

Decoupling Taste-Based versus Statistical Discrimination in Elections*

Amanda de Albuquerque Frederico Finan Anubhav Jha
Laura Karpuska Francesco Trebbi

May 2025

Abstract

We present a methodology for decoupling taste-based versus statistical discrimination in political behavior. We combine a flexible empirical model of voting, featuring vertical and horizontal candidate differentiation in gender, ability, and policy positions, with a large-scale micro-targeted electoral experiment aimed at increasing female candidate vote shares. Our structural econometric approach allows to separately identify preference parameters driving taste-based discrimination and beliefs parameters driving statistical discrimination through expectations about ability and policy positions of female politicians. Our application to Brazilian municipal elections uncovers substantial levels of taste-based and statistical discrimination. Counterfactual political campaigns show promise in reducing both.

JEL Codes: D72, P0, P16

Keywords: Discrimination; Elections; Political Behavior, Bayesian Updating; Identity; Gender; Ability; Policy; Brazil.

*de Albuquerque: Rede A Ponte (amanda@redeaponte.org). Finan: UC Berkeley and National Bureau of Economic Research (ffinan@berkeley.edu). Jha: Ashoka University (anubhavpcjha@gmail.com). Karpuska: INSPER (laukarpuska@gmail.com). Trebbi: UC Berkeley and National Bureau of Economic Research (ftrebbi@berkeley.edu). American Economic Association RCT Registry ID: AEARCTR-0014405, initial registration date September 20, 2024. We thank Daniel Deisenroth, Chris Hooton, Steve Tadelis, and Nils Wernerfelt for early conversations on the use of Instagram features relevant to the analysis. Vivek Venkata, Helena Arruda, Livia Almeida, Fernanda Peron, and Amanda Pena provided excellent research assistance. The analysis in this article was not performed with the financial support or in collaboration with Instagram or Meta Platforms, Inc. We also thank seminar participants at UC Berkeley and UC Davis for useful comments and suggestions.

1 Introduction

A persistent challenge in political representation is understanding why women and minorities remain underrepresented in elected office. One key explanation lies on the demand side: voter discrimination. Discrimination, however, can take two distinct forms: statistical discrimination, in which voters judge individuals based on group-level characteristics or averages; and taste-based discrimination, which stems from intrinsic biases against candidate identities, regardless of their actual behavior or qualifications.¹ While both forms of discrimination can produce similar patterns of underrepresentation, distinguishing between them is crucial for effective policy design. Economic theory suggests that voter beliefs –the basis for statistical discrimination– are more amenable to change through targeted information and learning than the deeply rooted preferences (Fang and Moro (2011); Bertrand and Duflo (2017)). If statistical discrimination is the primary barrier, well-designed information campaigns could meaningfully improve the descriptive representation of women and ethnic or religious minorities. Conversely, if taste-based animus is the dominant force, the potential for policy interventions to reduce underrepresentation may be more limited.

Empirically disentangling statistical from taste-based discrimination is difficult for two main reasons. First, statistical discrimination operates through voter beliefs used to form expectations about candidates' unobserved characteristics, whereas taste-based discrimination reflects preferences over observed candidate attributes. Since both voter beliefs and preferences are unobserved, and jointly determine vote choice, reduced-form analysis of vote shares has difficulty in separating the two mechanisms.² Second, while labor market discrimination models often focus on vertical attributes like productivity or competence, voter decisions also involve horizontal attributes, such as policy positions and identity. Because horizontal attributes affect different voters differently (e.g. in spatial models), they introduce non-monotonicities in voting behavior that can obscure the presence of discrimination, especially when fitting aggregate vote share data through linear models.

In this paper, we present a random utility voting model, estimated in combination with a large-scale electoral experiment, that enables us to address this complex empirical identification problem. In our framework, voters choose among candidates based on three key attributes: gender identity, ability, and policy position. Ability is modeled as a vertical attribute,

¹See Phelps (1972); Arrow (1973, 1998) and for a more recent discussion Fang and Moro (2011); Guryan and Charles (2013).

²See Altonji and Pierret (2001) for a canonical example of how statistical discrimination arises from unobserved characteristics, and the difficulty of separating it from other mechanisms.

a common-value component of preferences (i.e., voters uniformly prefer higher-quality candidates).³ Gender identity and policy position are modeled as horizontal attributes, or private-value components (i.e., voters evaluate these features differently depending on how closely they align with their own identity or policy ideal point).⁴

While we assume that voters observe gender identity with certainty (a realistic assumption in our context), both ability and policy position are subject to uncertainty. Voters hold subjective and heterogeneous beliefs about these attributes, and it is through these beliefs that one can formally incorporate statistical discrimination in the model. Importantly, the model also allows utility weights assigned to each dimension –gender, ability, and policy– to vary across voters. This state-dependent component of preferences captures the psychological salience of each attribute, a behavioral feature that can be shaped by campaign messages independently of belief updating. Dimension-specific utility weights provide a natural way to incorporate taste-based discrimination in voting, reflecting intrinsic preferences over observed candidate identity.

To estimate our model, we conducted a randomized experiment in partnership with a Brazilian nonpartisan non-profit organization that micro-targeted voters via Instagram one week prior to the 2024 municipal elections in Brazil. Brazil provides a compelling setting for this experiment. Gender representation remains a persistent challenge in Brazilian politics, particularly at the local level. Although women make up 51% of the population, they are severely underrepresented in municipal councils, which are the focus of our paper. In the 2024 municipal elections, only 18.2% of elected councilors were women.

Social media, and Instagram in particular, offer a highly effective channel for reaching Brazilian voters due to their popularity. As of March 2024, Instagram had approximately 141 million users in Brazil, representing 63.9% of the population. The platform boasts a diverse and engaged user base across age groups, and Brazilian politicians have increasingly leveraged Instagram to connect with constituents. Social networks and messaging apps are now among the most important tools in political campaigning in Brazil.

Our experiment, conducted in 1,000 municipalities across the country, is designed to identify statistical versus taste-based discrimination within the structure of our model of political behavior. We randomly exposed female and male voters to informative and uninformative messages about female candidates. The informative messages provided hard information about candidates’ attributes, aiming to influence both voters’ beliefs and the salience of these at-

³This follows the standard treatment in models that incorporate both common-value and private-value dimensions of candidate evaluation. See for example [Kawai and Sunada \(2025\)](#).

⁴For an example see [Kendall et al. \(2015\)](#). For a recent discussion of the role of identity see [Bonomi et al. \(2021\)](#); [Gennaioli and Tabellini \(2025\)](#).

tributes in their utility functions. In contrast, uninformative messages were designed to only increase salience/preference weights without altering beliefs. This combination of message types offers a novel empirical strategy to separately identify the roles of statistical and taste-based discrimination in voter decision-making.

Our analysis yields several findings relevant to the Political Economy literature on political behavior and the Economics literature on discrimination. Estimation of our model provides evidence of taste-based discrimination, but we also show it can be experimentally reduced. Specifically, our gender identity treatment micro-targeted to male voters increased support for female politicians by reducing their taste-based disutility from voting against their own gender. In contrast, female voters exhibited a higher baseline willingness to vote for female candidates and the treatment had no measurable effect on female voters.

With respect to statistical discrimination, our estimates suggest that voters are heterogeneous in their perception of whether female candidates have higher ability than male candidates. Quantitatively, our ability messages did not significantly increase the vote share for female candidates.

Policy messages appear more effective in our context. We uncover a significant mismatch between female voters' policy preferences and their beliefs about female candidates' policy positions. Specifically, female candidates are perceived as more conservative than female voters themselves. This perceived ideological distance helps explain the persistent gap in vote shares for female candidates, particularly among female voters. Encouragingly, this form of statistical discrimination –rooted in beliefs about policy positions– appears more amenable to correction than misperceptions about ability in the Brazilian context. Our informative policy message effectively reduced the perceived distance between female voters' ideal points and the positions they attributed to female candidates. Overall, informative policy messages substantially increased female vote shares and representation.

Our model also allows us to quantitatively assess the relative importance of statistical versus taste-based discrimination. Eliminating statistical discrimination consistently reduces female underrepresentation and can even lead to overrepresentation. In contrast, removing taste-based discrimination has mixed effects and may sometimes increase underrepresentation. This counterintuitive result arises because both male and female voters tend to prefer same-gender candidates, so removing female voters' preference for women can reduce support for female candidates in some municipalities.

Another key advantage of our approach is its ability to quantitatively evaluate the impact

of alternative campaign strategies on both taste-based and statistical discrimination. This includes the simulation of “optimal” campaigns, which tailor the mix of electoral messages based on municipality-level characteristics. Results show that such optimized campaigns may increase vote shares of female candidates on average by 2 percentage points. When translated into cost-effectiveness terms, this corresponds to approximately 1.6 additional votes for female candidates per dollar spent.

This paper contributes to an active area of research in Political Economy and Political Science on discrimination in electoral representation. It also contributes more broadly to the wider literature in Labor and Public Economics on gender and minority discrimination.

First, we add to the extensive literature within Economics and Political Science examining demand-side factors contributing to the underrepresentation of women in politics. Specifically, we focus on voter bias, a factor previously explored with mixed findings. For example, in Economics [Casas-Arce and Saiz \(2015\)](#) find little evidence of voter discrimination following Spain’s introduction of reforms requiring political parties in municipalities exceeding 5,000 inhabitants to field candidate lists containing at least 40 percent women. This result aligns with [Gonzalez-Eiras and Sanz \(2021\)](#), who also examine Spanish local elections and find no evidence of voter bias when comparing female vote shares across closed-list and open-list electoral systems. Similarly, in Political Science, [Broockman and Soltas \(2020\)](#) investigate voter discrimination during Illinois Republican presidential primary elections from 2000 to 2016. Although they identify substantial voter bias against non-white delegates, they find no evidence of discrimination against women. Stereotypes and statistical discrimination are central to [Anzia and Bernhard \(2022\)](#) for the case of U.S. elections, even if evidence on valence of female candidates points towards higher levels of effort and quality in that context ([Anzia and Berry, 2011](#); [Ashworth et al., 2024](#)).

Conversely, [Fréchette et al. \(2008\)](#) argues that male bias explains increased incumbency advantages for male candidates in French legislative elections following the implementation of gender quotas, especially under single-member majority systems rather than closed-list proportional representation. [Le Barbanchon and Sauvagnat \(2022\)](#) further illustrates voter bias through survey data on gender roles in politics, demonstrating that female candidates receive fewer votes in French municipalities characterized by larger local gender earnings gaps within the same electoral district and election year. [Beaman et al. \(2009\)](#), using random assignment of gender quotas for village council leadership positions in India, show that prior exposure to female leadership increases women’s electoral success and reduces stereotypes about women’s effectiveness in both public and domestic roles.

Our study advances this literature in two significant ways. First, while previous research primarily relies on changes to gender quotas or electoral systems, we introduce, to our knowledge, the first field experiment specifically designed to quantify the effect of voter bias on female vote shares. Second, our experimental design allows us not only to detect voter bias, but also to decouple the underlying sources of discrimination.

By distinguishing between taste-based and statistical discrimination, our research intersects with an extensive Labor Economics literature.⁵ We offer a methodological contribution to a body of work that predominantly uses audit or correspondence studies to differentiate these discrimination theories. Typically, these studies assume that provided applications include all relevant information about a candidate's productivity and observed outcome differences between minority and non-minority candidates indicate taste-based discrimination. However, this method requires strong assumptions: either all relevant productivity factors must be observed, or unobserved factors must not differ systematically between groups in both mean and variance, unless linearity is imposed (Heckman and Siegelman, 1993; Heckman, 1998). Related to the approach proposed by Neumark (2012), our study demonstrates how experimental variation in the salience and beliefs about candidates' attributes can allow one to decouple taste-based discrimination from statistical discrimination. Most correspondence studies lack such experimental variation concerning productive attributes independent of group identity.

Our research also relates to a substantial Political Economy literature on the effects of informational campaigns on voter behavior. This literature predominantly examines information provision in a reduced-form manner, highlighting its efficacy in increasing voter turnout (Gerber and Green, 2000), influencing how citizen's vote (DellaVigna and Gentzkow, 2010; Aker et al., 2011), and even affecting vote-buying behaviors (Vicente and Wantchekon, 2009; Fujiwara and Wantchekon, 2013; Vicente, 2014). As our experiments highlight, to understand how information impacts voter behavior requires structural estimation models combined with messaging that distinctly targets voter beliefs and preferences. In this respect, our approach is related to Kendall et al. (2015); Cruz et al. (2024), who demonstrate that campaign messages influence voter beliefs and electoral choices, finding valence-based appeals particularly effective. While these papers estimate belief updating and preference shifts based on ideology and valence, our framework extends to the study of discrimination. Additionally, our study employs a large-scale social media campaign, demonstrating the effectiveness of these micro-targeted online interventions in electoral settings.

⁵For comprehensive reviews of the discrimination literature, see Guryan and Charles (2013); Lang and Lehmann (2012); Bertrand and Duflo (2017).

Finally, our paper contributes to a growing Political Economy literature focused on structural models of voter demand. This literature covers elections in both high income (Berry et al., 2024; Cox, 2024; Longuet-Marx, 2025) and low income countries (Ujhelyi et al., 2021; Iaryczower et al., 2022; Montero, 2025) –the latter context being closer to ours. This literature typically does not integrate structural modeling and randomized experimental variation, nor does it make explicit the sources of discrimination in elections.⁶

The paper proceeds as follows. The next section presents the econometric model and our methodological approach. Section 3 discusses the institutional context, details of the electoral experiment, and the data. Section 4 presents the reduced-form results. Section 5 includes our main empirical results, model fit, robustness, and validation. Section 6 presents a series of electoral counterfactuals, relevant to future applications of this methodology and to current research on descriptive electoral representation. Section 7 concludes.

2 Empirical Model

2.1 Setup

A voter living in municipality $m = 1, \dots, M$ is choosing among a set of political candidates $j = 1, \dots, J$ for election to the state legislature by open list proportional representation. The voter is indicated as $i \in \{1, 2, \dots, N_m\}$, with binary sex at birth $G_i \in \{0, 1\}$.⁷ Each candidate j is represented by a set of features $k = 1, \dots, K$. In this analysis, the features are $K = 3$ and comprise gender of the candidate, G_j , ability in performing administrative tasks, A_j , and policy position on a uni-dimensional progressive-conservative scale, P_j . For the purposes of the exposition, we can focus on the choice between two candidates, a female and a male.

Voters enjoy utility from supporting candidates with certain features, for example candidates who share their own gender identity or who are closer to their policy views. Voters may not know features of candidates with certainty, or at least not for all candidates, and have their own beliefs over particular features. Specifically, while we assume voters know gender identity with certainty, ability and policy position are dimensions of j over which there is a degree of uncertainty, i.e. voter i has subjective and heterogeneous beliefs over features A_j and P_j . We define voter expectations as $\mathbb{E}_i[A_j] = A_{ij}$ and $\mathbb{E}_i[P_j] = P_{ij}$, where the subscript in the expectation operator indicates that expectations are taken over i 's subjective belief distribution about candidate

⁶For a related structural approach that emphasizes heterogeneity in voter response, see also Magesan et al. 2024.

⁷For the remainder of the paper, we use sex and gender as synonyms.

j 's features. This allows one to capture situations such as when voter i may support male candidate j over a female candidate j' because i expects female candidates to have on average lower ability than male ones ($A_{ij} > A_{ij'}$). In practice, i 's expectations about j allow to formally incorporate *statistical discrimination* into our problem, as implied by discrimination theory (Phelps, 1972; Arrow, 1973, 1998).⁸

Voter preferences are expressive and additively separable across the K dimensions. Given voter i 's expectations $\mathbb{E}_i[\cdot]$, gender G_i , and ideal policy position P_i , i 's utility from voting for politician j is assumed to take the form:

$$u_{i,j} = -w_i^G \times |G_j - G_i| + w_i^A \times A_{ij} - w_i^P \times |P_{ij} - P_i| + \varepsilon_{ij}, \quad (1)$$

where w_i^k are preference weights placed on the gender identity, ability, and policy dimensions (we will also use the term salience or salience weights) and are non-negative, $w_i^k \geq 0$.

In equation (1) preferences are spatial in policy and in gender identity, that is everyone likes candidates that are closer to themselves more than candidates that are farther away. Operator $|\cdot|$ represents a generic distance function determining losses (e.g., quadratic losses associated to voting for candidates far from the voter in terms of their policy position imply $|P_{ij} - P_i| = (P_{ij} - P_i)^2$). Preferences are vertical along the ability dimension, as everyone likes higher ability/competence in their elected officials, unconditionally of policy position or gender. In this sense, equation (1) accommodates a mix of private value horizontal dimensions (gender, policy) and common value vertical dimensions (ability). Weights w_i^k represent pure taste parameters and allow us to incorporate *taste-based discrimination* elements into our problem. For instance, a male voter i may consider female candidate j' less desirable than male candidate j because $w_i^G > 0$ and i does not share the same identity as candidate j' ($|G_j - G_i| = 1$). Finally, (1) includes a random utility component. ε_{ij} is an i.i.d. (across all i and j), idiosyncratic preference shock. It characterizes a stochastic element of voter utility realized at the time they cast their ballot.

2.2 Mechanisms and Estimation

Formalization via equation (1) makes the central mechanisms of the political behavior model explicit. We focus on two in particular. Voters are allowed to respond to electoral signals through: (i) Beliefs used in voter expectations, $\mathbb{E}_i[A_j] = A_{ij}$ and $\mathbb{E}_i[P_j] = P_{ij}$; (ii) Preference weights of the gender identity, ability, and policy dimensions, w_i^k , for $k = G, A, P$.

⁸For a discussion, see also Fang and Moro (2011).

With respect to mechanism (i), we assume that voter expectations can only be affected by information and that no belief updating takes place absent informative signals. As absence of information can be in itself an informative signal and induce updating, this can be considered a strong assumption in a context where voters and politicians are perfectly rational/Bayesian. We believe it is a less stringent assumption in our specific application. In Brazilian local elections, hard information is typically limited, and its absence may be relatively unsurprising to voters, hardly inducing any learning. We indicate informative signals about ability and policy in municipality m as $T_{i,m}^A$ and $T_{i,m}^P$, respectively. Voter expectations of candidate ability are defined as $A_{ij} = \xi^A G_j + \rho_A T_{i,m}^A G_j + \eta^A X_m$, and function of $T_{i,m}^A$, the candidate's gender (G_j), as well as vector of candidate and municipality voter characteristics $X_m = [X_{m,Cand}, X_{m,Voters}]$. Similarly, we parameterize voter policy preferences as $|P_{ij} - P_i| = (P_{ij} - P_i)^2$ with $P_{ij} = \xi^P G_j + \rho_P T_{i,m}^P G_j + X_{m,Cand}^\top \eta_{Cand}^P$, and voter ideal point $P_i = X_{m,Voter}^\top \cdot \eta_{Voter}^P + \mu \cdot G_i$, with $\eta^P = [\eta_{Cand}^P, -\eta_{Voter}^P]$. The set of control variables included in the estimation encompasses the relative wealth of female candidates compared to male candidates, the proportion of female candidates with college degrees relative to male candidates, marital status ratios among candidates by gender, GDP per capita, and average age and education levels of voters, among other characteristics.

With respect to mechanism (ii), w_i^k plays the role of salience weights along dimension k . These weights are designed to parsimoniously capture psychological components of choice beyond rational learning. Examples are shifts in awareness (or neglect) of issues, or changes in voter attention occurring during the campaign. As attention shifts may happen for a number of reasons, and w_i^k may be sensitive to multiple types of stimuli, we allow salience weights to respond to all signals, including uninformative ones. We indicate uninformative signals about ability and policy as $V_{i,m}^A$ and $V_{i,m}^P$. Signals about gender are indicated as $V_{i,m}^G$. We assume that weights $w_i^k = \exp(\omega^k + \sigma_k v_{i,m}^k + \lambda^k \max\{T_{i,m}^k, V_{i,m}^k\})$ are influenced by a dimension-specific intercept ω^k , unobserved heterogeneity preference shocks $v_{i,m}^k$ and signals $T_{i,m}^k$ or $V_{i,m}^k$. The exponential function ensures weights are positive and we experiment with other transformations as well in our empirical application.

With these assumptions in place, the parameterization of (1) is complete. i 's utility from

normally distributed with mean 0 and variance σ_k . Further, we fix the scale of each dimension by calibrating σ_k to 1. This restricts the dispersion of weights associated with each dimension, while allowing the scale of the utility to remain free and be estimated from the data. The resulting parameter estimates are relative to this normalized value.

Variation in the assignment of the signals $V_{i,m}^G$ and $\{T_{i,m}^k, V_{i,m}^k\}$, $k = A, P$, (discussed in the following section) allows to identify all parameters of the voter weight functions, particularly ω^k . Along the vertical dimension, voter expectations vary monotonically with the covariates, whereas expectations along the horizontal (spatial) dimension are inherently non-monotonic. This theory-driven distinction is central to identification of $\{\xi^k, \eta^k\}$, $k = A, P$. The linearity of the expected ability function A_{ij} in candidate and voter characteristics $X_m = [X_{m,Cand}, X_{m,Voters}]$ guarantees identification of the parameters $\{\xi^A, \eta^A\}$. The non-monotonic nature of the policy loss function, in combination with the interaction of candidate and voter characteristics, allows to identify the remaining parameters $\{\xi^P, \eta_{Cand}^P, \eta_{Voter}^P\}$.

Based on these assumptions, we model vote choice via a discrete choice, random utility framework. Let $v_{i,m}^s$ denote if i voted for female candidate or not. We assume that the decision to vote for a female candidate is determined by:

$$v_{i,m} = \mathbf{1}\{u_{i,1,m} > u_{i,0,m}\}. \quad (3)$$

The model-based total votes for female candidates in municipality m are then given by $v_m = \sum_{i=1}^{N_m} v_{i,m}$ and can be matched to empirical moments measured at the municipal level. The estimator that we employ is Generalized Method of Moments (GMM) and Appendix A provides the details of its construction. Appendix B reports Monte Carlo simulation exercises focused on probing the robustness of this estimator to measurement error and spillovers.

3 Design of the Experiment and Data

During Brazil's 2024 municipal elections, we collaborated with a non-governmental organization dedicated to reducing gender discrimination in politics to conduct a large-scale randomized experiment. The experiment involved 1,000 municipalities, with randomization implemented at the municipal level, and focused on the three dimensions central to our model: gender identity, ability, and policy. Messages were disseminated exclusively through Instagram and were tailored to emphasize these dimensions, by shifting salience and altering beliefs. This section details the institutional context, the experimental setting, and the data.

3.1 Institutional Context

In Brazil, municipal elections are held every four years to elect mayors, vice mayors, and city councilors in the country's 5,570 municipalities (*municípios*). The elections took place in October 2024, with voting mandatory for literate citizens aged 18 to 70 and optional for those aged 16 to 17, over 70, or illiterate. As a result of mandatory voting, turnout is high, typically above 80%. The electoral process is overseen by the Superior Electoral Court (TSE), which ensures transparency and the efficient use of electronic voting machines, a hallmark of Brazil's electoral system (Fujiwara, 2015). These machines enable quick vote counting and minimize the risk of fraud, enhancing trust in election outcomes.

Municipalities in Brazil are responsible for essential services, including education, health-care, urban planning, and infrastructure. Mayors wield significant executive power, managing budgets and local public service delivery, while city councilors serve as the legislative body, enacting local bylaws and overseeing the administration. These roles make municipal elections a vital mechanism for determining policies that directly impact citizens' daily lives.

Candidates for city council are elected through a proportional representation system based on the “*quociente eleitoral*” (electoral quotient) and the “*quociente partidário*” (party quotient). The electoral quotient is calculated by dividing the total number of valid votes in a municipality by the number of available council seats. Each party must meet or exceed this threshold to secure seats. Once a party achieves the electoral quotient, its share of the council seats is determined by the total number of votes received by its candidates. This system encourages smaller parties to form coalitions, but can also result in over-representation of larger political alliances. The structure ensures proportionality, but often limits individual candidates from smaller or independent parties from gaining traction.

Gender representation remains a significant challenge in Brazil's political landscape, particularly within city councils. Although women constitute 51% of the Brazilian population, their representation in municipal legislatures remains disproportionately low. In the 2024 municipal elections, only 18.2% of elected city councilors were women, totaling 10,654 female councilors out of 58,446 elected officials. City council positions often serve as the main entry point for women into politics in Brazil: 89% of women who pursue a political career begin as city councilors (Rede A Ponte, 2024). This first step is crucial, as it provides legislative experience, political visibility, and networking opportunities that are fundamental for advancing to higher political offices at the state and federal levels. Despite legislation mandating that 30% of party resources be allocated to support female candidates, enforcement has been inconsistent, and

structural barriers persist.

Municipal elections provide an opportunity to evaluate how targeted interventions, particularly those designed to address both statistical and taste-based discrimination, can influence voter beliefs about candidates' ability and policy positions, while also shifting the salience of gender in electoral decision-making. In the context of this study, the decentralized structure and diverse political dynamics of Brazil's municipalities create a rich setting to explore how micro-targeted campaign messages impact perceptions of candidates along dimensions such as gender, ability, and policy.

3.2 Experimental Design

To assess the impact of targeted messages on voter behavior, we conducted a randomized campaign 8 days before Brazil's 2024 municipal elections. This experiment, carried out in partnership with a non-governmental organization, covered 1,000 municipalities with randomization at the municipal level. The campaign utilized Instagram to deliver tailored messages focusing on gender identity, ability, and policy.

3.2.1 Digital Campaign Implementation via Instagram

The campaign messages were delivered exclusively via Instagram, a widely used platform in Brazil, making it an effective tool for targeting voters. As of March 2024, Instagram had approximately 141 million users in Brazil, representing 63.9% of the population (NapoleonCat). Its user base is diverse, with significant engagement across various age groups. Brazilian politicians have increasingly turned to Instagram to connect with constituents, leveraging its broad reach and engagement potential. Social media and messaging apps are among the most popular channels for political campaigns in Brazil.

Approximately 40% of city council candidates registered an Instagram account with Brazil's Superior Electoral Court (TSE). However, this figure likely underestimates actual usage, as candidates are not required to report their social media presence to the TSE, and many may have active Instagram accounts without officially registering them. This widespread adoption by both voters and politicians underscores Instagram's pivotal role in political messaging, making it an ideal platform for our experimental campaign.

Instagram's advertising algorithm enabled us to target voters based on their municipality, gender, and age group. The campaign leveraged the platform's geolocation tools to ensure that advertisements were displayed only to voters within the assigned treatment municipalities. Ad-

ditionally, we used Instagram's engagement metrics to track impressions (total views) and reach (unique users who saw the ad), ensuring effective message dissemination.

Table 1 presents the total number of impressions and unique views, broken down by gender. Overall, our campaign generated 9.9 million impressions and 5.2 million unique views. Given that the total population in our sample is 6.9 million, we were able to reach approximately 76% of the population –or nearly 90.2% of registered voters in our sample. Further details on the technical implementation of the digital campaign, including budget allocation, contingency measures, and performance evaluation, can be found in Appendix C.

3.2.2 Randomization and Treatment Assignment

To ensure a well-balanced sample and a robust first stage of our treatment, we implemented a three-step selection process. First, we focused on medium-sized municipalities with populations between 5,000 and 30,000. Larger municipalities, typically located near metropolitan areas, are more prone to substantial spillover effects due to regular commuting patterns. Additionally, achieving an adequate penetration rate in these areas could be challenging. Conversely, smaller municipalities may experience a high proportion of individuals traveling in and out, leading to significant noise in our penetration estimates. For instance, even a fluctuation of 100 individuals in a municipality of 2,000 voters could create considerable discrepancies between the actual percentage of voters treated and the measured penetration rate.

Second, within this set of municipalities, we calculated the minimum Instagram penetration rates and selected the top 1,000 municipalities with the highest minimum penetration. To estimate minimum penetration rates, we conducted an independent data collection effort simulating a campaign before the experiment. This allowed us to gauge the reach of Instagram ads at the municipal level. Finally, we assigned treatment arms to the selected municipalities, while ensuring balance across key observable characteristics, including gender composition, education levels, racial composition, internet availability, past voting patterns for female candidates, GDP per capita, and age distribution. To verify that our sample is representative of the country, we compared our experimental sample to the full set of Brazilian municipalities main characteristics (see Table 2). A comparison of summary statistics confirms that the selected municipalities are broadly similar in terms of demographic composition, economic development, and internet access.

While Instagram's geolocation tools allow us to target voters within specific municipalities, potential spillover effects remain a challenge. These spillovers can occur through two main channels: social sharing of ads and commuting patterns. The first type of spillover arises when

Instagram users share our ads with others via Instagram messages. This can result in our messages reaching voters outside the intended municipalities, leading to unintended exposure to treatment. To monitor this, we analyzed engagement data (e.g., number of times our post were shared) and found no systematic evidence of significant spillover effects between treatment and control municipalities. The second type of spillover occurs when voters commute between municipalities for work or other purposes, potentially exposing them to different treatment statuses. This can lead to two types of confounding forces. First, a voter assigned to one treatment status may travel to another municipality where a different treatment status applies, resulting in exposure to two messages instead of one. Second, voters from municipalities outside our sample may travel to those within our sample, inflating the estimated penetration rate of our messages. Since these voters ultimately vote in their home municipality, which is outside our sample, this could lead to overestimating actual reach. However, given the small and predominantly rural nature of the municipalities in our sample –averaging approximately 8,000 residents and located far from major metropolitan centers– regular intermunicipal commuting for work is limited. [Instituto de Pesquisa Econômica Aplicada \(2016\)](#) used national commuting data to show that only around 12.8% of employed Brazilians commute across municipalities, with much lower rates in rural and less connected areas. Therefore, the likelihood of voters from outside our sample regularly entering treated municipalities and inflating our estimated reach appears minimal.

Before randomization, municipalities were stratified based on demographic factors (population size, literacy rate, schooling levels, urbanization rate, proportion of racial groups, age distribution), economic indicators (per capita GDP) political characteristics (past electoral performance of female candidates), and digital infrastructure (internet coverage, minimum Instagram penetration rate, Meta’s small-municipality classification). They were then randomly allocated to one of seven experimental arms: i) a gender treatment targeted at male voters, ii) a gender treatment targeted at female voters, iii) an informative ability treatment, iv) a non-informative ability treatment, v) an informative policy treatment, vi) a non-informative policy treatment, and vii) a control group that did not receive any campaign messages. This randomization approach ensured that the assignment of municipalities across experimental conditions was balanced, allowing for identification of the effects of each message type on voter preferences and beliefs.

Appendix C contains all messages. Gender, uninformative ability and uninformative policy messages raise awareness of the electoral importance of candidate’s gender, ability, and policy, respectively. Informative ability messages add verifiable information on the average educational attainment and qualifications of female candidates and signal the high quality of fe-

male candidates. Informative policy messages add verifiable information about average social spending by female candidates and signal a progressive stance of female candidates.

In Table 3, we present summary statistics for municipalities in our sample by treatment assignment. For each characteristic, we also report the mean difference between each treatment arm and the control group. The average population size in our sample is approximately 8,000 inhabitants, with women representing slightly less than half of the population. Despite this near-equal demographic representation, women remain vastly underrepresented in politics. This is evident both in the share of votes female candidates received in the 2022 state and congressional elections and in the share of female candidates running for local council in the 2024 elections. Women comprised only 35% of local council candidates, barely exceeding the 30% legal requirement. Only about 21% of candidates have a college degree, but this proportion is significantly higher among female candidates (30%) compared to male candidates (16%).

When comparing all municipal characteristics across treatment and control groups, we observe minimal differences. Out of 126 comparisons, only 6 are statistically different from the control at the 10% significance level.⁹ Furthermore, a multinomial logit test fails to reject the null hypothesis that all variables are jointly significant ($\text{Prob} > \chi^2 = 0.9383$; $\text{Pseudo } R^2 = 0.0293$). Overall, the results in Table 3 indicate that the randomization was effective in achieving balance on covariates.

3.3 Data

Our analysis relies on multiple official data sources to construct a comprehensive dataset on electoral outcomes, demographic characteristics, and internet accessibility at the municipal level. Electoral data come from the TSE, which maintains official records publicly accessible online for all elections since 1994. This dataset includes detailed information on the total number of registered voters, votes for each candidate, and candidates' characteristics such as gender, race, education level and declared wealth. We construct key indicators such as the vote share for female candidates, turnout rate, the proportion of municipal council seats won by women, partisan distribution of elected officials, median campaign contributions received, and historical electoral data.

To account for demographic and economic differences across municipalities, we use data from the Brazilian Institute of Geography and Statistics (IBGE), which provides information on population size, age distribution, literacy rates, racial composition, per capita GDP, schooling

⁹When considering all 483 pairwise comparisons, only 35 are statistically significant at the 10% level.

levels, and the degree of urbanization. These variables are critical for assessing the representativeness of our experimental sample relative to the broader Brazilian population and serve as key controls in our empirical analysis. Differences in municipal characteristics, particularly in economic development and demographic composition, may shape how voters process political information, making it essential to include these factors in our modeling. Age distribution is captured through three categories, distinguishing younger, middle-aged, and older voters. Literacy rates serve as proxy for education, which may influence information processing and belief updating. Racial composition is incorporated by measuring the share of *pardos*, black, and indigenous residents in each municipality, given their relevance for analyzing disparities in political representation. Economic indicators such as per capita GDP and urbanization rates further allow us to control for variations in municipal wealth and the degree of rural-urban differences in electoral behavior.

Finally, since the campaign was conducted exclusively on Instagram, internet access plays a central role in determining potential voter exposure to our intervention. To account for differences in digital accessibility across municipalities, we incorporate data from the National Telecommunications Agency (Anatel), which provides indicators on internet coverage. The dataset includes measures of the population covered by broadband or mobile internet services and the percentage of households with internet access. These variables ensure that we can accurately assess the reach of our campaign, as exposure to digital ads depends on both individual internet access and broader infrastructural coverage. Municipalities with extremely low internet penetration were excluded from our sample to avoid biases arising from non-exposure.

4 Reduced-Form Estimation

In this section, we present results based on the RCT design outlined in Section 3.2. The reduced-form analysis performed in this section does not rely on the structure of Section 2, but it serves the purpose of highlighting elements of our empirical model that cannot be recovered from pure reduced-form analysis. To estimate the effects of the experiment, we analyze variants of the following linear regression model:

$$v_m = \beta_0 + \beta_1 V_m^{G,1} + \beta_2 V_m^{G,0} + \beta_3 V_m^A + \beta_4 T_m^A + \beta_5 V_m^P + \beta_6 T_m^P + X_m' \gamma + \delta_{s(m)} + \epsilon_m. \quad (4)$$

Our main dependent variable, v_m , measures the share of votes received by female candidates for local council. The treatment indicators are defined as follows: $V_m^{G,1} = 1$ if females in the municipality received the gender message, $V_m^{G,0} = 1$ if males in the municipality received the

gender message, $V_m^A = 1$ if the municipality received the uninformative ability message, $T_m^A = 1$ if the municipality received the informative ability message, $V_m^P = 1$ if the municipality received the uninformative policy message, and $T_m^P = 1$ if the municipality received the informative policy message. Since treatment assignment was stratified, our regression always includes strata fixed effects (δ_s). In some specifications, we also include a vector X_m of municipal controls. The reduced-form error term ϵ_m is robust to heteroskedasticity.

Table 4 presents the results of our experiment on female candidates' vote share. Column (1) reports results without controls, but includes strata fixed effects. In column (2), we add the share of votes female candidates received in the most recent past election (i.e., 2022). In column (3), we further include the baseline controls that were statistically significant across treatment arms in the balance tables. Finally, in column (4), we incorporate region fixed effects. Unsurprisingly given our experimental design, the results remain consistent across these different specifications.

Focusing on column (1), we find that gender messages increased the vote share for female candidates, primarily when targeted at male voters. A coefficient of 1.8 (robust standard error = 0.9) represents a substantial increase of 7.6% relative to the control group mean of 23.0%. Moreover, this estimate likely represents an intention-to-treat effect. Given that our ads reached only 90% of voters, assuming no spillovers, full exposure (100% of voters) would have led to an estimated 8.6% increase in female candidates' vote shares. In terms of our model, this result suggests that reminding men about the lack of female political representation reduced their disutility from voting for female candidates. Among women, the messaging slightly increased disutility for voting for male candidates, though with less statistical confidence.

The results for ability messaging are less clear-cut. The coefficient on the uninformative ability message, while sizable in magnitude, is statistically insignificant. However, as our model highlights, the effect of λ_A –the parameter this treatment targets– depends on individuals' subjective prior beliefs about female candidates' ability relative to male candidates. For individuals who believe female candidates are less capable, an uninformative message may reduce the likelihood of voting for a female candidate by increasing the salience of a negative attribute, while the opposite holds for those who believe female candidates are more capable. These opposing effects could explain the lack of a clear overall treatment effect in the aggregate. When we structurally estimate our empirical model, we will be able to account for these differences in voter's beliefs.

For the informative ability message, we also find an insignificant effect (coefficient = 0.3, robust standard error = 0.9). Note that this coefficient captures both the salience effect and any

learning about the expected ability of female candidates induced by the message. Removing salience effects using our estimated impact from the uninformative ability message, we estimate a small learning effect (coefficient = -0.7, p-value = 0.40), suggesting that voters either did not update their beliefs in response to the message or that any updating was too minor to affect aggregate vote shares meaningfully. A more subtle reason for the insignificant result may also stem from the unobserved heterogeneity in the salience weights of the three preference dimensions. If voters assign greater weight to gender than to ability, then, even if our informative treatment alters beliefs about female candidates' ability, this may not manifest in vote choices, as voters' decisions are primarily driven by gender identity. Therefore, the configuration of salience weights plays a crucial role in determining whether the informative treatment will lead to significant changes in vote choices. Our empirical model accounts for these effects by recovering the salience weight distribution.

The effects of the uninformative policy message are small and statistically insignificant. However, similarly to the uninformative ability message, this result may reflect an underlying heterogeneity in beliefs. As our model emphasizes, the impact of λ_p on vote choice depends on the beliefs about candidates' policy position relative to the voter's ideal point. Consequently, the uninformative policy messages may have different effects depending on whether voters perceive female candidates as more liberal or conservative than male candidates and on where voters' own policy preferences are located.

Informative policy messages focus on social spending of female candidates and generally point to a more progressive stance of female candidates (in addition to increasing the salience of policy in vote choice). When we remove salience effects, we find that the informative policy message had a substantial impact on female candidates' vote shares. With a coefficient of 1.7 (p-value = 0.08), the effect size is comparable to that of gender messages targeted at men. Within the context of our model, this finding suggests that the message shifted voters' beliefs about female candidates' policy positions closer to their ideal points and therefore boosted female vote share.

Overall, our reduced-form evidence indicates that a subset of micro-targeted ads increased the vote share for female candidates within our sample of municipalities. Crucially, however, our findings highlight the necessity of incorporating the heterogeneity of voters' beliefs about candidates' abilities and policy positions to accurately assess the overall impact of these messages.

4.1 Other electoral outcomes

For completeness, Table 5 examines the effects of our campaign messages on other electoral outcomes. In column (1) - (3), we investigate the effects on voter turnout and campaign expenditures. Although, in principle, our treatments could have influenced these outcomes, *ex ante* such effects were highly unlikely. In Brazil, voting is mandatory, and abstention typically occurs when individuals no longer reside in the municipality where they are registered. Additionally, our campaign began only eight days before the election, making it difficult for candidates to increase their spending. Nevertheless, we estimate the impact of our treatments on these outcomes, disaggregating campaign expenditures by candidate gender. As shown in the table, we find no evidence that our treatments affected these different electoral outcomes.

In column (4), we examine how our campaign messages affect the share of women elected to municipal councils. As discussed in Section 3.1, local councilors are elected using an open-list proportional representation system based on the D'Hondt method. Consequently, the relationship between vote share and election outcomes may not be strictly linear. Nevertheless, we find evidence that our campaign messages increased the share of women elected to the councils, especially in the case of the uninformative ability messages and informative policy messages. For instance, the uninformative policy message increased the share of elected women by 0.028 (s.e. = 0.014), representing an approximately 23.3% increase relative to the control group mean of 12 percentage points. Moreover, aside from the informative ability message, the treatment coefficients are generally large in magnitude, although some are estimated with less statistical precision. When we pool all treatment groups and estimate a single regression of the share of elected females on receiving any campaign message, the estimated treatment coefficient is 0.019 (s.e. = 0.012), significant at the 10 percent level. A treatment effect of this magnitude implies a 15.6% increase in the share of women elected to local councils in 2024.

5 Structural Estimation

In this section, we begin by presenting the empirical model's estimates. We then assess its quantitative implications and out-of-sample predictive performance and validate our approach. In the next section, we utilize the estimated model to evaluate several counterfactual scenarios, including message campaigns optimized specifically for each municipality with the goal of increasing female vote shares.

5.1 Parameter Estimates

Panel A of Table 6 reports GMM estimates for the main parameters of our model.¹⁰ To illustrate the magnitudes of these estimates in explaining voting decisions, we also present their average marginal effects in Panel B of Table 6. The average marginal effects are computed by individually setting each parameter to zero, simulating the resulting vote choices, and comparing these simulated choices with the baseline scenario (with parameters at their estimated values). The average differences across simulations yield the corresponding marginal effects, expressed in percentage points (p.p.).¹¹

In Table 6, Column (1) represents our preferred specification. It includes region fixed effects for ability and policy weights, it accounts for candidate- and municipality-level observable characteristics, and it employs exponential weights. The exponential function is utilized to constrain the weights to be non-negative. Column (2) differs by excluding region fixed effects. Column (3) differs by excluding candidate-level controls. Columns (4) and (5) use quadratic and absolute functions respectively, instead of the exponential function, for defining salience weights.

Several notable findings emerge from this analysis.

Baseline Preference Parameters. Gender identity parameters are precisely estimated and underscore gender’s significant role in shaping voter behavior. The estimated parameter reflecting female voters’ baseline distaste for voting against their gender is $\omega_{GF} = 4.088$ (s.e. = 0.010). For male voters, the corresponding parameter is $\omega_{GM} = 6.233$ (s.e. = 0.012). This implies that, holding all other candidate’s characteristics (ability, policy platforms) fixed, female and male voters both exhibit lower willingness to vote for opposite-gender candidates compared to candidates of their own gender. Estimates are both positive and statistically different from zero at standard significance levels, so that we can strongly reject the null hypothesis that no weight is given on the gender of the candidate relative to the gender of voter. Overall, the estimated ω_{GF} and ω_{GM} represent prima facie evidence of taste-based discrimination in our context.

Economic magnitudes of taste-based discrimination based on gender appear substantial for both groups. Setting the baseline female distaste parameter to zero ($\omega_{GF} = 0$) would decrease female candidates’ vote share by 19.5 p.p., a substantial drop, as reported in Panel B of Table 6.

¹⁰See Appendix Tables D1 and D2 for the remaining parameter estimates of the control variables and region fixed effects.

¹¹Formally, the average marginal effect of parameter θ_k is evaluated by computing $\frac{1}{M} \sum_m (s(\hat{\theta}; X_m) - s(\tilde{\theta}; X_m))$ where $\tilde{\theta}$ is defined as for all $l \neq k$ we have $\tilde{\theta}_l = \hat{\theta}_l$ and $\tilde{\theta}_k = 0$.

Conversely, setting male voters' ω_{GM} to zero would increase vote shares for female candidates by approximately 8 p.p. (s.e. = 4.130).¹²

The baseline salience parameter estimates for candidate ability ($\omega_{A0} = 2.15$, s.e. = 0.519) and policy ($\omega_{P0} = 1.343$, s.e. = 0.564) are both positive and statistically significant in Panel A of Table 6. Although the baseline ability parameter is numerically larger, it plays a less critical role than policy salience in determining vote choices for female candidates. Specifically, setting either the ability or policy salience to zero effectively removes these dimensions from influencing vote choices. As reported in panel B doing so increases female candidates' vote shares modestly by 1.19 p.p. (s.e. = 0.611) for ability salience, but substantially by 13.69 p.p. (s.e. = 6.785) for policy salience. The fact that removing either dimension boosts vote shares for female candidates highlights that female underrepresentation in politics arises not only from gender identity, but also from voters' subjective beliefs about relative candidate ability and policy positions. This is, indeed, a key marker of statistical discrimination in this context.

Baseline Belief Parameters. The baseline parameter estimate for voter beliefs about female candidates' relative ability, $\xi_A = 0.120$ (s.e. = 0.059), suggests voters, on average, perceive female candidates as having higher ability compared to male candidates. This estimate, however, is sensitive to the inclusion of candidate-level controls and the functional form used to model salience across columns (1)-(5).¹³

The baseline parameter for voter beliefs about female candidates' policy positions relative to male candidates is $\xi_P = 2.157$ (s.e. = 0.624), and the relative policy preference (bliss point) for female voters compared to male voters is $\mu_P = -3.193$ (s.e. = 0.876). The opposite signs of these parameters indicate a mismatch between beliefs about female candidates' policy positions and female voters' policy preferences. Specifically, female candidates are perceived as more conservative than female voters themselves. This is illustrated graphically in Figure 1. The figure

¹²The mismatch between the estimated coefficients (ω_{GF} , ω_{GM}) and their associated average marginal effects arises because the mapping from parameters to vote shares is nonlinear. Specifically, even if ω_{GF} is smaller in magnitude than ω_{GM} , the impact of changing ω_{GF} on vote shares may be larger depending on where the average voter's utility lies on the logistic curve. This curve is S-shaped, so the marginal effect of a parameter depends not just on its size but also on whether it falls near the steep or flat part of the curve. For example, a moderate change in ω_{GF} near the steep part can have a large effect on vote shares, while even a large change in ω_{GM} near the flat part might yield little impact. This suggests that the average utility for female voters lies near the steeper part of the logistic curve, while that for male voters is closer to the flatter region. As a result, the marginal effect of ω_{GF} can exceed that of ω_{GM} , despite its smaller magnitude. This is evident by the estimated gradients of vote share with respect to each parameter: 0.13 for ω_{GF} and -0.0052 for ω_{GM} , indicating much greater local sensitivity to the former.

¹³Omitting candidate controls obscures baseline beliefs regarding female candidates' ability relative to males. Additionally, quadratic and absolute functional forms are sub-optimal choices for modeling salience weights because their symmetry around zero interferes with accurate parameter estimation.

also illustrates the heterogeneity across different types of municipalities in terms of the policy positions of voters and candidates of different gender. This substantial mismatch helps explain a significant gap in vote shares for female candidates. Setting the perceived policy position parameter (ξ_p) to zero would increase female candidates' vote shares by 20.13 p.p. (s.e. = 0.968). Similarly, setting female voters' relative policy position parameter (μ_p) to zero would increase female candidates' vote shares by 13.29 p.p. (s.e. = 0.655).

Treatment Parameters. We now focus on how our experimental treatments affected both taste-based and statistical discrimination in this election. We find that female voters experience no significant change in their distaste for voting against their gender ($\lambda_{GF} = -0.018$, s.e. = 0.034) after receiving our gender message. At baseline, female voters already show a strong tendency to vote for female candidates (recall that removing this distaste entirely would decrease female candidate vote shares by 19.5 p.p.). Given this strong baseline preference, our intervention does not further enhance female voters' inclination toward female candidates. In contrast, male voters show a significant reduction in their distaste for voting against their gender following our intervention ($\lambda_{GM} = -0.533$, s.e. = 0.015). This reduction makes male voters more receptive to female candidates, resulting in an increased female candidate vote share of 0.356 p.p. (s.e. = 0.097), which corresponds to a 1.5% increase relative to the average municipality female vote share. In synthesis, our estimates show that the gender treatment reduced taste-based discrimination in male voters.

Both ability message did not produce strong effects.

Our uninformative ability message appeared to reduce the salience of candidate ability ($\lambda_A = -0.162$, s.e. = 0.022), yet translated into a marginal effect that is positive and insignificant (0.049, s.e. 0.05). While apparently counterintuitive, this result is due to λ_A 's estimate being a precise zero in driving vote choice. In Appendix Table D3 we perform an analysis of economic magnitudes for all treatments looking at the contribution of the parameter in terms of changes in total utility. For the case of λ_A , the change is only -0.09 p.p., indicating that total voter utility virtually does not change.¹⁴

In a similar vein, the informative ability message did not alter voters' beliefs about female candidates' ability ($\rho_A = -0.092$, s.e. = 0.064). As the informative ability message was designed to reduce this specific dimension of statistical discrimination, here we must report a null result.

¹⁴In addition, the underlying heterogeneity in ability beliefs is substantial in the data, as evident in Appendix Figure D1. The figure shows an even distribution between municipalities where voters believe female candidates are on average less able than male candidates and ones where voters believe the opposite. It is intuitive, then, that the average effect of any change in the salience of the ability dimension, absent a change in beliefs, may be ambiguous in terms of the support for female candidates.

Our uninformative policy message increased the salience of the policy dimension ($\lambda_p = 0.083$, s.e. = 0.012) and with credible economic magnitudes per Appendix Table D3. The effect of this message on voters varies depending on the alignment of their policy bliss points relative to male and female candidate policy positions. Specifically, voters whose bliss points are closer to male candidates' positions than female candidates' positions show a decrease in support for female candidates after receiving our uninformative message. Conversely, voters with bliss points closer to female candidates' positions increase their support for female candidates. However, since voter bliss points are, on average, closer to male candidates' positions, the overall effect of our uninformative policy message is a decrease in average support for female candidates. This is illustrated in Figure 1, which displays the three observed configurations of candidate platform positions relative to voter bliss points. We find that in 85.8% of municipalities, the male candidate's platform is closer to voters' ideal points than that of the female candidate. In 14.1% of municipalities, the female candidate's platform is closer to male voters' bliss points but further from female voters' bliss points. Only in 0.1% of municipalities is the female candidate's platform closer to the overall distribution of voter bliss points. Confirming this, Appendix Figures D2 and D3 report the distribution of beliefs about the policy position of female candidates for male and female voters. We elaborate further on these implications in the counterfactual exercises section.

Encouragingly, statistical discrimination arising from voters' beliefs about candidate policy positions appears more amenable to reduction than along the ability dimension, at least in our context. The informative policy message reduced the perceived distance between female voters' bliss points and female candidates' positions. At baseline, we have $\xi_p = 2.157$ and $\mu = -3.193$, with a treatment parameter of $\rho_p = -0.192$ (s.e. = 0.067). Since $\xi_p > 0$ and $\mu < \xi_p$, initially female candidates are perceived as being further away from female voters' bliss points compared to male candidates. After exposure to our informative message, this distance decreases (since $\rho_p < 0$), bringing the new perceived policy position ($\tilde{\xi}_p = \xi_p + \rho_p < \xi_p$) closer to μ . The same logic applies to comparisons with male voters, whose average bliss point is at 0. We illustrate this finding in Figure 2, which shows the leftward shift in the density of perceived female candidate platform positions in treated municipalities compared to control, consistent with voters updating their beliefs in response to the informative policy message. Consequently, as perceived female candidate positions shift closer to voter bliss points, the average marginal effects are positive in Panel B of Table 6. Note that the sign reversals across specifications for ξ_p , ρ_p , and μ are not problematic due to the quadratic loss function underlying policy preferences. Although candidate positions and voter bliss points can be mirrored on the policy axis without affecting voter choices, our results remain robust as long as the relative ordering (either $\xi_p >$

$\xi_p + \rho_p > \mu$ or $\xi_p < \xi_p + \rho_p < \mu$) holds true. This robustness is consistent across all five columns of Table 6.

5.2 Quantitative Implications

To give a clearer sense of the relative importance across the dimensions of gender identity, policy, and ability in voter preferences, Figure 3 illustrates the shares and distributions of municipalities for which each candidate dimension dominates on average across voters. Specifically, we first compute the product of weights and the relative utility voters receive from voting for female candidates for each dimension. For gender, this utility is given by $\Delta U_G = -\omega_{G,i} \times 1\{G_i \neq 1\} + \omega_{G,i} \times 1\{G_i \neq 0\}$; for ability, it is $\Delta U_A = \omega_{A_i} \times (A_{i1} - A_{i0})$; and for policy, it is $\Delta U_P = \omega_{P_i} \times ((P_{i1} - P_i)^2 - (P_{i0} - P_i)^2)$. On the y-axis, we plot $|\Delta U_A| - |\Delta U_G|$, and on the x-axis, we plot $|\Delta U_P| - |\Delta U_G|$. We find that 69.5% of vote choices were dominated by the gender dimension, 7.7% by the ability dimension, and the remaining 22.8% by the policy dimension. It is therefore unsurprising that it is especially along the gender identity and policy dimensions that we see economically significant treatment effects.

The model further allows to quantitatively assess the relative importance of statistical versus taste-based discrimination. This is presented visually in Figure 4.

The figure decomposes voter discrimination into its sources: taste-based discrimination and statistical discrimination. To do this, we ensure that our decomposition is not sensitive to the order in which each source is shut down. Because vote choices are non-linear functions of the degree of taste-based and statistical discrimination, the order in which each source is shut down matters in determining how much discrimination is due to each source. To avoid this issue, we consider first the effect of shutting down a source of discrimination relative to when both sources are active. Then, we also compute the effect of shutting down a source of discrimination relative to when the other source has already been shut down. We then compute the marginal change in underrepresentation attributable to each source of discrimination. We define the base level of underrepresentation as $\Delta_{base} = \text{Fem Pop Share} - s(\hat{\theta})$, where $s(\hat{\theta})$ is the model predicted female vote share. Consider the sequence ST , where statistical discrimination is shut down first, followed by taste-based discrimination. The marginal contribution of statistical discrimination is given by: $\Delta_{ST,stat} = s(\hat{\theta}) - s(\hat{\theta}; \text{no-stat})$, and for taste-based discrimination, this is given by: $\Delta_{ST,taste} = s(\hat{\theta}; \text{no-stat}) - s(\hat{\theta}; \text{no-taste, no-stat})$. Similarly, calculate $\Delta_{TS,stat}$ and $\Delta_{TS,taste}$. Then, our decomposition measures are given by: $\text{Decomp}_{stat} = \frac{\Delta_{ST,stat} + \Delta_{TS,stat}}{2 \cdot \Delta_{base}}$ and $\text{Decomp}_{taste} = \frac{\Delta_{ST,taste} + \Delta_{TS,taste}}{2 \cdot \Delta_{base}}$. Note that the measures satisfy $\text{Decomp}_{taste} + \text{Decomp}_{stat} = -100\%$ by construction, which is a desirable property for a decomposition measure.

We compute the above measure for each municipality and plot the densities in Figure 4. Shutting down statistical discrimination is shown to unambiguously reduce female underrepresentation, with an average reduction of 118% of the status quo ante. Effectively female representation increases to more than its proportional level given the correction. Appendix Figure D4 further decomposes statistical discrimination in its components coming from ability and policy.

The effect of shutting down taste-based discrimination is more dispersed around zero and the presence of mass above zero may appear counterintuitive. About half of the municipalities would appear to experience an increase in underrepresentation of women absent taste-based discrimination. The rationale for this result is intuitive, however. Recall that both male and female voters are characterized, by preferences for same-gender candidates. As a consequence, while shutting down discrimination against female candidates for male voters increases female vote shares, shutting down preference for female candidates by female voters decreases female vote shares. Depending on voter and candidate characteristics in a municipality (for example, a city where highly qualified male candidates compete with mediocre female candidates), this second component may dominate, increasing female underrepresentation. The average in the Figure 4 is 18% of the status quo ante.

5.3 Out-of-Sample Fit and Validation

This subsection presents our assessment of the out-of-sample goodness of fit for the model and performs model validation. We consider three exercises designed to assuage concerns of substantial mis-specification in our framework and to lend additional credibility to the structural estimates presented in the previous subsection.

We begin by evaluating model performance in all Brazilian municipalities not included in our RCT sample. The mean squared error (MSE) between model predictions and observed vote shares is 0.0049, compared to an in-sample MSE of 0.004. In panel (a) of Figure 5, we plot predicted versus observed vote shares. The fitted line (slope coefficient = 0.84, s.e. = 0.02) closely follows the 45-degree line, indicating a strong alignment. Panel (b) displays region-specific averages of predicted and observed vote shares, which also show a close correspondence to the real data, further supporting the model's out-of-sample validity.

Second, we examine all municipalities in Brazil that border those included in our RCT sample. For this subsample, the mean squared error is 0.0047. Panel (c) of Figure 5 presents a scatter plot of observed versus predicted vote shares. The fitted line (slope coefficient = 0.87, s.e. =

0.03) lies close to the 45-degree line, again indicating strong alignment. Panel (d) plots region-specific averages of predicted and observed vote shares, which again reassuringly shows a close correspondence, suggesting that the model performs well even in municipalities adjacent to the original RCT sample. These two model-fit exercises allow one to assess the predictive accuracy of the baseline parameters in municipalities outside the RCT sample. Together, they provide encouraging evidence that our model generalizes beyond the experimental context, supporting both an external validity argument for our analysis and that the model does not suffer from in-sample overfitting.

As a third exercise, we perform a standard two-fold validation analysis. To assess whether the treatment parameters consistently predict voting behavior, we split our RCT sample into two parts: (a) a training set comprising 80% of the observations, and (b) a validation set containing the remaining 20% of the observations. We estimate the model using the training sample and evaluate its fit on the validation sample. Importantly, the validation set includes municipalities that received experimental treatment assignments, allowing us to directly test the model's predictive performance in treated contexts.

In this exercise, we find that the mean squared error is 0.0041, closely matching the in-sample MSE of 0.004. Panel (e) of Figure 5 plots predicted versus observed vote shares, where the fitted line (slope coefficient = 0.88, s.e. = 0.099) lies close to the 45-degree line. In panel (f), we plot region-specific averages of predicted and observed vote shares and again observe a close correspondence, in line with the model's predictive power for treatment parameters.

6 Counterfactual Campaigns

An advantage of the structural methodology employed in this paper is its ability to quantitatively assess the value of alternative campaigns in reducing taste-based and statistical discrimination. We investigate a portfolio of counterfactual electoral campaigns and present the results of these exercises in Panel A of Table 7 for the set of municipalities in our RCT sample and in Panel B for the entire country. Geographic visualization of the counterfactuals is in Figures 6 and 7. For completeness, Appendix Table D4 also reports persuasion rates for each counterfactual, computed following [DellaVigna and Gentzkow \(2010\)](#).

We begin by focusing on targeted gender campaigns directed at female and male voters. As observed in the reduced-form analysis (first and second rows of Table 4), only gender identity messages targeted at male voters appear successful in increasing female vote shares, whereas the reduced-form slope is indistinguishable from zero for female voters. In Table 7, we report

a counterfactual effect of a 0.36 percentage points (s.e.=0.1) increase in female vote shares as a result of the gender campaign targeting male voters. This is a highly statistically significant result, but smaller in magnitude when compared to reduced-form estimate reported in column (1) of Table 4, which shows a 1.8 percentage points increase (s.e. =0.9). The difference between reduced-form OLS slopes and the structural-form counterfactual effects arises from treatment heterogeneity that the OLS estimator cannot capture. While our model is characterized by effect sizes that have distributions dictated by heterogeneity in beliefs along both the ability and policy dimensions, OLS in Table 4 forces its slopes to an average that may underweight or overweight certain municipalities, depending on the specifics of the distribution of the effects.¹⁵

Next we consider three blanket campaign counterfactuals that employ gender, informative ability, or uninformative ability messages. In these scenarios, we simulate vote choices under three distinct, un-targeted campaigns in which the respective messages are sent to all voters across all municipalities. We find that none of these three campaigns increases female vote shares significantly, either in the RCT sample (Panel A Table 7) or in the broader sample of Brazil (Panel B Table 7). In particular, the counterfactual effects of the ability messages mirror the estimates in Table 4, where we also observe imprecise null effects.

In contrast to the findings in Table 4, accounting for heterogeneity across voters and the ideological distance between voter and candidate positions reveals that some electoral messages can substantially backfire. To illustrate this, we conduct a counterfactual analysis focused on the *Uninformative Policy* message. Vote choices are simulated under a blanket campaign in which only this uninformative policy message is shown to all voters across all municipalities. We find that such a campaign leads to a decrease in female vote shares, both in the RCT sample (Panel A Table 7) and in the full Brazil sample (Panel B Table 7). Intuitively, this decline in support arises because, in a substantial subset of municipalities, though not all, voters believe the female candidate's platform positions to be further from their own policy bliss points than those of the male candidate (see Figure 1). As a result, increasing the salience of policy preferences amplifies the impact of these ideological distances on vote choice, ultimately reducing support for female candidates. While the structural model accommodates this heterogeneity, the reduced-form model in Equation (4) does not. Consequently, the reduced-form analysis fails to detect the negative effect of this treatment.

We can further contrast the previous results with a counterfactual that also operates along the policy dimension but incorporates information. In the *Informative Policy* counterfactual,

¹⁵A quantitative exploration of the difference between fitting OLS versus our structural model to simulated data is reported in Appendix Table D5. As it can be seen from the table, OLS bias due to misspecification of voter heterogeneity can be quantitative large.

vote choices are simulated under a blanket campaign in which only an informative policy message is shown to all voters across all municipalities. In this case, we find that such a campaign increases female candidates' vote shares. Again, the intuition is straightforward: when voter beliefs are corrected by showing voters that female candidates are closer to a substantial share of the electorate in terms of policy positions, making policy considerations more salient leads to electoral gains for female politicians. In effect, the treatment reduces statistical discrimination.¹⁶

6.1 Optimal Campaigns

The previous set of counterfactual experiments shows that not all messages succeed in increasing electoral support for female candidates. Moreover, voter-level heterogeneity appears to play an important role in mediating the effects of some of the campaign messages we employ. In the next set of counterfactual experiments, we analyze the potential of an optimal campaign to increase female vote shares. To do so, we first introduce notation to define two notions of optimal campaigns.

We define an ad campaign as a binary vector

$$D = (V^{G,0}, V^{A,0}, T^{A,0}, V^{P,0}, T^{P,0}, V^{G,1}, V^{A,1}, T^{A,1}, V^{P,1}, T^{P,1}) \in \{0, 1\}^{10}, \quad (5)$$

where $V^{G,g}$ indicates whether the gender message is shown to gender g ; $V^{A,g}$ indicates whether the informative ability message is shown to gender g ; $T^{A,g}$ indicates whether the uninformative ability message is shown to gender g ; $V^{P,g}$ indicates whether the informative policy message is shown to gender g ; and $T^{P,0}$ indicates whether the uninformative policy message is shown to gender g . We define $g = 0$ for male voters and $g = 1$ for female voters. The set of all feasible campaigns is therefore $\{0, 1\}^{10}$, comprising 1024 possible message combinations.

Let X_m denote the observable characteristics of municipality m , and let $s_m(D, X_m; \hat{\theta})$ be the simulated vote share for female candidates in municipality m under campaign D , based on es-

¹⁶In Table 7, we also demonstrate the usefulness of counterfactual analysis for message bundling and micro-targeting. Bundling informative messages on policy and ability (the *Beliefs* counterfactual) yields a small, statistically insignificant increase in female vote shares, as the ability message dilutes the impact of the policy message. In contrast, bundling uninformative messages (the *Salience* counterfactual) leads to a significant decline in female vote shares, underscoring how bundling can weaken or even reverse the effects of individual messages. Micro-targeting also matters: when all five messages are shown only to male voters, female candidates gain significant electoral support. However, targeting only female voters or all voters produces no significant effect, suggesting male voters are more responsive to these messages (see Panels (c) and (d) of Figure 7).

estimated model parameters $\hat{\theta}$. We assume that all voters are treated in a municipality whenever a treatment is assigned to the municipality.

First, we examine the *Aggregate Optimal Campaign*, which is defined as the campaign $D^{\text{agg-optimal}}$ that maximizes the average predicted vote share across all municipalities:

$$D^{\text{agg-optimal}} = \arg \max_{D \in \{0,1\}^{10}} \frac{1}{M} \sum_{m=1}^M s_m(D, X_m; \hat{\theta}).$$

This optimization focuses on a single, uniform campaign that is optimal at the national level. In this counterfactual scenario, the optimal female vote share is calculated as:

$$S^{\text{agg-optimal}} = \max_{D \in \{0,1\}^{10}} \frac{1}{M} \sum_{m=1}^M s_m(D, X_m; \hat{\theta})$$

For both the RCT sample and the full Brazil sample, the optimal message combination under this counterfactual is the same: for males, the optimal campaign includes the gender message and the informative policy message, that is $V^{G,0} = 1, V^{A,0} = 0, T^{A,0} = 0, V^{P,0} = 1, T^{P,0} = 0$; for females, the optimal combination consists of the uninformative ability message and the informative policy message, corresponding to $V^{G,1} = 0, V^{A,1} = 0, T^{A,1} = 1, V^{P,1} = 1, T^{P,1} = 0$.

We find that this campaign increases the vote share for female candidates by 1.05 percentage points (s.e. = 0.281) within the RCT sample, and by 1.01 percentage points (s.e. = 0.304) in the full Brazil sample. In absolute terms, this translates to an increase of approximately 72,000 votes for female candidates across all municipalities in the RCT sample, and about 975,000 additional votes nationwide. To assess the cost-effectiveness of this intervention, we perform a back-of-the-envelope calculation: within the RCT sample, the campaign persuades approximately 1.6 voters to support a female candidate per dollar spent; in the full Brazil sample, the corresponding figure is 0.96 voters per dollar.¹⁷

¹⁷Obtaining these costs figures requires a series of steps. First, we calculate the total costs for both samples (RCT municipalities only and whole Brazil). For the RCT sample the cost of running one ad across 859 municipalities is \$9,300. This implies a per-message cost per municipality in the RCT sample of \$10.95, calculated as total cost divided by the number of municipalities (\$9,300 / 859). The cost of running four ads, two for males and two for females, is approximately given by \$10.95 × 1000 × 4 = \$43,306.17. Given the total number of additional votes these ads is obtain is 72,000, the ratio is 1.66 votes per dollar. For the whole Brazil we need to take account of the fact that we focused on smaller municipalities where the average population was 8,026 while the average municipality population for the whole Brazil is 35,988. Therefore an estimate that accounts for the fact that one has to treat larger populations yields that the total cost of the campaign is approximately given by \$9300/(859 · 8026.21) × (5507 · 35988.19) × 4 = \$1,069,336. Given that the total additional votes this campaign will obtain is given by 975,000, the ratio is 0.96 additional votes per dollar.

The effects of this optimal campaign counterfactual can be further enhanced by tailoring the message campaign to each municipality. We refer to this as the *Municipality Optimal Campaign*. This counterfactual allows the campaign to vary across municipalities, optimally adjusting the mix of messages to the local context. For each municipality m , we select the campaign D_m^* that maximizes the predicted vote share locally:

$$D_m^{\text{municipality-optimal}} = \arg \max_{D \in \{0,1\}^{10}} s_m(D, X_m; \hat{\theta}).$$

This approach customizes the message combination based on each municipality’s specific characteristics and baseline beliefs, yielding greater gains by leveraging local heterogeneity, while also minimizing potential losses. Simulated vote shares under the municipal-level optimum are calculated as:

$$s^{\text{municipality-optimal}} = \frac{1}{M} \sum_{m=1}^M \left(\max_{D_m \in \{0,1\}^{10}} s_m(D_m, X_m; \hat{\theta}) \right)$$

The Municipality Optimal Campaign increases vote shares for female candidates by 1.463 percentage points (s.e. = 0.281) within the RCT sample and by 1.408 percentage points (s.e. = 0.293) across the full Brazil sample. These effects represent an improvement of approximately 0.4 percentage points (s.e. = 0.12) over the aggregate optimal campaign in both samples. Notably, around 24% of municipalities across Brazil experience a vote share increase of at least 2 percentage points under the Municipality Optimal Campaign (see Figure 8). In terms of total votes, this counterfactual increases support for female candidates by 99,727 votes within the RCT sample and by 1,598,754 votes nationwide. The Municipality Optimal Campaign persuasion rate is estimated at 1.92 percentage points (see Appendix Table D4), a result quantitatively lower than the literature. Gerber and Green (2000); Green and Gerber (2019) focusing on get-out-the-vote campaigns find 11.5-15.6 percentage points. Enikolopov et al. (2011), studying the role of independent media in Russia, reports rates of 7.7 percentage points. Gentzkow (2006) focuses on the introduction of the TV and reports rates of 4.4 percentage points. DellaVigna and Kaplan (2007), based on the entry of Fox News, reports persuasion rates of 11.6 percentage points. The reason for the lower persuasion rates in our setting is intuitively due to the lighter intensity and shorter exposure of our treatments.

For each message type, we can also calculate the proportion of municipalities in which it appears as part of the optimal message mix in the Municipality Optimal Campaign. For the RCT sample (Brazil sample, respectively), the proportions are as follows: gender–male: 73.9%

(75.5%), informative ability–male: 0% (0%), uninformative ability–male: 25.8% (27.8%), informative policy–male: 69.6% (71.9%), uninformative policy–male: 0% (0%), gender–female: 0% (0%), informative ability–female: 0% (0%), uninformative ability–female: 52.5% (52.7%), informative policy–female: 85.1% (85.7%), and uninformative policy–female: 0% (0%).

Given the proportion of municipalities in which each message appears in the optimal mix, we can estimate the monetary cost of implementing the Municipality Optimal Campaign. Using the average optimal message combination –calculated as the sum of all message inclusion proportions– we find that each municipality, on average, receives 3.07 messages. In comparison, the Aggregate Optimal Campaign delivers 4 messages per municipality. This indicates that the Municipality Optimal Campaign not only enhances effectiveness, but also improves cost efficiency. Based on our estimates, the votes-per-dollar ratio for the Municipality Optimal Campaign is 2.97 in the RCT sample. For Brazil as a whole, back-of-the-envelope calculations yield a ratio of 1.91 votes per dollar. These higher ratios reflect both increased vote shares and a more efficient allocation of campaign resources.

7 Conclusions

This paper develops an empirical model of political behavior that incorporates multidimensional candidate characteristics. Candidates are horizontally differentiated along policy and gender identity dimensions, and vertically differentiated by perceived ability. The model is designed to quantitatively assess sources of discrimination in elections by incorporating psychological factors –captured through salience weights in voter preferences– and belief updating mechanisms. Applying the model to local municipal elections in Brazil, we uncover evidence of both taste-based and statistical discrimination against female candidates.

The paper’s structural analysis of a multi-arm electoral experiment, implemented through a large-scale, micro-targeted social media campaign, demonstrates the potential for reducing both taste-based discrimination (by shifting male voters’ preferences regarding candidate gender) and statistical discrimination (by updating beliefs about the policy positions of female candidates). In the context of Brazil, our counterfactual analysis reveals that substantial improvements in gender representation are achievable through an optimally tailored mix of electoral messages across municipalities, with some configurations increasing female candidates’ vote shares by more than 2 percentage points. The analysis also identifies which specific messages may backfire, and explains the mechanisms behind their counterproductive effects.

Future research in Political Economy and Political Science can build on the flexibility of our

empirical framework by incorporating alternative dimensions of candidate identity, policy positions, and perceived ability. Our model is portable to elections in different contexts and political systems, making it well-suited for application beyond the Brazilian setting. Our quantitative methodology can also be applied to other large-scale interventions, particularly those aimed at correcting voter misperceptions or addressing various forms of discrimination. Notably, contemporary challenges such as mass polarization and voter backlash appear especially well-suited for analysis within the framework developed in this paper.

References

- Aker, Jenny, Paul Collier, and Pedro Vicente**, “Is Information Power? Using Cell Phones during an Election in Mozambique,” *World Development*, 2011, pp. 1–53.
- Altonji, Joseph G and Charles R Pierret**, “Employer learning and statistical discrimination,” *The Quarterly Journal of Economics*, 2001, 116 (1), 313–350.
- Anzia, Sarah F and Christopher R Berry**, “The Jackie (and Jill) Robinson effect: Why do congresswomen outperform congressmen?,” *American Journal of Political Science*, 2011, 55 (3), 478–493.
- **and Rachel Bernhard**, “Gender stereotyping and the electoral success of women candidates: New evidence from local elections in the United States,” *British Journal of Political Science*, 2022, 52 (4), 1544–1563.
- Arrow, Kenneth J**, “The Theory of Discrimination,” in O Ashenfelter and A Rees, eds., *Discrimination in Labor Markets*, Princeton University Press, 1973.
- , “What has economics to say about racial discrimination?,” *Journal of economic perspectives*, 1998, 12 (2), 91–100.
- Ashworth, Scott, Christopher R Berry, and Ethan Bueno de Mesquita**, “Modeling theories of women’s underrepresentation in elections,” *American Journal of Political Science*, 2024, 68 (1), 289–303.
- Barbanchon, Thomas Le and Julien Sauvagnat**, “Electoral competition, voter bias, and women in politics,” *Journal of the European Economic Association*, 2022, 20 (1), 352–394.
- Beaman, Lori, Raghavendra Chattopadhyay, Esther Duflo, Rohini Pande, and Petia Topalova**, “Powerful women: does exposure reduce bias?,” *The Quarterly journal of economics*, 2009, 124 (4), 1497–1540.
- Berry, Steve, Christian Cox, and Philip Haile**, “Voting in Two-Party Elections: An Exploration Using Multi-Level Data,” *Working Paper*, 2024.
- Bertrand, Marianne and Esther Duflo**, “Field experiments on discrimination,” *Handbook of Economic Field Experiments*, 2017, 1, 309–393.
- Bonomi, Giampaolo, Nicola Gennaioli, and Guido Tabellini**, “Identity, beliefs, and political conflict,” *The Quarterly Journal of Economics*, 2021, 136 (4), 2371–2411.

- Broockman, David E and Evan J Soltas**, “A natural experiment on discrimination in elections,” *Journal of Public Economics*, 2020, 188, 104201.
- Casas-Arce, Pablo and Albert Saiz**, “Women and power: unpopular, unwilling, or held back?,” *Journal of Political Economy*, 2015, 123 (3), 641–669.
- Cox, Christian**, “The Equilibrium Effects of Campaign Finance Deregulation on US Elections,” *mimeo University of Arizona*, 2024.
- Cruz, Cesi, Philip Keefer, Julien Labonne, and Francesco Trebbi**, “Making policies matter: Voter responses to campaign promises,” *The Economic Journal*, 2024, 134 (661), 1875–1913.
- DellaVigna, Stefano and Ethan Kaplan**, “The Fox News effect: Media bias and voting,” *The Quarterly Journal of Economics*, 2007, 122 (3), 1187–1234.
- **and Matthew Gentzkow**, “Persuasion: Empirical Evidence,” *Annual Review of Economics*, 2010, 2, 643–669.
- Enikolopov, Ruben, Maria Petrova, and Ekaterina Zhuravskaya**, “Media and political persuasion: Evidence from Russia,” *American economic review*, 2011, 101 (7), 3253–3285.
- Fang, Hanming and Andrea Moro**, “Theories of statistical discrimination and affirmative action: A survey,” *Handbook of social economics*, 2011, 1, 133–200.
- Fréchette, Guillaume R, Francois Maniquet, and Massimo Morelli**, “Incumbents’ interests and gender quotas,” *American Journal of Political Science*, 2008, 52 (4), 891–909.
- Fujiwara, Thomas**, “Voting technology, political responsiveness, and infant health: Evidence from Brazil,” *Econometrica*, 2015, 83 (2), 423–464.
- **and Leonard Wantchekon**, “Can informed public deliberation overcome clientelism? Experimental evidence from Benin,” *American Economic Journal: Applied Economics*, 2013, 5 (4), 241–255.
- Gennaioli, Nicola and Guido Tabellini**, “Presidential Lecture: Identity Politics,” *Econometrica*, 2025.
- Gentzkow, Matthew**, “Television and voter turnout,” *The Quarterly Journal of Economics*, 2006, 121 (3), 931–972.
- Gerber, Alan S and Donald P Green**, “The effects of canvassing, telephone calls, and direct mail on voter turnout: A field experiment,” *American political science review*, 2000, 94 (3), 653–663.

- Gonzalez-Eiras, Martin and Carlos Sanz**, “Women’s representation in politics: The effect of electoral systems,” *Journal of Public Economics*, 2021, 198, 104399.
- Green, Donald P and Alan S Gerber**, *Get out the vote: How to increase voter turnout*, Brookings Institution Press, 2019.
- Guryan, Jonathan and Kerwin Kofi Charles**, “Taste-based or statistical discrimination: The economics of discrimination returns to its roots,” *The Economic Journal*, 2013, 123 (572), F417–F432.
- Heckman, James J**, “Detecting discrimination,” *Journal of Economic Perspectives*, 1998, 12 (2), 101–116.
- **and P Siegelman**, “Clear and convincing evidence: Measurement of discrimination in America,” in M Fix and R Struyk, eds., *The Urban Institute audit studies: Their methods and findings*, Washington DC: Urban Institute Press., 1993.
- Iaryczower, Matias, Sergio Montero, and Galileu Kim**, “Representation failure,” Technical Report, National Bureau of Economic Research 2022.
- Instituto de Pesquisa Econômica Aplicada**, *City and Movement: Intermunicipal Commuting in Brazil*, IPEA, 2016.
- Kawai, Kei and Takeaki Sunada**, “Estimating candidate valence,” *Econometrica*, 2025, 93 (2), 463–501.
- Kendall, Chad, Tommaso Nannicini, and Francesco Trebbi**, “How Do Voters Respond to Information? Evidence from a Randomized Campaign,” *American Economic Review*, January 2015, 105 (1), 322–53.
- Lang, Kevin and Jee-Yeon K Lehmann**, “Racial discrimination in the labor market: Theory and empirics,” *Journal of Economic Literature*, 2012, 50 (4), 959–1006.
- Longuet-Marx, Nicolas**, “Party Lines or Voter Preferences? Explaining Political Realignment,” *mimeo Columbia University*, 2025.
- Magesan, Arvind, Andrea Szabó, and Gergely Ujhelyi**, “Candidate Selection by Parties: Crime and Politics in India,” Technical Report, mimeo University of Huston 2024.
- Montero, Sergio**, “Going It Alone? A Structural Analysis of Coalition Formation in Elections,” *The Journal of Politics*, 2025, 87 (2), 774–789.

Neumark, David, “Detecting discrimination in audit and correspondence studies,” *Journal of Human Resources*, 2012, 47 (4), 1128–1157.

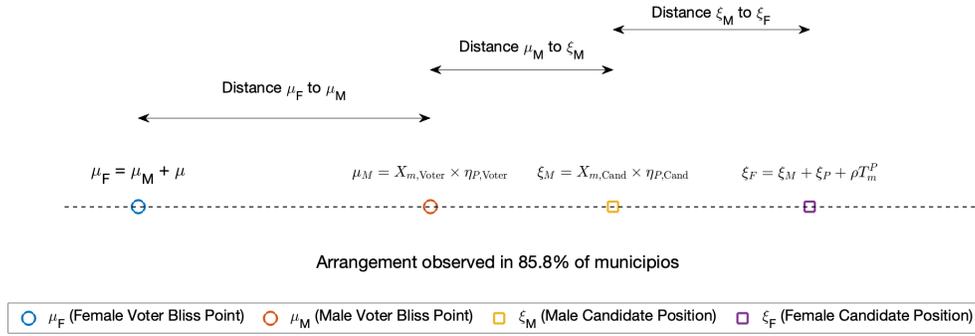
Phelps, Edmund, “The Statistical Theory of Racism and Sexism,” *American Economic Review*, September 1972, 62, 659–661.

Rede A Ponte, “Raio X da Diversidade: VereadorAs do Brasil,” 2024. Disponível em: <https://redeaponte.org/relatorios/#raioxdiversidade>.

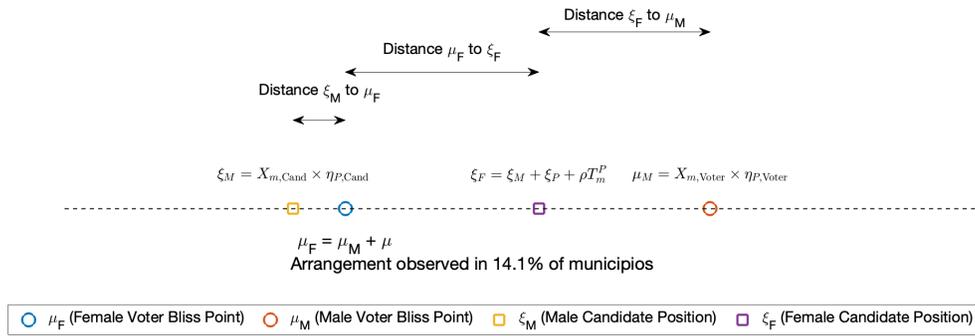
Ujhelyi, Gergely, Somdeep Chatterjee, and Andrea Szabo, “None of the Above: Protest Voting in the World’s Largest Democracy,” *Journal of the European Economic Association*, 2021, 19 (3), 1936–1979.

Vicente, Pedro and Leonard Wantchekon, “Clientelism and Vote Buying: Lessons from Field Experiments in African Elections,” *Oxford Review of Economic Policy*, 2009, 25 (2), 292–305.

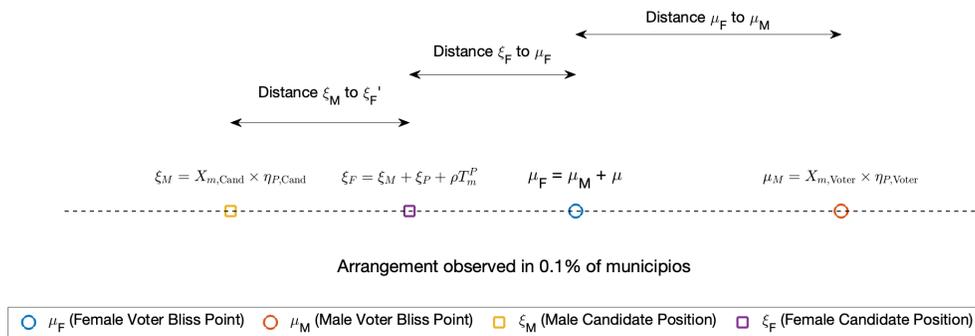
Vicente, Pedro C, “Is vote buying effective? Evidence from a field experiment in West Africa,” *The Economic Journal*, 2014, 124 (574), F356–F387.



(a) Arrangement 1



(b) Arrangement 2



(c) Arrangement 3

Figure 1: These figures illustrate three distinct configurations of male and female voter bliss points alongside the estimated platform positions of male and female candidates, as recovered from our structural model.

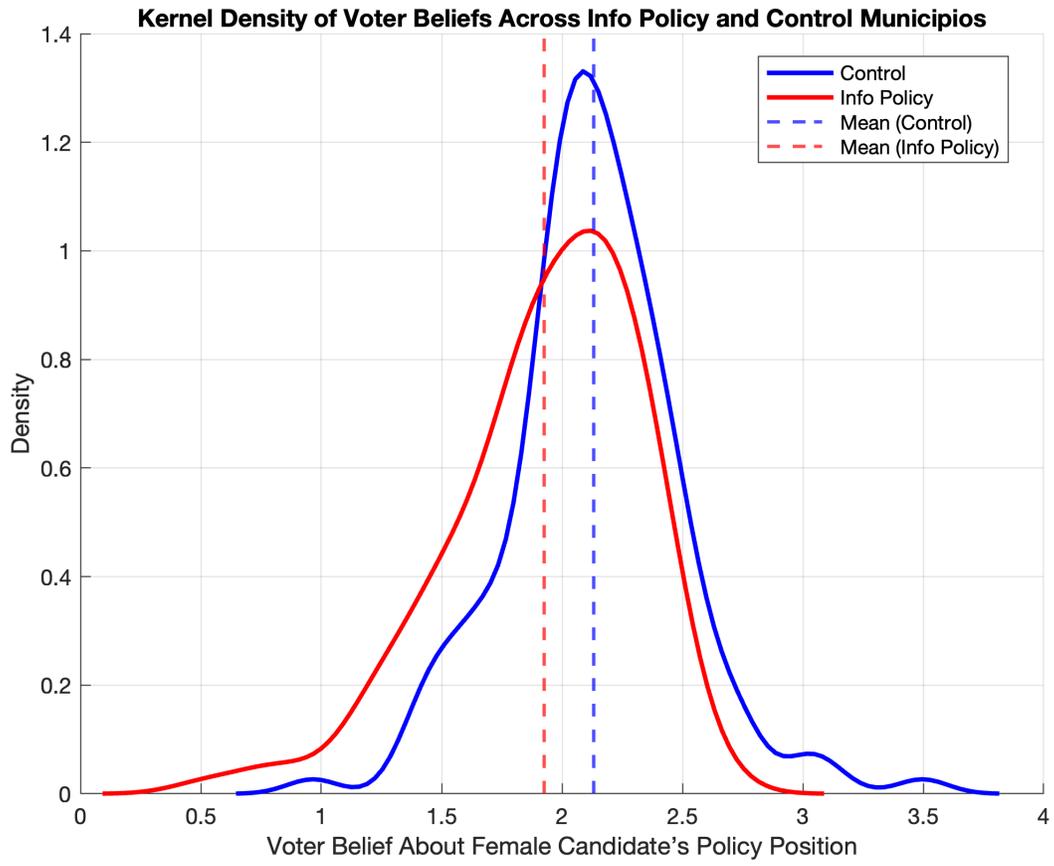


Figure 2: This figure plots voters' beliefs about the female candidate's policy position in control municipalities versus those in municipalities treated with the informative policy message.

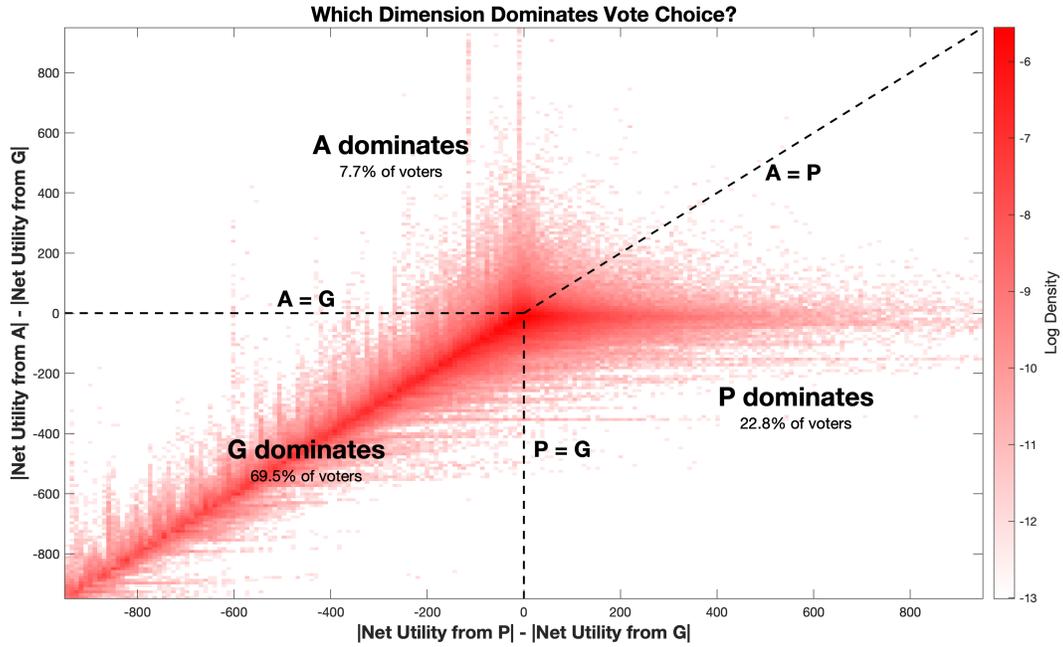


Figure 3: This figure illustrates which dimensions dominate vote choice for voters. To demonstrate this, we show the relative magnitudes of the Ability dimension versus the Gender dimension against the Policy dimension versus the Gender dimension. The regions where each dimension dominates vote choice are highlighted. To be precise, we first compute the product of weights and the relative utility voters receive from voting for female candidates for each dimension. For Gender, this utility is given by $\Delta U_G = -\omega_{G,i} \times 1 \{G_i \neq 1\} + \omega_{G,i} \times 1 \{G_i \neq 0\}$; for Ability, it is $\Delta U_A = \omega_{A_i} \times (A_{i1} - A_{i0})$; and for Policy, it is $\Delta U_P = \omega_{P_i} \times ((P_{i1} - P_i)^2 - (P_{i0} - P_i)^2)$. On the y-axis, we plot $|\Delta U_A| - |\Delta U_G|$, and on the x-axis, we plot $|\Delta U_P| - |\Delta U_G|$. We find that 69.5% of vote choices were dominated by the Gender dimension, 7.7% by the Ability dimension, and the remaining 22.8% by the Policy dimension.

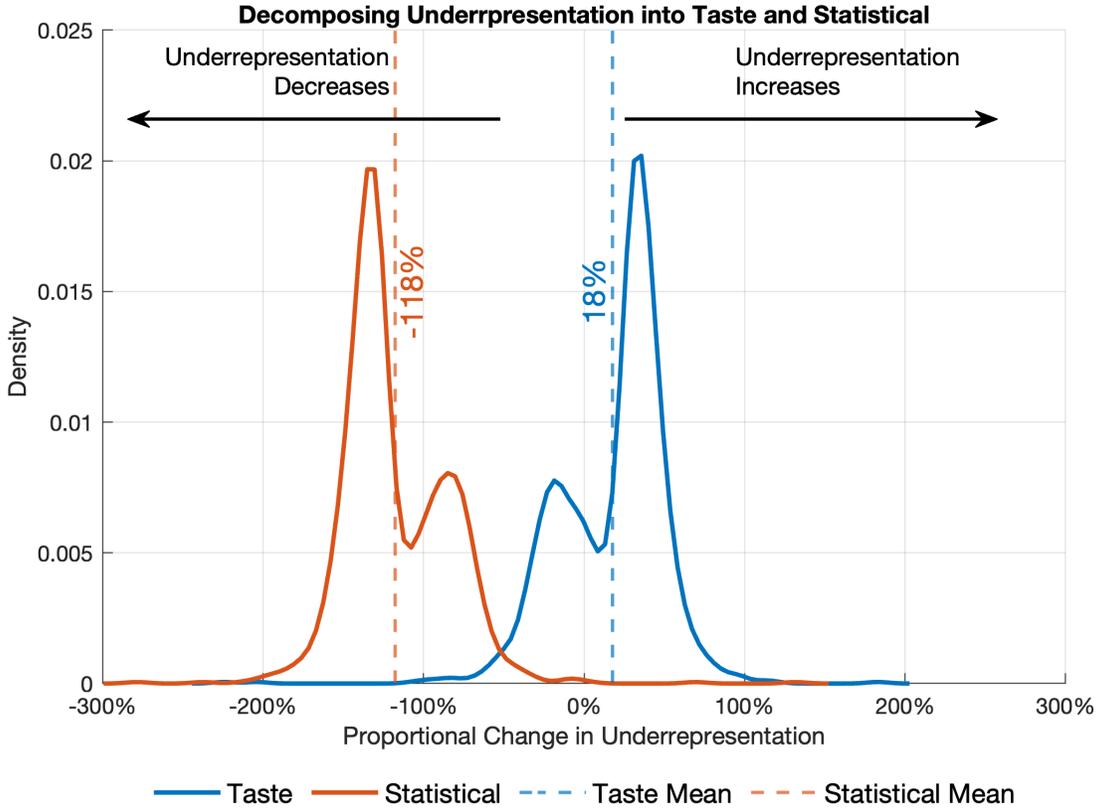


Figure 4: We decompose voter discrimination into its sources: taste-based discrimination vs. statistical discrimination. To do this, we ensure that our decomposition is not sensitive to the order in which each source is shut down. Note that, since vote choices are non-linear functions of the degree of taste-based and statistical discrimination, the order in which each source is shut down matters in determining how much discrimination is due to each source. To avoid this issue, we consider both possible orders and then compute the marginal changes in underrepresentation attributable to each source of discrimination. Define the base level of discrimination as $\Delta_{base} = \text{Fem Pop Share} - s(\hat{\theta})$. Consider the sequence ST , where statistical discrimination is shut down first, followed by taste-based discrimination. The marginal contribution of statistical discrimination is given by: $\Delta_{ST,stat} = s(\hat{\theta}) - s(\hat{\theta}; \text{no-stat})$, and for taste-based discrimination, this is given by: $\Delta_{ST,taste} = s(\hat{\theta}; \text{no-stat}) - s(\hat{\theta}; \text{no-taste, no-stat})$. Similarly, calculate $\Delta_{TS,stat}$ and $\Delta_{TS,taste}$. Then, our decomposition measures are given by: $\text{Decomp}_{stat} = \frac{\Delta_{ST,stat} + \Delta_{TS,stat}}{2 \cdot \Delta_{base}}$ and $\text{Decomp}_{taste} = \frac{\Delta_{ST,taste} + \Delta_{TS,taste}}{2 \cdot \Delta_{base}}$. Note that the measures satisfy $\text{Decomp}_{taste} + \text{Decomp}_{stat} = -1$ by construction, which is a desirable property for a decomposition measure. We compute the above measure for each municipality and plot the densities.

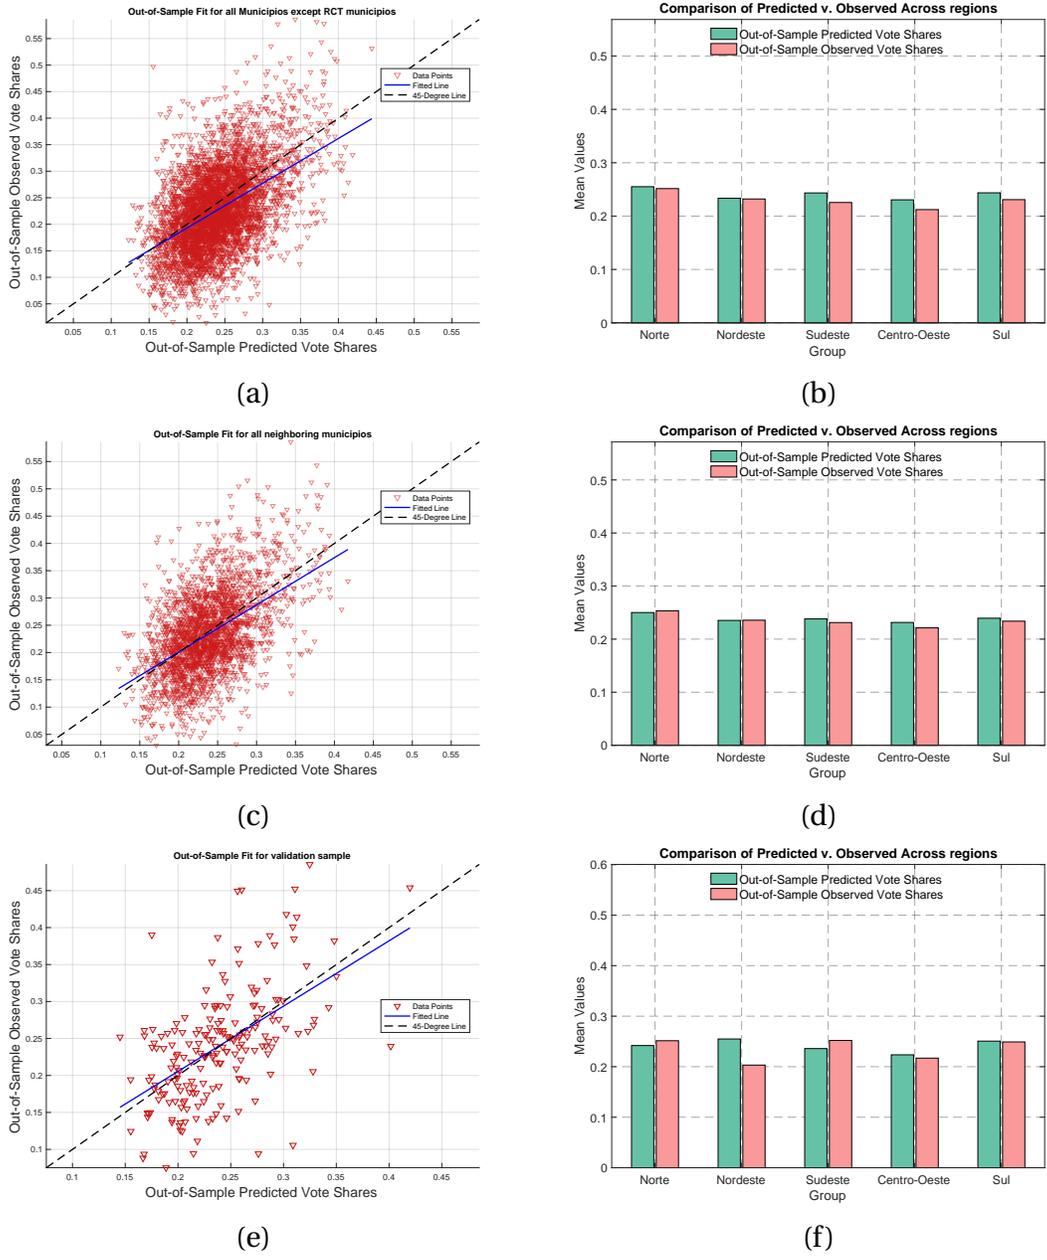


Figure 5: Out-of-Sample Fit: Panels (a) and (b) display model fit for all Brazilian municipalities not included in the RCT sample. Panels (c) and (d) show fit for municipalities neighboring those in the RCT. Panels (e) and (f) report fit for a hold-out validation set. For these last two panels, the RCT sample is randomly split into a training set (80% of RCT municipalities) and a validation set (20%). The model is re-estimated using only the training set, and predictive accuracy is evaluated on the validation set. While Panels (a) through (d) assess the predictive power of baseline parameters, Panels (e) and (f) provide insight into the model's ability to generalize treatment parameter effects.

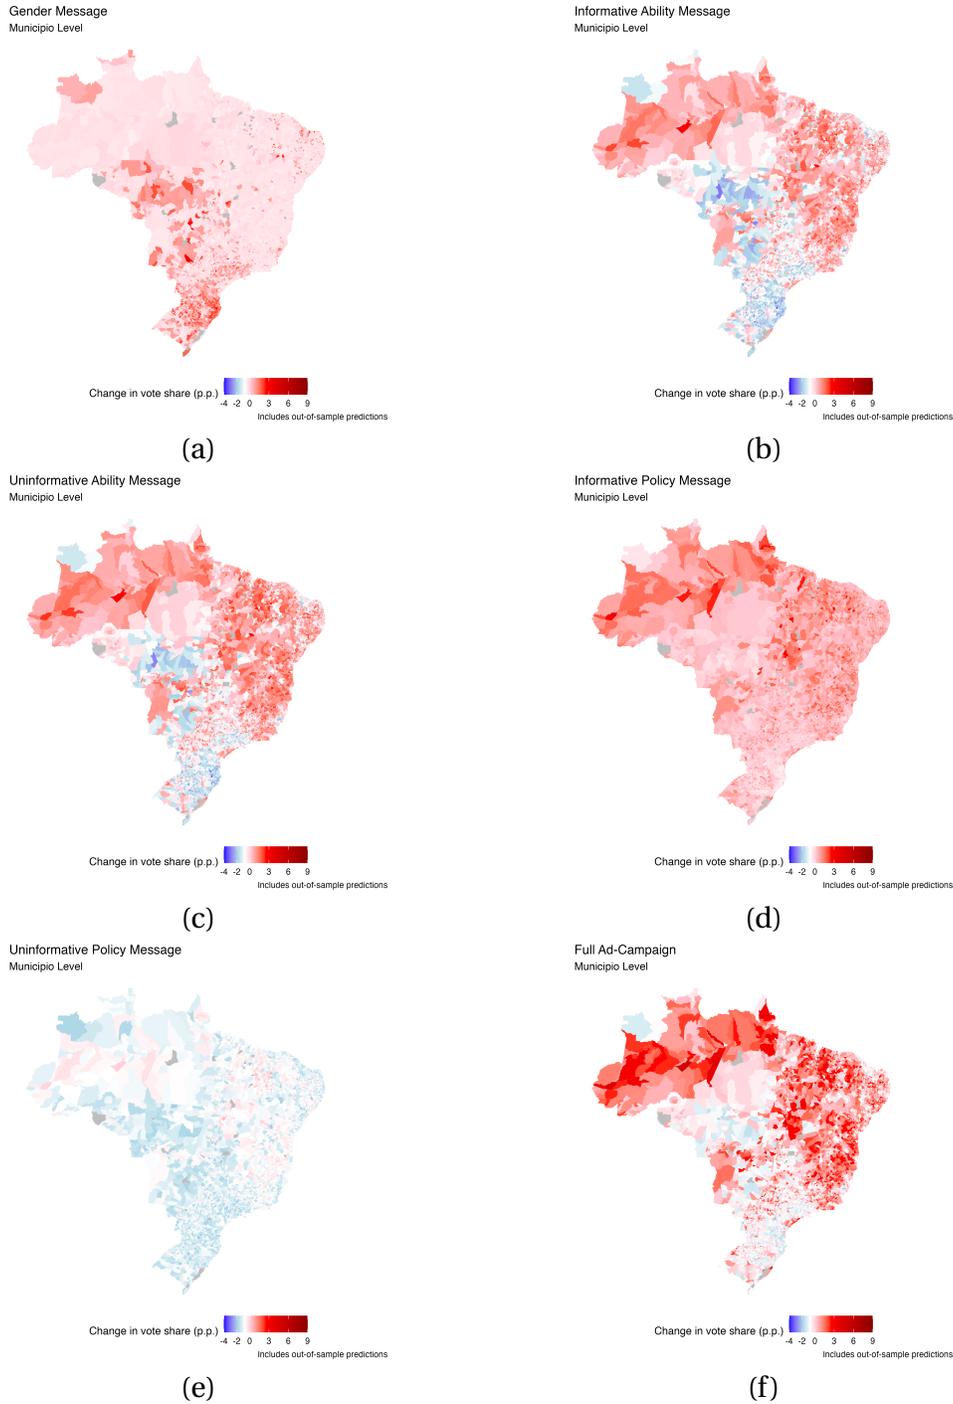
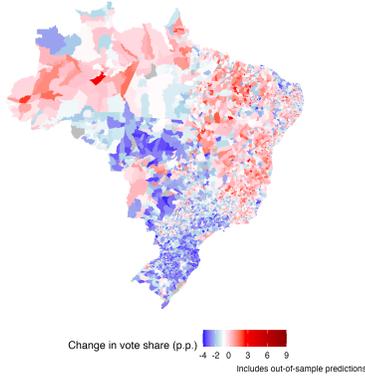


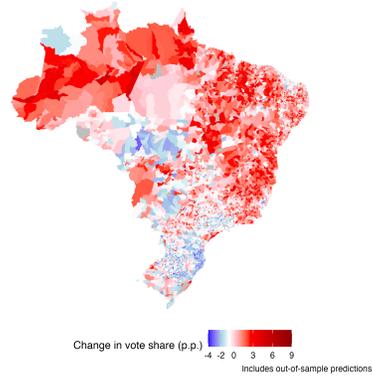
Figure 6: Choropleth plots at the municipality level depicting responses to ad campaigns. Panel (a) shows changes in vote shares when all voters are targeted with the Gender-Identity message. Panel (b) displays changes in vote shares when all voters are targeted with the Informative Ability message. Panel (c) illustrates changes in vote shares when all voters are targeted with the Uninformative Ability message. Panel (d) presents changes in vote shares when all voters are targeted with the Informative Policy message. Panel (e) shows changes in vote shares when all voters are targeted with the Uninformative Policy message. Finally, Panel (f) highlights changes in vote shares when all voters are targeted with all messages simultaneously.

Only Uninformative Messages for Ability and Policy
Município Level



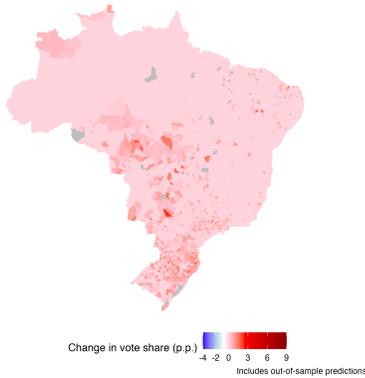
(a)

Only Informative Message for Ability and Policy
Município Level



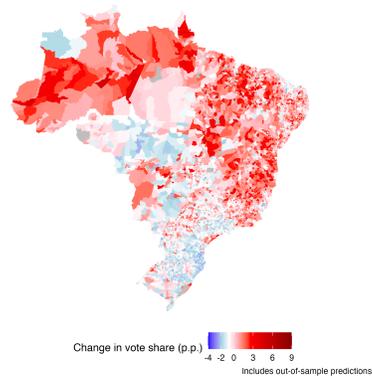
(b)

Only Males Voters Targetted with All messages
Município Level



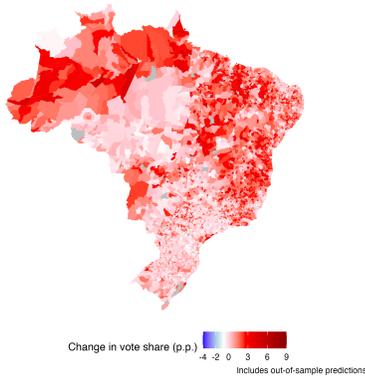
(c)

Only Females Voters Targetted with All messages
Município Level



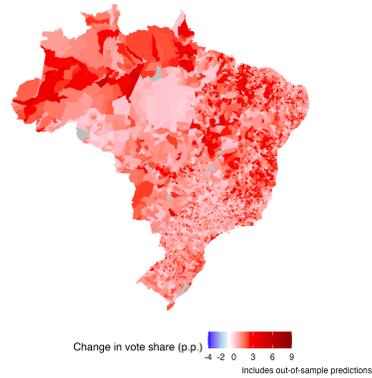
(d)

Optimal Ad-Campaign at National Level
Município Level



(e)

Optimal Ad-Campaign Message
Município Level



(f)

Figure 7: Choropleth plots at the municipality level illustrating responses to ad campaigns. Panel (a) shows changes in vote shares when all voters are targeted with the Uninformative Ability and Policy messages, highlighting the impact of targeting the salience of these two dimensions simultaneously. Panel (b) displays changes in vote shares when all voters are targeted with the Informative Ability and Policy messages, focusing on the effects of targeting beliefs about these dimensions simultaneously. Panel (c) illustrates changes in vote shares when only male voters are targeted with all messages simultaneously. Panel (d) presents changes in vote shares when only female voters are targeted with all messages simultaneously. Panel (e) depicts changes in vote shares when all voters are targeted with the national-level optimal ad campaign. Finally, Panel (f) highlights changes in vote shares when all voters are targeted with the municipality-specific optimal ad campaign.

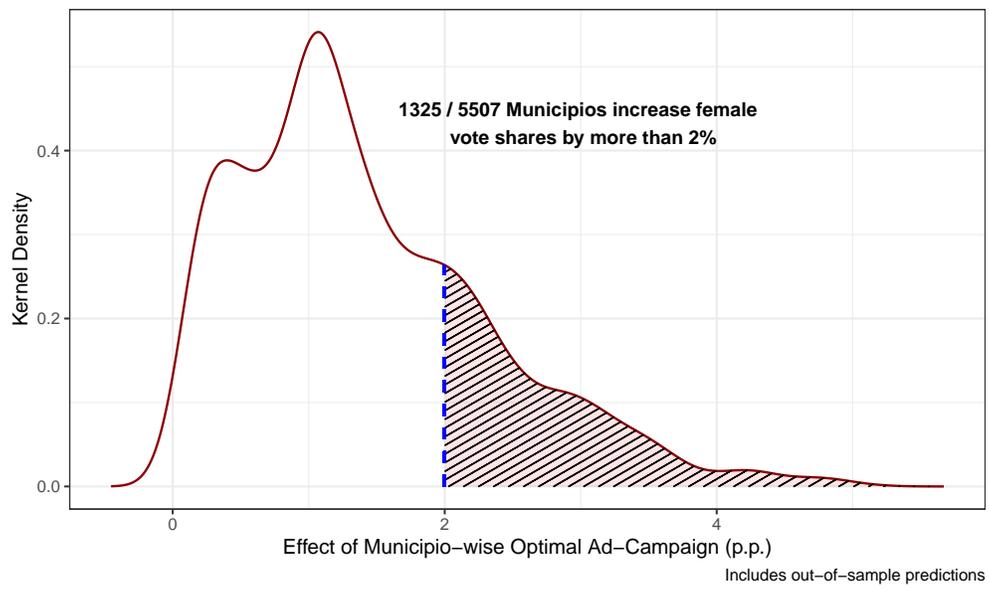


Figure 8: Distribution of change in vote share across municipalities under the optimal ad campaign

Table 1: First Stage

Variable	(1) C Mean/(SE)	(2) AI Mean/(SE)	(3) AU Mean/(SE)	(4) GI_F Mean/(SE)	(5) GI_M Mean/(SE)	(6) PI Mean/(SE)	(7) PU Mean/(SE)
Total Ad Impressions	0.000 (0.000)	12933.586 (274.168)	12828.257 (238.235)	8613.394 (209.645)	8162.650 (200.835)	13164.972 (268.357)	13094.028 (298.303)
Ad Impressions (Female)	0.000 (0.000)	7386.407 (174.559)	7330.382 (152.796)	8613.394 (209.645)	0.000 (0.000)	7536.732 (169.700)	7473.458 (188.240)
Ad Impressions (Male)	0.000 (0.000)	5505.041 (107.286)	5455.535 (92.051)	0.000 (0.000)	8162.650 (200.835)	5584.634 (105.374)	5579.718 (114.562)
Total Ad Reach	0.000 (0.000)	7025.497 (164.059)	6832.861 (140.312)	4469.585 (121.631)	4360.811 (123.260)	7031.120 (163.401)	6883.000 (176.126)
Ad Reach (Female)	0.000 (0.000)	3888.400 (99.904)	3781.500 (85.935)	4469.585 (121.631)	0.000 (0.000)	3899.401 (98.000)	3813.092 (106.146)
Ad Reach (Male)	0.000 (0.000)	3114.000 (67.669)	3029.576 (57.408)	0.000 (0.000)	4360.811 (123.260)	3108.155 (68.265)	3049.289 (72.157)
Ad Reach per Voter	0.000 (0.000)	1.059 (0.022)	1.044 (0.018)	0.675 (0.016)	0.656 (0.017)	1.055 (0.023)	1.035 (0.021)
Number of observations	140	145	144	142	143	142	142

Notes: Robust standard errors in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Table 2: How representative is the RCT Sample

Variable	RCT-Sample	All Municipios
Population	8026.21 (84.63)	35988.19 (2725.11)
% of Literate Population	84.13 (0.3)	84.01 (0.13)
Congressional female vote share in 2022 (p.p.)	15.14 (0.43)	15.9 (0.18)
% of Brown	43.34 (0.63)	47.45 (0.26)
% of Black	7.2 (0.17)	8.54 (0.08)
% of Indegenous	1.35 (0.22)	0.97 (0.07)
% covered with internet	82.82 (0.54)	84.59 (0.22)
GDP per capita	33702.95 (1060.4)	33837.15 (561.41)
% completed secondary school	88.28 (0.24)	88.19 (0.1)
% of Urban population	60.32 (0.65)	63.82 (0.3)
% of population b/w 20 to 44	35.6 (0.09)	36.15 (0.04)
% of population b/w 45 to 64	24.56 (0.11)	24.07 (0.05)
% of 65 plus population	12.65 (0.1)	12.23 (0.05)

Notes: This table compares socioeconomic characteristics between the RCT sample and all Brazilian municipalities. Variables include population size, literacy rate, racial composition, internet coverage, urbanization, GDP per capita, and age distribution. Standard errors are reported in parentheses. All values are calculated using the most recent available data from the Brazilian Institute of Geography and Statistics (IBGE) and the National Telecommunications Agency (Anatel).

Table 3: Balance

Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(2)-(1)	(3)-(1)	(4)-(1)	Pairwise t-test		(6)-(1)	(7)-(1)
	C	AI	AJ	GI_F	GI_M	PI	PJ				Mean difference	Mean difference		
Log Population	8.972 (0.025)	8.938 (0.023)	8.926 (0.021)	8.938 (0.026)	8.946 (0.023)	8.959 (0.023)	8.954 (0.024)	-0.034	-0.046*	-0.034	-0.026	-0.013	-0.018	
Log GDP per Capita	10.128 (0.063)	10.173 (0.058)	10.051 (0.059)	10.174 (0.062)	10.118 (0.059)	10.200 (0.063)	10.065 (0.067)	0.045	-0.078	0.045	-0.011	0.072	-0.063	
Internet Penetration Rate	81.717 (1.460)	84.926 (1.284)	83.344 (1.312)	83.212 (1.435)	84.108 (1.438)	82.382 (1.564)	79.999 (1.557)	3.209*	1.627	1.495	2.391	0.665	-1.718	
Black Population (%)	50.660 (1.941)	50.910 (1.809)	51.106 (1.893)	50.863 (1.923)	50.410 (1.924)	48.145 (1.990)	51.636 (1.926)	0.251	0.446	0.203	-0.250	-2.515	0.977	
Urban Population (%)	0.592 (0.016)	0.627 (0.016)	0.602 (0.016)	0.624 (0.018)	0.591 (0.018)	0.603 (0.017)	0.583 (0.018)	0.035*	0.010	0.033	-0.001	0.011	-0.009	
Literacy Rate	87.983 (0.656)	88.551 (0.622)	88.101 (0.631)	88.789 (0.643)	88.557 (0.658)	88.652 (0.628)	87.332 (0.641)	0.568	0.118	0.806	0.575	0.670	-0.651	
Population Aged 15-24 (%)	14.221 (0.157)	14.410 (0.170)	14.472 (0.173)	14.030 (0.165)	14.016 (0.155)	14.082 (0.187)	14.495 (0.175)	0.189	0.250	-0.191	-0.205	-0.139	0.274	
Population Aged 60+ (%)	17.877 (0.356)	17.736 (0.355)	17.893 (0.360)	17.994 (0.364)	18.517 (0.358)	18.123 (0.388)	17.956 (0.370)	-0.141	0.015	0.116	0.639**	0.245	0.079	
Female Population (%)	49.702 (0.096)	49.710 (0.134)	49.523 (0.177)	49.620 (0.159)	49.510 (0.206)	49.790 (0.084)	49.689 (0.098)	0.008	-0.179	-0.082	-0.192	0.088	-0.013	
Vote Share - Female - State (2022)	15.771 (1.065)	14.330 (1.018)	14.654 (1.024)	16.456 (1.104)	16.092 (1.078)	15.543 (1.028)	15.158 (1.201)	-1.440	-1.117	0.685	0.322	-0.228	-0.612	
Vote Share - Female - Congress (2022)	15.359 (1.180)	14.658 (1.140)	15.870 (1.192)	16.795 (1.283)	13.931 (1.048)	15.258 (1.081)	14.109 (1.056)	-0.701	0.511	1.436	-1.428	-0.101	-1.250	
Vote Share for Bolsonaro (2022)	0.440 (0.015)	0.456 (0.015)	0.436 (0.015)	0.449 (0.016)	0.443 (0.015)	0.463 (0.015)	0.409 (0.016)	0.016	-0.004	0.008	0.003	0.022	-0.031*	
Log Campaign Expenses - Female	6.987 (0.133)	6.807 (0.142)	6.897 (0.134)	6.920 (0.143)	6.815 (0.161)	7.092 (0.120)	7.078 (0.141)	-0.180	-0.090	-0.066	-0.171	0.105	0.092	
Log Campaign Expenses - Male	7.017 (0.086)	6.923 (0.095)	6.909 (0.109)	7.031 (0.087)	6.949 (0.099)	6.989 (0.105)	7.077 (0.107)	-0.094	-0.108	0.015	-0.068	-0.028	0.061	
College Degree - Female (%)	0.285 (0.012)	0.260 (0.011)	0.314 (0.013)	0.297 (0.013)	0.300 (0.014)	0.307 (0.013)	0.293 (0.014)	-0.025	0.029	0.012	0.015	0.022	0.008	
College Degree - Male (%)	0.159 (0.007)	0.162 (0.008)	0.168 (0.008)	0.164 (0.007)	0.164 (0.008)	0.170 (0.009)	0.168 (0.007)	0.003	0.009	0.005	0.005	0.011	0.009	
Black - Female (%)	0.459 (0.025)	0.450 (0.021)	0.448 (0.024)	0.450 (0.023)	0.437 (0.024)	0.447 (0.025)	0.459 (0.024)	-0.009	-0.011	-0.009	-0.021	-0.011	0.000	
Black - Male (%)	0.459 (0.025)	0.456 (0.023)	0.450 (0.024)	0.455 (0.024)	0.436 (0.024)	0.429 (0.025)	0.463 (0.023)	-0.003	-0.009	-0.004	-0.023	-0.030	0.004	
Married - Female (%)	0.534 (0.011)	0.540 (0.012)	0.558 (0.011)	0.539 (0.012)	0.548 (0.013)	0.545 (0.012)	0.531 (0.013)	0.006	0.024	0.005	0.013	0.011	-0.003	
Married - Male (%)	0.435 (0.013)	0.415 (0.014)	0.434 (0.014)	0.411 (0.015)	0.452 (0.015)	0.438 (0.015)	0.411 (0.015)	-0.020	-0.001	-0.024	0.017	0.003	-0.024	
Log Wealth - Female	11.711 (0.065)	11.622 (0.084)	11.706 (0.077)	11.728 (0.085)	11.853 (0.069)	11.775 (0.072)	11.737 (0.073)	-0.089	-0.006	0.017	0.142	0.064	0.026	
Log Wealth - Male	12.175 (0.056)	12.192 (0.058)	12.232 (0.058)	12.154 (0.060)	12.210 (0.064)	12.264 (0.055)	12.142 (0.055)	0.017	0.057	-0.021	0.036	0.089	-0.032	
Female Candidate (%)	0.353 (0.003)	0.357 (0.003)	0.355 (0.003)	0.357 (0.003)	0.358 (0.003)	0.354 (0.003)	0.354 (0.003)	0.004	0.002	0.004	0.005	0.001	0.001	
F-test of joint significance (F-stat)								0.948	1.113	0.657	0.791	0.407	0.647	
Number of observations	141	145	145	142	143	142	142	286	286	283	284	283	283	

Notes: Robust standard errors in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Table 4: The Effects of Campaign Message on Vote Share for Female Candidates

	Vote Share for Female Candidates			
	(1)	(2)	(3)	(4)
Gender - Female	1.141 (0.918)	0.940 (0.915)	0.948 (0.917)	1.008 (0.908)
Gender - Male	1.790* (0.933)	1.879** (0.922)	1.806* (0.924)	1.800* (0.928)
Ability Uninformative	1.085 (0.924)	1.045 (0.915)	1.031 (0.915)	1.100 (0.916)
Ability Informative	0.284 (0.923)	0.280 (0.908)	0.285 (0.912)	0.354 (0.919)
Policy Uninformative	-0.298 (0.961)	-0.218 (0.967)	-0.216 (0.966)	-0.267 (0.970)
Policy Informative	1.432 (0.959)	1.394 (0.946)	1.385 (0.948)	1.419 (0.946)
DV Control Mean	22.97	22.97	22.97	22.97
R^2	0.03	0.05	0.05	0.06
Number of Obs.	1000	1000	1000	1000
Strata FE	Y	Y	Y	Y
Lagged DV	N	Y	Y	Y
Controls	N	N	Y	Y

Notes: Robust standard errors in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Table 5: The Effects of Campaign Message on Other Electoral Outcomes

	Turnout	Campaign Spending Females	Campaign Spending Males	Share Elected Females
	(1)	(2)	(3)	(4)
Gender - Female	-0.002 (0.003)	-10.643 (174.848)	21.877 (153.136)	0.014 (0.013)
Gender - Male	-0.001 (0.003)	49.682 (180.144)	19.731 (155.374)	0.018 (0.014)
Ability Uninformative	0.000 (0.003)	31.604 (184.925)	31.722 (156.190)	0.038** (0.016)
Ability Informative	-0.001 (0.003)	-21.038 (176.264)	2.883 (156.901)	-0.003 (0.013)
Policy Uninformative	-0.001 (0.003)	298.856 (221.219)	182.369 (186.854)	0.023 (0.016)
Policy Informative	0.003 (0.003)	154.277 (183.438)	77.181 (163.118)	0.028** (0.014)
DV Control Mean	0.85	1969.23	1680.76	0.12
R ²	0.63	0.23	0.22	0.10
Number of Obs.	998	1000	1000	1000
Strata FE	Y	Y	Y	Y
Lagged DV	Y	N	N	N
Controls	Y	Y	Y	Y

Notes: Robust standard errors in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Table 6: Main Parameter Estimates and Their Marginal Effects

Panel A: Main Parameters		Models				
Parameter	Description	(1)	(2)	(3)	(4)	(5)
ω_{GF}	Baseline weight gender-identity (F)	4.088*** (0.010)	3.878*** (0.031)	4.847*** (0.015)	-6.772*** (0.382)	1498.302*** (8.599)
ω_{GM}	Baseline weight gender-identity (M)	6.233*** (0.012)	7.049*** (0.028)	6.613*** (0.023)	13.404*** (0.663)	931.560*** (0.000)
λ_{GF}	Effect of gender message (F)	-0.018 (0.034)	-0.000 (0.032)	-0.003 (0.009)	-0.068 (0.189)	48.147*** (1.351)
λ_{GM}	Effect of gender message (M)	-0.533*** (0.015)	-1.528*** (0.018)	-0.859*** (0.014)	-1.552*** (0.361)	-931.456*** (0.000)
ω_{A0}	Baseline weight ability	2.150*** (0.519)	1.995*** (0.412)	1.705*** (0.593)	0.312*** (0.036)	0.861*** (0.003)
λ_A	Effect of informative ability message	-0.162*** (0.022)	-0.168*** (0.026)	0.121*** (0.012)	-0.331*** (0.078)	-0.458*** (0.002)
ω_{P0}	Baseline weight policy	1.343** (0.564)	-0.077 (0.600)	1.008* (0.598)	0.663*** (0.028)	0.526*** (0.001)
λ_P	Effect of informative policy message	0.083*** (0.012)	0.089*** (0.015)	0.064*** (0.013)	0.095*** (0.015)	0.036*** (0.002)
ξ_A	Baseline relative ability F v. M candidates	0.120** (0.059)	0.206** (0.096)	-3.050* (1.847)	-16.142*** (2.231)	-397.479*** (2.558)
ρ_A	Effect of informative ability message	-0.092 (0.064)	-0.052 (0.185)	0.324 (0.218)	-1.562 (2.354)	-46.227*** (1.530)
ξ_P	Baseline relative policy position of F v. M candidates	2.157*** (0.624)	-3.078*** (0.857)	4.084*** (1.212)	-3.855*** (0.235)	22.653*** (0.071)
ρ_P	Effect of informative policy message	-0.192*** (0.067)	0.311*** (0.092)	-0.380*** (0.128)	0.470*** (0.089)	-0.953*** (0.056)
μ	Relative ideal point of F v. M voters	-3.193*** (0.876)	8.216*** (2.952)	-2.886*** (0.905)	4.471*** (0.245)	-23.343*** (0.121)
Region Fixed Effect		✓	-	✓	✓	✓
Candidate controls		✓	✓	✓	✓	✓
Municipio Characteristics		✓	✓	-	✓	✓
Saliency-Weight functional Form		Exponential	Exponential	Exponential	Quadratic	Absolute
N	Number of Observations	1000	1000	1000	1000	1000
Obj Fun	Objective Function Value (MSE)	0.004	0.0041	0.0042	0.004	0.004

Panel B: Average Marginal Effects		Models				
Parameter	Description	(1)	(2)	(3)	(4)	(5)
ω_{GF}	Baseline weight gender-identity (F)	19.470* (10.513)	19.138* (10.149)	22.030* (11.816)	19.856** (9.899)	23.070** (11.535)
ω_{GM}	Baseline weight gender-identity (M)	-7.978* (4.130)	-10.862** (5.404)	-2.943 (2.602)	-3.841* (1.977)	-0.693** (0.346)
λ_{GF}	Effect of gender message (F)	-0.227 (0.434)	-0.006 (0.391)	-0.040 (0.124)	0.270 (0.726)	0.845*** (0.023)
λ_{GM}	Effect of gender message (M)	0.356*** (0.097)	0.763*** (0.155)	0.186 (0.340)	0.083*** (0.023)	0.957*** (0.007)
ω_{A0}	Baseline weight ability	-1.190* (0.611)	-1.105* (0.600)	-2.927* (1.595)	-0.195 (0.120)	-0.171** (0.087)
λ_A	Effect of informative ability message	0.049 (0.050)	0.067 (0.045)	-0.284** (0.142)	0.363** (0.178)	0.916*** (0.010)
ω_{P0}	Baseline weight policy	-13.690** (6.785)	0.930 (4.108)	-10.396* (5.665)	-1.666** (0.829)	1.541** (0.770)
λ_P	Effect of informative policy message	-1.005*** (0.161)	-1.060*** (0.196)	-0.741*** (0.169)	-0.757*** (0.124)	-0.033*** (0.003)
ξ_A	Baseline relative ability F v. M candidates	0.293* (0.151)	0.498** (0.248)	-3.158* (1.678)	-4.864** (2.401)	-8.146** (4.072)
ρ_A	Effect of informative ability message	-0.159 (0.131)	-0.050 (0.579)	0.041 (0.154)	-0.053 (0.496)	0.132*** (0.026)
ξ_P	Baseline relative policy position of F v. M candidates	-20.131** (9.681)	-19.929** (9.626)	-19.049** (9.152)	-15.832** (7.911)	-24.858** (12.429)
ρ_P	Effect of informative policy message	0.536** (0.262)	0.629** (0.302)	0.882*** (0.236)	0.862*** (0.298)	1.049*** (0.058)
μ	Relative ideal point of F v. M voters	-13.294** (6.550)	-16.018** (7.788)	-8.946** (4.409)	-9.661** (4.840)	-18.157** (9.078)

*, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Notes: This table presents the estimated structural parameters (Panel A) and average marginal effects (Panel B) from our random utility voting model. Parameters are estimated using GMM and capture both statistical and taste-based discrimination. Column (1) is the baseline specification, which includes candidate- and municipality-level controls, region fixed effects, and exponential saliency weights. Columns (2)–(5) report estimates from alternative specifications that remove fixed effects, controls, or modify the functional form of saliency. In Panel B, average marginal effects for baseline parameters are computed as the average difference in predicted vote share when the parameter is set to zero, holding all other parameters fixed at their estimated values. This represents the ceteris paribus impact of each parameter on female candidate vote share. For treatment parameters, marginal effects are computed through counterfactual simulations that compare the predicted vote shares when all individuals in the RCT sample are treated versus when none are treated. Standard errors are in parentheses. Results are discussed in Section 5. Control estimates are in Table D1 and region fixed effect estimates are in Table D2.

Table 7: Counterfactual Experiments

Panel A: RCT-Sample							
Counterfactual Name	Description	Vote-Share			Votes		
		Est (p.p.)	Diff (p.p.)	p-value	Est	Diff	p-value
Baseline	No messages sent to voters	23.846 (2.266)			1618.866 (154.610)		
Gender - female	Gender messages sent to all female voters	23.620 (2.266)	-0.227 (0.434)	0.601	1603.431 (154.683)	-15.435 (29.589)	0.602
Gender - male	Gender messages sent to all male voters	24.202 (2.268)	0.356*** (0.097)	0.000	1642.324 (154.777)	23.458*** (6.464)	0.000
Gender	Gender messages sent to all voters	23.976 (2.267)	0.129 (0.438)	0.768	1626.889 (154.768)	8.023 (29.865)	0.788
Informative ability	Informative ability message sent to all voters	23.688 (2.305)	-0.159 (0.131)	0.226	1609.258 (157.247)	-9.608 (8.803)	0.275
Uninformative ability	Uninformative ability message sent to all voters	23.896 (2.290)	0.049 (0.050)	0.327	1623.423 (156.116)	4.557 (3.584)	0.204
Informative policy	Informative policy message sent to all voters	24.383 (2.188)	0.536** (0.262)	0.041	1655.860 (149.278)	36.994** (17.850)	0.038
Uninformative policy	Uninformative policy message sent to all voters	22.842 (2.328)	-1.005*** (0.161)	0.000	1550.894 (158.808)	-67.972*** (10.933)	0.000
Salience message	Only ability and policy uninformative messages sent	22.872 (2.351)	-0.975*** (0.177)	0.000	1554.015 (160.245)	-64.851*** (12.005)	0.000
Belief messages	Only ability and policy informative messages sent	24.238 (2.224)	0.391 (0.275)	0.155	1647.084 (151.693)	28.218 (18.668)	0.131
Only males	Only male voters are sent all messages	24.020 (2.269)	0.174*** (0.056)	0.002	1630.279 (154.788)	11.413*** (3.754)	0.002
Only females	Only female voters are sent all messages	24.106 (2.217)	0.259 (0.492)	0.599	1637.866 (151.320)	19.000 (33.542)	0.571
All treatments	All messages sent	24.280 (2.219)	0.433 (0.493)	0.380	1649.279 (151.467)	30.413 (33.588)	0.365
Aggregate optimal	Optimal campaign for the RCT-sample	24.898 (2.253)	1.051*** (0.281)	0.000	1691.042 (153.878)	72.176*** (19.090)	0.000
Municipio-wise optimal	Municipio-wise optimal campaign	25.309 (2.218)	1.463*** (0.281)	0.000	1718.593 (151.452)	99.727*** (19.150)	0.000
Municipio-wise v. aggregate	Difference b/w municipio-wise and national		0.411*** (0.116)	0.000		27.551*** (7.736)	0.000

Panel B: Full Brazil Sample							
Counterfactual Name	Description	Vote-Share			Votes		
		Est (p.p.)	Diff (p.p.)	p-value	Est	Diff	p-value
Baseline	No messages sent to voters	24.005 (2.266)			24.005 (2.266)		
Gender - female	Gender messages sent to all female voters	23.777 (2.264)	-0.228 (0.434)	0.600	6600.542 (701.237)	-65.232 (124.377)	0.600
Gender - male	Gender messages sent to all male voters	24.341 (2.269)	0.336*** (0.091)	0.000	6739.532 (705.808)	73.759*** (20.154)	0.000
Gender	Gender messages sent to all voters	24.114 (2.265)	0.109 (0.438)	0.804	24.114 (2.265)	0.109 (0.438)	0.804
Informative ability	Informative ability message sent to all voters	23.828 (2.305)	-0.177 (0.134)	0.187	23.828 (2.305)	-0.177 (0.134)	0.187
Uninformative ability	Uninformative ability message sent to all voters	24.038 (2.292)	0.033 (0.050)	0.513	24.038 (2.292)	0.033 (0.050)	0.513
Informative policy	Informative policy message sent to all voters	24.533 (2.190)	0.528** (0.263)	0.045	24.533 (2.190)	0.528** (0.263)	0.045
Uninformative policy	Uninformative policy message sent to all voters	22.993 (2.329)	-1.012*** (0.163)	0.000	22.993 (2.329)	-1.012*** (0.163)	0.000
Salience message	Only ability and policy uninformative messages sent	23.021 (2.351)	-0.984*** (0.178)	0.000	23.021 (2.351)	-0.984*** (0.178)	0.000
Belief messages	Only ability and policy informative messages sent	24.372 (2.225)	0.366 (0.275)	0.184	24.372 (2.225)	0.366 (0.275)	0.184
Only males	Only male voters are sent all messages	24.173 (2.268)	0.168*** (0.053)	0.002	24.173 (2.268)	0.168*** (0.053)	0.002
Only females	Only female voters are sent all messages	24.237 (2.216)	0.232 (0.494)	0.639	24.237 (2.216)	0.232 (0.494)	0.639
All treatments	All messages sent	24.405 (2.218)	0.400 (0.494)	0.418	24.405 (2.218)	0.400 (0.494)	0.418
Aggregate optimal	Optimal campaign for the full sample	25.015 (2.807)	1.011*** (0.304)	0.001	6842.793 (885.929)	177.019** (85.276)	0.038
Municipio-wise optimal	Municipio-wise optimal campaign	25.412 (2.763)	1.408*** (0.293)	0.000	6956.087 (877.416)	290.313*** (74.544)	0.000
Municipio-wise v. aggregate	Difference b/w municipio-wise and national		0.397*** (0.118)	0.001		113.294*** (36.178)	0.002

*, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.
Notes: This table reports results from counterfactual simulations evaluating the impact of alternative electoral messaging campaigns on female candidate vote shares and total votes. Panel A focuses on municipalities in the RCT sample, while Panel B presents results for the full set of Brazilian municipalities. Each counterfactual simulates voter behavior under a different campaign strategy using the estimated structural model. The "Est." columns report predicted average vote share and vote count outcomes under each campaign, while the "Diff." columns show the change relative to the baseline scenario with no messaging. Campaigns include gender-targeted, ability-targeted, and policy-targeted messages, as well as bundled salience- or belief-based messages. Standard errors, reported in parentheses, are computed via model-based simulation using 400 draws for all scenarios, except for the optimal campaigns, which use 200 draws. All simulations incorporate the model's estimated heterogeneity in voter preferences and salience weights. Results are discussed in Section 6. Persuasion rate estimates are in Table D4.

Online Appendix to
“Decoupling Taste-based versus Statistical
Discrimination in Elections”

Not for Publication

A Appendix: GMM Estimator

The estimator that we employ is a Generalized Method of Moments estimator defined as follows. The econometrician observes $\{v_m, W_m\}_{m=1}^M$ where $W_m = \{G_j, \{G_{i,m}, \mathcal{T}_{i,m}\}_{i=1}^{N_m}, X_m\}$, N_m is the number of voters in municipality m ¹⁸. Also, x_m is the number of votes female candidates receives and X_m is a vector of observable municipality and ballot characteristics. This is data.

The estimation process has two steps. The first step of the GMM estimator proceeds as follows. Fix a $\theta = (\omega_0^G, \omega_1^G, \sigma^G, \lambda_{G,0}, \lambda_{G,1}, \omega^A, \sigma^A, \lambda_A, \omega^P, \sigma^P, \lambda_P, \xi^A, \xi^P, \rho_A, \rho_P, \eta^A, \eta^P, \mu) \in \Theta$ and then for each $m \in \{1, 2, \dots, M\}$.

1. Consider the set $\{(G, \mathcal{T}) : G \in \{0, 1\} \text{ and } \mathcal{T} \in \{1, 2, \dots, 6\}\}$. This set has 12 elements. Define $N_{G,\mathcal{T},m}$, as the numbers of voters that have gender G and received the treatment \mathcal{T} , i.e.:

$$N_{G,\mathcal{T},m} = \sum_{i=1}^{N_m} \mathbf{1}\{G_{i,m} = G, \mathcal{T}_{i,m} = \mathcal{T}\}$$

2. For each (G, \mathcal{T}) , construct indicators $T^G = \mathbf{1}\{\mathcal{T} = 2\}$, $T^A = \mathbf{1}\{\mathcal{T} = 3\}$, $V^A = \mathbf{1}\{\mathcal{T} = 4\}$, $T^P = \mathbf{1}\{\mathcal{T} = 5\}$, $V^P = \mathbf{1}\{\mathcal{T} = 6\}$.
3. Do the following if $N_{G,\mathcal{T},m} > 0$:
 - (a) Draw shocks $v_{v,G,\mathcal{T}}^k \sim N(0, 1)$ for each $k \in \{G, A, P\}$ for $v = 1, 2, \dots, V$ total draws.
 - (b) Calculate $\tilde{u}_{v,G,\mathcal{T},m}$ using:

$$\begin{aligned} \tilde{u}_{v,G,\mathcal{T},j,m}(\theta) = & - \exp \left(\sum_{g \in \{0,1\}} \omega_g^G \cdot \mathbf{1}\{G = g\} + \sigma_G \cdot v_{v,G,\mathcal{T}}^G \right. \\ & \left. + \sum_{g \in \{0,1\}} \lambda_{G,g} \cdot V^G \cdot \mathbf{1}\{G = g\} \right) \times \mathbf{1}\{G \neq G_j\} \\ & + \exp \left(\omega^A + \sigma^A v_{v,G,\mathcal{T}}^A + \lambda^A \max\{T^A, V^A\} \right) \\ & \times \left(\xi^A G_j + \rho_A T^A G_j + \eta^A X_m \right) \\ & - \exp \left(\omega^P + \sigma^P v_{v,G,\mathcal{T}}^P + \lambda^P \max\{T^P, V^P\} \right) \\ & \times \left(\xi^P G_j + \rho_P T^A G_j + \eta^P X_m - \mu G \right)^2 \end{aligned}$$

- (c) Let $p_{v,i,m}$ denote the probability voter i that got a draw v votes for a female candidate:

¹⁸Please note that $V_{i,m}^G, T_{i,m}^A, V_{i,m}^A, T_{i,m}^P, V_{i,m}^P$ can be derived from $\mathcal{T}_{i,m}$ in the same fashion as in Section 2.2.

(note $G_j = 1$ corresponds to the female candidate, $G = 1$ female voter)

$$p_{v,G,\mathcal{T},m}(\theta) = \frac{\exp(\tilde{u}_{v,G,\mathcal{T},1,m}(\theta) - \tilde{u}_{v,G,\mathcal{T},0,m}(\theta))}{1 + \exp(\tilde{u}_{v,G,\mathcal{T},1,m}(\theta) - \tilde{u}_{v,G,\mathcal{T},0,m}(\theta))}$$

$$p_{G,\mathcal{T},m}(\theta) = \frac{1}{V} \sum_{v=1}^V \frac{\exp(\tilde{u}_{v,G,\mathcal{T},1,m}(\theta) - \tilde{u}_{v,G,\mathcal{T},0,m}(\theta))}{1 + \exp(\tilde{u}_{v,G,\mathcal{T},1,m}(\theta) - \tilde{u}_{v,G,\mathcal{T},0,m}(\theta))}$$

4. If $N_{G,\mathcal{T},m} = 0$ then $p_{G,\mathcal{T},m} = 0$. The value here does not matter as the product $N_{G,\mathcal{T},m} \times p_{G,\mathcal{T},m} = 0$.

5. Calculate $\lambda_m(\theta)$ (calculated for Female candidate only) as follows:

$$\lambda_m(\theta) = \sum_{G=0}^1 \sum_{\mathcal{T}=1}^6 N_{G,\mathcal{T},m} \times p_{G,\mathcal{T},m}$$

The second step of the GMM estimator is then to minimize

$$d(\theta) = \frac{1}{M} \times \sum_{m=1}^M \left(\frac{x_m - \lambda_m(\theta)}{N_m} \right)^2$$

B Appendix: Monte-Carlo Experiments

B.1 Baseline case

Here, we demonstrate the Monte Carlo performance of our model. To generate the data, we use the true observable characteristics of municipalities in our RCT sample, ensuring that the number of treated individuals matches the actual number observed in each municipality. We estimate the model for three potential sample sizes: $N=250$, $N=500$, and $N=1,000$. The samples for $N=250$ and $N=500$ are constructed by randomly selecting observations. The parameter values are set to the estimated values obtained from our model, allowing us to assess the bias and mean squared error (MSE) at these estimated values. For brevity, we report only the estimates for the main parameters, as the performance for controls and fixed effects closely aligns with that of the main parameters.

The Monte Carlo results indicate that bias and MSE consistently decrease as the sample size increases. For instance, the bias of the baseline weight ω_{GF} decreases from 0.0593 at $N=250$ to just 0.0255 at $N=1,000$, while the corresponding MSE remains low at 0.0128. Similarly, the bias for the effect of the G message on males (λ_{GM}) reduces from -0.0012 at $N=250$ to a negligible -0.0006 at $N=1,000$, with MSE decreasing from 0.0003 to effectively zero. The effect of uninformative ability messages (λ_A) shows nearly no bias across all sample sizes, with MSE consistently at 0.0000. Additionally, for the relative policy preference of female voters (μ), bias decreases from -0.0261 at $N=250$ to just 0.0011 at $N=1,000$, and MSE stabilizes at 0.0166. Overall, these results highlight the strong Monte Carlo performance of our model and estimator, with bias and MSE becoming negligible at $N=1,000$, demonstrating the robustness and reliability of our estimation approach.

Table B1: Monte-Carlo Experiments

		Avg. Est.			Bias			MSE			
		N=250	N=500	N=1000	N=250	N=500	N=1000	N=250	N=500	N=1000	
ω_{GF}	Baseline weight gender (F)	4.088	4.147	4.140	4.113	0.0593	0.0528	0.0255	0.0153	0.0104	0.0128
ω_{GM}	Baseline weight gender (M)	6.233	6.295	6.292	6.259	0.0620	0.0587	0.0263	0.0169	0.0118	0.0134
λ_{GF}	Effect of gender message (F)	-0.018	-0.017	-0.018	-0.018	0.0003	0.0001	-0.0001	0.0000	0.0000	0.0000
λ_{GM}	Effect of gender message (M)	-0.533	-0.535	-0.538	-0.534	-0.0012	-0.0044	-0.0006	0.0003	0.0002	0.0000
ω_{A0}	Baseline weight ability	2.150	2.205	2.205	2.174	0.0554	0.0553	0.0248	0.0155	0.0125	0.0113
λ_A	Effect of uninfo ability message	-0.162	-0.162	-0.162	-0.162	0.0004	0.0003	-0.0002	0.0000	0.0000	0.0000
ω_{P0}	Baseline weight policy	1.343	1.386	1.375	1.369	0.0426	0.0316	0.0255	0.0151	0.0113	0.0121
λ_P	Effect of uninfo policy message	0.083	0.083	0.083	0.083	0.0001	0.0001	0.0003	0.0000	0.0000	0.0000
ξ_A	Baseline net ability (F vs M)	0.120	0.122	0.120	0.121	0.0016	0.0005	0.0007	0.0004	0.0001	0.0001
ρ_A	Effect of ability info	-0.092	-0.088	-0.090	-0.092	0.0034	0.0017	-0.0005	0.0002	0.0001	0.0001
ξ_P	Baseline net policy (F vs M)	2.157	2.178	2.181	2.160	0.0212	0.0244	0.0031	0.0082	0.0082	0.0077
ρ_P	Effect of policy info	-0.192	-0.194	-0.195	-0.192	-0.0023	-0.0027	-0.0006	0.0001	0.0001	0.0001
μ	Relative policy (F voters)	-3.193	-3.219	-3.231	-3.192	-0.0261	-0.0376	0.0011	0.0186	0.0165	0.0166

B.2 Robustness to Spillovers

To assess the robustness of our parameter estimates to spillovers, we introduce a modification to our data-generating process. We retain the same covariates and "true reach" as in the baseline scenario. However, during model estimation, we construct the "observed reach" by incorporating stochastic deviations. Specifically, for each municipality and treatment status, we generate two random variables: a) A Bernoulli random variable that determines whether the observed reach is greater or smaller than the true reach. b) A Poisson random variable that quantifies the magnitude of the deviation between the observed reach and the true reach.

Formally, this approach allows us to systematically introduce spillover-induced measurement error, enabling us to evaluate how well our estimator performs under varying degrees of mismeasurement. $N_{m,Treat}^{obs} = N_{m,Treat}^{true} + (-1)^{B_{m,Treat}+1} \cdot P_{m,Treat}$ such that $B_{m,Treat} \sim Bern(0.5)$ and $P_{m,Treat} \sim Poisson(p \cdot N_{m,Treat}^{true})$. Here p controls the mean absolute deviation of observed reach from true reach.

Note that spillovers ultimately introduce measurement error in our reach variable, which defines the number of individuals treated by our message. When voters move from one municipality (likely outside our sample, given the spatially scattered nature of our dataset) to a municipality within our sample, they may receive a treatment message that was not intended for them. However, our data would still record that someone in the municipality was treated, even though the individual is not a registered voter there. This discrepancy introduces measurement error, which we explicitly model in our analysis.

We assess the Monte Carlo performance of our estimator at the estimated parameters under varying spillover rates: $p = 0.05, 0.10, 0.15, 0.20$. Overall, the estimates remain robust for spillover rates up to 15%, with minimal impact on bias and MSE. For instance, the bias in the baseline weight for female voters (ω_{GF}) remains stable at 0.0255 for spillover rates up to 15%, with a consistently low MSE of 0.0128. Similarly, the effect of the G message on female voters (λ_{GF}) shows an almost negligible bias (-0.0001) across all cases, with MSE remaining at 0.0000. However, at $p = 20\%$, some parameters exhibit notable deviations. The bias in the baseline weight for male voters (ω_{GM}) increases substantially to -0.6835, with a sharp rise in MSE to 5.7213. Likewise, the baseline net ability (ξ_A) experiences a jump in bias (0.2459) and a corresponding increase in MSE (0.8980).

While these higher spillover rates introduce noticeable increases in bias and MSE for some parameters, overall, the errors remain within acceptable limits. This suggests that our estimator performs well under moderate levels of spillover. When spillover rates exceed 15%, as they may

introduce nontrivial distortions in parameter estimates however these distortions still remain within some degree of acceptable range.

Table B2: Robustness to Spillovers

		Avg. Est.				Bias				MSE				
		0.05	0.10	0.15	0.20	0.05	0.10	0.15	0.20	0.05	0.10	0.15	0.20	
ω_{GF}	Baseline weight gender (F)	4.088	4.113	4.113	4.113	4.335	0.0255	0.0255	0.0255	0.2479	0.0128	0.0128	0.0128	0.0908
ω_{GM}	Baseline weight gender (M)	6.233	6.259	6.259	6.259	5.549	0.0263	0.0263	0.0263	-0.6835	0.0134	0.0134	0.0134	5.7213
λ_{GF}	Effect of gender message (F)	-0.018	-0.018	-0.018	-0.018	-0.022	-0.0001	-0.0001	-0.0001	-0.0042	0.0000	0.0000	0.0000	0.0001
λ_{GM}	Effect of gender message (M)	-0.533	-0.534	-0.534	-0.534	-0.492	-0.0006	-0.0006	-0.0006	0.0415	0.0000	0.0000	0.0000	1.3885
ω_{A0}	Baseline weight ability	2.150	2.174	2.174	2.174	2.201	0.0248	0.0248	0.0248	0.0515	0.0113	0.0113	0.0113	0.4163
λ_A	Effect of uninfo ability message	-0.162	-0.162	-0.162	-0.162	-0.136	-0.0002	-0.0002	-0.0002	0.0257	0.0000	0.0000	0.0000	0.0012
ω_{P0}	Baseline weight policy	1.343	1.369	1.369	1.369	1.604	0.0255	0.0255	0.0255	0.2604	0.0121	0.0121	0.0121	0.3486
λ_P	Effect of uninfo policy message	0.083	0.083	0.083	0.083	0.087	0.0003	0.0003	0.0003	0.0042	0.0000	0.0000	0.0000	0.0001
ξ_A	Baseline net ability (F vs M)	0.120	0.121	0.121	0.121	0.366	0.0007	0.0007	0.0007	0.2459	0.0001	0.0001	0.0001	0.8980
ρ_A	Effect of ability info	-0.092	-0.092	-0.092	-0.092	-0.173	-0.0005	-0.0005	-0.0005	-0.0810	0.0001	0.0001	0.0001	0.0176
ξ_P	Baseline net policy (F vs M)	2.157	2.160	2.160	2.160	2.431	0.0031	0.0031	0.0031	0.2740	0.0077	0.0077	0.0077	0.3112
ρ_P	Effect of policy info	-0.192	-0.192	-0.192	-0.192	-0.179	-0.0006	-0.0006	-0.0006	0.0132	0.0001	0.0001	0.0001	0.0015
μ	Relative policy (F voters)	-3.193	-3.192	-3.192	-3.192	-3.124	0.0011	0.0011	0.0011	0.0694	0.0166	0.0166	0.0166	1.6343

C Appendix: Digital Campaign Execution Details

This appendix provides supplementary details regarding the implementation, execution, and monitoring of the digital intervention, complementing the description in Section 3.2.1. The campaign was executed through Meta Ads, primarily on Instagram, with a focus on maximizing reach among voters in targeted municipalities. A specialized advertising firm was contracted to manage the technical execution of the campaign, ensuring proper implementation of ad targeting, budget allocation, and performance monitoring. The intervention was conducted between September 28 and October 3, 2024, ensuring compliance with electoral regulations that prohibit political ads within 48 hours of election day (October 6).

The campaign targeted users based on municipality, gender, and age group, using Meta Ads' geolocation tools to ensure that advertisements were displayed exclusively within designated treatment municipalities. The segmentation strategy employed municipality-level geofencing within the Meta Ads platform, relying on the first five digits of each postal code to define geographic boundaries. Age-based segmentation covered voters aged 18 to 65 years or older. The advertisements were delivered automatically across multiple placements within the Instagram platform, including feed, stories, and reels.

Ads were implemented using Meta Ads' algorithm designed to maximize exposure within the defined geolocation parameters. The total budget for the campaign was R\$52,000 (around USD 9,300), distributed across 859 treated municipalities, with daily spending limits varying according to population size and digital penetration levels in each municipality. The campaign creative strategy relied on five distinct video advertisements, which visual elements were designed to maintain consistency across different treatment groups while ensuring clarity and salience in message delivery.

Ethical and regulatory compliance was a central component of the intervention. The campaign adhered to Meta's policies regarding political advertising, ensuring transparency in message dissemination. To mitigate the risk of ad delivery disruptions or account restrictions, a contingency structure was implemented prior to campaign launch. This included obtaining the required political ad authorization from Meta, creating redundant advertising accounts and social media pages, and executing a pre-warming phase in which the newly created pages were actively maintained for 30 days before the campaign commenced. These measures minimized the likelihood of algorithmic flagging or restrictions due to unusual activity patterns.

The campaign setup followed a structured process, with ad configurations uploaded via spreadsheet integration into Meta Ads Manager. A subsequent quality control review was con-

ducted to verify the accuracy of ad placements, budget allocations, and geolocation targeting. This verification process was handled independently by a separate analyst to reduce the probability of configuration errors. Performance monitoring was conducted daily using a comprehensive set of metrics, with the primary indicator being the reach within each municipality. Additional metrics included total impressions, unique reach, and cost per thousand impressions, allowing for an evaluation of ad efficiency and engagement levels. Interaction data, including video views, reactions, shares, and comments, were also tracked to assess audience responsiveness.

C.1 Messages

Portuguese original:

- **Gender identity:** Sabia que mulheres são mais de 50% da população, mas são só 16% do Congresso Nacional? Ao redor do mundo, apenas 27% dos parlamentares são mulheres.¹⁹ Quem merece seu voto nestas eleições? Uma candidata mulher ou um candidato homem? Pense bem nisso.

- **Policy uninformative:** O que é importante para você nesta eleição? Educação, saúde, qualidade de vida para as crianças? Vote em candidatos que realmente defendem o que é importante para você todos os dias.

- **Policy informative:** O que é importante para você nesta eleição? Educação, saúde, qualidade de vida para as crianças? Você sabia que estudos mostram que parlamentares mulheres investem 77% mais em cuidados infantis,²⁰ educação e saúde²¹ do que políticos homens? Vote em candidatos que realmente defendem o que é importante para você todos os dias.

- **Ability uninformative:** O que é importante para você nesta eleição? Políticos competentes e capacitados que trabalham duro para melhorar seu governo local e sua comunidade. Vote em candidatos que têm a qualidade que você exige.

- **Ability informative:** O que é importante para você nesta eleição? Políticos competentes e capacitados que trabalham duro para melhorar seu governo local e sua comunidade. Você

¹⁹Inter-Parliamentary Union. Women in national parliaments, as of 1st January 2024

²⁰K. A. Bratton and L. P. Ray. 2002. "Descriptive representation: Policy outcomes and municipal day-care coverage in Norway," *American Journal of Political Science*, 46(2), pp. 428–437.

²¹R. Chattopadhyay and E. Duflo (2004). "Women as policy makers: Evidence from a randomized policy experiment in India," *Econometrica* 72(5), pp. 1409–1443;

Gerrity JC, Osborn T, Mendez JM. 2007. Women and representation: a different view of the district? *Polit. Gender* 3:2179–200

sabia que estudos mostram que, em média, políticas mulheres têm maior qualidade,²² são mais competentes e trabalham mais²³ do que os políticos homens? Vote em candidatos que têm a qualidade que você exige.

English translation:

- **Gender identity:** Did you know that women make up more than 50% the population, but they represent only 16% of the National Congress? Around the world only 27% of parliamentarians are women. Who deserves your vote in these elections? A female candidate or a male candidate? Think carefully about this.

- **Policy uninformative:** What is important to you in this election? Education, health care, child welfare? Vote for candidates who truly defend what is important for you every day.

- **Policy informative:** What is important to you in this election? Education, health care, child welfare? Did you know that studies show that female parliamentarians invest 77% more on childcare, education, and health care than male politicians? Vote for candidates who truly defend what is important for you every day.

- **Ability uninformative:** What is important to you in this election? Competent and qualified politicians who work hard to improve your local government and community. Vote for candidates who meet the quality you demand.

- **Ability informative:** What is important to you in this election? Competent and qualified politicians who work hard to improve your local government and community. Did you know that studies show that, on average, female politicians are of higher quality, more competent, and work harder than their male counterparts? Vote for candidates who meet the quality you demand.

²²Baltrunaite, Audinga, Piera Bello, Alessandra Casarico, and Paola Profeta. "Gender quotas and the quality of politicians." *Journal of Public Economics* 118 (2014): 62-74.

²³Anzia, Sarah E., and Christopher R. Berry. "The Jackie (and Jill) Robinson effect: Why do congresswomen outperform congressmen?" *American Journal of Political Science* 55, no. 3 (2011): 478-493.



(a)



(b)



(c)



(d)

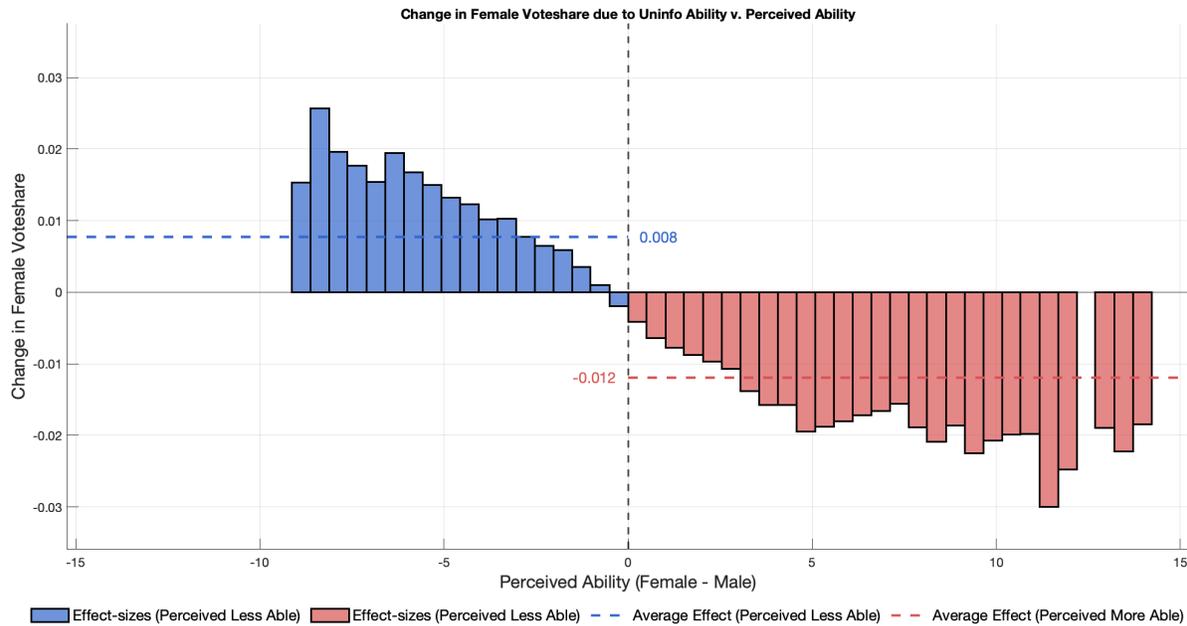


(e)

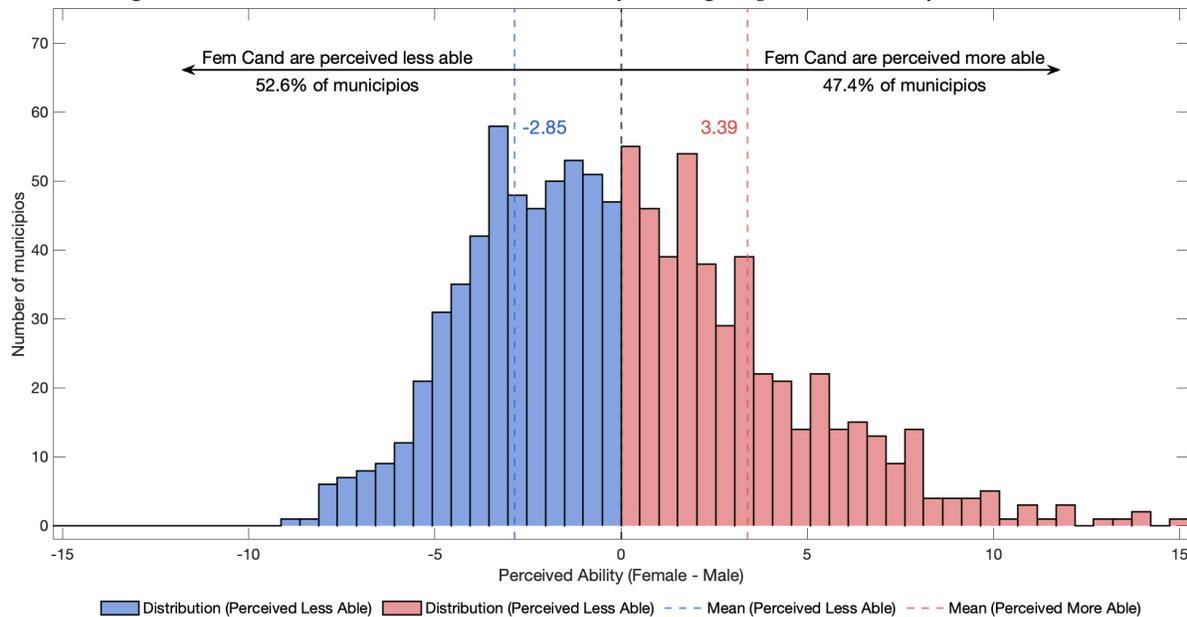
Figure C1: Screenshots of the Instagram treatment ads: Panel (a) shows the gender message. Panels (b) and (c) show the uninformed message and informative ability component respectively. Panels (d) and (e) show the uninformed message and informative policy component respectively.

D Supplemental Figures

This Appendix reports additional figures and tables supplementary to the main text.



(a) Change in vote-shares due to uninformative ability message v. perceived ability of female candidates



(b) Perceived relative ability of female candidates

Figure D1: This figures illustrate perceived relative ability of female candidates and the change in vote-shares due to uninformative ability message.

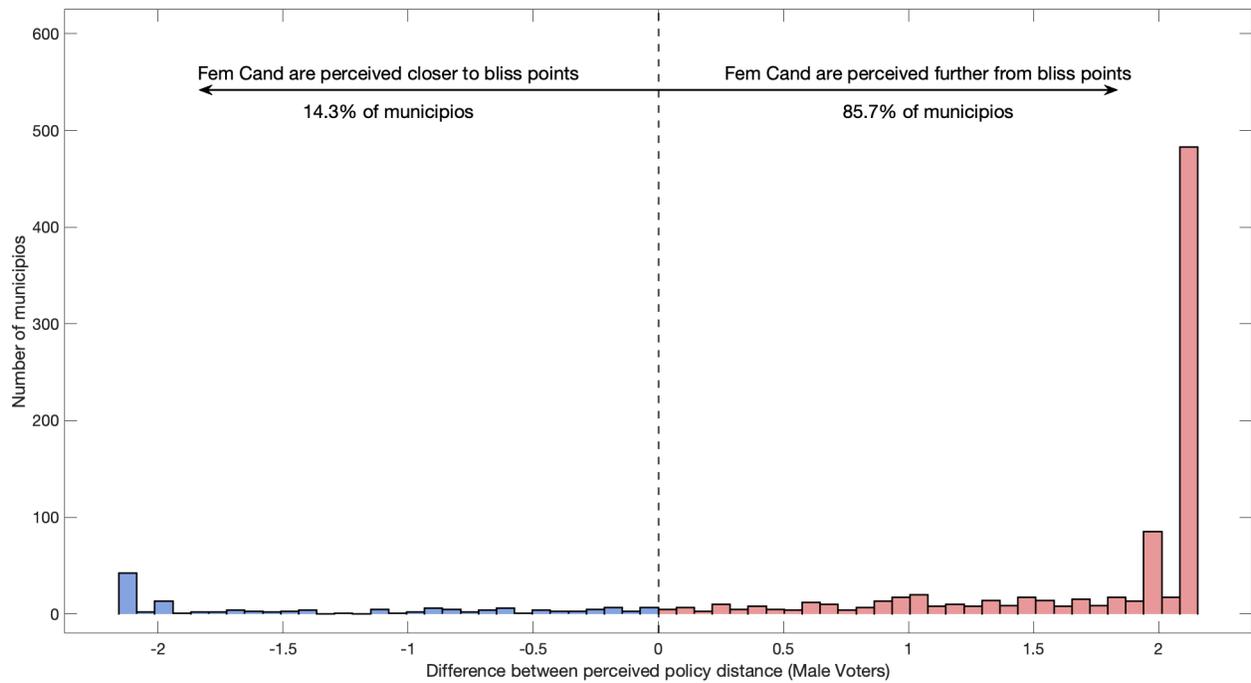


Figure D2: This figures illustrate difference in (a) distance between male voter blisspoint & female candidate platform positions and (b) distance between male voter blisspoint & male candidate platform positions.

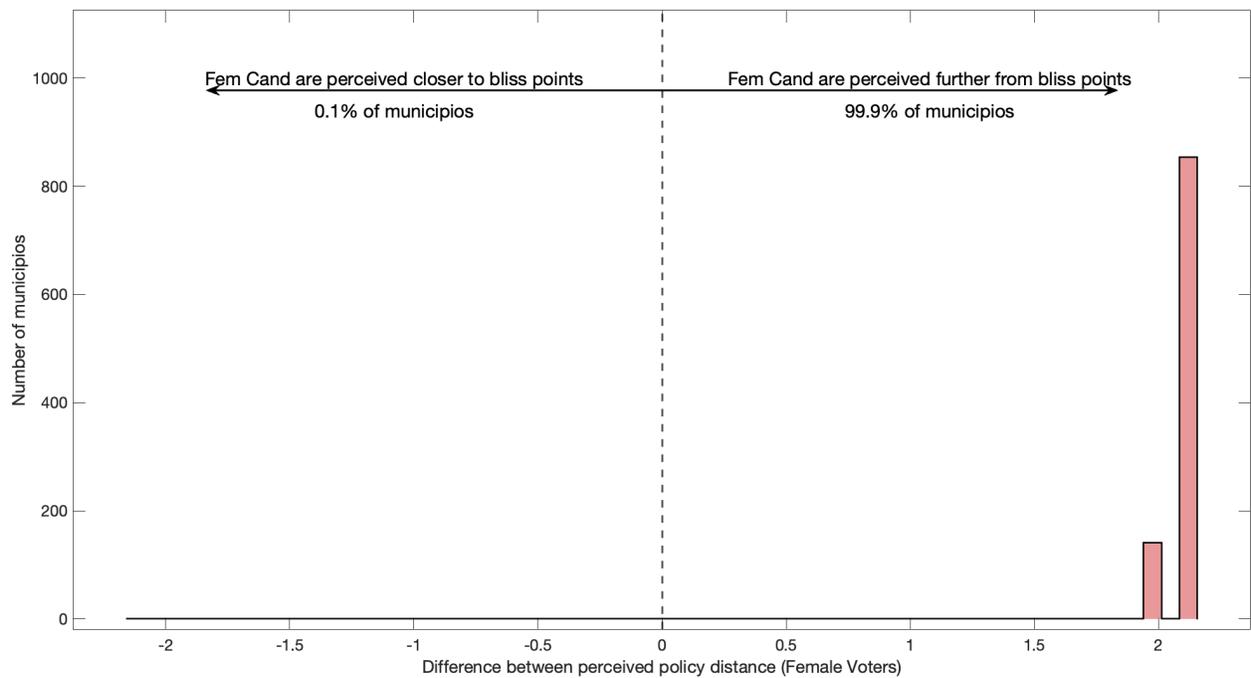


Figure D3: This figures illustrate difference in (a) distance between female voter blisspoint & female candidate platform positions and (b) distance between female voter blisspoint & male candidate platform positions.

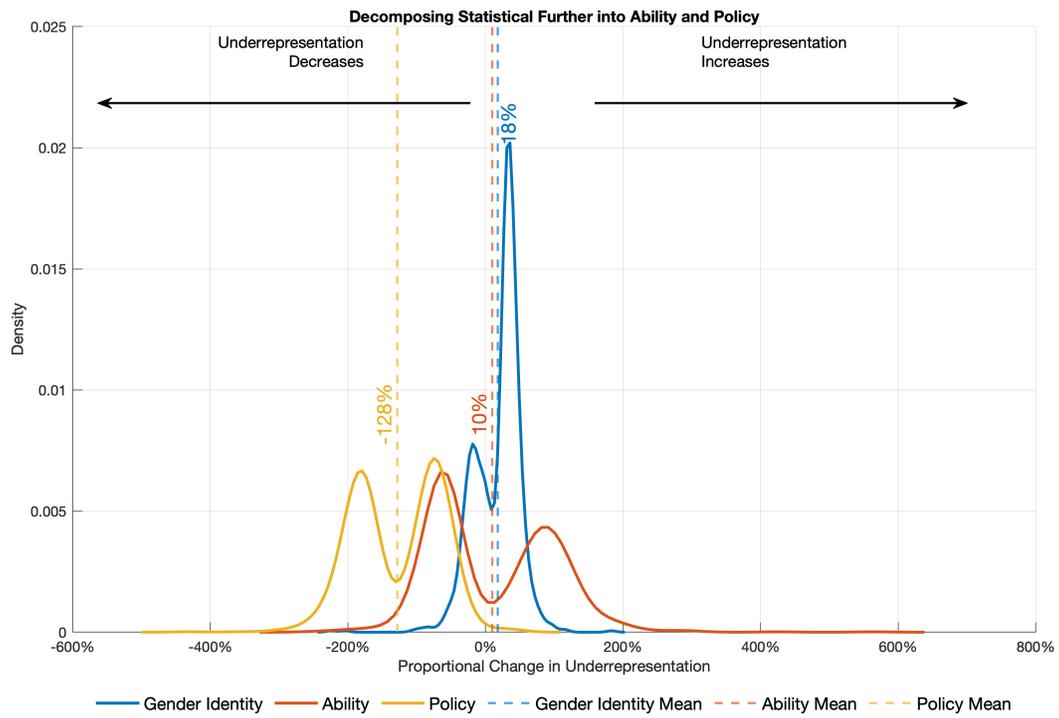


Figure D4: Here, we further decompose statistical discrimination into ability-based and policy-based statistical discrimination by following a similar set of steps as in Figure 4.

Table D1: Parameter Estimates for Controls

Parameter	Description	Models				
		(1)	(2)	(3)	(4)	(5)
η_{A1}	Ability coeff: relative wealth (Cand)	0.530* (0.276)	1.102** (0.466)	0.865 (0.535)	1.590** (0.620)	-53.098*** (0.286)
η_{A2}	Ability coeff: prop of F politicians (Cand)	1.669* (0.864)	1.195** (0.483)	4.294* (2.596)	12.413*** (1.027)	-178.171*** (1.533)
η_{A3}	Ability coeff: F vs M college degree (Cand)	-0.374* (0.198)	-1.000** (0.428)	1.949* (1.179)	5.451*** (1.933)	-127.577*** (0.643)
η_{A4}	Ability coeff: F vs M black (Cand)	-0.276* (0.147)	-0.291** (0.113)	-3.472* (2.101)	-2.485 (1.922)	141.948*** (0.717)
η_{A5}	Ability coeff: F vs M married (Cand)	-0.343* (0.184)	-0.337** (0.146)	3.574* (2.162)	8.454*** (0.936)	-220.818*** (1.023)
η_{A6}	Ability coeff: urban pop (Voters)	-1.020* (0.532)	-0.437** (0.170)	-	-6.564*** (1.028)	94.557*** (1.017)
η_{A7}	Ability coeff: lagged F vote share (Voters)	-0.195* (0.110)	-0.178* (0.100)	-	3.237*** (0.616)	18.146*** (0.342)
η_{A8}	Ability coeff: pop <=25 yrs (Voters)	-3.886* (2.037)	-5.180** (2.162)	-	-9.028*** (1.773)	221.220*** (0.768)
η_{A9}	Ability coeff: pop >60 yrs (Voters)	-4.733* (2.477)	-6.308** (2.626)	-	-0.236 (1.391)	141.972*** (0.758)
η_{A10}	Ability coeff: GDP per capita (Voters)	0.060 (0.038)	-0.541*** (0.207)	-	0.689 (1.194)	-151.718*** (1.063)
η_{A11}	Ability coeff: black voters (Voters)	-2.965* (1.538)	-2.829** (1.167)	-	-12.099*** (1.161)	-288.190*** (1.910)
η_{P1}	Policy coeff: relative wealth (Cand)	0.050 (0.038)	-0.381*** (0.108)	0.035*** (0.013)	0.149*** (0.043)	-1.583*** (0.008)
η_{P2}	Policy coeff: prop of F politicians (Cand)	0.124*** (0.043)	0.048** (0.023)	0.264*** (0.079)	-0.098*** (0.037)	-5.700*** (0.022)
η_{P3}	Policy coeff: F vs M college degree (Cand)	-0.077*** (0.028)	0.419*** (0.120)	0.553*** (0.170)	-0.541*** (0.087)	-1.752*** (0.007)
η_{P4}	Policy coeff: F vs M black (Cand)	-0.266*** (0.091)	0.424*** (0.118)	-1.237*** (0.368)	0.658*** (0.048)	0.823*** (0.006)
η_{P5}	Policy coeff: F vs M married (Cand)	-0.253*** (0.074)	0.360*** (0.106)	0.530*** (0.156)	-0.348*** (0.030)	-4.710*** (0.016)
η_{P6}	Policy coeff: urban pop (Voters)	-0.314*** (0.097)	0.187*** (0.053)	-	0.501*** (0.082)	1.577*** (0.018)
η_{P7}	Policy coeff: lagged F vote share (Voters)	-0.098*** (0.037)	0.112*** (0.027)	-	-0.184** (0.076)	-0.073*** (0.005)
η_{P8}	Policy coeff: pop <=25 yrs (Voters)	-1.172*** (0.348)	2.328*** (0.641)	-	0.656*** (0.088)	4.140*** (0.034)
η_{P9}	Policy coeff: pop >60 yrs (Voters)	-1.430*** (0.438)	2.823*** (0.775)	-	-0.078 (0.134)	2.809*** (0.033)
η_{P10}	Policy coeff: GDP per capita (Voters)	0.045** (0.022)	0.212*** (0.060)	-	-0.019 (0.166)	-2.506*** (0.010)
η_{P11}	Policy coeff: black voters (Voters)	-0.779*** (0.216)	1.211*** (0.343)	-	0.726*** (0.047)	-4.237*** (0.023)
Region Fixed Effect		✓	-	✓	✓	✓
Candidate controls		✓	✓	✓	✓	✓
Municipio Characteristics		✓	✓	-	✓	✓
Salience-Weight functional Form		Exponential	Exponential	Exponential	Quadratic	Absolute
N	Number of Observations	1000	1000	1000	1000	1000
Obj Fun	Objective Function Value (MSE)	0.004	0.0041	0.0042	0.004	0.004

Notes: This table presents the estimated parameters of controls used in our random utility voting model. Parameters are estimated using GMM and capture both statistical and taste-based discrimination. Column (1) is the baseline specification, which includes candidate- and municipality-level controls, region fixed effects, and exponential salience weights. Columns (2)–(5) report estimates from alternative specifications that remove fixed effects, controls, or modify the functional form of salience. Results are discussed in Section 5. Main parameter estimates are in Table 6 and region fixed effect estimates are in Table D2.

Table D2: Parameter Estimates for Fixed-Effects

Parameter	Description	Models				
		(1)	(2)	(3)	(4)	(5)
ω_{A1}	Ability weight FE (Norte)	-0.419*** (0.041)	-	-0.484*** (0.036)	-0.241*** (0.032)	-1.586*** (0.003)
ω_{A2}	Ability weight FE (Nordeste)	0.041** (0.021)	-	0.254*** (0.024)	0.465*** (0.061)	-0.265*** (0.003)
ω_{A3}	Ability weight FE (Sudeste)	-0.010 (0.019)	-	-0.038 (0.025)	-0.051 (0.032)	0.244*** (0.014)
ω_{A4}	Ability weight FE (Centro-Oeste)	0.183*** (0.027)	-	0.316*** (0.026)	-0.873*** (0.028)	-1.235*** (0.003)
ω_{P1}	Policy weight FE (Norte)	-0.017 (0.041)	-	0.057*** (0.011)	0.027 (0.038)	-0.324*** (0.001)
ω_{P2}	Policy weight FE (Nordeste)	-0.176*** (0.038)	-	-0.024** (0.009)	-0.607*** (0.039)	-1.108*** (0.001)
ω_{P3}	Policy weight FE (Sudeste)	-0.080*** (0.029)	-	-0.013 (0.012)	-0.083** (0.035)	-0.757*** (0.002)
ω_{P4}	Policy weight FE (Centro-Oeste)	-0.096** (0.046)	-	-0.058*** (0.011)	-0.480*** (0.018)	0.097*** (0.002)
Region Fixed Effect		✓	-	✓	✓	✓
Candidate controls		✓	✓	✓	✓	✓
Município Characteristics		✓	✓	-	✓	✓
Saliency-Weight functional Form		Exponential	Exponential	Exponential	Quadratic	Absolute
N	Number of Observations	1000	1000	1000	1000	1000
Obj Fun	Objective Function Value (MSE)	0.004	0.0041	0.0042	0.004	0.004

Notes: This table presents the estimated region fixed effects used in our random utility voting model. Parameters are estimated using GMM and capture both statistical and taste-based discrimination. Column (1) is the baseline specification, which includes candidate- and municipality-level controls, region fixed effects, and exponential saliency weights. Columns (2)–(5) report estimates from alternative specifications that remove fixed effects, controls, or modify the functional form of saliency. Results are discussed in Section 5. Main parameter estimates are in Table 6 and control parameter estimates are in Table D1.

Table D3: Changes in Weights and Utility

Treatment	Change in Weights			Change in Utility				
	Est	Diff	p value	Est	Diff	p value	$\frac{\text{Diff}}{\text{Dimension Utility}}$ (p.p.)	$\frac{\text{Diff}}{\text{Utility}}$ (p.p.)
Gender - female	95.017 (3.545)	-1.667 (3.190)	0.601	95.017 (3.545)	-1.667 (3.190)	0.601	-1.724	-0.5307
Gender - male	-409.307 (3.376)	174.625*** (4.880)	0.000	-409.307 (3.376)	174.625*** (4.880)	0.000	29.905	35.6093
Informative ability	12.325 (7.748)	-2.168 (1.540)	0.159	0.664 (0.563)	-1.444* (0.762)	0.058	-68.512	-0.4593
Uninformative ability	12.325 (7.748)	-2.168 (1.540)	0.159	1.792 (0.497)	-0.315*** (0.098)	0.001	-14.959	-0.0999
Informative policy	-6.310 (4.940)	-0.501 (0.369)	0.174	-72.090 (9.110)	2.493* (1.290)	0.053	3.343	0.7834
Uninformative policy	-6.310 (4.940)	-0.501 (0.369)	0.174	-81.022 (10.390)	-6.438*** (1.239)	0.000	-8.632	-2.0814

Notes: The table shows how model objects change in response to treatments. First, we examine whether significant parameter estimates of λ_{GF} , λ_{GM} , λ_A , and λ_P lead to significant changes in actual salience weights (recall that the weights are given by an exponential functional form). Only the gender treatment for males results in significant changes in weights. We also examine whether the treatments induce significant changes in voter utility. Except for the gender treatment for females, all treatments lead to significant changes in utility. However, not all changes are economically significant. In particular, for the uninformative ability treatment, although it induces a 15% change in utility from ability, this translates to only a 0.1% change in overall utility. The informative policy treatment has an effect roughly eight times larger, while the uninformative policy treatment induces a change approximately twenty times greater.

Table D4: Persuasion Rates

Counterfactual Name	Description	Vote-Share			Persuasion Rates		
		Est (p.p.)	Diff (p.p.)	p-value	Est (p.p.)	Diff (p.p.)	p-value
Baseline	No messages sent to voters	23.846 (2.266)			23.846 (2.266)		
Gender - Female	Gender messages sent to all female voters	23.620 (2.266)	-0.227 (0.434)	0.601	23.620 (2.266)	-0.298 (0.583)	0.610
Gender - Male	Gender messages sent to all male voters	24.202 (2.268)	0.356*** (0.097)	0.000	24.202 (2.268)	0.467*** (0.133)	0.000
Gender	Gender messages sent to all voters	23.976 (2.267)	0.129 (0.438)	0.768	23.976 (2.267)	0.169 (0.587)	0.773
Info Ability	Informative ability message sent to all voters	23.688 (2.305)	-0.159 (0.131)	0.226	23.688 (2.305)	-0.209 (0.179)	0.243
Uninformative Ability	Uninformative ability message sent to all voters	23.896 (2.290)	0.049 (0.050)	0.327	23.896 (2.290)	0.065 (0.076)	0.392
Informative Policy	Informative policy message sent to all voters	24.383 (2.188)	0.536** (0.262)	0.041	24.383 (2.188)	0.704** (0.358)	0.049
Uninformative Policy	Uninformative policy message sent to all voters	22.842 (2.328)	-1.005*** (0.161)	0.000	22.842 (2.328)	-1.319*** (0.209)	0.000
Salience Message	Only ability and policy uninformative messages sent	22.872 (2.351)	-0.975*** (0.177)	0.000	22.872 (2.351)	-1.280*** (0.232)	0.000
Belief Messages	Only ability and policy informative messages sent	24.238 (2.224)	0.391 (0.275)	0.155	24.238 (2.224)	0.514 (0.374)	0.170
Only Males	Only male voters are sent all messages	24.020 (2.269)	0.174*** (0.056)	0.002	24.020 (2.269)	0.228*** (0.077)	0.003
Only Females	Only female voters are sent all messages	24.106 (2.217)	0.259 (0.492)	0.599	24.106 (2.217)	0.340 (0.665)	0.609
All Treatments	All messages sent	24.280 (2.219)	0.433 (0.493)	0.380	24.280 (2.219)	0.569 (0.665)	0.392
Aggregate Optimal	Optimal campaign for the RCT-sample	24.898 (2.253)	1.051*** (0.281)	0.000	24.898 (2.253)	1.381*** (0.374)	0.000
Municipio-wise Optimal	Municipio-wise optimal campaign	25.309 (2.218)	1.463*** (0.281)	0.000	25.309 (2.218)	1.921*** (0.369)	0.000

*, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Notes: The table reports persuasion rates for counterfactual ad campaigns. The persuasion rate in this setting is calculated by allowing for full exposure/penetration and turnout since voting is mandatory. In our setting the persuasion rate takes the form : $\frac{\text{CF Female Vote Share} - \text{Baseline Female Vote Share}}{1 - \text{Baseline Female Vote Share}}$. Standard errors are calculated by simulating vote shares across 400 simulations for all counterfactuals except for the optimal campaign ones. For the optimal campaigns we rely on 200 simulations due to computational constraints.

Table D5: Structural Model v. OLS Simulations

Treatment	Data		Simulated			Simulated v. Data	
	OLS	Structural Model	OLS	Logit	Structural	OLS v. OLS	OLS v. Logit
Gender-F	0.011 (0.009)	-0.0023 (0.004)	0.0032	0.0038	-0.0022	-0.0078 (0.009)	-0.0072 (0.009)
Gender-M	0.018 (0.009)	0.0036 (0.001)	0.0114	0.0138	0.0030	-0.0066 (0.009)	-0.0042 (0.009)
Info-Ability	0.011 (0.009)	-0.0016 (0.001)	0.0049	0.0065	0.0006	-0.0061 (0.009)	-0.0045 (0.009)
Uninfo-Ability	0.003 (0.009)	0.0005 (0.001)	-0.0009	-0.0011	-0.0012	-0.0039 (0.009)	-0.0041 (0.009)
Info-Policy	-0.003 (0.010)	0.0054 (0.003)	-0.0093	-0.0132	-0.0101	-0.0063 (0.010)	-0.0102 (0.010)
Uninfo-Policy	0.014 (0.010)	-0.0100 (0.002)	0.0078	0.0093	0.0046	-0.0062 (0.010)	-0.0047 (0.010)

Notes: Standard errors are reported in parentheses. Column (1) reproduces estimates from Column (1) of Table 4. Column (2) presents average marginal effects from the structural model. Columns (3)–(5) show average estimates across 200 simulations where data are generated from the estimated model. Column (6) compares the OLS estimate in the real data to its simulated counterpart. Column (7) compares the real-data OLS estimate to the average marginal effect from a logit model estimated on simulated data.