetary gain from submitting the conspicuous rather than the neutral profile to a real evaluation, computed as the difference in positive evaluation probabilities between the two profiles (as realized in the evaluations experiment) multiplied by the bonus payment of TZS 10,000 ($15 PPP). Light blue bars show mean willingness to accept (WTA) for the conspicuous over the neutral profile, elicited through an incentive-compatible multiple price list ($N = 502$). WTA is signed positive when participants prefer the conspicuous profile and negative when they prefer the neutral profile. All values are expressed as TZS. Error bars show 95% confidence intervals.

evaluates each domain but not the exact composition of these groups. If senders pictured a different set of evaluators than the ones actually recruited, their beliefs could be wrong for that reason alone. This would be a misperception created by the experimental design, with little external validity. Senders typically know their audience. Two pieces of evidence speak against this concern.

First, consider credit, the domain where senders know their evaluators most precisely. Loan officers and credit professionals are defined by a single, clearly communicated characteristic, their job, and senders have direct personal experience with this group: 74% of participants in pilot data have taken loans from microfinance institutions or banks before, and have thus been

assessed by precisely this group before. If misperceived evaluators were driving misperceived returns, the misperception should therefore be smallest for credit. Instead it is largest, at 28.8 pp.

Second, consider friendship and help, where the concern is most plausible. Senders are told these domains are evaluated by community members in Dar es Salaam. If each sender pictures people like those around them, the average sender should picture roughly the right composition of evaluators. This comes with one caveat: as Table A1 shows, the community evaluator samples differ somewhat from the sample of senders, most notably in income. Suppose, then, the worst case for this concern: every sender imagines evaluators who look exactly like themselves or the community around them. I construct evaluations of this counterfactual evaluator sample directly. I reweight the friendship and help evaluator samples to match the senders' distributions of gender, age, district, occupation sector, and income, and re-estimate the realized returns. Figure A6 shows the result: the reweighted returns differ from the original ones by less than 0.5 percentage points, so the misperceptions are essentially unchanged. Even if senders pictured evaluators just like themselves, they would still carry the same misperceptions about signaling returns.

Order Effects. Senders complete belief elicitation and choice tasks in the four domains sequentially, which raises the concern of cross-domain spillovers. Senders' beliefs and choices in later domains may be contaminated by what came before: having already reported beliefs, preferences, and WTA for one domain, senders might anchor on their earlier answers, update across domains, or simply grow confused. This concern warrants a clean test: the contamination does not apply to the first round of elicitations for the first domain of each sender. This takes place before any other domain has even been mentioned to senders. Since the domain order is randomized, restricting to first-domain responses yields a random quarter of the sample for each domain. I therefore re-estimate the misperceptions, preferences, and WTA results on first-domain responses only. Figures A7, A8, and A9 in the Appendix show the results: the patterns of overestimation and overspending are substantively unchanged, though estimates are necessarily less precise with a quarter of the sample per domain. Two differences are worth noting. First, belief returns are somewhat lower in first elicitations throughout — significantly so for professional prestige — consistent with mild anchoring in later rounds; as a consequence, prestige and credit beliefs are no longer statistically distinguishable, so the perceived ranking of domains is recovered less precisely. Second, the first-domain WTA for help moves onto its potential-earnings benchmark. Neither difference affects the central findings: overestimated returns and overspending on conspicuous attire are present already in each sender's very first elicitation.

5 Decomposing Willingness to Accept

This paper has so far shown that the same signal, conspicuous attire, can pay off in some domains and not in others. Senders overestimate how much it pays off in every domain and invest more in the signal than it earns them. But does this mean that signaling is wasteful? The size of the belief error makes it a natural suspect for this overspending, but WTA may also reflect intrinsic utility from a taste for conspicuous self-presentation, independent of how evaluators respond to it. Only if misperceptions drive the bulk of WTA is the investment wasteful: money spent on wrong beliefs buys nothing, whereas money spent on taste buys utility. In this section, I structurally decompose senders' utility, as measured by WTA, into a belief-driven and an intrinsic utility component. I then quantify the share of WTA driven by each component and estimate the welfare losses due to misperceptions.

5.1 Setup

To separate these two forces, I set up a sender utility function with an instrumental component, driven by beliefs about evaluation returns, and an intrinsic component driven by taste for the signal. I first show that beliefs and WTA co-move, which warrants a decomposition in which beliefs enter utility defining WTA. I then structurally decompose WTA to quantify the share of WTA implied by each of the two forces and quantify welfare losses due to misperceptions. Finally, I discuss two complementary experimental approaches to establishing the link between beliefs and WTA causally.

A sender can choose between sending a signal C or sending no signal N such that $s \in \{C, N\}$ equivalent to choosing conspicuous versus neutral attire in the preceding experiment. For notational simplicity, I do not indicate individual senders by i, instead I will refer to a representative sender. As for the experiment, the sender can choose to send a separate signal in each domain $d \in D$. I begin by setting up the utility of the sender. The key assumption is that the utility of choosing signal s for evaluation in domain d consists of two additively separable components.

The first component is the *instrumental* utility component, $B^d(s)$. It is the expected evaluation payoff from choosing profile s, computed under the sender's beliefs of successful evaluation of the signal in domain d, $\hat{b}^d(s)$. The instrumental component of the sender's utility is assumed to take a standard expected utility form:

$$\hat{B}^d(s) = \hat{b}^d(s) \cdot u(\bar{\omega} + \pi) + (1 - \hat{b}^d(s)) \cdot u(\bar{\omega}), \tag{3}$$

π is the bonus earned from a positive evaluation and $\bar{\omega}$ is the wealth endowment of the sender. $u(\cdot)$ is the Bernoulli utility function which I assume to be CRRA with risk aversion parameter ρ:

$$u(y,\rho) = \begin{cases} y & \rho = 0 \\ \dfrac{y^{1-\rho}}{1-\rho} & \rho \neq 0 \end{cases} \tag{4}$$

The belief $\hat{b}^d(s)$ summarizes everything the sender expects about how evaluators in domain d respond to her signal. Conceptually, it combines the two margins of signaling discussed in the introduction: (i) what traits the sender believes the signal conveys about them, and (ii) how they believe evaluators in domain d value these traits. Section **??** makes these two margins explicit; for now, it suffices to treat $\hat{b}^d(s)$ as a single term.

The second component is the *intrinsic* utility component, $\phi^d(s)$. It captures the utility the sender derives from the signal itself, such as a taste for conspicuous self-presentation, regardless of how evaluators respond. Because it does not depend on evaluation payoffs, it is invariant to beliefs. The complete utility of the sender with risk aversion ρ, for sending signal s for evaluation in domain d can therefore be written as:

$$U^d(s,\rho) = \hat{B}^d(s,\rho) + \phi^d(s) = \underbrace{\hat{b}^d(s) \cdot u(\bar{\omega}+\pi,\rho) + (1-\hat{b}^d(s)) \cdot u(\bar{\omega},\rho)}_{\text{instrumental}} + \underbrace{\phi^d(s)}_{\text{intrinsic}}. \tag{5}$$

The separability assumption is warranted here precisely as I seek to isolate anything that is belief-driven from anything that is not.

5.2 Correlations

Before proceeding to structurally decompose WTA into a beliefs driven and an intrinsic utility component, I first establish that beliefs belong in the utility defining WTA equation at all. The evidence comes from the correlation between elicited beliefs and WTA across senders and domains. If miscalibrated beliefs drive WTA, individuals who overestimate the return to conspicuous attire in a given domain should also exhibit higher WTA in that domain. Figure A10 plots individual-level WTA against belief returns for each of the four domains. The correlation is strong and positive across all domains. Domain-specific correlations range from 0.52 for prestige to 0.73 for help. Figure A11 shows that the relationship also holds *within* senders: demeaning both beliefs and WTA by sender, the domains in which a sender expects a larger evaluation return than her own average are the domains in which her WTA is higher than her own average ($r = 0.65$). This relationship addresses the concern that the raw correlation merely reflects a sender-level confound that raises both variables at once: a

sender who is generally optimistic, for instance, would report high belief returns *and* high WTA in every domain. Such uniform level differences are absorbed by the sender means in this latter regression.

These patterns are consistent with beliefs being an important driver of WTA in all domains and motivate the assumption that beliefs enter the utility function defining WTA, but it cannot quantify how much of the WTA is attributable to miscalibrated beliefs rather than intrinsic motives. I address this in the next section.

5.3 Structural decomposition

To decompose WTA into a belief- and an intrinsic preference implied component, I compute for each individual the WTA they would exhibit if they were a purely belief-driven agent with

$$\phi^d(C) = \phi^d(N) = 0.$$

For $\hat{b}^d(N) < \hat{b}^d(C)$ such a belief-driven agent is indifferent between the conspicuous and neutral profile at the transfer $W^d_{\hat{B}}$ which gives rise to the following indifference equation:

$$\hat{b}^d(C)\cdot u(\bar{\omega}+\pi,\rho)+(1-\hat{b}^d(C))\cdot u(\bar{\omega},\rho) = \hat{b}^d(N)\cdot u(\bar{\omega}+\pi+W^d_{\hat{B}},\rho)+(1-\hat{b}^d(N))\cdot u(\bar{\omega}+W^d_{\hat{B}},\rho). \quad (6)$$

Under risk neutrality ($\rho = 0$), Equation (6) simplifies to:

$$W^d_{\hat{B}} = (\hat{b}^d(C) - \hat{b}^d(N)) \cdot \pi, \quad (7)$$

which means that WTA equals the effective belief return of sending $s = C$ rather than $s = N$.[21] Let $\bar{W}^d$ denote the average WTA measured in this experiment, in domain d, and $\bar{W}^d_X$ the WTA implied by intrinsic preferences alone such that $\bar{W}^d = \bar{W}^d_{\hat{B}} + \bar{W}^d_X$. Then the belief-driven share of WTA in domain d is computed as:

$$s^d_B = \frac{\bar{W}^d_{\hat{B}}}{\bar{W}^d}, \quad (8)$$

[21] Equation (6) always has a unique solution on the compact interval $W^d_{\hat{B}} \in [0, \pi]$ as expected utility is monotonically increasing in $W^d_{\hat{B}}$ and at $W^d_{\hat{B}} = 0$ $LHS < RHS$, whereas at $W^d_{\hat{B}} = \pi$, the agent gets an unconditional transfer of π when submitting the neutral profile which makes this the strictly preferable option: $RHS < LHS$. The conclusion follows from the intermediate value theorem. For participants with $\hat{b}^d(N) > \hat{b}^d(C)$, the transfer is added to the conspicuous profile side and the indifference condition is symmetric to Equation (6).

I evaluate equation (8) under $\rho \in \{0, 0.5, 2.0, 5.0\}$, spanning risk neutrality to substantial risk aversion. Beliefs are taken from the beliefs elicitation directly and wealth endowment, $\bar{\omega}$ is based on the monthly income of a household member TZS 142,548 ($215 PPP), π =TZS 10,000 as specified through the experimental incentives. For each value of ρ I also adjust the potential earnings benchmark to account for the respective degree of risk aversion.

Figure 8 plots the decomposition under $\rho = 0.5$: the belief-implied WTA co-moves closely with observed WTA for prestige, credit, and friendship, accounting for 81.8%, 84.6% and 74.2% respectively. For help the two diverge in sign. Beliefs for help are pulling toward conspicuous attire and observed WTA towards neutral attire. Non-instrumental preferences therefore appear to partially offset the miscalibratied beliefs in this domain, through not statistically significantly so. These patterns are robust to other specifications for ρ as shown in Figure A12 in the Appendix.

Welfare losses from misperceptions can be quantified by comparing belief-implied WTA to its correct-belief benchmark. The value of choosing $s = C$ rather than $s = N$ to a sender with beliefs $\hat{b}$ can be approximated by $W^d = W^d_{\hat{B}} + W^d_X$, and she is willing to give up to this amount for the conspicuous profile. A sender with correct beliefs values the same choice at $W^d_B + W^d_X$. The welfare loss due to misperceptions is therefore $W^d_{\hat{B}} - W^d_B$: the distance between the light blue and the dark blue bars in Figure 8. These losses amount to around TZS 2,800–2,900 ($4 PPP) in prestige, credit, and friendship, and around TZS 950 ($1.4 PPP) in help.

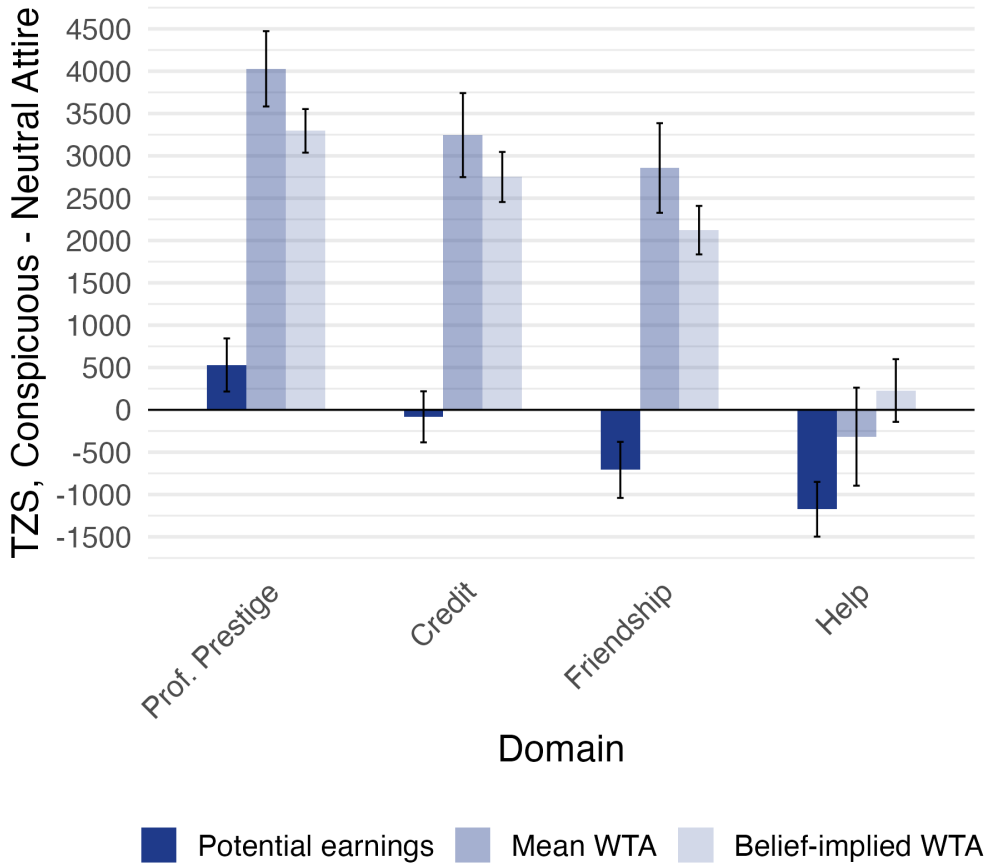
Taken together, conspicuous consumption in this setting is largely instrumental and misperceptions over its returns reduce welfare in every domain, leaving scope for welfare-improving policy.

5.4 Experimental Approaches

The structural decomposition of WTA holds under maintained structural assumptions about the utility function. To go beyond these assumptions, I now discuss two experimental approaches that aim at causally identifying how much of WTA is driven by beliefs.

One experimental approach to causally identifying the effect of beliefs on WTA is to isolate the intrinsic component by removing all evaluative stakes from the WTA choice. I attempted this by embedding a photo book control condition in the beliefs experiment. Participants were told that for one randomly chosen domain, one of their profiles will be displayed in a non-evaluative photo book. No evaluations or bonus payment were associated with this condition. To avoid contaminating the primary belief and WTA elicitation in the present data collection, the photo book condition always succeeded the primary outcome

Figure 8: Willingness to accept for conspicuous attire versus neutral attire profile, with belief implied WTA



Notes: Mean WTA for the conspicuous relative to the neutral profile (CC minus Neutral) across the four outcome domains, alongside imputed potential earnings and belief-implied WTA. Dark blue bars show imputed potential earnings — the expected monetary return from submitting the conspicuous rather than neutral profile to a real evaluation, computed as the difference in evaluation probabilities multiplied by the bonus payment of TZS 10,000 ($15 PPP). Medium blue bars show mean WTA from the MPL, signed positive when participants prefer the conspicuous profile and negative when they prefer the neutral profile. Light blue bars show the non-instrumental component of WTA inferred from the structural analysis. All bars show means with 95% confidence intervals. TZS values are in levels (CC minus Neutral).

tasks. Participants faced an identical preference and MPL choice using the same format as in the main evaluation condition. Results are reported in Figure A13 in the Appendix. Examination of the data, however, reveals two sources of contamination that prevent these results from being interpreted as reliable measures of the non-instrumental component. First, a non-trivial share of participants describe reasoning for the WTA task that reflects the evaluation context, suggesting that the preceding evaluation tasks influence choices even when instrumental stakes are nominally removed. Residual correlations between photo book WTA

and elicited beliefs between 0.18-0.46 corroborate that beliefs about evaluations continue to shape choices in the photobook condition. Second, a substantial share of participants describe the photo book as being viewed by a large audience, rather than the single photo book reader described in the instructions. Both sources of contamination have the potential to inflate photo book WTA above the true non-instrumental component.

A second approach attempts to circumvent these obstacles by isolating the intrinsic motive. In a follow-up visit to the participants of the beliefs experiment, one outcome domain is randomly assigned to be the subject of this survey, for each participant. The procedure unfolds in three steps. First, I re-elicit baseline beliefs and WTA for the assigned domain, establishing a within-visit reference point unaffected by any treatment. Second, for the assigned domain, I randomly draw 100 evaluations (separately for conspicuous and neutral attire profiles) from the pool of evaluations collected in the first experiment, and provide information on the average earnings associated with each profile type to the participant.[22] The random draw of 100 evaluations ensures that the information varies exogenously across participants. Third, I re-elicit beliefs and WTA for the same domain. This design identifies the causal effect of beliefs on WTA using the randomly drawn evaluation results as an instrument to exogenously shift beliefs. I defer the econometric details of this design to the appendix, Section D.1. It is the goal to collect data on this latter approach in Summer 2027.

6 Sources of Misperception

The preceding section establishes that beliefs are the quantitatively dominant driver of senders' willingness to accept and that misperceptions can therefore lead to welfare-diminishing signaling choices. This section asks: how can the welfare losses from misperceived signaling returns be reduced? Misperceptions can lie in one of the three layers of signaling (i) which traits the signal conveys, (ii) how these traits are being valued by evaluators and (iii) how both differ by domain. I show that misperceptions in each of these warrants a different policy. Using complementary experimental data, I locate the misperceptions and conclude on the optimal policy.

[22]The information is communicated in the following form: "We have shown profiles of people from Dar es Salaam, just like yours, to [domain-specific evaluators]. On average, people earned TZS X with their conspicuous attire profile in this evaluation and TZS Y with their neutral attire profile."

6.1 Layers of signaling

Recall the sender utility from sending signal s in domain d as defined in Section 5.1:

$$U^d(s,\rho) = \hat{B}^d(s,\rho) + \phi^d(s) = \hat{b}^d(s) \cdot u(\bar{\omega}+\pi,\rho) + (1-\hat{b}^d(s)) \cdot u(\bar{\omega},\rho) \; + \; \phi^d(s). \tag{5}$$

This section focuses on $\hat{b}^d(s)$, the sender's belief of successful evaluation of their signal in domain d. Forming a belief about how successful a signal is in a given domain requires forming a belief about three layers of signaling: (i) which traits the signal conveys, (ii) how these traits are being valued by evaluators, (iii) how both of these differ from the average in domain d; both (i) and (ii) may differ by domain, different evaluators may infer different traits from the same signal and different evaluators may value the same traits differently. A sender has type $\theta \in \mathbb{R}^K$, where θ_k is the $k-$th trait of the sender (e.g. wealth, trustworthiness, kindness). Then formally beliefs can be written as the product:

$$\hat{b}^d(s) = \hat{v}^d \cdot \widehat{\mathbb{E}}(\theta|s,d) \equiv \underbrace{\hat{v}^d}_{\text{Valuation of traits}} \cdot \underbrace{\hat{x}^d(s)}_{\text{Traits conveyed}} \tag{9}$$

$\hat{v}, \hat{x}$ are both $K \times 1$ vectors, defined component wise for each trait k such that the operation is a dot product. I therefore assume that the belief of receiving a successful evaluation for signal s is linear in the posterior mean of the sender's traits. Similarly, the realized probability of receiving a successful evaluation for signal s in domain d is:

$$b^d(s) = v^d \cdot x^d(s) \tag{10}$$

Misperceptions arise whenever senders' perceptions deviate from their true counterparts at either stage: $\hat{x}^d \neq x^d$ (senders misjudge which traits the signal conveys) or $\hat{v}^d \neq v^d$ (senders misjudge how evaluators value these traits). Both produce the same symptom, a miscalibrated $\hat{b}^d$, but they are different errors with different remedies.

The third layer, how traits conveyed and their valuations differ by domain, does not appear to be a major source of misperceptions. Figure 5 shows that senders' beliefs preserve the ranking of realized payoffs across domains: they correctly identify where conspicuous attire pays off most and where it pays off least, and overestimate the payoff in every domain. Misperceptions about how domains differ from one another would distort this ranking. An upward shift that leaves the ranking intact instead points to misperceptions in the parts of the belief that all domains share: which traits the signal conveys, and how these traits are valued relative to one another. In what follows, I therefore focus on misperceptions in the domain-invariant

layers of signaling.

6.2 Policies

For the remaining layers (i) and (ii) consider two policies, each targeting the misperception in one part of the belief:

1. **Content policy:** reduce misperceptions about what the signal conveys ($\hat{x}^d \to x^d$). This can take two forms: (a) substituting attire with a signal whose content is better understood[23] — in the credit domain, lenders assess documented transaction or income histories instead of conducting in-person interviews in which attire is visible — or (b) informing senders about which traits evaluators actually infer from conspicuous attire.

2. **Valuation policy:** reduce misperceptions about how evaluators value the conveyed traits ($\hat{v}^d \to v^d$), by publishing evaluation criteria and committing to them — in the credit domain, lenders publish the criteria by which applications are assessed, including the weight placed on traits such as trustworthiness, and bind loan officers to these weights.

By construction, each policy eliminates the misperception at one part of the belief and leaves the other untouched. Recall that the misperception driving the welfare losses of Section **??** is

$$\hat{b}^d - b^d \;=\; \hat{v}^d \cdot \hat{x}^d - v^d \cdot x^d.$$

The content policy ($\hat{x}^d = x^d$) reduces this misperception to $x^d \cdot (\hat{v}^d - v^d)$; the valuation policy ($\hat{v}^d = v^d$) reduces it to $v^d \cdot (\hat{x}^d - x^d)$. Which policy should be deployed therefore depends on where the misperception sits. If it sits in one part only, the matched policy eliminates the welfare loss entirely, while the mismatched policy achieves nothing; deploying both is then unnecessarily costly. If it sits in both parts, no single policy eliminates the loss in general:[24] each leaves a residual belief error and, when the two misperceptions partially offset each other, removing only one may increase the total error, as the following proposition shows.

Proposition 1 *Suppose misperceptions are present in both parts of the belief.*

[23] Substituting signals may change the cost of signaling; the welfare comparison below assumes signaling costs to remain.

[24] The exception is the knife-edge in which the truth revealed by one policy makes the remaining misperception irrelevant: if evaluators place no value on the trait ($v^d = 0$), the valuation policy alone eliminates the loss, since it no longer matters what senders believe the signal conveys.

(i) If the two misperceptions push the perceived return in the same direction,

$$[x^d \cdot (\hat{v}^d - v^d)][\hat{v}^d \cdot (\hat{x}^d - x^d)] > 0,$$

then each single policy strictly reduces the belief error $|\hat{b}^d - b^d|$.

(ii) If the two misperceptions push the perceived return in opposite directions,

$$[x^d \cdot (\hat{v}^d - v^d)][\hat{v}^d \cdot (\hat{x}^d - x^d)] < 0,$$

then a single policy increases *the belief error* $|\hat{b}^d - b^d|$ *whenever the misperception it removes is smaller in magnitude than twice the misperception it leaves behind.*

(iii) Deploying both policies eliminates the belief error entirely.

The proof follows from comparing absolute errors before and after each policy was implemented. Intuitively, a sender who overestimates both how much wealth her attire conveys to the loan officer and how much the officer values wealth is brought closer to the optimal signaling choice by either correction: assessing documented income instead of appearance, or publishing the lending criteria, each removes part of her inflated return. By contrast, publishing lending criteria to a sender who overestimates how much wealth her attire conveys but underestimates how much the officer values wealth then reveals that wealth matters more than she thought, while her content misperception persists. Her perceived return to conspicuous attire rises, and she overspends more than before. In either case, only the combined policy eliminates the welfare loss. Locating and characterizing the misperceptions is therefore the final step before concluding on the optimal policy.

6.3 Locating Misperceptions

To locate misperceptions, I first investigate whether senders correctly infer (i) which traits the signal conveys and to what extent, using complementary experimental data. Second, I quantify misperceptions in layer (ii), the valuation of traits, using a back-of-the-envelope calculation. Finally, I conclude on the optimal policy.

6.3.1 Content Layer

Traits conveyed by the signal are stable across domains. That is, the trait posterior is invariant to domains:

$$x^d(s) = x(s) \quad \forall d \in D.$$

As discussed in Section 3.4, in the Evaluations experiment, evaluators in all four outcome domains rate a subset of two profiles with respect to four traits: competence, assertiveness, trustworthiness, and kindness. In additional rounds of data collection, they also evaluate wealth. Because each trait was elicited identically across domains, I can test directly whether evaluators from different domains draw different trait inferences from conspicuous versus neutral attire. Figure A15 shows no evidence that they do. All sixteen domain–trait estimates agree in sign within traits, and estimates cluster by trait rather than by domain. The F-statistics reported in the figure test the joint equality of the conspicuous-attire effect across all four evaluator domains, trait by trait; I cannot reject equality for any trait ($p \geq 0.45$). In the next part, I therefore treat evalutions and beliefs about effects on traits as domain invariant.

Misperceptions in signal content concern a single trait: wealth. In particular I find that

$$\hat{x}_{wealth}(s) > x_{wealth}(s).$$

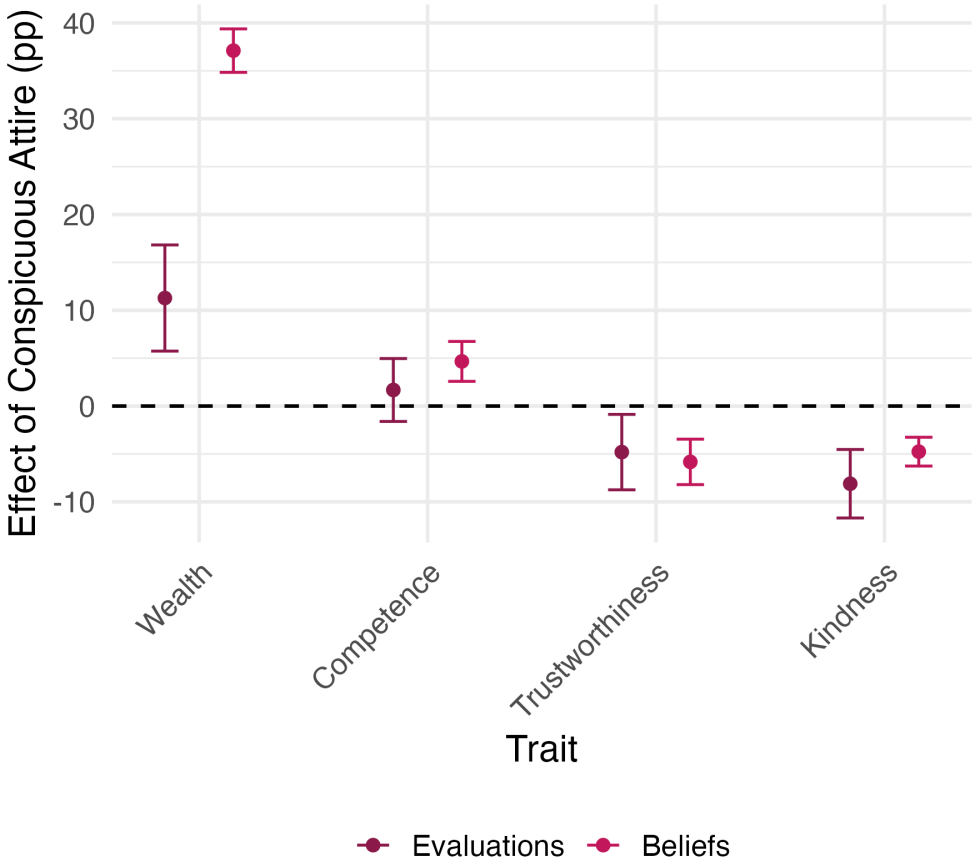
To show this I first need to bring two objects on the same scale: senders' beliefs about how conspicuous attire changes trait perceptions, and the realized changes in trait perceptions measured in the evaluations experiment. Eliciting senders' beliefs about every trait directly was not feasible: a full elicitation would have required each participant to state, for both profiles and each of the traits, the share of evaluators who would ascribe the relevant trait to them. This is a task whose length and repetitiveness would have compromised response quality. I therefore combine two lighter elicitation approaches that together recover the same object and bring sender's perceptions on the same scale as evaluations. I outline these elicitation approaches in Appendix D.2.

Figure 9 presents the result of this exercise. For wealth, competence, trustworthiness and kindness, the figure plots the imputed sender belief (light pink) against the realized effect of conspicuous attire on trait evaluations (dark pink) as previously shown in Figure 3.[25] Perceived effects on assertiveness were not re-elicited from senders. The figure shows that sender's beliefs provide an accurate account of elicited trait evaluations for all traits except for wealth. Sender's beliefs over the effect of conspicuous attire on wealth is more than three-fold the evaluations effect. For all three other traits, the imputed belief lies within the 95 percent confidence interval of the realized evaluation effect, with gaps of two to four percentage points. Senders therefore appear to correctly perceive that conspicuous attire

[25] From senders, the evaluation instrument elicited selfishness. I treat the selfishness association differential, sign-reversed, as the kindness measure.

makes them appear less trustworthy, and less kind as well as slightly more competent, but they perceive conspicuous attire to boost perceptions of wealth much more than it actually does.

Figure 9: Perceived vs. realized effct of conspicuous attire on evaluations, by trait



Notes: Believed versus realized trait effects of conspicuous attire. The figure compares senders' beliefs about how conspicuous attire changes trait perceptions (light pink) with the realized effects measured in the Evaluations experiment (dark pink). Realized effects are estimated by linear probability models of a positive trait rating on the conspicuous-attire indicator with profile fixed effects, standard errors clustered at the profile-by-condition level (Equation 1). The belief for wealth is directly elicited; beliefs for the remaining traits are imputed from an association task. Whiskers show 95% confidence intervals; $N = 502$ belief-side participants, $N = 1,701 - 1,708$ evaluations per trait.

6.4 Valuation Layer

Misperceptions in valuations overweight wealth relative to other traits. I begin by asking how well sender's beliefs are tracked by a *wealth-only* return rule, that is whether sender's beliefs can be approximated by a benchmark that sets valuations of all traits but wealth to

zero:

$$\hat{b}^d_{Wealth}(s) = \hat{v}^d_{Wealth} \cdot \hat{x}^d_{Wealth}(s). \tag{11}$$

To estimate this benchmark, I elicit both ingredients to Equation (11) directly in the beliefs experiment: the perceived wealth gain $\Delta x_{Wealth}(s)$ is measured through an analogous elicitation to the beliefs of the four domains at 37.1 pp. The domain weights $\hat{v}_d$ are measured as participants' stated correlations between being seen as wealthy and success in each of the four domains.[26] Figure 10 plots the wealth-only benchmark return alongside elicited beliefs and evaluation returns for each domain. Beliefs track the benchmark closely: professional prestige, credit and friendship lie within 8 pp of the prediction, while help falls below it by 13.3 pp. All beliefs are located closer to the wealth-only benchmark than to the evaluation returns. Senders' stated beliefs therefore appear to be largely based on an inference of the signals' return to wealth. Notably, the wealth-only benchmark shares the cross domain variation of both the beliefs and evaluation returns, which suggests that wealth is the main trait dimension along which domains differ and this is perceived correctly.

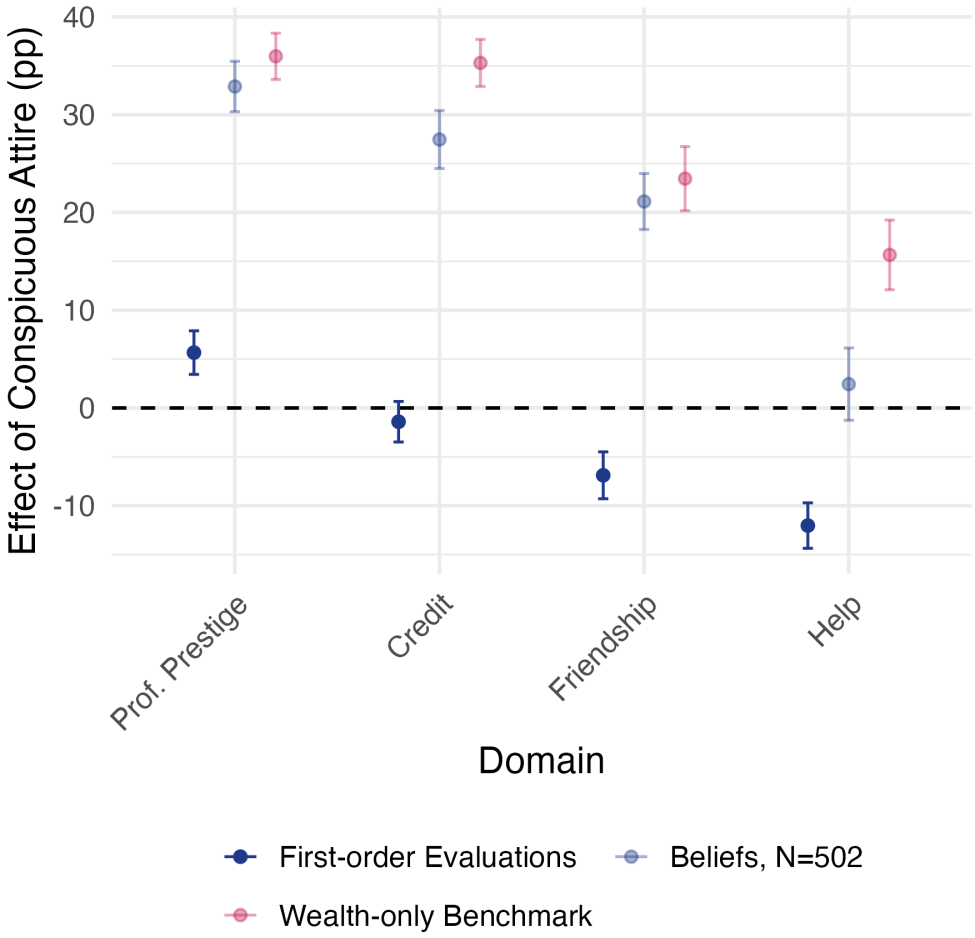
The fact that beliefs appear to mainly rest on inference of the signals' return to wealth does not by itself establish a misperception in valuations. The gap between beliefs and evaluations could be driven entirely by the content error documented in Section 6.3.1 above. To estimate whether misperceptions of the valuation of wealth relative to other traits adds to the misperception, I conduct a back-of-the envelope calibration. I write both realized and perceived payoffs as weighted sums of a wealth component and a composite of all other conveyed traits. For the trait components, I use the realized and perceived trait effects from Figure 9. Because these already contain the content error, the weights recovered from evaluations and from beliefs reflect misperceptions in valuations alone. I find that the weight senders place on wealth relative to other traits is roughly 1.4 times the weight it carries in realized evaluations. Appendix D.3 details the calibration. Senders thus overstate the valuation of wealth on top of overestimating how much wealth conspicuous attire conveys.

6.4.1 Optimal Policy

This section thus concludes that senders exhibit substantial misperceptions in both parts of the belief: in (i) the traits conveyed by conspicuous relative to neutral attire, and in (ii) how much these traits are valued by evaluators. Both misperceptions inflate the role of the same trait, wealth: senders overestimate the effect of conspicuous attire on perceived wealth more than threefold, and they overweight its valuation relative to other traits by a factor of

[26] The elicited correlations average 0.97 (professional prestige), 0.95 (credit), 0.59 (friendship), and 0.35 (help).

Figure 10: Perceived vs. realized effect of conspicuous attire on evaluations, with wealth-only belief benchmark



Notes: Elicited domain-specific belief returns versus predictions constructed from the elicited wealth belief return and the elicited wealth-domain correlation. Dark blue circles show realized first-order evaluation returns (as in Figure 5); light blue circles show elicited domain-specific belief returns; dark pink shows the elicited wealth belief return; light pink shows the predicted belief return for each domain, constructed as the wealth belief return scaled by the elicited wealth-domain correlation. Error bars show 95% confidence intervals. The dashed horizontal line marks the level of the wealth belief return.

1.4. Because the two misperceptions push perceived returns in the same direction, part (i) of Proposition 1 applies: either the content or the valuation policy individually reduces the welfare losses from misperceptions, but only their combination eliminates them.

6.5 Causes of Misperception

Where misperceptions come from matters for policy. Their source determines what a correction has to target; their underlying cause determines how far-reaching it must be. I therefore

close by weighing candidate causes of the misperceptions documented in this paper. The discussion is not exhaustive, and a thorough investigation of each mechanism is beyond the scope of this paper.

Pluralistic ignorance. I rule out pluralistic ignorance as the main driver of misperceptions. Under pluralistic ignorance, individuals hold accurate private views but misperceive what others in their group think (Bursztyn et al., 2020). Evaluators' second-order beliefs (Figure A3), however, closely replicate their first-order evaluations (Figure 2): evaluators know how their peers respond to conspicuous attire. The misperceptions are therefore located on the sender side rather than in a shared misreading of the signal among evaluators.

Scarce feedback. A lack of feedback is unlikely to explain misperceptions of this size. The gaps are largest where feedback is arguably most available: credit decisions are communicated directly and have immediate consequences, and gossip spreads quickly in social circles (Banerjee et al., 2019). If misperceptions persisted only for lack of feedback, they should be smallest in exactly these domains. The depth or quality of feedback on the other hand is precisely what the policies proposed in the preceeding section aim to improve.

Ingroup versus outgroup. Senders might project the evaluations of their ingroup, those in their immediate circles, onto evaluators outside those circles, such as loan officers and influential business people. This is unlikely to account for the pattern. Evaluators in the help and friendship domains are sampled on the same observables as the senders themselves, yet friendship produces among the largest misperceptions. Figure A6 further confirms that smaller deviations in this sampling leave the measured misperceptions unchanged.

Motivated beliefs. The structural decomposition in Section 5 assumes that preferences and beliefs are separable. Through motivated beliefs (Bénabou and Tirole, 2016), the two could instead interact: senders who intrinsically value conspicuous attire may persuade themselves that it also pays off socially. Evidence from the help domain speaks against this being the dominant force. In this domain, the belief component and the residual intrinsic-utility component have opposite signs. As in the other domains, beliefs pull toward conspicuous attire, while intrinsic utility pulls against it. If beliefs were motivated, they should instead be swayed in the direction of the intrinsic preference.

Salience and stereotypes. Salience and representativeness-based stereotypes (Bordalo et al., 2016, 2012) offer a more likely explanation of the misperceptions documented in this paper. Conspicuous attire is first and foremost a signal of wealth (Veblen, 1899). It is the trait this signal is most representative of. Senders who focus on this salient attribute form beliefs that track the wealth channel and neglect the other inferences their audience draws. The

evidence from Section 6.2 is consistent with this account. Senders' beliefs about the returns to conspicuous attire lie close to a benchmark that operates entirely through wealth (Figure 10), whereas evaluators also penalize conspicuous attire on decreased trustworthiness and kindness (Figure 3).

Sender–receiver differences. Finally, misperceptions may arise from the difference in roles senders take on when signaling relative to receivers. Senders choose their attire knowing their own trustworthiness and intentions, while evaluators observe only the outward signal. A sender who knows they are trustworthy may not anticipate that others read the same attire as a sign of the opposite. Such gaps between how people expect to be perceived and how they are actually perceived are a common source of misperceptions about others (Bursztyn and Yang, 2022).

Taken together, the evidence points toward causes rooted in how senders form beliefs, through salient stereotypes and the gap between the sender's and the receiver's perspective. This matters for two reasons. First, it leaves senders highly vulnerable to signaling mistakes: they invest in a signal whose returns they systematically misjudge, in domains where the stakes are high. Second, it shapes what policy can achieve. If misperceptions stemmed from scarce feedback, a one-off information intervention could correct them durably. Because they are instead regenerated by the way senders reason about signals, a one-time correction is likely to fade as senders revert to the salient wealth channel. Effective policies therefore need to be persistent, for example by being repeated or embedded in the settings where signaling decisions are made.

7 Conclusion

This paper identifies individual-level welfare losses from signaling: senders invest in a signal whose returns they misperceive, and pay for returns that do not materialize. Signaling can thus cause be privately, not just socially, wasteful.

In two linked experiments in Dar es Salaam, I show that conspicuous attire generates countervailing returns across socio-economic domains: it raises professional-prestige evaluations, leaves credit approval unchanged, and lowers friendship formation and informal help. Senders rank these returns correctly but overestimate their levels uniformly. For friendship, senders believe the signal helps when it hurts. Beliefs carry real stakes: senders give up considerable amounts of money for conspicuous attire even where it backfires. A structural decomposition attributes at least 74% of these amounts to beliefs rather than taste in domains of prestige, credit and friendship. Overspending therefore reflects misperceptions,

not preference. A conceptual framework locates the misperception in both parts of signaling: senders overestimate how much wealth conspicuous attire conveys, threefold, and how much evaluators value wealth relative to other traits, by a factor of 1.4. Because these two errors reinforce each other, correcting either the signal's content or its valuation reduces the welfare losses; only correcting both eliminates them.

More broadly, the paper points to a vulnerability of low-income settings. Where formal institutions are weak, socio-economic life runs on informal signals — and this paper shows that relying on them is expensive and hard to get right. Anticipating what a signal returns requires knowing both which traits it conveys and how these traits are valued, and senders who misjudge either invest in signals that work against them. The very settings that depend most on informal signaling are thus the ones most exposed to its welfare costs. Institutional reform like credit bureaus, employment records, published evaluation criteria, does more than improve information: it lifts the burden of signaling off individuals who can least afford to carry it.

Bibliography

Akerlof, G. (1976). The Economics of Caste and of the Rat Race and Other Woeful Tales*. *The Quarterly Journal of Economics*, 90(4):599–617.

Akerlof, G. A. and Kranton, R. E. (2000). Economics and identity. *The quarterly journal of economics*, 115(3):715–753.

Bagwell, K. and Riordan, M. (1991). High and declining prices signal product quality. *The American Economic Review*, pages 224–239.

Bagwell, L. S. and Bernheim, B. D. (1996). Veblen effects in a theory of conspicuous consumption. *The American Economic Review*, pages 349–373.

Banerjee, A., Chandrasekhar, A. G., Duflo, E., and Jackson, M. O. (2019). Using Gossips to Spread Information: Theory and Evidence from Two Randomized Controlled Trials. *The Review of Economic Studies*, 86(6):2453–2490.

Banerjee, A. V. and Duflo, E. (2007). The economic lives of the poor. *Journal of economic perspectives*, 21(1):141–167.

Baseler, T. (2023). Hidden Income and the Perceived Returns to Migration. *American Economic Journal: Applied Economics*, 15(4):321–352.

Becker, G. M., DeGroot, M. H., and Marschak, J. (1964). Measuring utility by a single-response sequential method. *Behavioral science*, 9(3):226–232.

Beekman, G., Gatto, M., and Nillesen, E. (2015). Family networks and income hiding: Evidence from lab-in-the-field experiments in rural Liberia. *Journal of African Economies*, 24(3):453–469.

Bellezza, S., Gino, F., and Keinan, A. (2014). The Red Sneakers Effect: Inferring Status and Competence from Signals of Nonconformity. *Journal of Consumer Research*, 41(1):35–54.

Bénabou, R. and Tirole, J. (2016). Mindful Economics: The Production, Consumption, and Value of Beliefs. *Journal of Economic Perspectives*, 30(3):141–164.

Bloch, F., Rao, V., and Desai, S. (2004). Wedding celebrations as conspicuous consumption: Signaling social status in rural India. *Journal of Human Resources*, 39(3):675–695.

Bordalo, P., Coffman, K., Gennaioli, N., and Shleifer, A. (2016). Stereotypes. *The Quarterly Journal of Economics*, 131(4):1753–1794.

Bordalo, P., Gennaioli, N., and Shleifer, A. (2012). Salience Theory of Choice Under Risk. *The Quarterly Journal of Economics*, 127(3):1243–1285.

Bordalo, P., Gennaioli, N., and Shleifer, A. (2013). Salience and Consumer Choice. *Journal of Political Economy*, 121(5):803–843.

Boston Consulting Group (2021). Understanding Consumers in the "Many Africas". https://www.bcg.com/publications/2014/understanding-consumers-in-the-many-africas.

Burger, R., Louw, M., Pegado, B. B. I. d. O., and van der Berg, S. (2015). Understanding consumption patterns of the established and emerging South African black middle class. *Development Southern Africa*, 32(1):41–56.

Bursztyn, L., Ferman, B., Fiorin, S., Kanz, M., and Rao, G. (2018). Status goods: Experimental evidence from platinum credit cards. *The Quarterly Journal of Economics*, 133(3):1561–1595.

Bursztyn, L., González, A. L., and Yanagizawa-Drott, D. (2020). Misperceived Social Norms: Women Working Outside the Home in Saudi Arabia. *American Economic Review*, 110(10):2997–3029.

Bursztyn, L. and Yang, D. Y. (2022). Misperceptions about others. *Annual Review of Economics*, 14(1):425–452.

Carranza, E., Donald, A., Grosset-Touba, F., and Kaur, S. (2025). The Social Tax: Redistributive Pressure and Labor Supply. *Econometrica*, 93(6):2273–2308.

Dizon-Ross, R. (2019). Parents' Beliefs about Their Children's Academic Ability: Implications for Educational Investments. *American Economic Review*, 109(8):2728–2765.

Dupas, P., Fafchamps, M., and Hernandez-Nunez, L. (2024). Keeping up appearances: An experimental investigation on social pressure and relative ranks. *National Bureau of Economic Research*.

Farrell, J. and Gibbons, R. (1989). Cheap Talk with Two Audiences. *The American Economic Review*, 79(5):1214–1223.

Feltovich, N., Harbaugh, R., and To, T. (2002). Too cool for school? Signalling and countersignalling. *RAND Journal of Economics*, pages 630–649.

Financial Sector Deepening Tanzania (FSDT) (2023). FinScope Tanzania 2023: Insights that Drive Innovation. Survey report.

Fiske, S. T. (2018). Stereotype Content: Warmth and Competence Endure. *Current Directions in Psychological Science*, 27(2):67–73.

Frank, R. H. (1985). The demand for unobservable and other nonpositional goods. *The American Economic Review*, 75(1):101–116.

Garcia, S. M., Weaver, K., and Chen, P. (2019). The Status Signals Paradox. *Social Psychological and Personality Science*, 10(5):690–696.

Hopkins, E. and Kornienko, T. (2004). Running to keep in the same place: Consumer choice as a game of status. *American Economic Review*, 94(4):1085–1107.

Jack, B. K., McDermott, K., and Sautmann, A. (2022). Multiple price lists for willingness to pay elicitation. *Journal of Development Economics*, 159:102977.

Jaikumar, S. and Sharma, Y. (2021). Consuming beyond means: Debt trap of conspicuous consumption in an emerging economy. *Journal of Marketing Theory and Practice*, 29(2):233–249.

Jaikumar, S., Singh, R., and Sarin, A. (2018). 'I show off, so I am well off': Subjective economic well-being and conspicuous consumption in an emerging economy. *Journal of Business Research*, 86:386–393.

Jakiela, P. and Ozier, O. (2016). Does Africa need a rotten kin theorem? Experimental evidence from village economies. *The Review of Economic Studies*, 83(1):231–268.

Jensen, R. (2010). The (Perceived) Returns to Education and the Demand for Schooling. *The Quarterly Journal of Economics*, 125(2):515–548.

Macchi, E. (2023). Worth your weight: Experimental evidence on the benefits of obesity in low-income countries. *American Economic Review*, 113(9):2287–2322.

McKenzie, D., Gibson, J., and Stillman, S. (2013). A land of milk and honey with streets paved with gold: Do emigrants have over-optimistic expectations about incomes abroad? *Journal of Development Economics*, 102:116–127.

National Bureau of Statistics (NBS), Tanzania (2018). Household Budget Survey 2017/18. Technical report, Dodoma, Tanzania (NBS headquarters).

National Bureau of Statistics (NBS), Tanzania (2021). Integrated Labour Force Survey 2020/21: Analytical Report. Technical report.

National Bureau of Statistics (NBS), Tanzania (2022a). NBS | Administrative Reports. Technical report.

National Bureau of Statistics (NBS), Tanzania (2022b). NBS | Age and Sex Distribution Reports. Technical report.

Nelissen, R. M. and Meijers, M. H. (2011). Social benefits of luxury brands as costly signals of wealth and status. *Evolution and human behavior*, 32(5):343–355.

Squires, M. (2024). Kinship taxation as an impediment to growth: Experimental evidence from Kenyan microenterprises. *The Economic Journal*, 134(662):2558–2579.

Srna, S., Barasch, A., and Small, D. A. (2022). On the value of modesty: How signals of status undermine cooperation. *Journal of Personality and Social Psychology*, 123(4):676.

Statista (2024). Luxury goods - Tanzania market analysis. https://www.statista.com/outlook/cmo/luxury-goods/tanzania.

Tanzania National Bureau of Statistics (2021). Tanzania - Tanzania Panel Survey 2020-2021, Wave 5. Technical report.

Tanzania National Bureau of Statistics (2024). Tanzania - Integrated Labour Force Survey 2024. Technical report.

Tanzania Ports Authority (2026). Dar es Salaam Ports. Technical report, Tanzania Posrts Authority.

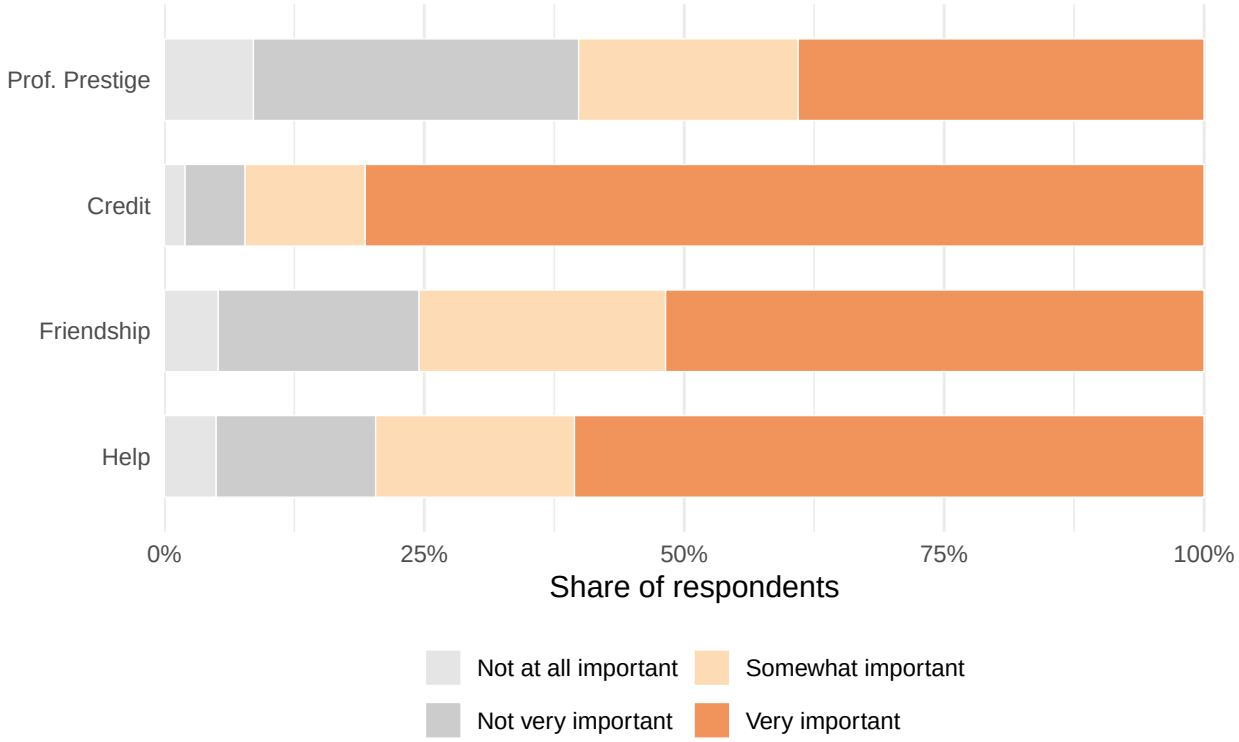
Townsend, R. M. (1994). Risk and Insurance in Village India. *Econometrica*, 62(3):539–591.

Udry, C. (1994). Risk and Insurance in a Rural Credit Market: An Empirical Investigation in Northern Nigeria. *The Review of Economic Studies*, 61(3):495–526.

Veblen, T. (1899). *The Theory of the Leisure Class: An Economic Study of Institutions*. Aakar Books.

A Supplementary Figures

Figure A1: Importance of domains for senders

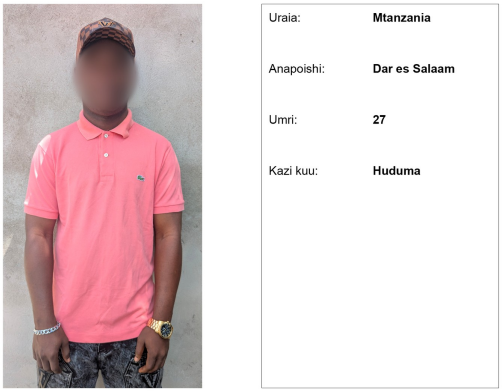


Notes: Importance senders attach to the four outcome domains. The figure shows the distribution of responses to the question how important each domain is to participants' own lives, elicited from all senders in the Beliefs experiment on a four-point scale from "not at all important" to "very important" ($N = 502$). Bars sum to 100% within each domain. Across all four domains, responses concentrate in the upper two categories, indicating that the domains in which signaling returns are measured are ones participants themselves regard as important.

Figure A2: Example profile pairs, neutral and conspicuous attire profiles



(a) Neutral attire profile



(b) Conspicuous attire profile



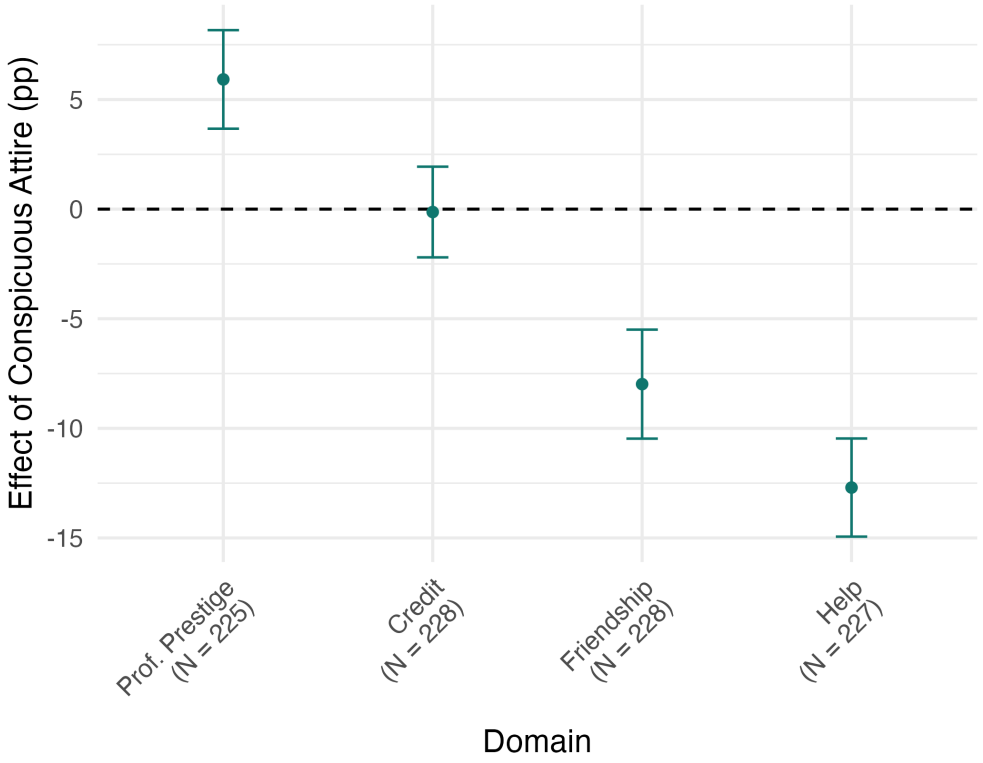
(c) Neutral attire profile



(d) Conspicuous attire profile

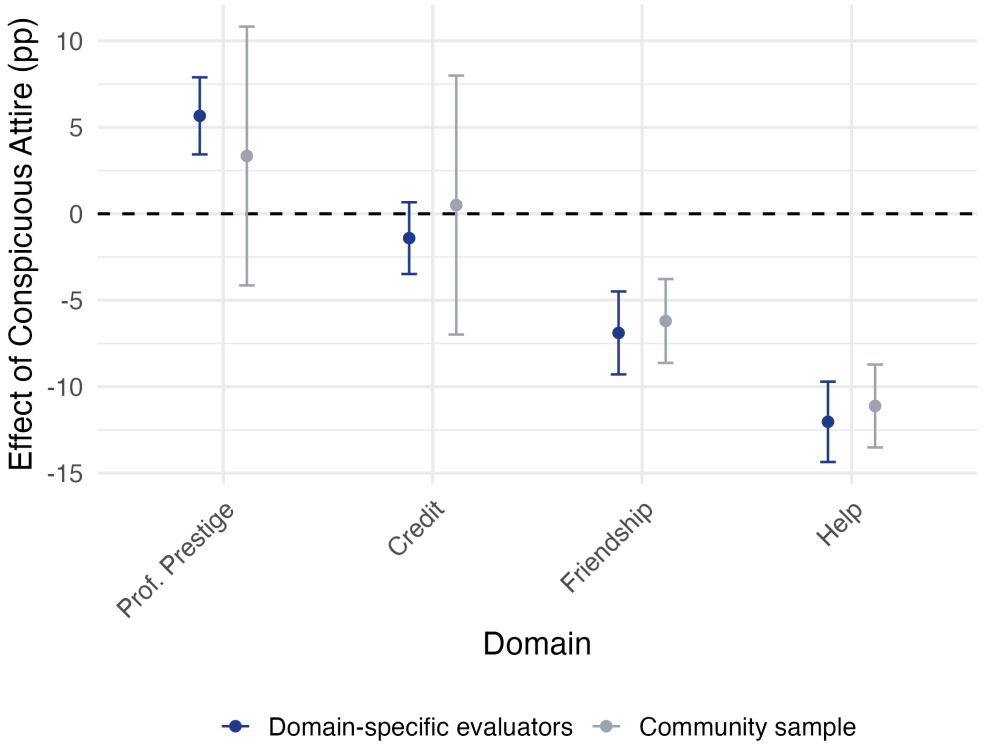
Notes: Example profile pairs (Swahili original). Panels (a) and (b) show the neutral and conspicuous attire profiles of one male profiled individual; Panels (c) and (d) show those of one female profiled individual. Within each pair, the two profiles are identical in all respects except attire: the same person is photographed against the same background, and demographic information — nationality, city of residence, age, and occupation sector — is held constant across profiles. Faces are redacted for anonymity. The conspicuous attire profile features visually prominent, branded clothing and accessories; the neutral profile features plain, unbranded clothing. Profile labels are in Swahili: *Uraia* (Nationality), *Anapoishi* (City of Residence), *Umri* (Age), *Kazi kuu* (Main Occupation Sector). Occupation sectors *Huduma* and *Biashara* translate to Services and Business, respectively. Standardized outfits were selected from a pre-determined set of garments to ensure consistency across profiles and to minimize variation in attire beyond the conspicuous versus neutral distinction.

Figure A3: Effect of conspicuous attire on evaluations, Second-order beliefs



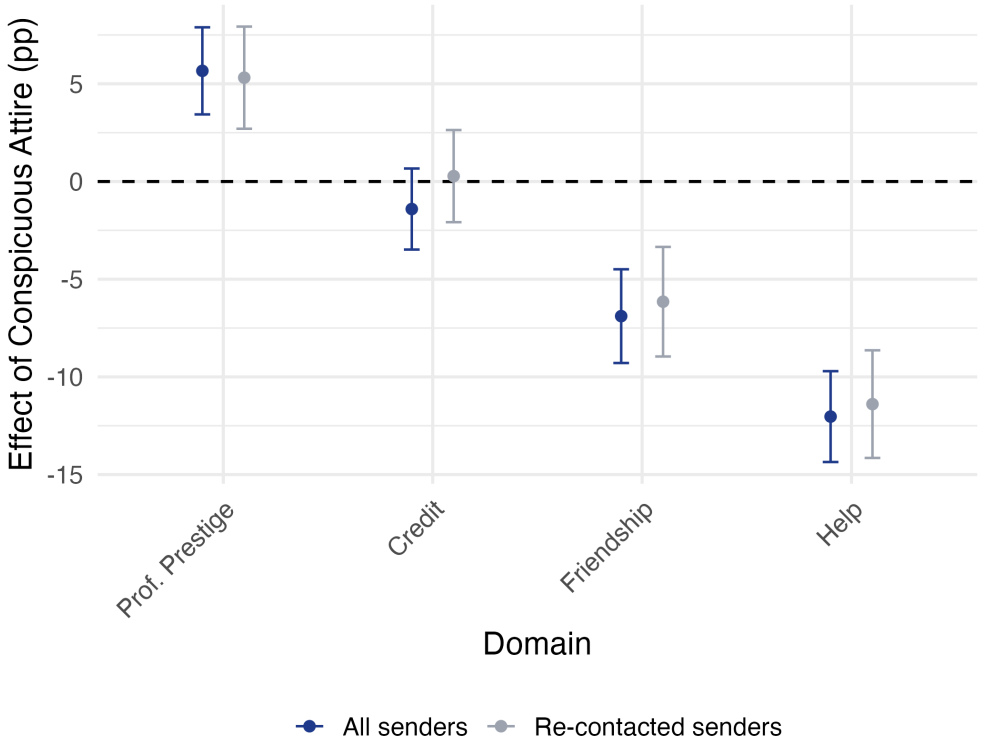
Notes: Effect of conspicuous attire on second-order (beliefs about others) evaluations across outcome domains. Coefficients are from domain-specific linear probability models (Equation 1), estimated separately for each outcome. The dependent variable is a binary indicator equal to one if the second-order evaluation falls in the positive range of the seven-point Likert scale. The treatment indicator equals one for profiles featuring conspicuous attire. All specifications include profiled individual fixed effects and evaluator fixed effects. Standard errors are clustered at the profiled individual × attire level. Horizontal bars indicate 95% confidence intervals.

Figure A4: Effect of conspicuous attire on evaluations, with community sample evaluations only



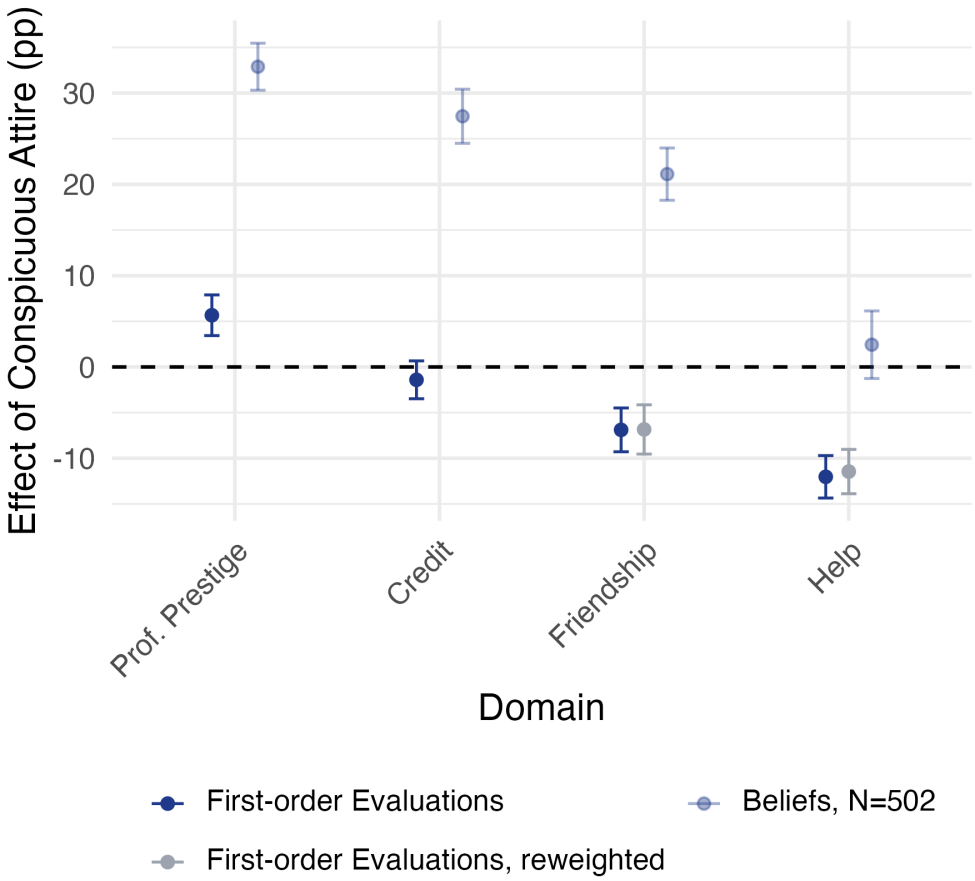
Notes: The figure reports LPM estimates of the effect of conspicuous attire on the probability of a positive first-order evaluation, estimated separately by domain on evaluations by the pooled community sample (the friendship and help evaluators, $N = 455$). Specifications include profiled-individual fixed effects; standard errors are clustered at the individual × attire level; whiskers show 95% confidence intervals. For friendship and help, estimates draw on the community samples' main evaluation task together with cross-evaluations by the other community sample ($N = 4{,}102$ and 4,085 evaluations); for professional prestige and credit, estimates rest solely on the two cross-evaluations each evaluator completes at the end of the survey ($N = 603$ and 653 evaluations after excluding profiles cross-evaluated only once, which the fixed effects absorb) and are accordingly less precise. Prestige cross-evaluations collected before the revision of the prestige question are excluded, mirroring the main prestige evaluator sample.

Figure A5: Effect of conspicuous attire on evaluations, with re-contacted sender profiles only



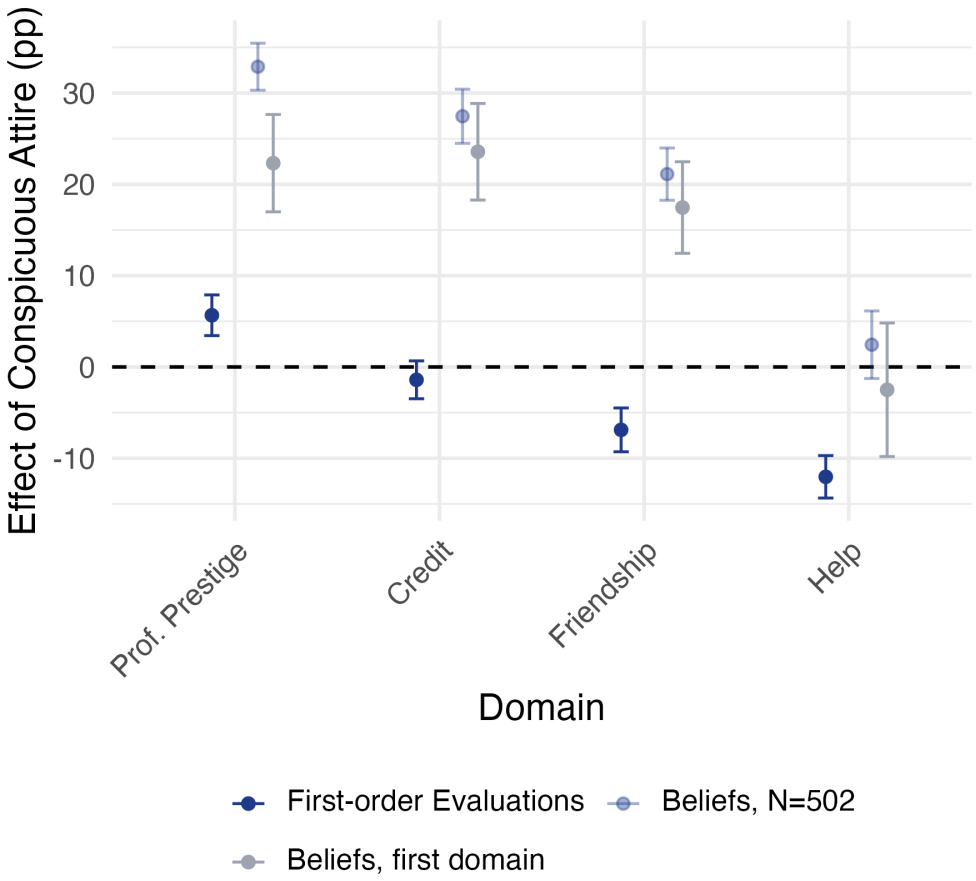
Notes: Effect of conspicuous vs. neutral attire on the probability of a positive first-order evaluation, by outcome domain. Grey estimates restrict the sample to evaluations of profiles of the 502 senders re-contacted for the beliefs experiment (Section 4.1.1); blue estimates reproduce the full-sample results of Figure 2 (655 senders). Coefficients are from domain-specific linear probability models (Equation 1). The dependent variable is a binary indicator equal to one if the evaluation falls in the positive range of the seven-point Likert scale. The treatment indicator equals one for profiles featuring conspicuous attire. All specifications include profiled-individual and evaluator fixed effects. Standard errors are clustered at the profiled-individual × attire level. Horizontal bars indicate 95% confidence intervals.

Figure A6: Perceived vs. realized effect of conspicuous attire, with reweighted evaluators



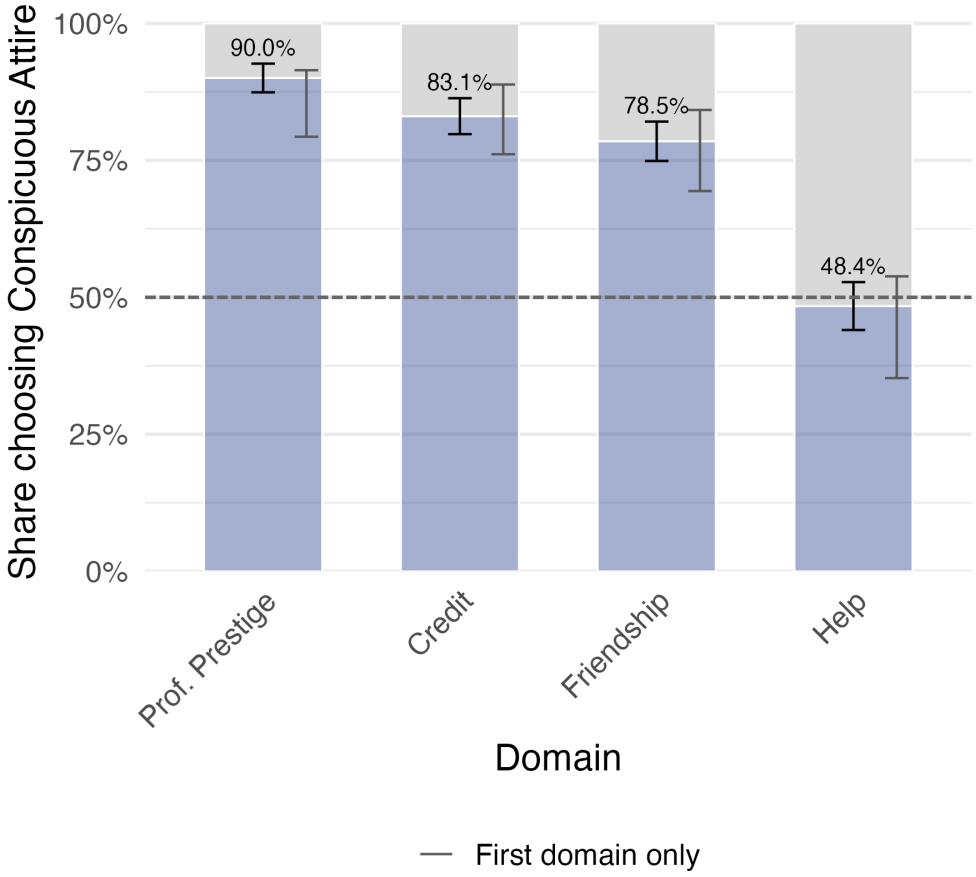
Notes: Robustness of Figure 5 to the composition of the community evaluator samples. Dark blue and light blue circles reproduce the realized first-order evaluation effects and the mean belief returns from Figure 5. Grey markers show realized evaluation effects for friendship and help after reweighting the respective evaluator samples to match the senders who completed the beliefs experiment ($N = 502$) on gender, age bracket, district, occupation sector, and household-income bracket. Weights are constructed by iterative proportional fitting (raking) of evaluator-level characteristics to the sender marginal distributions; effective sample sizes after weighting are X evaluators for friendship and Y for help. Eight friendship and four help evaluators in the top household-income bracket, which contains no beliefs-experiment sender, receive zero weight and are excluded. Professional-prestige and credit evaluators are professional populations about whose identity senders were explicitly informed, and are therefore not reweighted. All estimates are expressed as percentage-point differences (conspicuous minus neutral); error bars show 95% confidence intervals.

Figure A7: Perceived vs. realized effect of conspicuous attire, with first-domain beliefs only



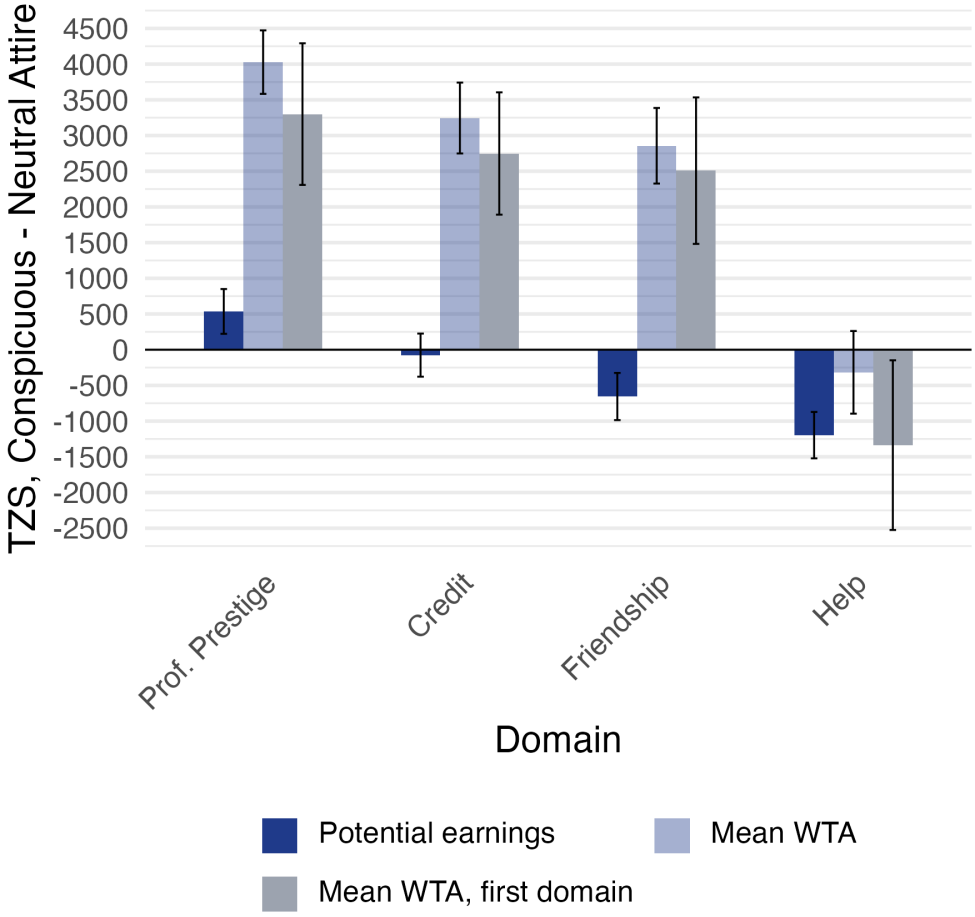
Notes: Robustness of Figure 5 to the order of elicitation. Dark blue and light blue circles reproduce the realized first-order evaluation effects and the full-sample belief returns from Figure 5. Grey circles show belief returns estimated only from each sender's first elicited domain, before any other domain was elicited ($N \approx 125$ senders per domain; the domain order is randomized). All estimates are expressed as percentage-point differences (conspicuous minus neutral); error bars show 95% confidence intervals.

Figure A8: Preference for conspicuous attire, with first-domain choices only



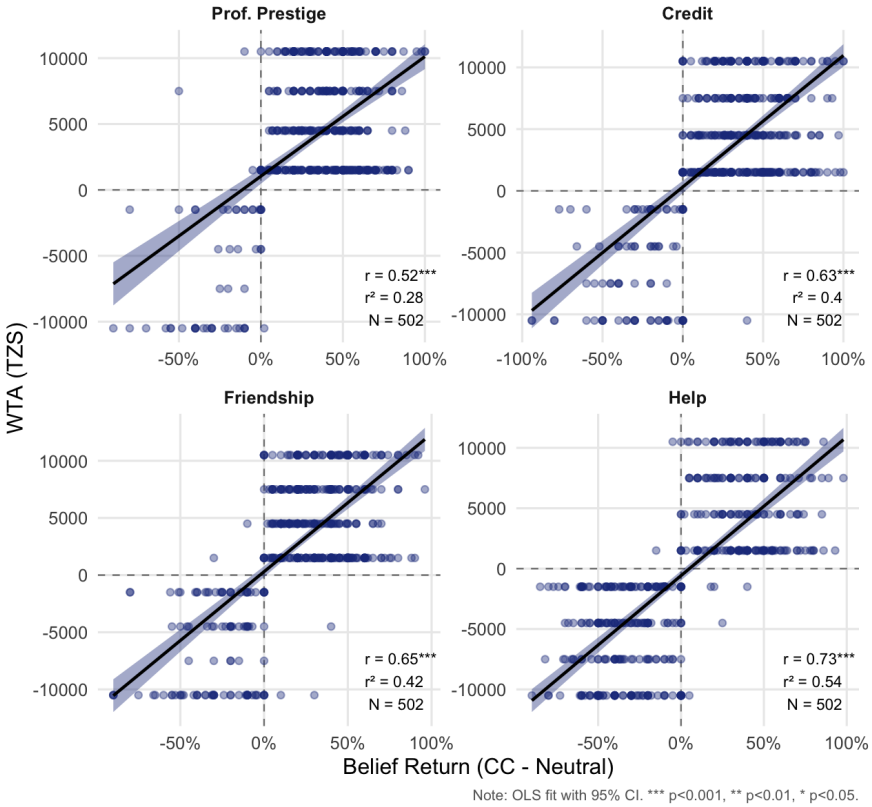
Notes: Robustness of Figure 6 to the order of elicitation. Bars and percentage labels reproduce the full-sample shares of senders choosing to submit their conspicuous rather than their neutral profile ($N = 502$); black bars indicate their 95% confidence intervals. Grey bars show the 95% confidence interval of the corresponding share computed only from each sender's first elicited domain ($N \approx 125$ senders per domain; the domain order is randomized). The dashed line marks indifference at 50%.

Figure A9: Willingness to accept and potential earnings, with first-domain choices only



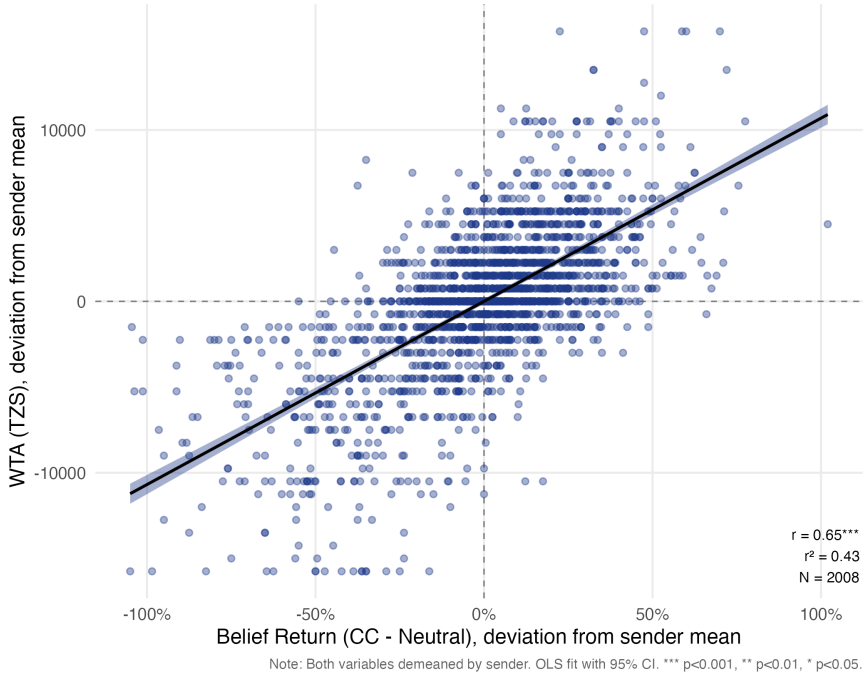
Notes: Robustness of Figure 7 to the order of elicitation. Dark blue bars show imputed potential earnings and light blue bars the full-sample mean WTA for the conspicuous over the neutral profile, as in Figure 7. Grey bars show mean WTA computed only from each sender's first elicited domain ($N \approx 125$ senders per domain; the domain order is randomized). All values are expressed in TZS (conspicuous minus neutral); error bars show 95% confidence intervals.

Figure A10: Correlation between perceived beliefs and Willingness to accept



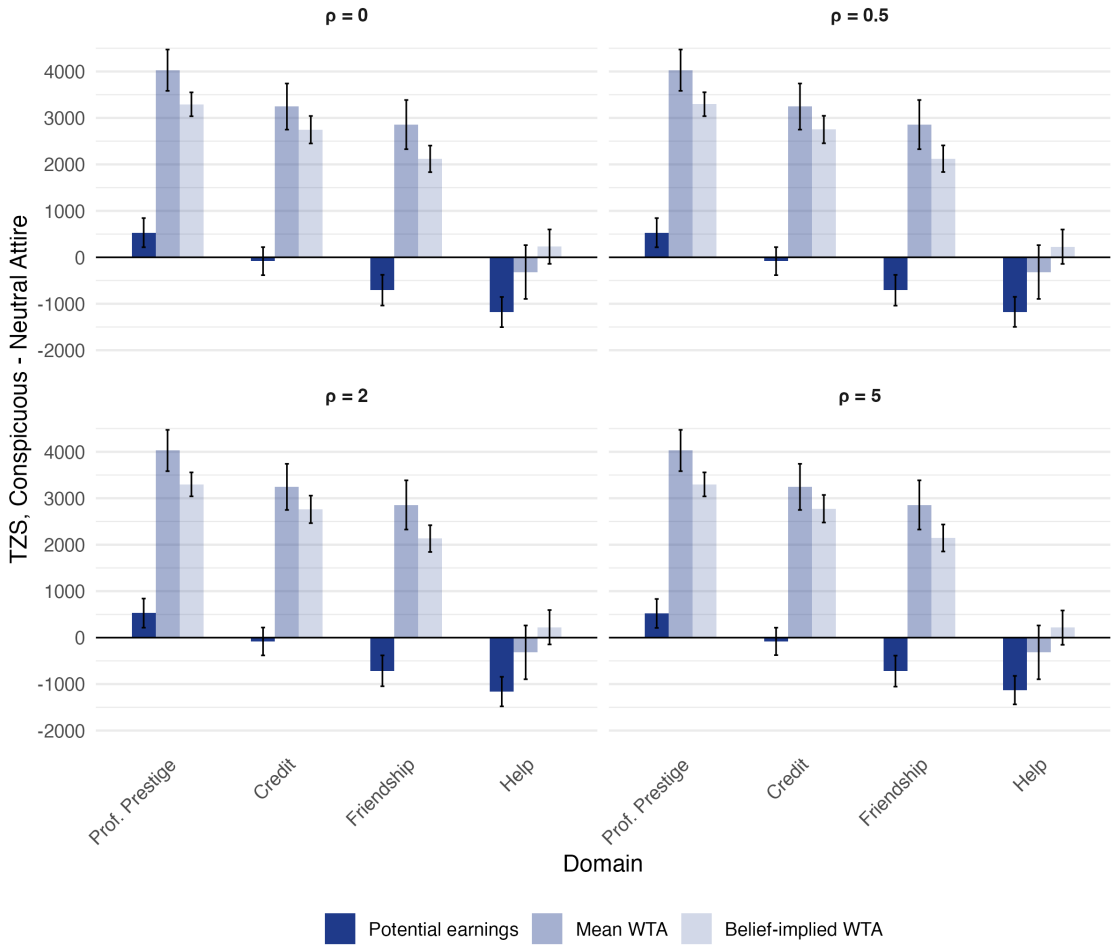
Notes: Individual-level scatter of WTA against belief returns for each outcome domain. Each dot represents one participant. The solid line shows the OLS fit with 95% confidence band. Belief returns are computed as the difference in elicited belief probabilities between the conspicuous and neutral profile (CC minus Neutral). WTA is the MPL midpoint at switching, in TZS, signed positive for conspicuous and negative for neutral preference. Reported statistics: Pearson correlation coefficient (r), coefficient of determination (r^2), and number of observations (N). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Figure A11: Correlation between perceived beliefs and Willingness to accept, in deviations from sender means



Notes: Within-sender scatter of WTA against belief returns, pooled across the four outcome domains. Both variables are demeaned by sender, so each dot represents one sender–domain observation expressed as a deviation from that sender's own mean, and the figure isolates within-sender variation: whether a given sender is willing to pay more in the domains where she believes the return to conspicuous attire is larger. The solid line shows the OLS fit with 95% confidence band; its slope equals the coefficient from a regression of WTA on belief returns with sender fixed effects. Belief returns and WTA are constructed as in Figure A10. Reported statistics: Pearson correlation coefficient (r), coefficient of determination (r^2), and number of sender–domain observations (N). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Figure A12: Structural decomposition of WTA under alternative risk-aversion parameters



Notes: Each panel replicates the decomposition in Figure 8 under a different coefficient of relative risk aversion, $\rho \in \{0, 0.5, 2, 5\}$; the main text uses $\rho = 0.5$. Dark bars show potential earnings, medium bars mean WTA, and light bars the belief-implied WTA, computed for each sender from her elicited beliefs under the respective curvature and averaged across senders. Potential earnings and mean WTA do not depend on ρ and are identical across panels. All values in TZS, conspicuous minus neutral attire. Bars show means with 95% confidence intervals. Across the four values of ρ, the mean belief-implied WTA varies by less than TZS 30 ($0.05 PPP) in every domain, so the decomposition is not sensitive to the assumed degree of risk aversion.

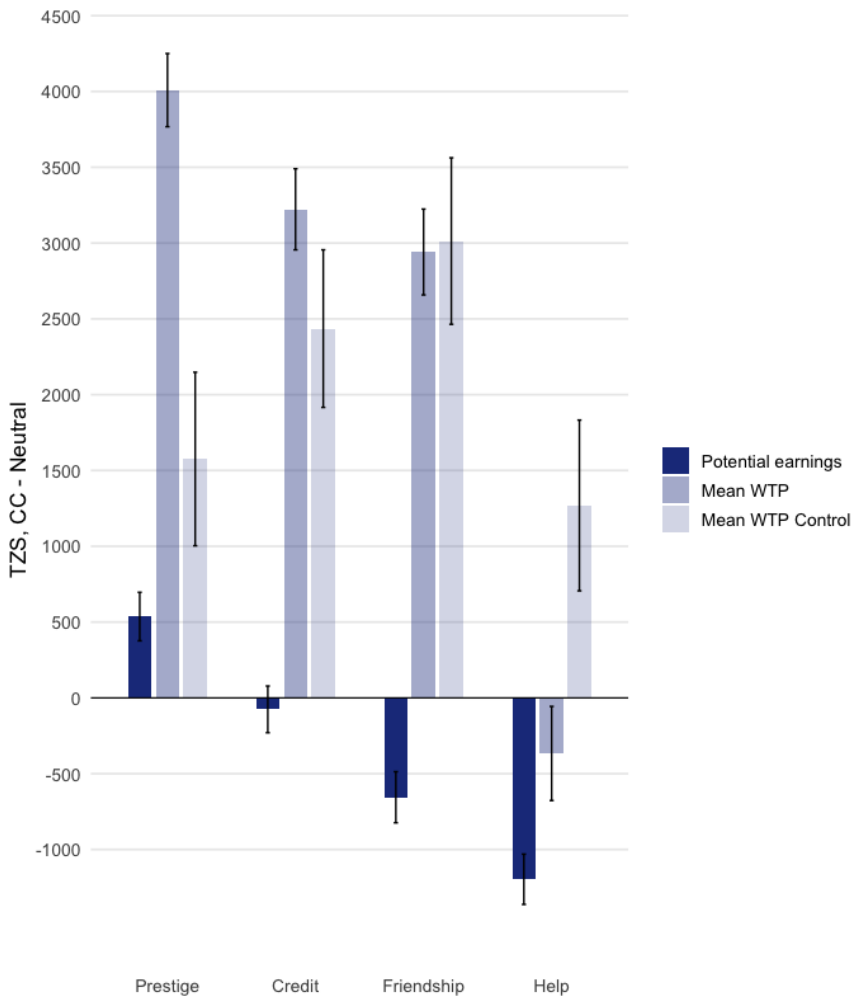


Figure A13: Mean WTA for the conspicuous relative to the neutral profile (CC minus Neutral) across the four outcome domains, alongside imputed potential earnings and photo book WTA. Dark blue bars show imputed potential earnings — the expected monetary return from submitting the conspicuous rather than neutral profile to a real evaluation, computed as the difference in evaluation probabilities multiplied by the bonus payment of TZS 10,000 ($15 PPP). Medium blue bars show mean WTA from the MPL, signed positive when participants prefer the conspicuous profile and negative when they prefer the neutral profile. Light blue bars show mean WTA in the photo book control condition, in which profiles are displayed in a non-evaluative photo book rather than submitted to real evaluators. All bars show means with 95% confidence intervals. TZS values are in levels (CC minus Neutral).

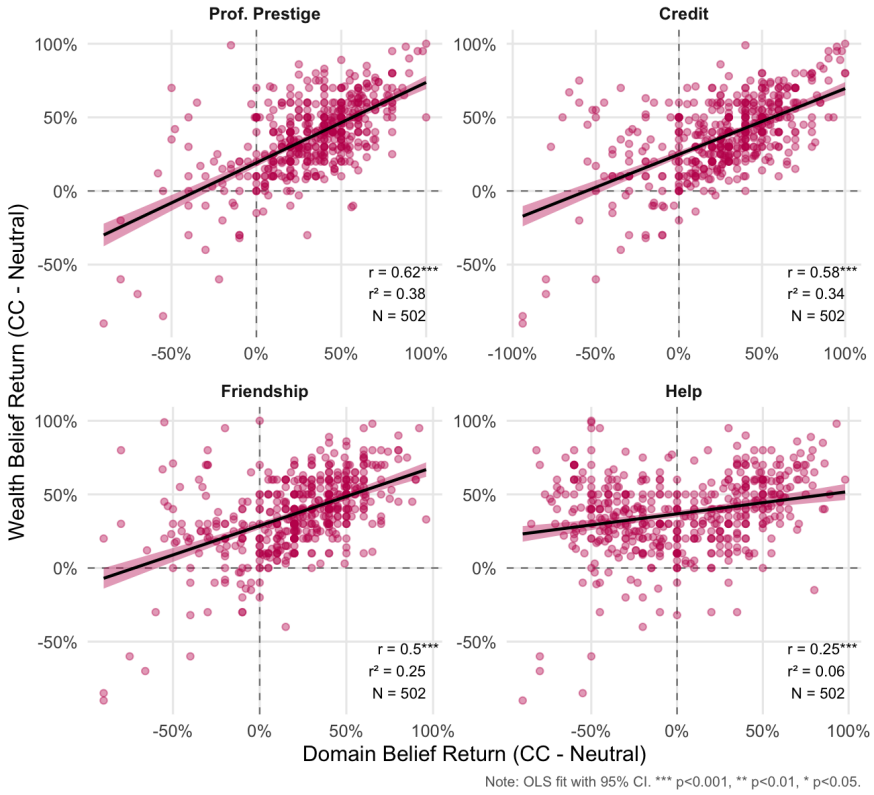


Figure A14: Individual-level scatter of wealth belief returns against domain belief returns for each outcome domain. Each dot represents one participant. The solid line shows the OLS fit with 95% confidence band. Belief returns are computed as the difference in elicited belief probabilities between the conspicuous and neutral profile (CC minus Neutral), separately for the wealth belief and for each domain-specific belief. Reported statistics: Pearson correlation coefficient (r), coefficient of determination (r^2), and number of observations (N). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

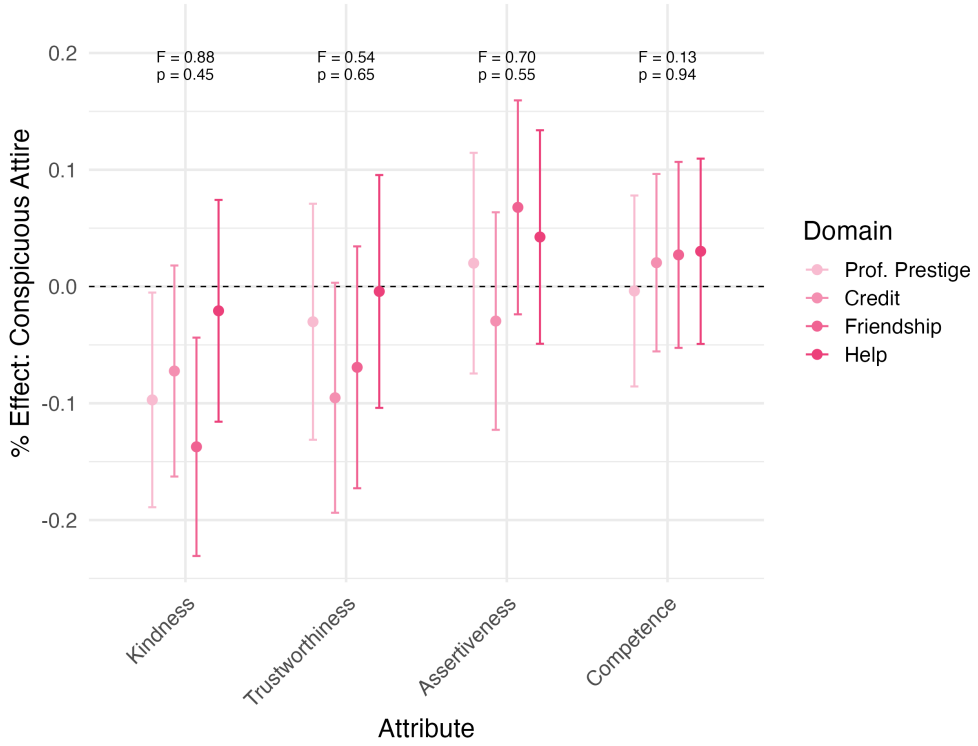


Figure A15: The figure plots the effect of conspicuous attire on the probability of a positive attribute rating, estimated separately for evaluators from each of the four domain samples (professional prestige, credit, friendship, and help). Estimates are from linear probability models of the form given in equation 1, regressing each attribute rating on the conspicuous-attire indicator interacted with evaluator-domain indicators, including model fixed effects and clustered at the profile level. Whiskers show 95% confidence intervals based on standard errors clustered at the profile-by-condition level; $N_{Kindness} = 1,701$, $N_{Trustworthinss} = 1,707$, $N_{Competence} = 1,708$, $N_{Assertiveness} = 1,708$. Displayed intervals are based on marginal standard errors; tests of equality across domains use the joint covariance matrix and can be more precise than the visual overlap suggests. Attribute inferences are stable across evaluator domains: all sixteen estimates agree in sign within attribute, and joint tests reject equality of effects across domains.

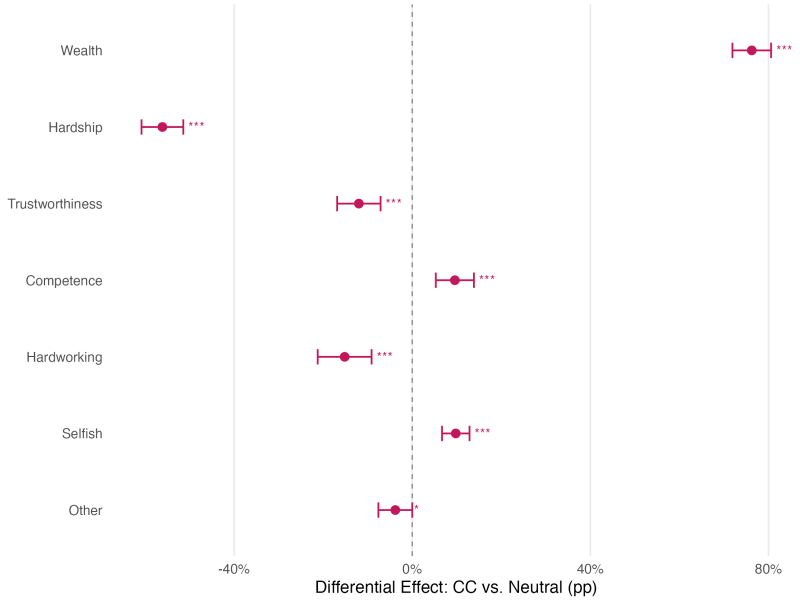


Figure A16: Differential effect of conspicuous versus neutral attire on trait associations. Coefficients show the change in the probability that a trait is selected (out of a maximum of two) when a profile is shown in conspicuous rather than neutral attire, estimated separately for each trait. Traits are elicited independently of any specific outcome domain. Horizontal bars indicate 95% confidence intervals. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

B Supplementary Tables

Table A1: Sample composition compared with the population of Dar es Salaam

		Senders		Evaluators			
	Population	All	Beliefs Sample	Friendship	Help	Prestige	Credit
N		655	502	228	227	225	228
Gender							
Male	48.7	49.2	48.4	50.0	49.8	60.0	59.6
Female	51.3	50.8	51.6	50.0	50.2	40.0	40.4
District							
Kinondoni	18.2	18.2	18.7	18.4	18.1	32.4	39.9
Ubungo	20.2	18.9	18.1	20.2	20.3	20.0	33.3
Ilala	30.6	30.8	31.3	30.7	30.8	27.1	16.2
Temeke	25.0	25.6	24.1	25.0	25.1	13.8	7.0
Kigamboni	5.9	6.4	7.8	5.7	5.7	6.7	3.5
Age							
18–24	25.9	23.2	21.3	24.6	24.2	1.8	14.5
25–29	18.2	20.9	17.5	19.3	18.5	16.4	28.5
30–39	25.9	29.5	32.9	28.1	28.2	41.3	43.4
40+	30.1	26.4	28.3	28.1	29.1	40.4	13.6
Monthly household income (TZS)							
Up to 200,000	14.9	22.7	16.7	11.0	16.7	0.0	3.9
200,001–400,000	28.2	27.8	32.3	28.9	27.8	0.0	18.4
400,001–800,000	32.5	29.6	35.1	31.1	30.0	0.0	43.0
800,001–2,000,000	21.2	18.5	14.5	25.4	23.8	14.2	23.2
Above 2,000,000	3.3	1.4	1.4	3.5	1.8	85.8	11.4
Occupation sector							
Agriculture & fishing	7.8	4.7	4.8	10.1	4.4	0.9	0.0
Industry	17.3	13.3	12.2	11.4	12.3	4.9	0.4
Services	18.7	26.1	24.9	18.4	24.2	8.4	70.6
Trade	32.3	34.0	37.1	32.0	32.6	84.9	11.8
Public administration	5.8	1.1	1.0	4.4	4.8	0.0	0.0
Transport	9.7	12.2	12.5	14.5	15.4	0.9	0.0
Other	8.4	3.8	3.8	2.6	2.2	0.0	17.1
Currently not working	–	4.7	3.8	6.6	4.0	0.0	0.0

Notes: Shares in percent. Senders: "All" refers to the full profile-creation sample (Section **??**); "Beliefs" to the subsample of senders who completed the beliefs experiment (Section 4). Senders and friendship and help evaluators were recruited on population quotas for gender, district, age, household income, and occupation sector. Population benchmarks for Dar es Salaam: gender and age (adults 18+) from the 2022 Census (National Bureau of Statistics (NBS), Tanzania, 2022b); district from the 2022 Census (National Bureau of Statistics (NBS), Tanzania, 2022a); occupation among employed persons from the ILFS 2024 (Tanzania National Bureau of Statistics, 2024); household consumption from a log-normal fit to the mean and median in the NPS 2020/21 (Tanzania National Bureau of Statistics, 2021), in monthly 2025 prices (CPI factor 1.15). Industry: mining, manufacturing, utilities, construction. Services: food service, education, other services.

Table A2: Evaluation returns to conspicuous attire

	Prof. Prestige		Credit		Friendship		Help	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Positive evaluation (LPM)								
Conspicuous attire	0.054***	0.057***	−0.015	−0.014	−0.068***	−0.069***	−0.120***	−0.120***
	(0.013)	(0.011)	(0.012)	(0.011)	(0.013)	(0.012)	(0.013)	(0.012)
Romano–Wolf p-value	[0.001]	[0.001]	[0.299]	[0.261]	[0.001]	[0.001]	[0.001]	[0.001]
Mean, neutral attire	0.45	0.45	0.47	0.47	0.62	0.62	0.47	0.47
Observations	3,599	3,599	3,645	3,645	3,647	3,647	3,628	3,628
Wald test, H_0: equal effects across domains: $p < 0.001$ (profile FE only), $p < 0.001$ (two-way FE)								
Panel B: Seven-point Likert scale								
Conspicuous attire	0.265***	0.279***	−0.089*	−0.085*	−0.332***	−0.335***	−0.611***	−0.611***
	(0.058)	(0.052)	(0.053)	(0.047)	(0.061)	(0.058)	(0.061)	(0.054)
Romano–Wolf p-value	[0.001]	[0.001]	[0.144]	[0.111]	[0.001]	[0.001]	[0.001]	[0.001]
Mean, neutral attire	−0.23	−0.23	−0.14	−0.14	0.61	0.61	−0.30	−0.30
Observations	3,599	3,599	3,645	3,645	3,647	3,647	3,628	3,628
Wald test, H_0: equal effects across domains: $p < 0.001$ (profile FE only), $p < 0.001$ (two-way FE)								
Profiled-individual FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Evaluator FE	No	Yes	No	Yes	No	Yes	No	Yes

Notes: Each column reports the coefficient on the conspicuous-attire indicator from a separate regression. In Panel A the dependent variable is an indicator equal to one if the evaluation falls in the positive range of the seven-point Likert scale; in Panel B it is the full seven-point score. Odd columns include profiled-individual fixed effects, even columns profiled-individual and evaluator fixed effects only. Standard errors, in parentheses, are clustered at the profiled-individual × attire level. Romano–Wolf p-values, in brackets, adjust for multiple hypothesis testing across the four domains within each outcome and specification, computed by a cluster bootstrap resampling profiled-individual × attire clusters (999 replications). Wald tests of equality are computed from a stacked regression interacting treatment with domain indicators. Means refer to evaluations of profiles in neutral attire. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table A3: Heterogeneous returns to conspicuous attire: evaluator characteristics

		Test of effect homogeneity			
Characteristic	Categories	Prof. Prestige	Credit	Friendship	Help
Gender	2	0.398 [1.000]	0.310 [1.000]	0.720 [1.000]	0.621 [1.000]
Age bracket	3–4	0.219 [0.991]	0.260 [0.998]	0.302 [1.000]	0.678 [1.000]
Schooling	2–3	0.486 [1.000]	0.489 [1.000]	0.552 [1.000]	0.318 [1.000]
High income	2	—	**0.071** [0.779]	0.899 [1.000]	0.760 [1.000]
Occupation sector	3–5	—	0.144 [0.954]	0.323 [1.000]	0.888 [1.000]
District	3–4	**0.053** [0.702]	0.667 [1.000]	0.903 [1.000]	**0.010** [0.235]
Marital status	2	0.530 [1.000]	**0.047** [0.674]	0.249 [0.995]	0.472 [1.000]
Children	4–5	0.299 [1.000]	0.276 [0.999]	0.383 [1.000]	0.938 [1.000]

Notes: Each cell corresponds to one regression in which the attire treatment is interacted with the listed evaluator characteristic, using the preferred specification: the dependent variable is the positive-evaluation indicator, with profiled-individual and evaluator fixed effects; standard errors are clustered at the profiled-individual × attire level. Main effects of evaluator characteristics are absorbed by the evaluator fixed effects. The first entry is the p-value of a Wald test that all interaction coefficients of the cell regression are jointly zero. The p-value in brackets adjusts for the fact that the table tests 30 such joint hypotheses, following the Romano–Wolf stepdown procedure with a standard cluster bootstrap resampling profiled-individual × attire clusters (999 replications). Raw p-values below 0.10 are in bold. Categories with fewer than 20 evaluators are excluded from the respective test, and tests with fewer than two remaining categories after this exclusion are not computed ("—").

Table A4: Heterogeneous returns to conspicuous attire: profiled-individual characteristics

		Test of effect homogeneity			
Characteristic	Categories	Prof. Prestige	Credit	Friendship	Help
Gender	2	**0.039** [0.791]	**0.002** [0.273]	**0.027** [0.724]	**0.023** [0.689]
Age bracket	4	0.855 [1.000]	0.886 [1.000]	0.429 [1.000]	0.721 [1.000]
Occupation sector	7	0.780 [1.000]	0.183 [0.981]	**0.084** [0.914]	0.462 [1.000]
Homophily	4	0.808 [1.000]	0.848 [1.000]	0.597 [1.000]	0.691 [1.000]

Notes: Each cell corresponds to one regression in which the attire treatment is interacted with the listed characteristic, using the preferred specification: the dependent variable is the positive-evaluation indicator, with profiled-individual and evaluator fixed effects; standard errors are clustered at the profiled-individual × attire level. Gender, age bracket, and occupation sector refer to the profiled individual, as displayed on the profile; their main effects are absorbed by the profiled-individual fixed effects. Homophily is a characteristic of the evaluator–individual pair: the number of the three displayed characteristics they share (0–3). The first entry is the p-value of a Wald test that all interaction coefficients of the cell regression are jointly zero. The p-value in brackets adjusts for the fact that the table tests 16 such joint hypotheses, following the Romano–Wolf stepdown procedure with a standard cluster bootstrap resampling profiled-individual × attire clusters (999 replications). Raw p-values below 0.10 are in bold. Categories with fewer than 20 profiled individuals are excluded from the respective test, and tests with fewer than two remaining categories after this exclusion are not computed ("—").

Table A5: Perceived vs. realized returns to conspicuous attire

	Prof. Prestige (1)	Credit (2)	Friendship (3)	Help (4)
Perceived return (β_B)	0.329***	0.275***	0.211***	0.024
	(0.013)	(0.015)	(0.015)	(0.019)
Realized return (β)	0.057***	−0.014	−0.069***	−0.120***
	(0.013)	(0.012)	(0.014)	(0.013)
Misperception ($\beta_B - \beta$)	0.273***	0.289***	0.281***	0.143***
	(0.021)	(0.023)	(0.023)	(0.028)
Romano–Wolf p-value	[0.001]	[0.001]	[0.001]	[0.001]
Mean belief, neutral attire	0.36	0.43	0.50	0.52
Observations, beliefs	1,002	1,002	1,002	1,002
Observations, evaluations	3,601	3,647	3,649	3,630
Wald test, H_0: equal perceived returns across domains: $p < 0.001$				
Wald test, H_0: no misperception in any domain: $p < 0.001$				
Profiled-individual FE	Yes	Yes	Yes	Yes

Notes: Each column stacks senders' incentivized beliefs, rescaled to probabilities, and first-order evaluations for one domain. Perceived returns are estimated from Equation 2; realized returns and the misperception are estimated jointly in the stacked regression, where the misperception is the interaction of the treatment with an indicator for belief observations. Standard errors, in parentheses, are clustered at the sender level for beliefs and the profiled-individual × attire level for evaluations. Romano–Wolf p-values, in brackets, adjust for multiple hypothesis testing across the four domains (cluster bootstrap, 999 replications). The first Wald test is of equal perceived returns across domains, from a stacked regression with standard errors clustered at the sender level; the second is of the joint null of no misperception in any domain. Means refer to beliefs about profiles in neutral attire. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

C Supplementary Survey Materials

C.1 Verbatim Vignette Questions, Evaluations

- **Prof. Prestige**
  - Question: *"Imagine you are organizing an exclusive dinner for influential business people. You are considering whether to invite this person, whom you have not met before. How likely are you to invite them as one of the influential business people to the dinner?"*
  - Evaluators, $N = 225$: Evaluators are purposively recruited local business decision-makers in Dar es Salaam. This group includes business owners or co-owners, individuals holding senior managerial roles with decision-making authority (e.g.

hiring, partnerships, procurement), members of business associations or professional networks, and individuals working in business incubators, accelerators, investment, or recruitment firms. High-income individuals (reporting monthly income above TZS 2 million ($3,023 PPP)) may also be included.

- **Credit Access**
  - Question: *"Imagine this person, whom you have never met before, is applying for a business loan of 1 Million Shilling. If you are asked to assess them, how likely are you to approve this person for the loan?"*
  - Evaluators, $N = 228$: Loan Officers, credit analysts, other loan professionals from Dar es Salaam.
- **Friendship**
  - Question: *"Imagine you are meeting this person for the first time at a social gathering. How likely are you to begin a friendship with this person after the first meeting?"*
  - Evaluators, $N = 228$: Evaluators are recruited to broadly reflect variation in gender, district, income, and age among adults in Dar es Salaam. Participation is restricted to individuals who report regular public exposure (spending time outside their home most days of the week).
- **Help**
  - Question: *"Imagine you are running late for an appointment, and you see this person, whom you have never met before, struggling to carry several heavy bags. How likely are you to stop and offer this person your help?"*
  - Evaluators, $N = 227$: Evaluators are recruited to broadly reflect variation in gender, district, income, and age among adults in Dar es Salaam. Participation is restricted to individuals who report regular public exposure (spending time outside their home at least five days per week).

D Conceptual Framework — Appendix

D.1 Experimental Approaches to Causal Identification of Effect of Beliefs on Willingness to Accept

Below, I outline the econometric details of the second experimental approach to identifying the causal link between beliefs and WTA as discussed in Section **??**.

Let $t \in \{pre, post\}$ index the two elicitation rounds, $b^o_{i,t}$ denote individual i's belief return ($b^o_{CC,i,t} - b^o_{N,i,t} \in [-1, 1]$) in domain o at time t, $W^o_{i,t}$ denote WTA, and $r^o_{i,drawn}$ denote the average realized return implied by the randomly drawn evaluations shown to participant i. The first and second stage are:

$$b^o_{i,t} = \alpha_i + \gamma \cdot Post_t + \beta_1 \cdot Post_t \cdot r^o_{i,drawn} + \varepsilon_{i,t} \tag{12}$$

$$W^o_{i,t} = \alpha_i + \gamma \cdot Post_t + \beta_2 \cdot \hat{b}^o_{i,t} + \nu_{i,t} \tag{13}$$

where α_i are individual fixed effects absorbing time-invariant baseline beliefs and WTA, $Post_t$ captures any common time trend between the two elicitation rounds, and $\hat{b}^o_{i,t}$ in equation (13) is instrumented by $Post_t \cdot r^o_{i,drawn}$. The key parameter, β_2 identifies the causal effect of a one percentage point change in belief return on WTA. Under purely belief-driven WTA, $\beta_2 = \pi/100$, meaning a one percentage point increase in belief return shifts WTA by the corresponding expected monetary value of the bonus. If β_2 is substantially smaller, non-instrumental motives account for a meaningful share of WTA that is unresponsive to belief updating.

D.2 Bringing Evaluator and Sender Perceptions of Traits on the same Scale

First, I conduct an association task with senders covering a broad set of traits: at the end of the Beliefs experiment, I ask all 502 participants which kinds of people come to mind when they think of each of the two attire styles, selecting up to two traits from a list of seven. [27] Traits are presented in randomized order. For both neutral and conspicuous attire style, the task yields the share of participants who associate it with each trait. The difference in these shares between conspicuous and neutral attire measures the direction and relative magnitude of the perceived signal content. The results are shown in Figure A16. The figure plots, for

[27] Traits included: wealthy, struggling financially, trustworthy, competent, hardworking, selfish, and other (open text response).

each of seven traits, the difference in the share of participants associating the relevant trait with conspicuous rather than neutral attire. Conspicuous attire is overwhelmingly associated with wealth (+76 pp) and correspondingly dissociated from its inverse, financial hardship (−56 pp). The remaining traits carry smaller but significant differentials: conspicuous attire is positively associated with competence and with selfishness — the inverse indicator of kindness — and negatively with trustworthiness and being hardworking. These differential associations however live on a different scale than the trait evaluations data examined under Assumption 1. Comparing association differentials to these therefore requires a conversion onto a common scale.

Second, I conduct a direct belief elicitation for a single trait: wealth. On the same screens previously used to elicit beliefs on returns for the four main outcomes, participants now state, for each of the two profiles, how many out of 100 people in Dar es Salaam would rate them as wealthy on the same seven-point Likert Scale as before. Unlike the return beliefs, this elicitation was not incentivized, as no corresponding evaluator-side measure exists against which to score it. This yields wealth as one trait for which the perceived signal content is measured both on the association and directly on the probability scale. The result is shown in Figure 9: conspicuous attire raises the perceived probability of being considered wealthy by 37.1 pp.

I link the two elicitations by a simple observation: for wealth, I observe both the association differential (76.2 pp) and the belief differential (37.1 pp). Their ratio, 0.487, converts association differentials into the belief scale, assuming uniform conversion across traits. I therefore multiply the trait differentials measured in the association task by the same conversion factor to convert them to the probability scale. This yields the simplest possible out-of-sample prediction.

D.3 Quantifying Misperceptions in Valuations - Back-of-the Envelope Calculation

Wealth appears to dominate senders' beliefs relative to other traits; a simple back-of-the envelope calibration shows by how much. Denote $\hat{b}^d_{Wealth} = (\hat{v}_{Wealth} + \hat{\varepsilon}^d_{Wealth})\hat{x}^d$ the market-specific return based on the wealth-only benchmark. The market specific correlations elicited in the beliefs experiment serve as the weights, with $\hat{v}_{Wealth} = 0.716$ its average and $\hat{\varepsilon}^d_{Wealth}$ the market-specific deviations. $\hat{x}_{Wealth}$ is approximated by the elicited return to wealth from conspicuous vs. neutral attire; $\hat{x}_{Wealth} = 0.371$.

Let $(\hat{v}_O + \hat{\varepsilon}^d_O)g_O = \sum_k (\hat{v}_k + \hat{\varepsilon}^d_k)g_k$ be a normalized composite measure of all traits signaled by conspicuous attire other than wealth. I can then write perceived and realized returns in

every market d as the convex combination of realized (perceived) returns to wealth and to all other traits, where I impose that the average weights of wealth and other attributes sum to one. An over-weighting factor can be computed by taking each component in the resulting system of equations to the data and computing the ratio of weights applied to wealth in the perceived returns over the weights applied to wealth in the realized returns equation.