ISSN 2282-6483

Alma Mater Studiorum - Università di Bologna
DEPARTMENT OF ECONOMICS

Beliefs about Gender Inequalities, Narratives and Support for Gender Quotas

Luca Di Corato
Federica Esposito
Natalia Montinari

Quaderni - Working Paper DSE N°1220

Beliefs about Gender Inequalities, Narratives and Support for Gender Quotas*

Luca Di Corato† Federica Esposito‡ Natalia Montinari§

Abstract

Gender quotas remain controversial despite evidence of their effectiveness in reducing labor-market gender inequality. We study how informational narratives about quotas affect support, and how effects depend on pre-existing causal beliefs about inequality. In a pre-registered survey experiment with 2,404 Italian workers and managers, we compare demand-side (discrimination, bias) versus supply-side (participation, confidence, role models) framings. All information increases unincentivized stated support, most strongly under demand-side narratives, but none affects the extensive margin of an incentivized donation, revealing a clear say–do gap. Conditional on donating, however, supply-side framing significantly raises amounts given. Open-ended responses show narratives reshape reasoning primarily among those with diffuse priors (generic cultural explanations). We formalize this in a simple model featuring misalignment costs and tail-driven effects: narrative success depends on the distribution of prior beliefs, which acts as a state variable determining optimal framing across contexts.

Keywords: gender quotas, gender inequality, support for policy, survey experiment
JEL codes: D63, D83, J16, J22, J31, J71.

*The experiment was pre-registered in the AEA RCT Registry (AEARCTR-0003252). Ethical approval was obtained from the Bioethical Committee of the University of Bologna. We gratefully acknowledge financial support from Ca' Foscari University of Venice (Fondi di Primo Insediamento) and from the PRIN project *"Advancing Gender Equality through Experiments: New Developments and Applications (AGEENDA)"* (CUP P2022R248L_002). We also thank the participants of the Lund University Microeconomics Seminar (September 2024), the Berlin Workshop on Empirical Public Economics: Gender Economics (26-27 September 2024), the VEM - Verona Experimental Meeting (28–29 November 2024), the CEPR Workshop on People's Understanding of and Support for Economic Policies (St. Gallen, Switzerland, 21–22 March 2025), the BEEN meeting, the 2025 Florence Workshop on Behavioural a Experimental Economics (8–10 May 2025) and the Gender Issue and Development Workshop at the University of Reading (3-4 November 2025) for their helpful comments and suggestions. We also thank our research assistants Maria Grazia Pagliarulo, Vlad C. Rotariu, and Elena Barbieri for their excellent work.

†Ca'Foscari University of Venice
‡University of Bologna
§University of Bologna

Non-Technical Abstract

Gender inequality in the labor market remains a persistent feature of most advanced economies. Women are highly educated and increasingly present in professional careers, yet they remain underrepresented in top managerial and leadership positions and continue to face substantial wage gaps. To address these disparities, many countries have introduced gender quotas in corporate boards and leadership positions. Although a large body of evidence suggests that quotas can be effective in increasing women's representation, they remain politically controversial and often face mixed public support.

This paper studies how public support for gender quotas responds to different informational narratives about gender inequality. In particular, we examine whether the way gender inequality is explained—either as the result of discrimination and structural barriers or as the outcome of constraints related to participation, confidence, or career choices—affects individuals' willingness to support quota policies. More importantly, we ask whether such informational messages are equally persuasive for everyone, or whether their impact depends on people's prior beliefs about why gender inequality exists.

To address these questions, we conducted a large pre-registered survey experiment with 2,404 Italian workers and managers. Participants were first asked about their beliefs regarding gender inequality in the labor market, including its magnitude and its underlying causes. They were then randomly exposed to one of several short informational videos. Some videos emphasized demand-side explanations for inequality, such as discrimination and biased evaluations in hiring and promotion. Others focused on supply-side explanations, such as family responsibilities, lack of role models, or lower participation in competitive environments. A control group received no information.

After viewing the information, participants reported their support for gender quotas and were given an opportunity to engage in a real behavioral decision: they could donate part of a potential lottery prize to a nonprofit organization promoting gender equality policies. This design allows us to distinguish between **expressed attitudes** toward quotas and **costly actions** that reflect stronger forms of support.

Our findings reveal an important difference between these two forms of support. Providing information increases stated approval of gender quotas, particularly when inequality is framed as the result of discrimination. However, these informational messages do not increase the likelihood that individuals engage in costly support through donations. In other words, information changes what people say about quotas but does not necessarily change whether they are willing to take action in support of them. This gap between expressed attitudes and costly behavior highlights the limits of purely informational interventions.

At the same time, we find that informational narratives do influence behavior along another dimension. Among those who choose to donate, individuals exposed to the

supply-side narrative donate significantly larger amounts. This suggests that narratives can affect the **intensity** of support even when they do not change the decision to participate.

To understand these patterns, we analyze respondents' prior beliefs about the causes of gender inequality using open-ended survey responses. These beliefs are highly heterogeneous. Many respondents attribute inequality primarily to work–life balance constraints or family responsibilities, while fewer mention discrimination or institutional barriers. We show that informational messages are most effective for individuals whose prior beliefs are relatively diffuse or uncertain. In contrast, individuals who hold strong prior views about the causes of inequality are much less responsive to new information.

Taken together, the results suggest that the effectiveness of policy communication depends crucially on how well the message aligns with the audience's existing beliefs. Informational campaigns that ignore these beliefs may change attitudes superficially but fail to generate meaningful behavioral support. More broadly, our findings highlight that public support for policies aimed at reducing inequality depends not only on the policies themselves but also on how people understand the causes of inequality. Designing effective communication strategies therefore requires understanding the distribution of beliefs in the population and tailoring messages accordingly.

1 Introduction

Despite prolonged efforts over the past decades, gender gaps in labor market outcomes persist in most advanced economies. Women often outperform men in education and enter the labor market better prepared than ever before (Goldin et al., 2006; Kuzmina and Melentyeva, 2021; Bertrand, 2011). Yet, within a few years after graduation and well before motherhood, men are more likely to be employed full-time, earn higher wages, experience steeper career trajectories, and attain top positions (Francesconi and Parey, 2018; Manning and Swaffield, 2008). Gender pay gaps are particularly pronounced at the top of the earnings distribution, and women remain heavily underrepresented in high-status and high-income sectors, occupations, and leadership roles (Blau and Kahn, 2017; Bertrand, 2018). Consistent with this pattern, women hold only about one-third of leadership positions in OECD countries—a figure that has remained strikingly stable over time despite a wide range of policy interventions (OECD, 2023).

The persistence of gender gaps at the top of the labor market has motivated a broad set of policy responses aimed at promoting women's participation and advancement. Guided by international organizations such as the UN, OECD, and the European Union, countries have adopted measures ranging from equal employment opportunity and pay-equality legislation to affirmative action policies, including gender quotas, as well as pay-transparency regulations. Among these initiatives, the OECD 2013 report explicitly calls on member states to address gender inequalities in employment and leadership, including through the use of gender quotas in both the public and private sectors.[1] While implementation varies across countries -ranging from binding quotas (e.g., Norway, France, Spain, and Italy) to softer targets or voluntary guidelines (e.g., the United Kingdom and the United States)- gender quotas have become a central and highly visible policy instrument in the debate over gender equality.

This paper studies how support for gender quotas responds to informational narratives and, critically, how these responses depend on individuals' prior beliefs about the causes of gender inequality. We show that information about how quotas work easily shifts expressed approval but only translates into costly support when narratives are aligned with the distribution of prior causal beliefs. We examine when expressed approval translates into concrete, costly action—and when it does not.

While gender quotas remain politically contested, a large body of empirical work shows that they are highly effective at changing who holds corporate power: quotas substantially increase the share of women on boards and improve observable board characteristics without systematically harming firm performance or shareholder value, with recent meta-analytic evidence documenting high compliance and no average negative effects on firm outcomes (Ferrari et al., 2016; Acutis et al., 2024). More recent studies also document spillovers beyond boards to connected firms and executive leadership, suggesting

[1] More recently, the European Parliament adopted a directive to improve gender balance among directors of listed companies, requiring large EU firms to achieve at least 40% representation of the under-represented sex among non-executive directors, or 33% among all directors, by June 30, 2026 (European Commission, 2022).

indirect and informational channels (Zaccaria et al., 2024; Falconieri et al., 2025). Yet public support often weakens when policies move from abstract principles to concrete instruments. Proponents view quotas as tools to counter discrimination and stereotypes, enhancing women's empowerment and firm performance, while critics stress individual choice, potential stigma, and threats to meritocratic selection. These competing narratives suggest that quota sustainability depends not only on effectiveness but also on how they are framed, understood, and evaluated by the public.

We conduct a pre-registered survey experiment to study how informational narratives about gender quotas shape individual support for such policies, accounting for prior beliefs about the magnitude and causes of gender inequalities in the labor market. The sample consists of 2,404 Italian workers and managers—individuals directly exposed to potential quota effects and likely to hold informed views on affirmative action. The experiment features three information treatments. The *demand* treatment frames quotas as tools to counter discrimination and biases; the *supply* treatment emphasizes constraints such as underconfidence and the lack of role models; and the *info-causes* treatment presents both types of causes without referencing specific policies. A fourth group serves as a control and receives no information.

Our analysis focuses on two main outcomes: unincentivized self-reported support for gender quotas and an incentivized measure of behavioral support, proxied by participants' willingness to donate part of a potential 500-euro lottery prize to an NGO promoting quotas, with a 50 percent matching contribution by the research team. All information treatments increase stated support relative to the control group, most strongly under the demand-side narrative, but none affects the decision to donate, revealing a clear say–do gap. Conditional on donating, however, exposure to the supply-side narrative significantly increases the amount given.

Beyond shifting levels of support, information also affects how respondents reason about quotas. Using coded open-ended responses, we show that narratives reshape the content and structure of justifications, broadening the set of arguments invoked, particularly among respondents with diffuse prior beliefs about the causes of gender inequality, whereas respondents with specific priors—such as strong views centered on work-life balance, women's ability, or discrimination—exhibit weaker or no responses. Our central contribution is to show that narrative effectiveness depends on the structure of prior causal beliefs: frames aligned with the distribution and specificity of such beliefs are the only ones that translate into costly engagement, and the prevailing belief distribution in the population thus acts as a state variable for the impact of informational campaigns. These patterns are consistent with a simple conceptual framework developed in Section 8, in which donors experience a disutility when the communicated narrative is misaligned with their prior beliefs and the agency's optimal message is determined by donors in the tails of the belief distribution whose contributions are most sensitive to ideological alignment and misalignment. Together, these mechanisms generate the observed say–do gap and imply that, in our context, a supply-side narrative uniquely raises donated amounts among those willing to contribute. A key challenge for extrapolating experimental evidence is understanding where prior beliefs come from. We show that causal beliefs about

gender inequality are systematically related to demographic characteristics and regional labor-market outcomes, indicating that they reflect structured normative environments rather than idiosyncratic noise. This allows us to interpret narrative effectiveness as a function of the distribution of priors in a given context and to derive predictions about which frames should be most effective across environments.

Our results are robust to a range of pre-treatment checks, including attrition and self-selection, respondents' prioritization of gender equality relative to other strategic challenges, and awareness of the existence and magnitude of gender inequalities. Our study complements recent work linking beliefs to policy demand, most notably Settele (2022), while extending it along several dimensions. Whereas Settele examines how information about the magnitude of the gender wage gap affects support for affirmative action in a relatively abstract policy setting, we study a concrete and contested instrument—gender quotas for corporate boards—and focus on narratives about why inequalities persist. By combining unincentivized attitudes with an incentivized behavioral outcome—a costly donation to an NGO promoting quotas—we assess not only whether information shifts approval, but also how it affects the strength of support and the resulting say–do gap in a context where respondents are directly exposed to potential quota effects.

The remainder of the paper proceeds as follows. Section 2 briefly reviews the relevant literature on beliefs, narratives, and support for gender quotas. Section 3 describes the experimental design; Section 4 presents our hypotheses; Section 5 details the data; Section 6 reports the main findings; Section 7 discusses potential mechanisms; Section 8 develops the conceptual framework underlying these mechanisms; and Section 9 concludes.

2 Review of the literature

This paper connects two strands of research. The first examines how beliefs and information shape support for public policies, emphasizing that attitudes toward redistribution and affirmative action depend not only on self-interest but also on fairness perceptions and causal beliefs about inequality (Stantcheva, 2022; Haaland and Roth, 2023; Alesina et al., 2021; Settele, 2022). The second explores the legitimacy and effectiveness of gender quotas, highlighting that their success hinges on public acceptance and perceptions of justice (Arve, 2023; Faniko et al., 2017; Bush, 2020; Latura and Weeks, 2022; Muriaas and collaborators, 2024; Zehnter and Nater, 2024). While the first literature identifies beliefs as key drivers of policy demand, the second reveals that fairness framing and context strongly condition quota support. Yet, few studies bridge these perspectives by examining how causal beliefs about inequality interact with information about how quotas work, or by linking attitudinal and behavioral measures of support. We next discuss these literatures in greater detail.

Information, beliefs, and policy demand. A large body of work in behavioral and political economics shows that what people believe about *why* inequality exists crucially shapes their preferences for redistribution and corrective policies. Settele 2022 provides

direct evidence in the context of affirmative action. In a survey experiment, she examines general affirmative-action policies and shows that individuals who perceive wider gender gaps in opportunities express stronger willingness to endorse such interventions. By combining self-reported attitudes with an incentivized donation measure, she documents a "say–do" gap between verbal approval and costly behavior, demonstrating that beliefs about the *causes* of inequality are proximate drivers of genuine policy demand.

Beyond affirmative action, several studies explore how narratives and information influence support for redistribution. Alesina et al. 2021 show that framing inequality as a product of luck rather than effort increases demand for redistribution, while Stantcheva 2022 models how "causal narratives" structure policy preferences. Haaland and Roth 2023 show that information about discrimination leads to convergence in beliefs about racial discrimination, but find little change in support for pro-black or equality-enhancing policies. Collectively, this work highlights the causal role of beliefs and information but primarily relies on abstract policy domains and self-reported measures, leaving open how such mechanisms operate when applied to specific, contested interventions such as gender quotas.

Who supports quotas, and why. A second strand of research focuses directly on public attitudes toward gender quotas and seeks to explain who supports them and through which mechanisms. Zehnter and Nater 2024 show that gender differences in quota attitudes arise because women perceive higher levels of discrimination and thus greater policy necessity, whereas men's opposition is driven by concerns about reverse discrimination and stigma. Muriaas and collaborators 2024 extend this work through a preregistered survey experiment in Norway that varies both the *domain* of quota application—politics, business, or religion—and its *justification*—social-utility versus group-justice arguments. They find that fairness-based frames are more persuasive than efficiency-based ones and that effects depend on respondents' prior attitudes, suggesting that framing and prior beliefs interact in shaping support. Evidence from organizational contexts reinforces these findings: Lefley et al. 2024 document widespread skepticism toward corporate-board quotas in the Czech Republic, particularly among men who view them as violating meritocratic norms. At the same time, Bush 2020 show that quotas can enhance external perceptions of legitimacy: in survey experiments in Sweden and the United States, higher women's representation abroad increased perceived democratic quality and willingness to provide foreign aid. Together, this literature evidences that quota support is multi-dimensional—shaped by fairness beliefs, ideology, gender identity, and contextual framing—but has relied mainly on self-reported attitudes, rarely considering how underlying causal beliefs interact with informational cues or translate into behavioral support.

Gap and contribution. Despite these advances, both literatures leave critical questions open. Research on policy support demonstrates that beliefs and narratives strongly influence preferences, but typically in abstract settings. Work on gender quotas identifies fairness framing and ideology as central correlates of acceptance but seldom examines how prior causal beliefs about inequality condition the response to information, nor whether

such beliefs predict real, costly support.

Our study bridges these literatures by focusing on a concrete and contested instrument -gender quotas- and combining three design elements rarely implemented jointly. First, we elicit respondents' *ex-ante causal beliefs* about why gender inequality persists, capturing both magnitude assessments and open-ended explanations. Second, we randomize *competing informational narratives* that highlight either demand-side (employer discrimination) or supply-side (self-selection) mechanisms. Third, we measure both *stated* and *incentivized* support through attitudinal responses and a real monetary donation task. This design allows us to test whether *narrative -prior alignment*-the consistency between individuals' initial beliefs and the information they receive-is pivotal for translating abstract endorsement into concrete behavioral action. In doing so, we contribute to understanding how beliefs mediate the effectiveness of information in shaping support for equity-enhancing policies.

3 Experimental Design

This section outlines the experimental design. We first provide an overview of the survey structure and timeline, followed by details on each component.

3.1 Timeline and Overview

Data were collected between March and June 2023 through an online survey (CAWI) administered in collaboration with Scenari Srl[2], which recruited participants from its national panel via email invitations[3]. Before random assignment to one of the four treatment conditions, respondents completed a background questionnaire and a set of belief-elicitation tasks. These included incentivized measures of (i) perceived gender gaps in the labor market, as well as open-ended questions eliciting (ii) beliefs about the causes of gender inequalities and (iii) perceptions of the most effective policies to address them. These pre-treatment measures allow us to account for heterogeneity in pre-existing beliefs

[2]Scenari Srl (`https://www.scenari.it/`) is an independent Italian company specialized in market research and opinion surveys, with more than thirty years of experience and a consolidated record of collaborations with public institutions and universities, including the Alma Mater Studiorum – University of Bologna, Ca' Foscari University of Venice, the University of Naples Federico II, the University of Catania, the University of Innsbruck, the University of Turin, and the University of Rome Tor Vergata. The company holds several professional and quality certifications, including ASSIRM, ESOMAR, MSPA, UNI EN ISO 9001:2015, UNI PdR 125:2022, ISO 45001:2018, and EN ISO/IEC 27001:2023. To encourage panel participation, Scenari adopts a point-based incentive system whereby respondents earn points for each completed survey, which can be redeemed for Amazon gift cards.

[3]To motivate members of the panel to take part in surveys, the company offers incentives by adopting a point-based system. Participants receive points for each survey they complete, depending on the survey length. Every 50 points they can get a 10 Euros Amazon gift card. In the prior beliefs elicitation we provided additional incentives: participants who correctly guessed the answer to the two questions were rewarded with 5 Euros per correct answer with an Amazon gift card paid by the company. Participants were informed that, after the completion of the data collection, 25 participants would be randomly selected and rewarded. Selected participants received the earnings associated with their correct guesses through the company in addition to the participation points awarded to all respondents for completing the survey.

about the magnitude and the causes of gender inequalities when estimating treatment effects.

Figure 1 summarizes the overall survey structure, which is discussed in detail in the next subsection.

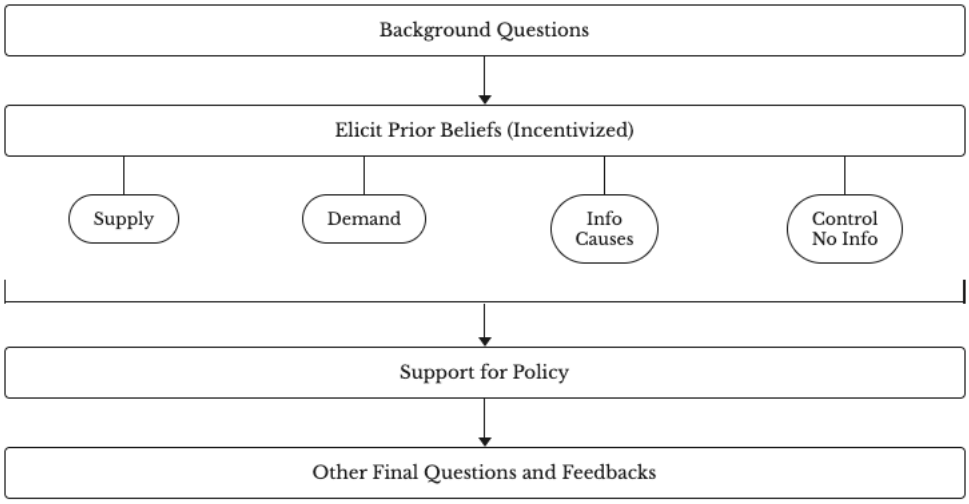

Figure 1: Outline of the survey structure and experimental timeline. The figure illustrates the sequence of stages: background questions, elicitation of prior beliefs, random assignment to one of four treatments, and measurement of post-treatment outcomes.

3.2 The Survey

This subsection describes the structure and content of the survey instrument, detailing the sequence of pre-treatment elicitation, information provision, and post-treatment outcome measurement. Appendix G provides full details on the survey design. Table G1 summarizes the survey structure and lists the variables elicited at each stage.

3.2.1 Pre-treatment elicitation

Importance assigned to gender equality. After completing background questions on demographics, employment, and firm characteristics, participants ranked six goals in order of priority for companies over the next five years: (i) promoting innovation, (ii) increasing productivity, (iii) increasing diversity and gender equality, (iv) promoting technological advancement, (v) enhancing teamwork, and (vi) reducing environmental impact. This ranking captures respondents' baseline prioritization of gender equality before any exposure to the survey content and allows us to test whether this prioritization correlates with survey completion.

Prior belief elicitation. Before treatment assignment, we elicited respondents' beliefs about three key aspects of gender inequality in managerial positions: its magnitude, its

underlying causes, and potential policy remedies. This step is crucial, as participants are likely to hold pre-existing views that may interact with the information treatments.

To assess the perceived magnitude of inequality, respondents answered two incentivized questions: (i) the average wage earned by female managers in Italy for every 100€ earned by male managers aged 30–49 working full-time, and (ii) the share of women CEOs among the 50 largest Italian companies. These measures are precise and difficult to verify online, ensuring that answers reflect genuine beliefs rather than researched values (Settele, 2022). Incentives were awarded to a randomly selected 2% of participants after data collection. Respondents who estimated the gender wage ratio within 2€ or the female CEO share within 2 percentage points of the latest Eurostat and EIGE data received a 5€ bonus per correct answer. This design aims to maintain attention and reduce politically motivated response bias by rewarding accuracy (Bullock et al., 2013; Prior et al., 2015).

To capture prior beliefs about the causes of gender inequality and preferred policy interventions, participants answered two open-ended questions: (i) what they believe causes gender gaps in wages and managerial representation, and (ii) whether they support stronger state intervention to reduce these gaps. Those in favor were asked to specify which policies they consider most effective. These responses allow us to characterize pre-existing beliefs and examine how treatment effects vary with prior views (Stantcheva, 2022).

3.2.2 Information treatments

After completing the pre-treatment elicitation, participants were randomly assigned to one of four groups, three information treatments: *Supply*, *Demand*, and *Info Causes* and one control group: *No Info*.

Participants in the two main treatment groups (*Supply* and *Demand*) watched one of two short, researcher-produced animated videos (Appendix B) explaining how gender quotas can reduce gender inequalities. Each video presented mechanisms through which quotas operate, focusing either on supply-side or demand-side factors.

The *Info Causes* group viewed a video of identical length describing the main causes of gender inequality in managerial positions, but without reference to quotas or policy interventions. This treatment mitigates potential demand effects arising from exposure to information or video content. The *No Info* group received no information and proceeded directly to the outcome questions, allowing us to isolate any priming effects from watching a video.

The two main treatment videos draw on established research on labor market inequalities. Demand-side explanations emphasize discrimination and bias in hiring or promotion, such as taste-based and statistical discrimination, or implicit bias in evaluating women's leadership (Beaman et al., 2009; Reuben et al., 2014; Bordalo et al., 2019). Supply-side explanations focus on women's characteristics and constraints, including heavier family responsibilities, limited leadership aspirations due to few role models, and preferences for less competitive environments (Niederle and Vesterlund, 2007; Günther et al., 2010; Porter and Serra, 2020; Cortés and Pan, 2023).

Both videos share a common introduction (Panels 1–5, Appendix Figure B1) presenting the concept of gender quotas and examples of their adoption across Europe. They then diverge (Panels 6–11, Appendix Figures B2–B3). The *Supply* video links quotas to mechanisms that empower women—such as increasing representation, creating role models, and encouraging investment in human capital—while the *Demand* video highlights how quotas can reduce discrimination by shifting social norms, correcting biased beliefs about competence, and improving fairness in selection processes. The *Info Causes* video covers both demand- and supply-side causes of gender inequality, omitting any discussion of quotas (Appendix Figure B5).

All three videos conclude with a statement noting that the information provided is based on academic research and include a list of references for interested viewers. Participants were required to remain on the video screen for at least 90 seconds (i.e., the full duration of the videos) before proceeding.

3.2.3 Post-treatment elicitation

Self-reported policy demand. After treatment exposure, we measured respondents' attitudes toward gender quotas in leadership positions, asking whether they consider quotas an effective tool and whether they would support their implementation in their own workplace. We also assessed support for other gender-equality measures, including gender-neutral language, flexible working hours to improve work-life balance, extended paternity leave, and wage transparency. Participants rated the importance of each policy on a four-point scale ranging from "extremely important" to "not important."

Additional post-treatment questions. To investigate the mechanisms underlying treatment effects, we included several post-treatment measures. First, participants answered an open-ended question explaining the reasons behind their support for or opposition to gender quotas in leadership positions, allowing us to assess how treatments influenced not only levels of support but also the rationale behind respondents' views.

Next, we introduced a set of mechanism-focused questions. The framing of quotas was randomly varied, presenting them either as increasing the number of women or decreasing the number of men in organizations. Participants then indicated their agreement with statements suggesting that quotas might lead women (or men) to work less, aspire to higher positions, or reduce teamwork. We also asked respondents to evaluate whether four groups—working men, working women, non-working men, and non-working women—would be advantaged or disadvantaged by quotas.

Finally, to capture potential belief updating induced by the treatments, respondents rated their agreement with statements attributing gender inequalities either (i) to women's individual effort or (ii) to social norms and historical factors.

Incentivized behavior. In studies on sensitive topics such as gender equality, experimenter demand and social desirability biases are common concerns, as respondents may overstate support for quotas to appear open-minded or to please the experimenter.

Although recent evidence suggests such effects may be limited (De Quidt et al., 2018; Mummolo and Peterson, 2019), we included a behavioral measure to mitigate these risks.

Participants were informed that by completing the survey, they were automatically entered into a lottery for a 500€ prize.[4] Before the draw, respondents indicated (i) whether they would donate any portion of a potential win to an NGO promoting gender quotas, and (ii) if so, the amount (0–500€). Each donated euro was matched by an additional 0.5€ from the research team, creating an incentive to commit to donations immediately rather than hypothetically.

Because even behavioral measures can be affected by self-presentation concerns, we adopted a *forced-response design* (Blair et al., 2015) that allowed participants to decline donation without explicitly saying "no." Before the donation question, respondents were instructed to consult the third-to-last digit of a banknote or their phone number: if it was 0, 1, or 2, they were required to answer "No"; otherwise, they could freely choose "Yes" or "No."[5]

Final questions. At the end of the survey, participants answered a short set of concluding questions covering six topics: (i) the perceived importance of gender for their own career prospects; (ii) the presence of workplace initiatives promoting gender equality; (iii) awareness of the Uni Pdr 125:2022 gender-equality certification and whether their company holds it; (iv) perceptions of potential political bias in the study; (v) willingness to donate any bonus earned from the belief-elicitation task if selected for payment; and (vi) interest in receiving additional information about the academic sources cited in the information treatments.

4 Hypotheses

This section presents the hypotheses tested in the experiment. All hypotheses were pre-registered prior to data collection on February 15, 2023.[6] The hypotheses below closely follow the preregistered statements, with minor edits to improve clarity and consistency with the experimental design.

H1. Information provision and support for gender quotas. *Treatments providing information increase support for gender quotas compared to the No Info condition.* This hypothesis builds on prior evidence showing that factual information can affect policy preferences, including in the context of gender inequality and affirmative action (Settele, 2022; Alesina et al., 2021). By providing information about gender inequalities and the functioning of quotas, the treatments are expected to lower informational and psychological

[4]The winner was randomly selected by the survey company after data collection using a standard lottery system.

[5]Exact wording: *"In the next question, you will need to answer YES or NO. However, we ask that you follow these instructions: take a banknote or, if you do not have one, think of the third-to-last digit of your phone number. If this digit is 0, 1, or 2, you must answer NO to the next question. In all other cases, you are free to answer YES or NO as you wish."*

[6]Pre-registration DOI: `https://doi.org/10.17605/OSF.IO/FG2XR`.

barriers, thereby increasing both stated and behavioral support for gender quotas relative to a no-information baseline.

H2. Demand-side versus supply-side framing. *The treatment focusing on demand-side sources of gender inequality increases support for gender quotas more than the treatment focusing on supply-side sources.* This hypothesis reflects the idea that framing gender inequalities as the result of discrimination and structural barriers elicits stronger support for corrective policies than framing them as stemming from supply-side factors, such as women's preferences or choices (Alesina et al., 2021). Because supply-side explanations are often perceived as reflecting outcomes that are at least partly "chosen" by women, a demand-side narrative is expected to generate higher support for gender quotas.

H3. Role of prior beliefs. *Prior beliefs about the magnitude of gender inequalities affect support for gender-equality policies.* Following Settele 2022, this hypothesis posits that individuals' pre-existing beliefs shape their support for policy interventions. In our design, prior beliefs are measured along multiple dimensions, including beliefs about the magnitude of gender inequalities, beliefs about their underlying causes, and views on appropriate policy responses. Respondents who perceive larger gender inequalities—and who attribute them to structural or discriminatory factors—are expected to display stronger support for gender quotas than those who perceive inequalities as smaller or primarily driven by individual choices.

H4. Heterogeneity by individual characteristics. *Individual characteristics, such as gender and age, are correlated with treatment effects and with support for gender quotas.* This hypothesis reflects the notion that personal characteristics and lived experiences influence both baseline attitudes toward gender equality policies and responsiveness to information. As a result, treatment effects are expected to vary systematically across demographic groups.

These preregistered hypotheses generate clear empirical predictions about how information provision, framing, prior beliefs, and individual characteristics shape support for gender quotas. We next describe the data and empirical strategy used to test these predictions.

5 Data

This section describes the dataset used in the analysis and assesses whether survey completion is affected by selective attrition related to pre-existing interest in gender-related issues.

5.1 Sample and summary statistics

The final sample consists of 2,404 respondents, all employed individuals aged between 20 and 87. Table 1 reports summary statistics for key demographic characteristics by treatment group, while Table 2 summarizes respondents' job-related characteristics.

As shown in Table 1, women represent slightly more than half of the sample, and respondents are broadly distributed across age groups, regions of residence, and family status. Table 2 indicates substantial heterogeneity in occupational roles, firm types, and job tenure, with representation from both managerial and non-managerial positions and from public, private, and non-profit organizations.

	Demand	**Supply**	**Info Causes**	**No Info**	**Total**
Women	370	350	334	268	1322
Has children	380	418	354	310	1462
Bachelor's degree or more	320	318	266	256	1160
Age range					
Age 20–25	6	4	13	9	32
Age 26–35	108	63	110	74	354
Age 36–44	193	183	165	140	681
Age 45–54	187	236	189	168	780
Age 55–64	104	125	90	96	415
Age 65–87	31	51	29	31	142
Geographical Area					
Center	115	117	94	66	392
South	183	225	142	203	753
North	325	315	357	244	1241
Observations	629	661	596	518	2404

Table 1: Demographic characteristics of respondents.

5.2 Balance across treatment groups

Appendix Table A1 reports pairwise balance tests across treatment groups for demographic characteristics, while Appendix Table A2 presents corresponding tests for job-related variables. Appendix Table A3 reports balance tests for the pre-treatment ranking of gender equality as a workplace priority. Overall, randomization produced broadly balanced samples. While a small number of pairwise differences emerge for specific characteristics, no systematic patterns are observed. To account for these minor imbalances, all subsequent analyses include standard demographic and job-related controls.

5.3 Attrition and self-selection

We next examine whether survey completion is affected by selective attrition. The overall attrition rate is 12% (N = 330 out of 2,734 initial respondents) and is significantly higher among women (16.75%) than among men (5.50%) ($t = -9.05$, $p < 0.001$). Table A4 in the Appendix report attrition rates by treatment group and present pairwise t-tests. These tests reveal no statistically significant differences in attrition across treatments, including when stratifying the sample by gender. Appendix Table A5 further examines attrition as a function of respondents' pre-treatment prioritization of gender equality, measured through the ranking task administered before treatment assignment. We find

	Demand	Supply	Info Causes	No Info	Total
Job Role					
Manager	75	74	86	85	320
Executive	65	77	27	24	193
Middle Manager	149	161	125	101	536
HR Manager	59	62	75	86	282
Other Managerial Role	28	36	24	19	107
Employees (White Collar)	178	166	158	123	625
Self-Employed	10	12	4	8	34
Employees (Blue Collar / Other)	58	71	88	58	275
Experience					
Average Job Tenure (years)	11.20	12.55	10.89	11.29	11.51
Average Firm Tenure (years)	13.01	14.49	12.21	14.49	13.10
Firm Type					
Public Firm	150	166	156	131	603
Private Firm	432	441	385	322	1580
Non-Profit Firm	47	54	55	65	221
Observations	629	661	596	518	2404

Table 2: Job characteristics of respondents.

no significant association between the priority assigned to gender equality and survey completion.

The results presented in this sections indicate that survey completion is not driven by self-selection based on pre-existing interest in gender-related issues, supporting the internal validity of the experimental design.

6 Results

We present the results in three steps. First, Section 6.1 examines respondents' perceptions of the magnitude of gender inequalities, their perceived causes, and preferred policy solutions, elicited prior to the information treatments. Second, Section 6.2 analyzes the effects of the information treatments on post-treatment outcomes, focusing first on self-reported support for gender quotas and then on incentivized behavior, with the donation decision and the amount donated considered last. Finally,

6.1 Pre-treatment beliefs about gender inequality

Informational narratives do not operate in a vacuum: their impact depends on the prior causal beliefs individuals hold about why inequality persists. We therefore begin by documenting the distribution and structure of these beliefs before treatment exposure. We show that priors are systematically related to demographic and regional characteristics and, importantly, that supply-side explanations are substantially more prevalent than discrimination-based accounts in our setting. Specifically, we elicited participants' beliefs on: (i) the magnitude of gender inequalities in managerial positions (incentivized measure);

(ii) the perceived causes of such inequalities; and (iii) potential policy interventions to address them (see Section 6.1.3).

6.1.1 Prior beliefs about the *magnitude* of gender inequalities

We begin by examining respondents' beliefs about the magnitude of gender inequalities in managerial and leadership positions, focusing on the gender wage gap and women's representation at the top of firms.

Gender wage gap. Respondents were asked how much a female manager earns for every 100 € earned by a comparable male manager. The correct value, based on Eurostat data, is 55.64 €. The average estimate in our sample is 71.5 €, implying a substantial underestimation of the true gap. Figure C1 in the Appendix displays the full distribution of responses by treatment group. Overall, fewer than 5% of respondents provide an answer within a narrow tolerance window around the true value. To summarize individual misperceptions, we construct an index (*indexGWG*) equal to the normalized distance between the respondent's estimate and the true value when the gender gap is underestimated, and zero otherwise (details in Appendix C). Beliefs about the wage gap are balanced across treatment groups. Appendix Tables C1–C5 report pairwise balance tests for raw answers and constructed misperception indices.

Women's representation in top positions. We next consider beliefs about the share of female CEOs among the 50 largest listed Italian firms. The true share is 2.9%. Respondents' average estimate is 19.8%, indicating a marked overestimation. Figure C2 reports the distribution of beliefs about women's representation among CEOs. Overall, fewer than 8% of respondents report a value close to the true figure. Analogously to the wage-gap measure, we construct a normalized misperception index (*indexCEO*) capturing the extent of overestimation (Appendix C). Again, pre-treatment beliefs are balanced across experimental groups.

Overall underestimation of gender inequalities. To summarize misperceptions across the two dimensions, we construct an *Underestimation Index*, defined as the average of *indexGWG* and *indexCEO*. Each component captures the normalized distance from the true value in the direction that attenuates the severity of inequality (i.e., underestimating the wage gap or overestimating female representation), and is set to zero when the respondent does not understate inequality. The combined index therefore ranges from 0 (accurate perception or no downward bias) to 1 (maximum underestimation of existing gender inequalities). Balance tests confirm that the index is evenly distributed across treatment groups, consistent with random assignment occurring after belief elicitation. The index exhibits systematic heterogeneity: women display significantly higher levels of underestimation, while respondents residing in Southern Italy tend to underestimate less. Employees and individuals working in large firms also exhibit greater misperceptions relative to managers (Figure C3). We use this index throughout the analysis as a summary

measure of prior beliefs about the magnitude of gender inequality, in particular to examine how informational treatments interact with ex-ante perceptions and how these beliefs correlate with both stated and incentivized support for quotas. When aggregated at the regional level, the Underestimation Index is not correlated with female labor-market outcomes (Appendix Table C6), suggesting that quantitative misperceptions are only weakly disciplined by local labor-market conditions: respondents may have limited ability to translate the environment they observe into a precise numerical assessment. This, in turn, highlights the need for richer belief elicitations that go beyond a single statistical estimate. We return to this point below when we consider alternative, more qualitative belief measures.

6.1.2 Prior Beliefs about the *Causes* of Gender Inequality

Attitudes toward gender quotas depend not only on perceptions of the magnitude of gender gaps, but also on beliefs about their underlying causes. Because informational treatments may resonate differently depending on these ex-ante views, characterizing prior causal beliefs is central to our analysis.

We elicit such beliefs through an open-ended question, following Ferrario and Stantcheva 2022. This approach captures the multidimensional narratives that structure policy preferences and avoids constraining respondents to pre-defined categories. Responses are typically substantive (median length: 9 words; 25th percentile: 5; 75th percentile: 18).

Text processing and classification. We preprocess responses using standard text-cleaning procedures (lowercasing, punctuation removal, stemming) and map them to topic dictionaries derived from the gender-inequality literature. For each category, we construct a binary indicator equal to one if at least one associated keyword appears in the response. Because respondents may refer to multiple mechanisms, categories are not mutually exclusive[7]. Topic definitions were manually validated to ensure conceptual coherence. Details are provided Appendix D.1. Table 3 reports the distribution of coded responses. Cultural causes is the only *generic* category; it captures broad references to norms or culture without identifying a specific mechanism. It is coded as mutually exclusive from all specific categories and it accounts for a substantial share of responses. The remaining explanations fall into *specific* categories, among which work–life balance is the most frequently cited (33.47%), followed by women's ability (28.30%). Institutional constraints (5.83%) and explicit discrimination (10.37%) are mentioned substantially less often. This asymmetry is economically meaningful: participation and caregiving constraints are markedly more salient in respondents' priors than discrimination-based accounts. The pattern becomes even clearer when focusing on respondents who articulate a single specific mechanism (75%, N=1,077). Although small differences across treatment groups emerge, pre-treatment causal beliefs are balanced overall, and we control for them

[7] 75% of respondents cite exactly one specific mechanism; 22% cite two and fewer than 4% cite three or more (Appendix Table D5). In some specifications, we restrict attention to single-mechanism responses excluding the generic category (Appendix Table D5).

in all analyses (Appendix Table D2). Causal beliefs display systematic heterogeneity across respondents. Women are significantly more likely to cite work–life balance constraints and women's ability, whereas men more frequently invoke generic cultural explanations (Appendix Table D4). These gender differences remain when controlling for demographics and job characteristics (Appendix Table D6). Other individual characteristics are also associated with recalling distinct categories. Married respondents are more likely to attribute inequality to work–life balance constraints. Relative to respondents younger than 25, older individuals are more likely to invoke taste discrimination or backlash. Respondents with at least a bachelor's degree are less likely to cite discrimination-based explanations. Work-related characteristics further correlate with causal beliefs. Compared to respondents in managerial positions, employees and the self-employed are more likely to refer to work–life balance constraints. Moreover, respondents working in very small (10–49 employees) or very large firms (more than 250 employees) are less likely to attribute inequality to women's ability relative to those in firms with fewer than 10 employees. Finally, respondents who invoke generic cultural explanations—and, separately, those citing discrimination—tend to exhibit greater underestimation of existing gender gaps (Appendix Tables D7 and Table D8). Moreover, when aggregated at the regional level, the distribution of belief categories is systematically related to female labor-market outcomes (Appendix D9). Regions where work–life balance explanations are more prevalent exhibit systematically different patterns of female part-time employment. This suggests that prior beliefs reflect broader economic and normative environments rather than idiosyncratic survey noise.

ID	Category	% of responses	Definition
Generic			
1	Cultural causes	28.39%	Acknowledges stereotypes or traditional social norms disadvantaging women, without specifying whether they concern work abilities or social roles.
Specific			
3	Work–Life balance	33.47%	Points to caregiving roles and family responsibilities
2	Women's ability	28.30%	Mentions women's (perceived) work capabilities
5	Taste discrimination and reaction	10.37%	Mentions explicit discrimination by men in power or defensive reactions to perceived threats from women.
4	Institutional problems	5.83%	Refers to insufficient institutional or organizational support (e.g., childcare, leave policies).
7	No idea	3.08%	States explicitly not knowing the causes.
6	Denial	1.29%	Denies the existence or salience of gender inequalities.
8	Not classifiable	5.62%	Does not fit any category above.

Table 3: Categories for causes of gender inequality in leadership positions

Notes: Percentages for the four specific categories are computed on the sample with non-missing open-ended responses ($N = 2{,}402$) and reflect "any mention" (categories are not mutually exclusive). See Appendix Table D1 for examples and keywords.

Importantly, the distribution of categories is not neutral with respect to the two narratives we experimentally vary. Work–life balance and participation-related explanations map naturally into supply-side accounts, whereas discrimination-based explanations align with demand-side narratives—a distinction we exploit in the analysis of treatment effects. Caregiving constraints and participation barriers are substantially more salient in respondents' ex-ante views than explicit discrimination or institutional failures, a pattern consistent with salient features of the Italian labor market Saraceno, 2020; Eurostat, 2023; Del Boca and Locatelli, 2022; OECD, 2023. As we show below, this asymmetry in prior beliefs plays a central role in determining which informational frames translate into costly support.

6.1.3 Views on Policies

After eliciting respondents' perceived *causes* of gender inequality, we assess their prior views on *policy* responses. We ask whether stronger state intervention is needed and, if so, which instruments respondents consider most effective. Using the same coding approach as above, we classify responses into nine literature-based categories (Table 4; examples in Appendix Table E1). Categories are not mutually exclusive. When quotas are explicitly mentioned, we code the stance as pro- or anti-quota.

The most frequently advocated policies target workplace discrimination and cultural change. Explicit pre-treatment support for gender quotas is limited (about 10%), and opposition is rare. Policy preferences vary systematically across respondents—for instance, women are more likely to endorse family-support measures (Appendix Table E5). Full distributions and balance checks are reported in Appendix E.

ID	Category	Definition
1	Family support policy	Support for motherhood and families.
2	Pro-quota	Favourable reference to gender quotas.
3	Anti-quota	Opposing reference to gender quotas.
4	Anti-discrimination policy	Policies to reduce workplace discrimination.
5	Incentives (non-quota)	Incentives other than quotas (e.g., bonuses, targets).
6	Cultural change	Changing stereotypes via education, media, and social norms.
7	Merit	Emphasis that pay/promotions should depend on merit/skills, not gender.
8	No idea	Not specified / respondent reports not knowing.
9	Not classifiable	Does not fit any category above.

Table 4: Categories for beliefs about policies to reduce gender inequalities

Notes: Categories coded from open-ended responses (binary: 1 if mentioned). Multiple categories may be flagged per response. Pro/anti-quota stance is coded when quotas are mentioned (see Appendix Table E1 for examples/keywords).

To examine the alignment between causal narratives and policy preferences, we regress

indicators for invoking a given cause on dummies for the policy categories mentioned in the same response (Figure E1). Policy proposals co-occur systematically with distinct causal beliefs. Respondents advocating family-support measures are more likely to frame inequality in work–life balance terms, whereas those emphasizing anti-discrimination or merit-based policies more often reference ability-related explanations. Similarly, proposals centered on anti-discrimination and cultural change are positively associated with generic cultural explanations. Overall, policy preferences and causal narratives are tightly linked, reinforcing the structured nature of respondents' prior beliefs.

6.2 Narratives and Support for Gender Quotas: Causal Evidence

As shown above, treatment groups are broadly balanced with respect to pre-treatment beliefs and characteristics. Where minor differences arise, we account for them by including standard controls in all specifications. We now present the main causal results of the information treatments. First, we examine the effects of the treatments on unincentivized, self-reported support for gender quotas (Section 6.2.1). We then turn to incentivized measures of support, considering in turn (i) respondents' willingness to donate to an NGO advocating quotas and (ii) the amount donated.

6.2.1 Unincentivized self-reported support for gender quotas

We begin by examining unincentivized self-reported support for gender quotas, measured immediately after the information treatments and prior to the donation task. While self-reported outcomes may reflect expressive motives, expressed opinions and the narratives underlying them play a central role in shaping workplace and public attitudes toward quotas, and constitute the primary outcome in much of the existing literature (see Section 2). Participants answered two open-ended questions on the perceived effectiveness of quotas and their willingness to support their introduction. We interpret these results alongside the behavioral outcomes. Table 5 reports estimates for the question "*Do you think that quota systems could be a useful tool for reducing inequalities?*" (DV = 1 if "Yes," 0 otherwise). Across specifications, all information treatments increase self-reported support relative to the No Info condition. The largest effect is observed for the *Demand* treatment (+14.1 percentage points), followed by *Info Causes* (+7.4 pp) and *Supply* (+5.7 pp). Women are also significantly more likely to support gender quotas. Models (1)–(6) progressively introduce additional controls: demographics (2), job characteristics (3), incentivized beliefs about the magnitude of gender gaps and issue prioritization (4), beliefs about causes (generic versus specific) (5), and detailed causes and policy views derived from open-ended responses (6). Focusing on Model (5), pairwise comparisons indicate that the effect of the *Demand* treatment is significantly larger than that of *Supply* ($p = 0.010$) and *Info Causes* ($p = 0.006$), while the latter two do not differ from each other

($p = 0.969$). These patterns are consistent with the preregistered hypotheses regarding the relative effectiveness of demand- versus supply-side narratives. In the most saturated specification (Model (6)), the indicator for *taste discrimination / reaction* is positive and marginally significant at the 10% level, and indicators for anti-discrimination policies and pro-quota stances are positive and statistically significant. Treatment effects remain stable across specifications, with the exception of *Info Causes*, whose coefficient is no longer statistically significant in the fullest model.

Narratives and ex-post reasoning

To better understand the reasoning underlying respondents' self-reported support for gender quotas, we analyze the open-ended justifications provided *immediately after treatment exposure*. After indicating whether quotas could address existing inequalities ("Yes/No"), respondents were asked to explain their answer in their own words. Using the coding procedure described in Section 6.1.3 and 6.1.2, we classify these free-text responses into a set of thematic categories, detailed in Appendix Tables F1–F2. In particular, Tables 6 and 7 report estimates of the effects of the information treatments on these *ex post* rationales. Appendix Table F3 reports the distribution of ex-post motivations by treatment group together with pairwise tests. For respondents answering "Yes," we code three main themes: i) *For Women*: quotas are viewed as addressing supply-side constraints (e.g., participation, confidence, role models); ii) *Signaling*: quotas are perceived as helping to reveal women's abilities that might otherwise be overlooked; iii) *Fairness*: quotas are framed as restoring equality by counteracting long-standing systemic disadvantages. For respondents answering "No," we code three corresponding themes: i) *Not the Right Instrument*: respondents argue that other policy tools would be more effective than quotas; ii) *Reverse Discrimination*: quotas are viewed as unfairly disadvantaging men; iii) *Unneeded / Backlash*: quotas are considered unnecessary or likely to generate adverse consequences, potentially harming women. Tables 6 and 7 report regression results in which each specification uses one ex-post justification category—elicited from the open-ended "Yes/No" question on quota effectiveness—as the dependent variable. Regressors include treatment indicators, the *Underestimation Index* from the incentivized belief-elicitation tasks, an indicator for whether the respondent views gender equality as a key future challenge for the firm (*Gender Priority*), and dummies capturing prior beliefs about effective policy responses. This empirical setup allows us to assess how information treatments reshape the rationales respondents use to justify their stated positions. Turning first to explanations offered in support of quotas (Table 6), Models (1) and (2) focus on justifications linking quota effectiveness to *supply-side* constraints. Exposure to the *Supply* treatment is associated with a higher likelihood of invoking these arguments, indicating that this narrative shifts perceptions toward viewing quotas as instruments that address

	Support for gender quotas					
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment (omitted category: No Info)						
Info Causes	0.074***	0.067**	0.062**	0.061**	0.062**	0.046
	(0.027)	(0.030)	(0.030)	(0.030)	(0.030)	(0.030)
Supply	0.057**	0.060**	0.064**	0.062**	0.062**	0.063**
	(0.027)	(0.028)	(0.028)	(0.028)	(0.028)	(0.028)
Demand	0.141***	0.127***	0.129***	0.131***	0.132***	0.107***
	(0.027)	(0.029)	(0.029)	(0.030)	(0.030)	(0.029)
Woman		0.047**	0.051**	0.050**	0.049**	0.041**
		(0.019)	(0.021)	(0.021)	(0.021)	(0.021)
Pre-Treatment measures						
Gender Priority				0.001	0.001	-0.002
				(0.006)	(0.006)	(0.006)
Underestimation Index				0.000	0.001	-0.021
				(0.027)	(0.027)	(0.027)
Pre-treatment beliefs: causes						
Cultural Causes					-0.011	
					(0.021)	
Women's Ability						0.010
						(0.020)
Work-Life Balance						0.002
						(0.021)
Taste Discrimination Reaction			-			0.052*
						(0.029)
No Idea (Causes)						-0.211***
						(0.061)
Pre-treatment beliefs: policy responses						
Against Discrimination Policy						0.099***
						(0.020)
Pro-Quota						0.267***
						(0.021)
Family Support Policy						0.039
						(0.025)
Know Quota						0.013
						(0.020)
No Idea (Policy)						0.031
						(0.062)
Demographics		✓	✓	✓	✓	✓
Job Characteristics			✓	✓	✓	✓
Observations	2,404	2,403	2,403	2,370	2,370	2,370

Notes: Linear probability models. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Dependent variable equals 1 if the respondent answers "Yes" to "Do you think a quota system could be a useful tool for reducing inequalities?" and 0 otherwise. *Controls:* Demographics — age group, born abroad, macro-region of residence (North, Center, South), college degree, marital status, has children, has a female child. Job characteristics — sector, firm type (public/private/non-profit), firm size (number of employees), job role (manager; entrepreneur; other managerial role; employee—white collar; self-employed; employee—blue collar/other).

Table 5: Effect of information treatments on self-reported support for gender quotas (LPM)

supply-side barriers. Models (3) and (4) examine *signaling* rationales—namely, the idea that quotas help reveal women's abilities that might otherwise be overlooked. Across these specifications, all treatment coefficients are positive and statistically significant. We next consider explanations given for opposing quotas (Table 7). Here, the *Demand* treatment reduces the probability that respondents cite *reverse discrimination* as a justification for opposition. By contrast, the *Info Causes* treatment—providing information about the sources of gender inequality without reference to quotas—raises the likelihood that respondents justify opposition by arguing that quotas are unnecessary or may generate backlash.

Overall, Tables 6 and 7 indicate that information treatments reshape ex-post rationales in systematic ways: supply-side messaging increases the salience of supply-based and signaling justifications in support of quotas, while demand-side messaging reduces the likelihood that opposition is framed in terms of reverse discrimination.

Beyond the *topics* respondents mention, we examine whether the information treatments affect the *way* respondents reason about gender quotas, as reflected in the structure of their open-ended justifications. As a proxy for argumentative breadth, we use the number of distinct thematic categories mentioned in each response. Using the coding scheme described above, we count how many justification categories are invoked within a message and estimate separate regressions for respondents who support quotas and for those who oppose them. Results are reported in Appendix Table F4. Among supporters, exposure to the information treatments—especially the *Supply* and *Demand* narratives—is associated with a higher number of distinct arguments, suggesting that these narratives broaden the set of considerations invoked in support of quotas. Among opponents, we find no systematic evidence that treatments affect the number of categories mentioned. Results are robust to allowing for interactions between treatments and the *Cultural Causes* indicator (Appendix Tables F4 -F5).

As a complementary proxy for the articulation of reasoning, we examine the length of respondents' justifications, measured by the number of words used in the free-form messages provided after the "Quota useful: Yes/No" question. Word counts are computed and regressed on treatment indicators and the same set of pre-treatment controls used in the main specifications. We run separate regressions for those who support quotas and those who do not. The corresponding results are reported in Appendix Table F6. Overall, we find limited evidence that treatments systematically change the length of respondents' reasoning, with one notable exception: exposure to the *Demand* narrative is associated with longer messages, both among respondents who report supporting or opposing quotas. The results are unchanged when allowing for heterogeneous effects by prior cultural explanations through interactions with the *Cultural Causes* indicator (Appendix Table F7).

These findings suggest that information treatments primarily affect the *content and*

	Ex-post rationales for those who support quotas					
	ForWomen		Signalling		Fairness	
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment (omitted category: No Info)						
Info Causes	0.001	0.005	0.098***	0.098***	-0.058*	-0.057*
	(0.019)	(0.019)	(0.025)	(0.025)	(0.031)	(0.031)
Supply	0.072***	0.076***	0.115***	0.119***	0.045	0.041
	(0.019)	(0.019)	(0.024)	(0.024)	(0.029)	(0.029)
Demand	0.012	0.014	0.189***	0.186***	0.056	0.060*
	(0.019)	(0.019)	(0.027)	(0.027)	(0.032)	(0.032)
Woman	0.004	0.001	0.042**	0.037*	0.054**	0.058***
	(0.014)	(0.014)	(0.019)	(0.019)	(0.022)	(0.022)
Pre-Treatment measures						
Gender Priority	0.004	0.004	0.009	0.008	0.000	0.001
	(0.004)	(0.004)	(0.006)	(0.006)	(0.006)	(0.006)
Underestimation Index	0.015	0.011	-0.051*	-0.060**	0.019	0.033
	(0.018)	(0.018)	(0.026)	(0.026)	(0.029)	(0.029)
Pre-treatment beliefs: causes						
Cultural Causes	-0.041***		-0.005		-0.027	-0.027
	(0.014)		(0.020)		(0.023)	(0.023)
Women's Ability		0.037**		0.039*	0.039*	0.036
		(0.015)		(0.020)	(0.023)	(0.023)
Work–Life Balance		0.062***		0.042**	0.030	0.030
		(0.016)		(0.020)	(0.023)	(0.023)
Taste Discrimination Reaction		-0.030		0.037	0.059*	0.059*
		(0.020)		(0.030)	(0.034)	(0.034)
No Idea (Causes)		-0.027		-0.144***	-0.204***	-0.147***
		(0.028)		(0.033)	(0.044)	(0.054)
Pre-treatment beliefs: policy responses						
Against Discrimination Policy	0.016	0.015	0.081***	0.079***	0.063***	0.067***
	(0.014)	(0.013)	(0.018)	(0.018)	(0.021)	(0.021)
Pro-Quota	0.043*	0.042*	0.097***	0.093***	0.178***	0.184***
	(0.025)	(0.025)	(0.034)	(0.034)	(0.036)	(0.036)
Family Support Policy	0.065***	0.050***	0.022	0.005	-0.007	0.003
	(0.018)	(0.018)	(0.022)	(0.023)	(0.026)	(0.025)
Know Quota	0.010	0.007	0.019	0.015	-0.018	-0.013
	(0.014)	(0.014)	(0.019)	(0.019)	(0.022)	(0.022)
No Idea (Policy)	-0.076***	-0.066***	-0.050	-0.022	-0.108**	-0.147***
	(0.019)	(0.019)	(0.043)	(0.044)	(0.054)	(0.054)
Demographics	✓	✓	✓	✓	✓	✓
Job Characteristics	✓	✓	✓	✓	✓	✓
Observations	2,370	2,370	2,370	2,370	2,370	2,370

Notes: Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.All specifications contains a set of controls. *Demographics:* age group; foreign-born; region (North, Center, South); degree; marital status; has children; has daughter. *Job characteristics:* sector; firm type (public, private, non-profit); firm size; job role. Cultural Causes is mutually exclusive; all other belief categories may co-occur.

Table 6: Explanations for supporting quotas

	Ex-post rationales for those who oppose quotas					
	Not Right Instrument		Reverse Discrimination		Unneeded / Backlash	
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment (omitted category: No Info)						
Info Causes	-0.029	-0.028	0.004	0.004	0.033**	0.033**
	(0.022)	(0.022)	(0.025)	(0.025)	(0.014)	(0.014)
Supply	-0.018	-0.015	-0.028	-0.026	0.002	0.003
	(0.021)	(0.021)	(0.023)	(0.023)	(0.011)	(0.011)
Demand	-0.032	-0.032	-0.054**	-0.055**	0.009	0.009
	(0.022)	(0.022)	(0.024)	(0.024)	(0.012)	(0.012)
Woman	-0.026*	-0.030**	-0.039**	-0.042**	-0.010	-0.011
	(0.015)	(0.015)	(0.018)	(0.018)	(0.010)	(0.010)
Pre-Treatment measures						
Gender Priority	0.003	0.003	0.002	0.002	0.004	0.004
	(0.005)	(0.005)	(0.005)	(0.005)	(0.003)	(0.003)
Underestimation Index	0.029	0.029	0.072***	0.071***	0.012	0.013
	(0.019)	(0.019)	(0.022)	(0.022)	(0.012)	(0.012)
Pre-treatment beliefs: causes						
Cultural Causes	0.026		0.004		-0.002	
	(0.016)		(0.018)		(0.009)	
Women's Ability		0.022		0.028		0.016
		(0.016)		(0.018)		(0.010)
Work–Life Balance		0.017		0.022		0.009
		(0.017)		(0.019)		(0.010)
Taste Discrimination Reaction		-0.021		0.005		-0.016
		(0.021)		(0.026)		(0.013)
No Idea (Causes)		-0.074**		-0.022		0.014
		(0.033)		(0.042)		(0.025)
Pre-treatment beliefs: policy responses						
Against Discrimination Policy	-0.043***	-0.042***	-0.033*	-0.034*	-0.021**	-0.021**
	(0.015)	(0.015)	(0.017)	(0.017)	(0.010)	(0.010)
Pro-Quota	-0.115***	-0.116***	-0.199***	-0.201***	-0.051***	-0.051***
	(0.015)	(0.015)	(0.017)	(0.017)	(0.009)	(0.009)
Family Support Policy	0.021	0.011	-0.032	-0.040*	-0.004	-0.006
	(0.018)	(0.019)	(0.020)	(0.022)	(0.011)	(0.012)
Know Quota	0.018	0.016	0.049***	0.048***	0.020**	0.020**
	(0.015)	(0.015)	(0.017)	(0.017)	(0.009)	(0.009)
No Idea (Policy)	0.008	0.020	-0.115***	-0.110***	-0.047***	-0.049***
	(0.046)	(0.047)	(0.040)	(0.040)	(0.018)	(0.016)
Demographics	✓	✓	✓	✓	✓	✓
Job Characteristics	✓	✓	✓	✓	✓	✓
Observations	2,370	2,370	2,370	2,370	2,370	2,370

Notes: Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. All specifications contains a set of controls. *Demographics:* age group; foreign-born; region (North, Center, South); degree; marital status; has children; has daughter. *Job characteristics:* sector; firm type (public, private, non-profit); firm size; job role. Cultural Causes is mutually exclusive; all other belief categories may co-occur.

Table 7: Explanations for opposing quotas

breadth of respondents' reasoning about gender quotas, while having more limited effects on the overall length of the arguments they provide. These analyses speak to how respondents *articulate* and justify their views, rather than to the behavioral mechanisms through which information affects policy support, which we examine next.

6.2.2 Incentivized measures of support

We next examine two incentivized behaviors: the decision to donate to an NGO advocating gender quotas and the amount pledged. Relative to self-reported attitudes, these outcomes are less prone to experimenter-demand concerns because they entail a tangible (though contingent) cost: lottery winners must donate the amount they stated. This setting allows us to assess the "say–do" gap—the extent to which stated support translates into costly action—and to distinguish extensive- from intensive-margin responses.

Donation decision (extensive margin).

We begin by examining whether exposure to the information treatments affects the likelihood of engaging in any costly support for gender quotas, measured by the decision to donate. Table **??** reports modified-logit estimates for the donation decision, where the dependent variable equals one if the respondent is willing to donate any positive amount to an NGO supporting quotas in the event of a lottery win; the estimator adjusts for the forced–response design used in the donation item.[8] All specifications include treatment indicators (with No Info as the omitted category) and sequentially add the same sets of controls used in Table 5.

In contrast to the results for unincentivized self-reported support, information treatments have virtually no effect on the willingness to donate. *Gender Priority* is positive and statistically significant in the richer specifications. Conditioning on prior beliefs, women appear less likely to donate, although the estimate is not consistently precise across models. Prior views play an important role: indicators for anti-discrimination policy preferences and pro-quota stances are positive and statistically significant, as is *KnowQuota*, highlighting the role of familiarity and policy preferences in incentivized support. Results are robust across specifications and in subgroup analyses by gender, job role, and age.

Amount donated (intensive margin)

We next restrict the sample to respondents who choose to donate and examine the causal effect of information on the *amount* donated. Among donors, the average contribution is €87 (SD = 113; range €1–€500), compared to €64 in the control group. Table 8 reports OLS estimates with ln(donation) as the dependent variable. Each specification includes

[8] Estimates obtained using the `RRreg` package: Heck et al. 2019.

treatment indicators (with No Info as the omitted category) and sequentially adds the same sets of controls used in the previous tables.

Across specifications, the *Supply* treatment exhibits a large and robust effect: conditional on donating, contributions increase by about 0.66–0.72 log points relative to the control group. The *Demand* treatment displays a smaller positive effect in the baseline specification but becomes statistically indistinguishable from zero once controls are included. In the most saturated specification (Model (6)), respondents who attribute gender inequalities to *women's ability* donate more (marginally significant), and those endorsing *anti-discrimination* policies also contribute higher amounts. Underestimation of the gender gap is negatively associated with donation levels, though not statistically significant. Finally, the coefficient on *Woman* is negative and statistically significant across models, indicating substantially lower donation amounts among women conditional on donating.

Summary and interpretation.

Looking across outcomes, treatment effects differ sharply across expressive and costly measures of support. The *Demand* and *Supply* narratives increase unincentivized self-reported support for gender quotas (Table 5). By contrast, none of the treatments affects the incentivized *decision* to donate (Table **??**), indicating that framing does not shift the extensive margin of costly engagement. However, conditional on donating, the *Supply* treatment significantly increases the *amount* donated (Table 8), revealing that the intensive margin is sensitive to narrative–prior alignment rather than to framing per se. This contrast suggests a behavioral distinction between expressive approval and aligned costly action. Stated support appears responsive to framing in general, consistent with expressive or low-cost attitude adjustment. In contrast, costly behavior is largely insensitive to framing at the participation stage, but becomes responsive once narratives align with individuals' prior causal beliefs. Informational frames therefore do not expand the set of donors; rather, they intensify engagement among those whose priors are congruent with the message. This distinction highlights the importance of combining expressive and incentivized measures: while self-reports capture framing-sensitive attitudes, donation behavior isolates alignment-sensitive costly engagement.

7 Narrative–Prior Alignment and Behavioral Responses

This section tests the central mechanism underlying our findings: narrative effectiveness depends on the alignment between informational frames and individuals' prior causal beliefs about gender inequality. As documented in Section 6.1.2 and Appendix D, prior beliefs about the causes of gender inequality are systematically structured. They vary predictably

	Amount donated (log)					
	(1)	(2)	(3)	(4)	(5)	(6)
Treatment (omitted category: No Info)						
Info Causes	0.271	0.009	0.018	0.001	0.001	-0.058
	(0.271)	(0.320)	(0.339)	(0.339)	(0.339)	(0.329)
Supply	0.660***	0.638**	0.683**	0.715***	0.715***	0.655**
	(0.248)	(0.253)	(0.271)	(0.273)	(0.274)	(0.269)
Demand	0.450*	0.214	0.218	0.193	0.193	0.185
	(0.251)	(0.302)	(0.323)	(0.325)	(0.326)	(0.317)
Woman		-0.445**	-0.507**	-0.496**	-0.497**	-0.527**
		(0.180)	(0.197)	(0.200)	(0.200)	(0.202)
Pre-Treatment measures						
Gender Priority				0.020	0.020	0.018
				(0.054)	(0.054)	(0.055)
Underestimation Index				-0.065	-0.065	-0.040
				(0.247)	(0.247)	(0.248)
Pre-Treatment beliefs: causes						
Cultural Causes					-0.008	
					(0.212)	
Women's Ability						0.374*
						(0.197)
Work–Life Balance						-0.062
						(0.214)
Taste Discrimination Reaction						-0.320
						(0.351)
No Idea (Causes)						0.532
						(0.614)
Pre-Treatment beliefs: policy responses						
Against Discrimination Policy						0.397*
						(0.211)
Pro-Quota						0.148
						(0.312)
Family Support Policy						0.000
						(0.221)
Know Quota						0.035
						(0.226)
No Idea (Policy)						-1.616**
						(0.648)
Demographics		✓	✓	✓	✓	✓
Job Characteristics			✓	✓	✓	✓
Observations	416	415	415	412	412	412

Notes: Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. *Demographics:* age group; born abroad; macro–region of residence (North, Center, South); college degree; marital status; number of children; has daughter (indicator). *Job characteristics:* sector; firm type (public, private, non-profit); firm size (number of employees); job role (manager; entrepreneur; other managerial role; employee—white collar; self-employed; employee—blue collar/other).

Table 8: Effect of treatments on amount donated (log)(range: 1–500 €)

with respondents' demographics, job characteristics, and regional labor-market outcomes, suggesting that they reflect stable normative environments rather than idiosyncratic survey noise. This motivates our focus on narrative–prior alignment as the key mechanism linking information to policy support. The previous section documented a divergence between unincentivized and incentivized outcomes: information treatments increase stated support for gender quotas but do not affect the decision to donate, while the *Supply* narrative increases donation amounts among donors. We now examine whether these patterns are driven by heterogeneous responses across prior-belief types. In particular, we ask whether the intensive-margin effects documented above emerge precisely among individuals whose priors are weak or aligned with the transmitted narrative. To test this mechanism, we estimate models that interact treatment indicators with respondents' pre-treatment belief categories about the causes of inequality. We estimate these models for all three outcomes: self-reported support (*QuotaUseful*), the donation decision (*Decision to Donate*), and the amount donated (*Donation*). The resulting interaction coefficients are plotted in Figures 2, 3, and 4.

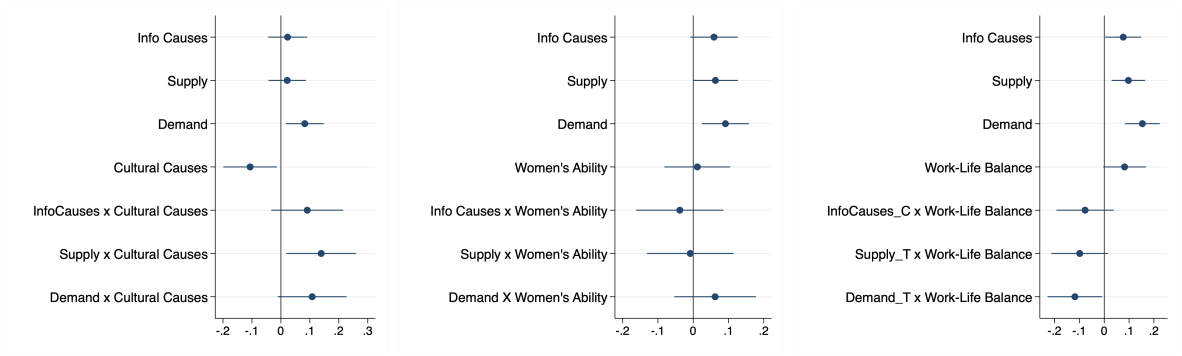

Figure 2: Comparison of belief interactions on self-reported support (QuotaUseful)

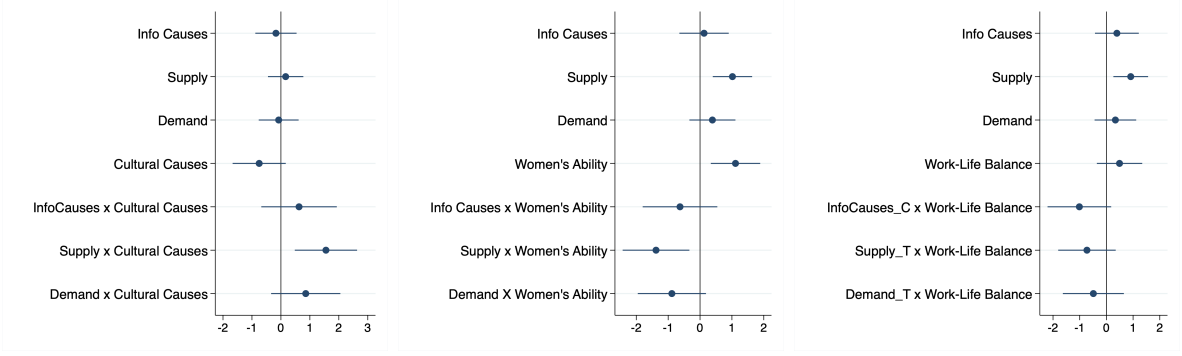

Figure 3: Comparison of belief interactions on ln(amount donated)

Each figure corresponds to one outcome. For each outcome, we estimate regressions that interact the main prior-belief categories about the causes of inequality with the treatment indicators, controlling for demographics, job characteristics, the Underestimation

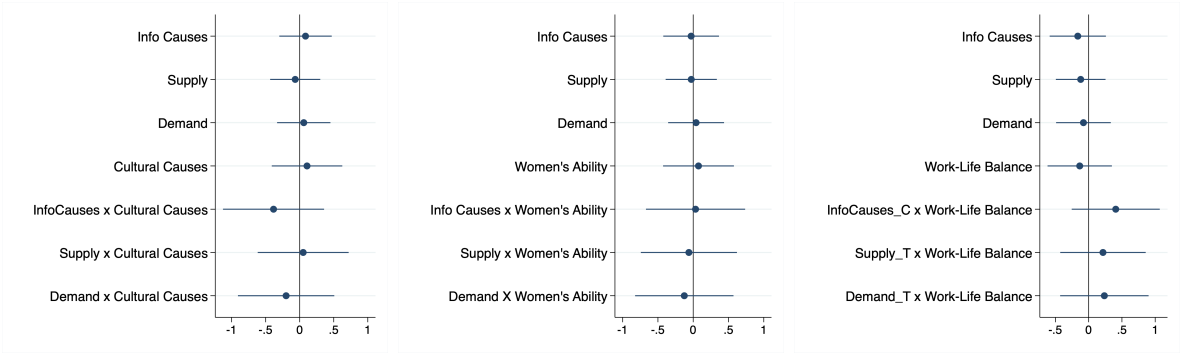

Figure 4: Comparison of belief interactions on decision to donate

Index, and prior beliefs about effective policy responses. Consistent with earlier results, we find no significant interactions for the *donation decision*. By contrast, prior beliefs systematically moderate treatment effects for both *self-reported support* and the *amount donated*. This absence of interaction effects on the extensive margin reinforces the earlier finding that framing does not expand participation in costly action. Heterogeneous effects instead emerge on the intensive margin. When interacting treatments with *Cultural Causes*—the generic prior-belief category capturing diffuse or non-committal views—the *Supply* narrative is the only one that increases support. This effect is visible for both *QuotaUseful* and *Donation*. In the *No Info* condition, respondents with generic beliefs exhibit lower baseline support (marginally so for donation amounts), indicating that individuals with diffuse priors are initially less supportive but more responsive to clarifying information. By contrast, when interacting treatments with more specific prior beliefs—such as *Women's Ability* or *Work–Life Balance*—treatment effects are weak or negative. When respondents hold well-defined causal narratives, informational messages do not increase support and may even backfire. Because the Women's Ability category captures a mix of perceptions about competence, stereotypes, and social norms, we interpret its interactions cautiously and focus primarily on the contrast between diffuse cultural beliefs and work–life balance explanations when distinguishing supply- and demand-side mechanisms.

Taken together, the results confirm that narrative effectiveness is governed by prior specificity and alignment, rather than by framing alone. Informational messages primarily affect respondents whose causal beliefs are diffuse rather than firmly articulated. Among competing narratives, only the *Supply* frame translates into higher donation amounts, consistent with the idea that clarifying how quotas address participation constraints reduces misalignment costs for a broader set of individuals.

Importantly, prior beliefs in the sample are not symmetrically distributed. As documented in Section 6.1.2, supply-side explanations—especially work–life balance constraints—are more prevalent than discrimination-based accounts. The interaction results therefore imply that, given the observed belief distribution, the *Supply* narrative minimizes

belief–message misalignment for a larger share of respondents, especially those located near the center of the belief distribution, where misalignment costs are relatively small. Motivated by this pattern, the next section introduces a conceptual framework in which donors experience disutility from narrative–belief misalignment and in which the agency's optimal message is shaped by the distribution of priors, especially in the tails. This framework formalizes why, given the prevailing distribution of beliefs in our sample, the supply-side narrative uniquely increases costly support, and clarifies how alternative belief distributions would generate different optimal narratives and treatment effects.

8 The Conceptual Framework

The framework formalizes the empirical patterns documented in Sections 6 and 7, where we show that treatment effectiveness depends on the distribution and specificity of prior causal beliefs. The data show two key patterns: (i) informational narratives primarily affect respondents with diffuse prior beliefs, while individuals with well-defined causal views exhibit little responsiveness; and (ii) only the Supply narrative increases donation amounts on the intensive margin. We develop a simple framework that rationalizes both findings through narrative-prior alignment and tail-driven responsiveness.

We model the interaction between an agency that communicates about gender inequality and a heterogeneous population of potential donors. The agency chooses which narrative to transmit in order to maximize expected donations.

There is a unit mass of donors, normalized to one. Donors hold heterogeneous prior beliefs about the origins of gender inequality. When exposed to a message that conflicts with their priors, they experience ideological disutility from misalignment.[9] At the same time, information may provide intrinsic value by clarifying the issue. Donors therefore face a trade-off between informational benefits and misalignment costs. If ideological disutility exceeds informational benefits, donors optimally ignore the message and rely solely on their prior beliefs when deciding how much to donate (see, e.g., Bénabou and Tirole 2011).[10]

Let the variable $m \in \{0, 1\}$ denote the message transmitted by the agency, with $m = 0$ representing the supply-side narrative and $m = 1$ the demand-side one. Each donor has a type θ drawn from a publicly known continuous distribution function $F(\theta)$ over the support $[0, 1]$, with density $f(\theta) \geq 0$. The interval $[0, 1]$ captures the heterogeneity in

[9]This assumption is consistent with the literature on belief utility and cognitive consistency, which highlights that individuals derive utility not only from outcomes but also from holding beliefs that align with their priors (Bénabou and Tirole 2011; Yariv 2002; Bendersky et al. 2022).

[10]In the experimental literature, Momsen and Ohndorf 2023 and Nam et al. 2013, among others, show that individuals sometimes deliberately avoid costless information about payoffs and the consequences of donations when it conflicts with their political predispositions. Moreover, patterns of information avoidance are systematically related to ex ante ideology.

the strength of donors' beliefs about the two competing narratives. The setup parallels a Hotelling 1929 spatial model: beliefs correspond to locations along an ideological line, and the distance between a donor's belief and the agency's message shapes engagement.[11]

When receiving message m, a donor of type θ chooses a donation D to maximize net utility:

$$U(D) = V(D) - C(D; \theta, m), \tag{1}$$

where $V(D)$ denotes the utility from giving and $C(D; \theta, m)$ the associated cost.

We assume that:

$$V(D) = \frac{D^\alpha}{\alpha}, \text{ with } \alpha \in (0, 1) \tag{1.1}$$

and

$$C(D; \theta, m) = \left[1 - \lambda \cdot |\theta - m| + t_m \cdot (\theta - m)^2\right] \cdot D, \ \lambda \in (0, 1), \ t_m > 0. \tag{1.2}$$

In Equation (1.2), the term $\lambda \cdot |\theta - m|$ captures the informational benefit associated with exposure to a message. This benefit reduces the marginal cost of donating and increases linearly (at rate λ) with the distance between prior beliefs and the message received. Intuitively, the larger the gap between the donor's prior and the message, the greater the potential for belief revision. When the message coincides with prior beliefs, it conveys little new insight; when it differs, it forces the donor to process discrepant information, potentially refining or reconsidering her views. The parameter λ therefore measures the marginal value of belief revision induced by ideological distance.

Conversely, the term $t_m \cdot (\theta - m)^2$ captures the ideological disutility that grows quadratically with misalignment between the donor's belief and the message's content. The quadratic form implies that small discrepancies are relatively easy to accommodate, but larger deviations generate disproportionately higher psychological or identity costs. While information may be more valuable when it is surprising, misalignment becomes increasingly costly as the gap widens. Beyond a certain threshold, the discomfort associated with holding or engaging with discrepant views outweighs the informational benefit, leading donors to disengage from the message altogether.

If ideological disutility dominates informational benefits, the donor optimally ignores the message and bases her donation decision solely on her prior belief. Ignoring the message is equivalent to assuming that the perceived narrative coincides with the donor's prior, so that no informational gain or ideological misalignment arises. Formally, this corresponds to evaluating Eq. (1.2) at $m = \theta$, which yields $C(D; \theta, \theta) = D$. As a benchmark, consider the case in which no communication occurs.

In the absence of any message, donors make their decision based solely on their prior

[11] Our model shares a Hotelling-type structure with Mullainathan and Shleifer 2005, where media outlets choose ideological positions to appeal to heterogeneous audiences. Unlike in their framework, where news consumption generates only costs, information in our setting carries intrinsic value but may also entail disutility when misaligned with prior beliefs.

beliefs, and neither informational benefits nor ideological misalignment arise. The marginal cost of donating therefore reduces to one, implying

$$D^{NI} = 1, \text{ for any } \theta \in [0,1]\,. \tag{2}$$

Maximizing $U(D)$ leads to the following result:

Proposition 1. *A donor exposed to message $m \in \{0,1\}$ donates:*

$$D^*(\theta; m) = \begin{cases} \left[1 - \lambda \cdot |\theta - m| + t_m \cdot (\theta - m)^2\right]^{-\frac{1}{1-\alpha}}, & \text{if } |\theta - m| \leq \lambda / t_m, \\ 1, & \text{if } |\theta - m| > \lambda / t_m. \end{cases} \tag{3}$$

According to (3), donations peak at $\left|\overline{\theta} - m\right| = \frac{\lambda}{2 \cdot t_m}$, and all donors contribute at least as much as under no communication. When beliefs and the message are perfectly aligned ($\theta = m$), $D^*(m; m) = 1$: the message yields neither informational benefit nor disutility. If $\lambda \geq t_m$, the informational benefit dominates for all other types, and $D^*(\theta; m) > 1$. Conversely, when $\lambda < t_m$, only donors sufficiently close to the message ($|\theta - m| \leq \lambda / t_m$) respond to it; others ignore the communication and donate $D^*(\theta; m) = 1$.

In the empirical analysis, respondents invoking generic cultural explanations provide a proxy for relatively diffuse priors, whereas those citing specific mechanisms correspond to individuals with more sharply defined belief locations along the ideological continuum. In our framework, "diffuse" priors are those for which the perceived narrative is less sharply pinned down, making donors more likely to fall within the region where informational gains outweigh misalignment costs. While this mapping is necessarily imperfect, it captures the key dimension emphasized by the model: heterogeneity in the structure and strength of prior beliefs. The heterogeneity patterns in Section 7 can thus be interpreted as empirical counterparts of variation in belief distance.

8.1 The Agency's Problem

The agency chooses the message m that maximizes expected total donations:

$$\max_{m \in \{0,1\}} \Pi(m) = \int_0^1 D^*(\theta; m) \cdot f(\theta) \cdot d\theta \tag{4}$$

where

$$D^*(\theta; 0) = \begin{cases} \left(1 - \lambda \cdot \theta + t_0 \cdot \theta^2\right)^{-\frac{1}{1-\alpha}}, & \theta \leq \lambda / t_0, \\ 1, & \theta > \lambda / t_0. \end{cases} \tag{4.1}$$

and

$$D^*(\theta; 1) = \begin{cases} \left[1 - \lambda \cdot (1 - \theta) + t_1 \cdot (1 - \theta)^2\right]^{-\frac{1}{1-\alpha}}, & \theta \geq 1 - \lambda / t_1, \\ 1, & \theta < 1 - \lambda / t_1. \end{cases} \tag{4.2}$$

Solving Problem (4) leads to the following result:

Proposition 2. *The agency's optimal message is*

$$m^* = \begin{cases} 0, & \text{if } \Delta \geq 0, \\ 1, & \text{otherwise}, \end{cases} \tag{5}$$

where $\Delta = \int_0^1 (\frac{dD^*(\theta,0)}{d\theta} - \frac{dD^*(\theta,1)}{d\theta}) \cdot (1 - F(\theta)) \cdot d\theta$.

The intuition behind Proposition (2) is immediate. The relative appeal of the two narratives depends on how individual donations vary across donor types, as reflected in the term $\frac{dD^*(\theta,0)}{d\theta} - \frac{dD^*(\theta,1)}{d\theta}$ within Δ, weighted by the upper tail of the donor distribution, $1 - F(\theta)$. The weight $1 - F(\theta)$ implies that changes in donations among donors with extreme priors have disproportionate leverage on expected fundraising. Consequently, the optimal narrative is not determined by the average belief but by those belief types whose donations are most responsive to ideological alignment.

When $\lambda \geq t_0, t_1$, this comparison spans the entire interval $[0, 1]$. When $\lambda < t_0, t_1$, only donors for whom $D^*(\theta; m) > 1$ respond to communication. In that case, the condition $\Delta \geq 0$ reduces to:

$$\int_0^{\frac{\lambda}{t_0}} \frac{dD^*(\theta,0)}{d\theta} \cdot (1 - F(\theta)) \cdot d\theta \geq \int_{1-\frac{\lambda}{t_1}}^1 \frac{dD^*(\theta,1)}{d\theta} \cdot (1 - F(\theta)) \cdot d\theta$$

In this case, the relative size and responsiveness of the active intervals—where donations exceed the baseline—determine the optimal communication strategy.

The structure of Δ yields clear comparative statics.

First, if the distribution of priors shifts toward demand-side beliefs (i.e., the density increases near $\theta = 1$), the demand-side narrative ($m = 1$) becomes more likely to maximize expected donations; conversely, when beliefs are concentrated toward supply-side explanations (θ closer to 0), the supply narrative ($m = 0$) is optimal.

Second, ideological costs shape the optimal narrative. If misalignment is less costly under one message (e.g., $t_0 < t_1$), that narrative is optimal for a wider range of belief distributions. To see this, consider the effect of increasing t_1, the misalignment penalty associated with the demand-side message. Differentiating expected donations under $m = 1$ yields

$$\frac{d\Pi(1)}{dt_1} = \begin{cases} \int_0^1 \frac{dD^*(\theta,1)}{dt_1} \cdot f(\theta) \cdot d\theta < 0 & \text{if } t_0 < t_1 \leq \lambda, \\ \int_{1-\frac{\lambda}{t_1}}^1 \frac{dD^*(\theta,1)}{dt_1} \cdot f(\theta) \cdot d\theta < 0 & \text{if } \lambda < t_0 < t_1. \end{cases}$$

In both regimes, a higher misalignment cost under the demand-side message reduces aggregate donations. Consequently, as the ideological penalty associated with $m = 1$ increases relative to that under $m = 0$, the demand-side narrative becomes less likely to maximize expected fundraising.

Third, when priors are widely dispersed—so that many donors are not closely aligned with either narrative—misalignment costs dominate informational benefits for most types, leading to muted average treatment effects. The effectiveness of informational campaigns therefore depends jointly on the location and the dispersion of prior beliefs: persuasion is inherently context-dependent.

The model provides a structural interpretation of the experimental findings. When donors hold well-defined prior beliefs, misalignment costs dominate informational gains, leading to muted or even negative treatment effects—consistent with the weak or backfiring effects observed among respondents with specific priors.

Moreover, because aggregate donations are effectively weighted by the upper tail of the belief distribution, the agency's optimal message is shaped by the prevailing orientation of those extreme priors. In our data, prior beliefs are predominantly oriented toward supply-side explanations, particularly work-life balance constraints (Section 6.1.2). Under the empirically observed distribution of priors, a supply-side narrative minimizes ideological misalignment for a larger share of high-response donors, thereby increasing donation amounts on the intensive margin among responsive donors. This mechanism explains why only the Supply treatment generates significant behavioral effects.

9 Conclusion

We ran a survey experiment with 2,404 Italian workers to study what shapes support for a highly debated affirmative-action tool: gender quotas. Treatments provided short informational narratives framed around either *demand*-side barriers (biases, discrimination) or *supply*-side barriers (participation constraints, confidence, self-selection). Before treatment exposure, we elicited respondents' priors about both the magnitude and—crucially—the *causes* of gender inequality, allowing us to examine how narrative effectiveness depends on pre-existing causal beliefs.

Across outcomes, we document a clear divergence between unincentivized attitudes and incentivized behavior. Information increases *stated* support for quotas, while the donation decision is largely unchanged; by contrast, among donors, the *Supply* narrative robustly increases the *amount* donated. This divergence highlights that narratives can shift approval without necessarily inducing costly action on the extensive margin, while still affecting the intensity of engagement. Our main contribution is to show that the effectiveness of informational framing hinges on the structure of individuals' prior beliefs about the *causes* of inequality. Treatments affect not only *whether* respondents express support, but also *how* they reason about quotas. Using coded open-ended responses, we document that narratives reshape the content and breadth of justifications, particularly among respondents with diffuse prior beliefs. Mechanism evidence indicates that pre-existing causal beliefs are the central moderator of treatment effects. Informational narratives

primarily influence respondents with diffuse or non-committal priors, whereas individuals holding specific causal explanations exhibit little responsiveness and may even react negatively. Moreover, among competing narratives, only the *Supply* frame translates into higher donation amounts. We formalize these patterns in a simple conceptual framework featuring two complementary mechanisms: (i) narrative–belief misalignment generates ideological disutility that limits responsiveness among individuals with strong priors, and (ii) the agency's optimal message is shaped by donors in the tails of the belief distribution, whose contributions are most sensitive to alignment. Because prior beliefs in our sample are predominantly oriented toward supply-side explanations, a supply-oriented narrative minimizes misalignment for a larger share of high-response donors, thereby increasing costly support on the intensive margin. In discrimination-focused contexts, demand-side narratives would instead optimize costly support. Our findings bridge the literature on belief-based policy demand and the study of quota legitimacy. They show that informational campaigns are most effective when they engage with the specificity and orientation of existing causal narratives. More broadly, the results highlight that durable support for equity-enhancing policies depends not only on communicating facts about inequality, but on addressing how individuals understand *why* inequality persists. This belief distribution can be measured from survey data, enabling context-specific narrative design. More broadly, our results suggest that the success of informational campaigns depends on their alignment with the distribution of beliefs in the population, implying that effective policy communication must be tailored to the prevailing narrative environment.

Declaration of generative AI and AI-assisted technologies in the writing process

During the preparation of this work, the authors used ChatGPT 5 to improve language and readability, with caution. After using this tool, the authors reviewed and edited the content as needed and took full responsibility for the content of the publication.

References

C. D. Acutis, A. Weber, and E. Wurm. The effects of board gender quotas: A meta-analysis. *Labour Economics*, 91:102634, 2024. doi: 10.1016/j.labeco.2024.102634. Open access.

A. Alesina, M. F. Ferroni, and S. Stantcheva. Perceptions of racial gaps, their causes, and ways to reduce them. Technical report, National Bureau of Economic Research, 2021.

M. Arve. Underrepresentation, quotas and quality: A dynamic argument for reform. *unknown journal / working paper*, 2023. June 2023 version available.

F. Barigozzi and N. Montinari. Social norms in survey experiments: Personal beliefs versus normative expectations. *Journal of Economic Behavior & Organization*, 239: 107292, 2025.

L. Beaman, R. Chattopadhyay, E. Duflo, R. Pande, and P. Topalova. Powerful women: does exposure reduce bias? *The Quarterly journal of economics*, 124(4):1497–1540, 2009.

R. Bénabou and J. Tirole. Identity, morals, and taboos: Beliefs as assets. *Quarterly Journal of Economics*, 126(2):805–855, 2011. doi: 10.1093/qje/qjr002.

C. Bendersky, M. Zeelenberg, and E. Van Dijk. Motivated beliefs and information processing. *American Economic Review: Insights*, 4(2):123–138, 2022. doi: 10.1257/aeri.20200415.

M. Bertrand. 17-new perspectives on gender. handbook of labor economics, volume 4b, 4 (11), 1543–1590, 2011.

M. Bertrand. Coase lecture–the glass ceiling. *Economica*, 85(338):205–231, 2018.

G. Blair, K. Imai, and Y.-Y. Zhou. Design and analysis of the randomized response technique. *Journal of the American Statistical Association*, 110(511):1304–1319, 2015.

F. D. Blau and L. M. Kahn. The gender wage gap: Extent, trends, and explanations. *Journal of economic literature*, 55(3):789–865, 2017.

P. Bordalo, K. Coffman, N. Gennaioli, and A. Shleifer. Beliefs about gender. *American Economic Review*, 109(3):739–773, 2019.

J. G. Bullock, A. S. Gerber, S. J. Hill, and G. A. Huber. Partisan bias in factual beliefs about politics. Technical report, National Bureau of Economic Research, 2013.

S. S. Bush. Gender quotas and international reputation. *American Journal of Political Science*, 2020. fill in volume/issue/pages.

P. Cortés and J. Pan. Children and the remaining gender gaps in the labor market. *Journal of Economic Literature*, 61(4):1359–1409, 2023.

J. De Quidt, J. Haushofer, and C. Roth. Measuring and bounding experimenter demand. *American Economic Review*, 108(11):3266–3302, 2018.

D. Del Boca and M. Locatelli. Gender gaps, care responsibilities, and labour supply in italy. *XXX Journal*, XX:XXX–XXX, 2022.

European Commission. Directive (eu) 2022/2381 of the european parliament and of the council on improving the gender balance among directors of listed companies and related measures. `https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32022L2381`, 2022. Official Journal of the European Union, L315/44, 7 December 2022.

Eurostat. Labour force survey main indicators by sex and age. `https://ec.europa.eu/eurostat/web/lfs/data`, 2023. Accessed: DATE.

S. Falconieri, M. Ortiz, F. Urzua, and P. Volpin. Board gender quotas and female ceos in private firms. Technical report, 2025.

K. Faniko, D. Bourguignon, O. Sarrasin, and S. Guimond. "when meritocracy opposes quota policy: How education and policy strength impact opinions". *Personality and Social Psychology Bulletin*, 43(5):638–651, 2017.

G. Ferrari, V. Ferraro, P. Profeta, and C. Pronzato. Gender quotas: Challenging the boards, performance, and the stock market. IZA Discussion Paper 10239, IZA Institute of Labor Economics, Bonn, 2016.

B. Ferrario and S. Stantcheva. Eliciting people's first-order concerns: Text analysis of open-ended survey questions. In *AEA Papers and Proceedings*, volume 112, pages 163–169. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, 2022.

N. M. Fortin. Gender role attitudes and the labour-market outcomes of women across oecd countries. *oxford review of Economic Policy*, 21(3):416–438, 2005.

M. Francesconi and M. Parey. Early gender gaps among university graduates. *European economic review*, 109:63–82, 2018.

C. Goldin, L. F. Katz, and I. Kuziemko. The homecoming of american college women: The reversal of the college gender gap. *Journal of Economic perspectives*, 20(4):133–156, 2006.

C. Günther, N. A. Ekinci, C. Schwieren, and M. Strobel. Women can't jump?—an experiment on competitive attitudes and stereotype threat. *Journal of Economic Behavior & Organization*, 75(3):395–401, 2010.

I. Haaland and C. Roth. Beliefs about racial discrimination and support for pro-black policies. *Review of Economics and Statistics*, 105(1):40–53, 2023.

D. W. Heck, M. Moshagen, and M. D. W. Heck. Package 'rrreg'. 2019.

H. Hotelling. Stability in competition. *The Economic Journal*, 39(153):41–57, Mar. 1929.

ISTAT. I nuovi indicatori sulle posizioni lavorative dipendenti per tempo di lavoro, 2021. URL `https://www.istat.it/`. Retrieved online from ISTAT.

O. Kuzmina and V. Melentyeva. Gender diversity in corporate boards: Evidence from quota-implied discontinuities. *Available at SSRN 3904886*, 2021.

E. Latura and W. Weeks. Corporate board quotas and gender equality policies in the workplace. *American Journal of Political Science*, 2022. fill in volume/issue/pages.

F. Lefley, G. Trnková, and H. Vychová. Views on corporate board gender quotas: Evidence from the czech republic. *Corporate Communications: An International Journal*, 29(2): 123–138, 2024. doi: 10.1108/CCIJ-10-2023-0137.

A. Manning and J. Swaffield. The gender gap in early-career wage growth. *The Economic Journal*, 118(530):983–1024, 2008.

K. Momsen and M. Ohndorf. Information avoidance: Self-image concerns, inattention, and ideology. *Journal of Economic Behavior & Organization*, 211:386–400, 2023.

S. Mullainathan and A. Shleifer. The market for news. *American economic review*, 95(4): 1031–1053, 2005.

J. Mummolo and E. Peterson. Demand effects in survey experiments: An empirical assessment. *American Political Science Review*, 113(2):517–529, 2019.

R. Muriaas and collaborators. Attitudes to gender quotas: Why and where to adjust gender imbalance in leadership. *European Journal of Political Research*, 2024. fill in volume/issue/pages.

H. H. Nam, J. T. Jost, and J. J. Van Bavel. "not for all the tea in china!" political ideology and the avoidance of dissonance-arousing situations. *PloS one*, 8(4):e59837, 2013.

M. Niederle and L. Vesterlund. Do women shy away from competition? do men compete too much? *The quarterly journal of economics*, 122(3):1067–1101, 2007.

OECD. Oecd recommendation on gender equality in education, employment and entrepreneurship, 2013. URL https://www.oecd.org/gender/OECD-Recommendation-Gender-Equality.pdf.

OECD. Family database: Pf2.1 parental leave. https://www.oecd.org/social/family-database.htm, 2023. Accessed: DATE.

OECD. Gender equality in a changing world, 2023. URL https://www.oecd.org/en/publications/gender-equality-in-a-changing-world_e808086f-en/.

C. Porter and D. Serra. Gender differences in the choice of major: The importance of female role models. *American Economic Journal: Applied Economics*, 12(3):226–254, 2020.

M. Prior, G. Sood, K. Khanna, et al. You cannot be serious: The impact of accuracy incentives on partisan bias in reports of economic perceptions. *Quarterly Journal of Political Science*, 10(4):489–518, 2015.

E. Reuben, P. Sapienza, and L. Zingales. How stereotypes impair women's careers in science. *Proceedings of the National Academy of Sciences*, 111(12):4403–4408, 2014.

C. Saraceno. Gender cultures and care regimes in southern europe. In XXX and XXX, editors, *Handbook on Gender and Social Policy*. Edward Elgar, 2020.

S. Settele. How do beliefs about the gender wage gap affect the demand for public policy? *American Economic Journal: Economic Policy*, 14(2):475–508, 2022.

S. Stantcheva. Understanding of trade. Technical report, National Bureau of Economic Research, 2022.

L. Yariv. I'll see it when i believe it: A simple model of cognitive consistency. Cowles foundation discussion paper no. 1352, Cowles Foundation for Research in Economics, Yale University, 2002. URL https://cowles.yale.edu/sites/default/files/files/pub/d13/d1352.pdf.

L. Zaccaria, F. Schivardi, and L. Guiso. Changing the board game: Horizontal spillovers of gender quotas. *Available at SSRN*, 2024.

M. K. Zehnter and C. Nater. Beyond being beneficiaries: Two mechanisms explain why women have more favourable attitudes towards gender quotas than men. *European Journal of Social Psychology*, 2024. fill in volume/issue/pages.

A Summary statistics, sample balance and attrition

Appendix A presents summary statistics and tests of sample balance and attrition. Tables A1–A2 report pairwise balance tests for demographic and job-related characteristics. Table A3 examines balance in the pre-treatment prioritization of gender equality. Tables A4-A5 document attrition rates and test for selective survey completion.

	(1) Demad_T	(2) Supply	(3) Info Causes	(4) No Info	(5) Total	(6) p-value (1) = (2)	(7) p-value (1) = (3)	(8) p-value (1) = (4)	(9) p-value (2) = (3)	(10) p-value (2) = (4)	(11) p-value (3) = (4)
Women	370	350	334	268	1322	0.034	0.325	0.016	0.272	0.679	0.151
Age 20-25	6	4	13	9	32	0.476	0.082	0.245	0.016	0.065	0.596
Age 26-35	108	63	110	74	354	0.000	0.557	0.184	0.000	0.009	0.062
Age 36-44	193	183	165	140	681	0.237	0.249	0.175	1.000	0.802	0.806
Age 45-54	187	236	189	168	780	0.022	0.453	0.325	0.135	0.241	0.797
Age 55-64	104	125	90	96	415	0.265	0.493	0.375	0.073	0.869	0.126
Age 65-87	31	51	29	31	142	0.040	0.956	0.432	0.039	0.247	0.410
Center	115	117	94	66	392	0.786	0.243	0.010	0.361	0.020	0.151
South	183	225	142	203	753	0.056	0.037	0.000	0.000	0.068	0.000
North	325	315	357	244	1241	0.150	0.004	0.124	0.000	0.851	0.000
Has Children	380	418	354	310	1462	0.297	0.717	0.845	0.163	0.235	0.879
Bachelor Degree or More	320	318	266	256	1160	0.321	0.029	0.625	0.217	0.655	0.110
Observations	629	661	596	518	2404						

Table A1: Demographics: Number of observations and balance test, p-value from t-tests

	(1) Demad_T	(2) Supply	(3) Info Causes	(4) No Info	(5) Total	(6) p-value (1) = (2)	(7) p-value (1) = (3)	(8) p-value (1) = (4)	(9) p-value (2) = (3)	(10) p-value (2) = (4)	(11) p-value (3) = (4)
Manager	75	74	86	85	320	0.683	0.195	0.029	0.086	0.009	0.361
Executive	65	77	27	24	193	0.451	0.000	0.000	0.000	0.000	0.935
Middle Mannager	149	161	125	101	536	0.779	0.255	0.087	0.153	0.046	0.542
HR Manager	59	62	75	86	282	1.000	0.073	0.000	0.069	0.000	0.057
Other Managerial Role	28	36	24	19	107	0.411	0.713	0.506	0.239	0.151	0.757
Employees_white collar	178	166	158	123	625	0.196	0.484	0.081	0.572	0.588	0.290
Self Employed	10	12	4	8	34	0.755	0.131	0.951	0.071	0.721	0.159
Employees_blue collar and Other	58	71	88	58	275	0.363	0.003	0.270	0.032	0.804	0.079
Mean of Years in Same Job Role	11.200	12.549	10.892	11.289	11.513	0.018	0.593	0.885	0.005	0.045	0.535
Mean of Years in Same Firm	13.079	14.490	12.212	14.490	13.096	0.027	0.188	0.282	0.000	0.002	0.849
Public Firm	150	166	156	131	603	0.598	0.347	0.572	0.667	0.945	0.737
Private Firm	432	441	385	322	1580	0.452	0.130	0.021	0.430	0.105	0.400
Non-Profit Firm	47	54	55	65	221	0.642	0.267	0.004	0.509	0.013	0.075
Observations	629	661	596	518	2404						

Table A2: Work characteristics: Number of observations and balance test, p-value from t-tests

N	Rank Priority: Gender Equality		t-test	p-value
	Supply	Demand		
	3.4909	3.455	0.405	0.685
1286				
	Supply	Info Causes		
	3.491	3.451	-0.451	0.652
1254				
	Supply	No Info		
	3.491	3.514	-0.246	0.806
1174				
	Demand	Info Causes		
	3.455	3.451	-0.045	0.964
1224				
	Demand	No Info		
	3.455	3.513	-0.612	0.541
1144				
	Info Causes	No Info		
	3.451	3.514	-0.655	0.513
596	516			

Notes: Sample mean and t-test by treatment for gender equality priority values (1-6). The value indicates in which position the participants ranked gender equality among the other five important goals for their workplace in the next five years. The other five issues were: to promote innovation, increase productivity, promote technological progress, increase teamwork, and promote projects aimed at reducing the environmental impact

Table A3: Rank Priority: Gender Equality by treatment

	Attrition		Total	t-test	p-value
	Supply	Demand			
	90	81	171	0.342	0.736
N	751	710	1461		
	Supply	Info Causes			
	90	81	171	-0.011	0.991
N	751	677	1428		
	Supply	No Info			
	90	78	168	-0.608	0.543
N	751	596	1347		
	Demand	Info Causes			
	81	81	162	0.322	0.747
N	710	677	1387		
	Demand	No Info			
	81	78	159	-0.924	0.356
N	710	596	1306		
	Info Causes	No Info			
	81	78	159	-0.604	0.546
N	677	596	1273		

Notes: Attrition represents the number of respondents who abandoned the survey before completing it. The last two columns show the t-test statistic and the p-value for the differences in mean by treatment

Table A4: Attrition by treatment

	Rank Priority: Gender Equality		Total	t-test	p-value
	Complete	Incomplete			
All Sample					
	3.477	3.514	3.479	-0.276	0.783
N	2398	142	2540		
Demand					
	3.456	3.667	3.468	-0.805	0.421
N	628	39	667		
Supply					
	3.491	3.538	3.494	-0.188	0.851
N	658	39	697		
Info Causes					
	3.451	3.326	3.442	0.524	0.601
N	596	46	642		
No Info					
	3.514	3.611	3.517	-0.608	0.543
N	516	18	534		

Notes: The table contains the sample means for the variable Rank Priority: Gender Equality and the corresponding t-tests and p-values. The value indicates in which position (1-6) the participants ranked gender equality among the other five important goals for their workplace in the next five years. The other five issues were: to promote innovation, increase productivity, promote technological progress, increase teamwork, and promote projects aimed at reducing the environmental impact The mean at the top represents the entire sample, while the means on the lower rows are computed separately for each treatment and control group.

Table A5: Attrition by Rank Priority: Gender Equality

B Treatment Videos

Appendix B provides the full content of the informational treatments administered in the experiment, including the visual panels and scripts of the Supply, Demand, and Info Causes videos. Figures B1–B3 reproduce the Supply and Demand treatments. These share a common introduction and structure (Figure B1) and differ only in the narrative framing of the causes and mechanisms through which quotas operate (Figure B2 for the Supply treatment and Figure B3 for the Demand treatment). Figures B4–B5 reproduce the Info Causes treatment. This appendix enables full replication of the informational interventions.

(a) Panel 1

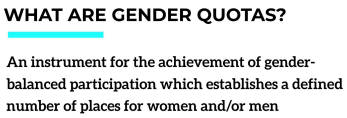

(b) Panel 2

(c) Panel 3

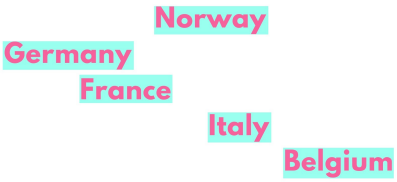

(d) Panel 4

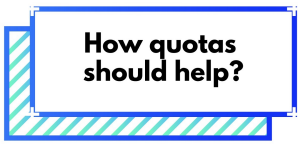

(e) Panel 5

Figure B1: First five panels for both Demand and Supply treatments

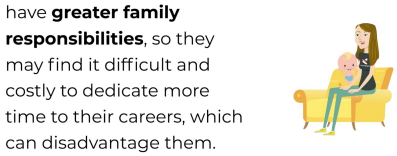

(a) Panel S_6

(b) Panel S_7

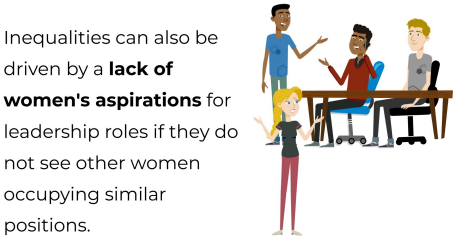

(c) Panel S_8

(d) Panel S_9

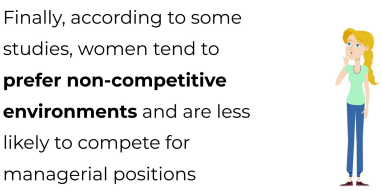

(e) Panel S_10

(f) Panel S_11

Figure B2: Supply Treatment: Panel 6-11

(a) Panel D_6

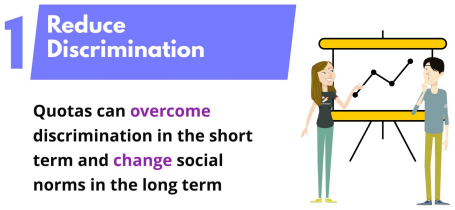

(b) Panel D_7

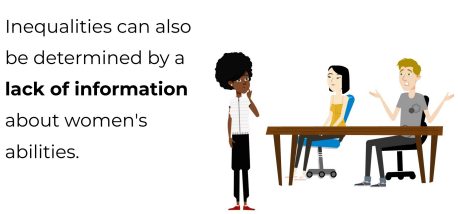

(c) Panel D_8

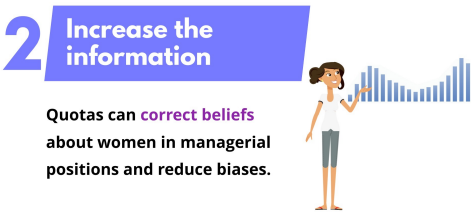

(d) Panel D_9

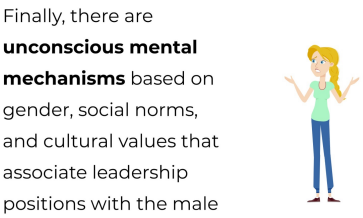

(e) Panel D_10

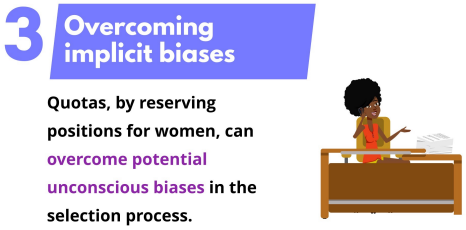

(f) Panel D_11

Figure B3: Demand Treatment: Panel 6-11

(a) Panel CI_1

(b) Panel CI_2

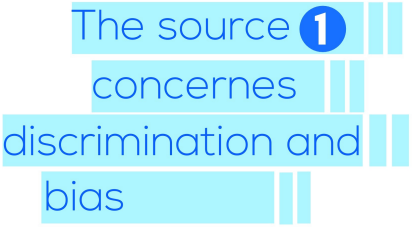

(c) Panel CI_3

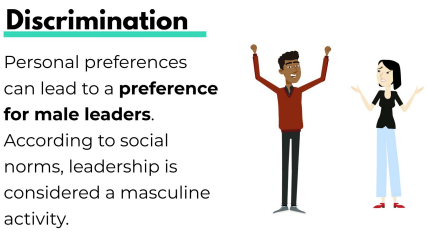

(d) Panel CI_4

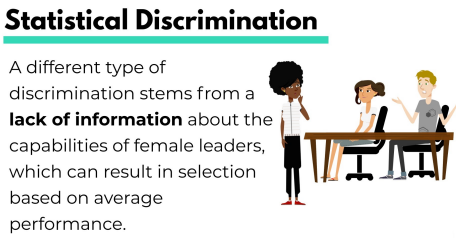

(e) Panel CI_5

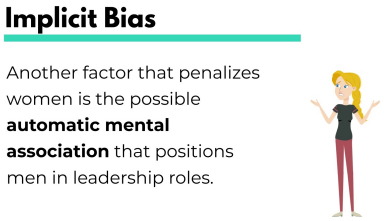

(f) Panel CI_6

Figure B4: Control Info Cauese: Panel 1-6

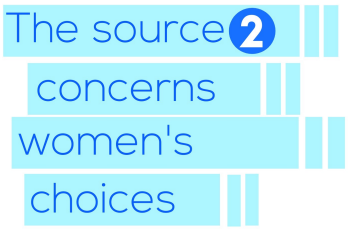

(a) Panel CI_7

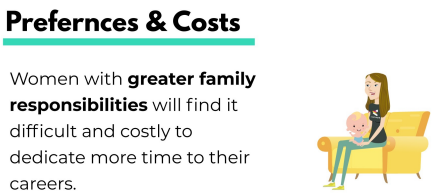

(b) Panel CI_8

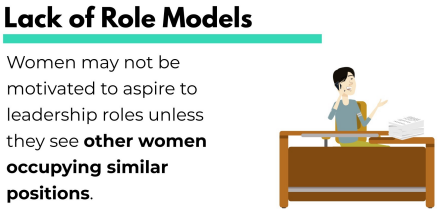

(c) Panel CI_9

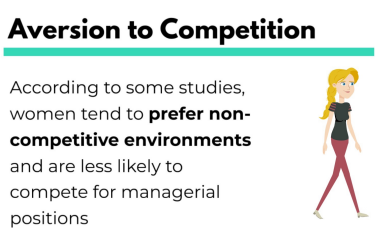

(d) Panel CI_10

(e) Panel CI_11

Figure B5: Control Info: Panel 7-11

C Incentivized Belief Elicitation and Construction of Misperception Indices

This appendix provides additional details on the incentivized belief-elicitation tasks used to measure respondents' prior perceptions of the *magnitude* of gender inequalities in managerial positions. We report the exact wording of the questions, describe the incentive scheme, and formally define the misperception indices used throughout the analysis.

C.1 Incentivized Questions

Before exposure to any informational treatment, respondents answered two incentivized questions concerning (i) the gender wage gap among managers and (ii) women's representation in top corporate positions.

Gender wage gap. Participants were asked:

> *"Think about managers (full-time employees aged 30–49) working in Italian firms with more than 10 employees. In your opinion, on average, how much does a female manager with these characteristics earn for every €100 earned by a male manager with the same characteristics?"*

Responses were given on a 0–100 scale. The correct value, based on Eurostat data, is 55.64.

Women's representation among CEOs. Participants were asked:

> *"In your opinion, what percentage of CEOs of the 50 largest listed companies in Italy are women?"*

Responses were given on a 0–100 scale. The correct value is 2.9 percent (EIGE, 2018).

C.2 Incentive Scheme

To encourage accurate reporting and reduce expressive responding, answers were incentivized. After data collection, 2% of respondents were randomly selected. For each of the two questions, participants received a €5 bonus if their answer fell within a pre-specified tolerance window:

- ± 2 euros for the gender wage ratio;
- ± 2 percentage points for the CEO share.

Bonuses were paid through the survey company in addition to standard participation rewards. This design follows established approaches to incentivized belief elicitation in politically sensitive contexts.

C.3 Construction of Misperception Indices

We focus on *underestimation* of gender inequality, defined as responses that attenuate the true magnitude of disparities.

Index of Gender Wage Gap Underestimation (indexGWG). Let y_i denote respondent i's answer and y^* the true value (55.64). We define:

$$indexGWG_i = \begin{cases} \frac{y_i - y^*}{100 - y^*} & \text{if } y_i > y^* \\ 0 & \text{otherwise} \end{cases}$$

The index ranges from 0 (accurate perception or overestimation of the gap) to 1 (maximum possible underestimation). Overestimation occurs in 19.25% of cases and is coded as zero.

Index of CEO Representation Overestimation (indexCEO). Let z_i denote respondent i's estimate and $z^* = 2.9$ the true share. We define:

$$indexCEO_i = \begin{cases} \frac{z_i - z^*}{100 - z^*} & \text{if } z_i > z^* \\ 0 & \text{otherwise} \end{cases}$$

The index captures overestimation of female representation, which corresponds to underestimating inequality.

Observations reporting values equal to or above 100% (N = 26) are excluded from this measure.

C.4 Underestimation Index

To obtain a summary measure of magnitude misperceptions, we define:

$$UnderestimationIndex_i = \frac{indexGWG_i + indexCEO_i}{2}.$$

The index ranges from 0 (no downward bias in either dimension) to 1 (maximum underestimation on both dimensions).

This composite index is used throughout the analysis to capture respondents' prior beliefs about the magnitude of gender inequality and to study heterogeneous treatment effects.

C.5 Balance Across Treatment Groups

Appendix Tables C1–C5 report pairwise t-tests of raw answers and index values across treatment groups. Consistent with random assignment occurring after belief elicitation,

we find no systematic differences across treatments.

Figures C1–C2 display the full distribution of responses by treatment group. Figure C3 reports correlates of the Underestimation Index with demographic and job characteristics.

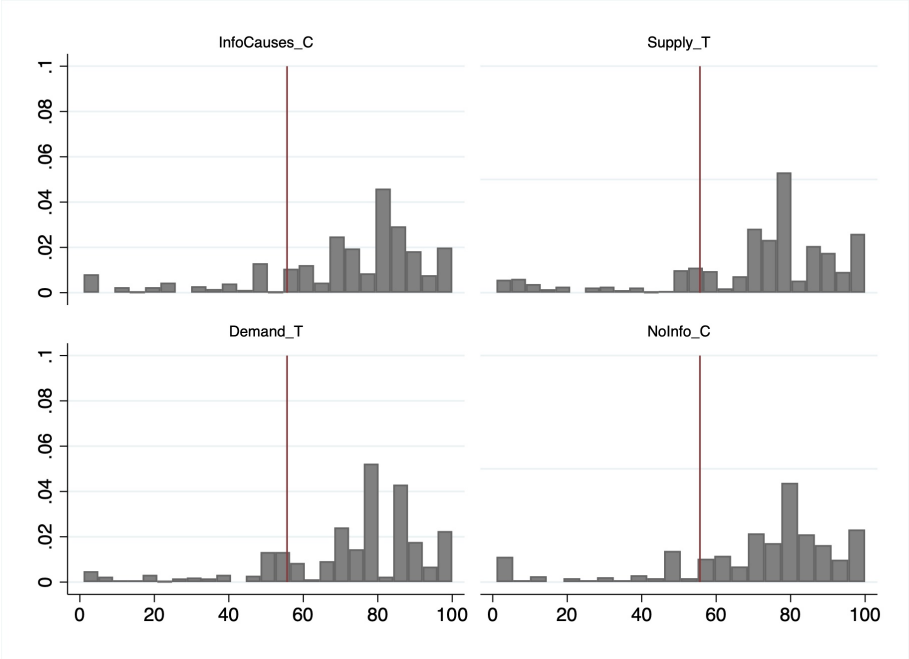

Figure C1: Distributions of prior beliefs about the Gender Wage Gap (as the wage of a female manager for every €100 of a male manager). The red line identifies the right answer:€55.64

Question: *Think about managers (full-time employees aged 30–49) working in Italian firms with more than 10 employees. In your opinion, on average, how much does a female manager with these characteristics earn for every €100 earned by a male manager with the same characteristics?*

N	Gender Wage Gap (GWG)		t-test	p-value
	Supply	Demand		
	71.106	72.951	-1.463	0.144
1289				
	Supply	Info Causes	t	p
	71.106	71.118	0.009	0.993
1257				
	Supply	No Info	t	p
	71.106	70.596	0.363	0.716
1178				
	Demand	Info Causes	t	p
	72.951	71.1178	-1.466	0.143
1224				
	Demand	No Info	t	p
	72.951	70.596	1.765	0.078
1145				

Avarage answer to the question: *Think about managers (full-time employees aged 30–49) working in Italian firms with more than 10 employees. In your opinion, on average, how much does a female manager with these characteristics earn for every €100 earned by a male manager with the same characteristics?*

Table C1: Mean and t-test for values of the answers to GWG question (from 1 to 100) by treatment

N	% CEO Women		t-test	p-value
	Supply	Demand		
	18.927	20.020	-1.095	0.274
1272				
	Supply	Info Causes		
	18.927	20.279	1.379	0.168
1244				
	Supply	No Info		
	18.927	20.185	-1.238	0.216
1168				
	Demand	No Info		
	20.019	20.185	-0.155	0.877
1134				
	Info Causes	Demand		
	20.279	20.019	0.253	0.801
1210				

Avarage answer to the question: *In your opinion, what percentage of CEOs of the 50 largest listed companies in Italy are women?*

Table C2: Mean and t-test for values of the answers to CEO question (from 1 to 100) by treatment

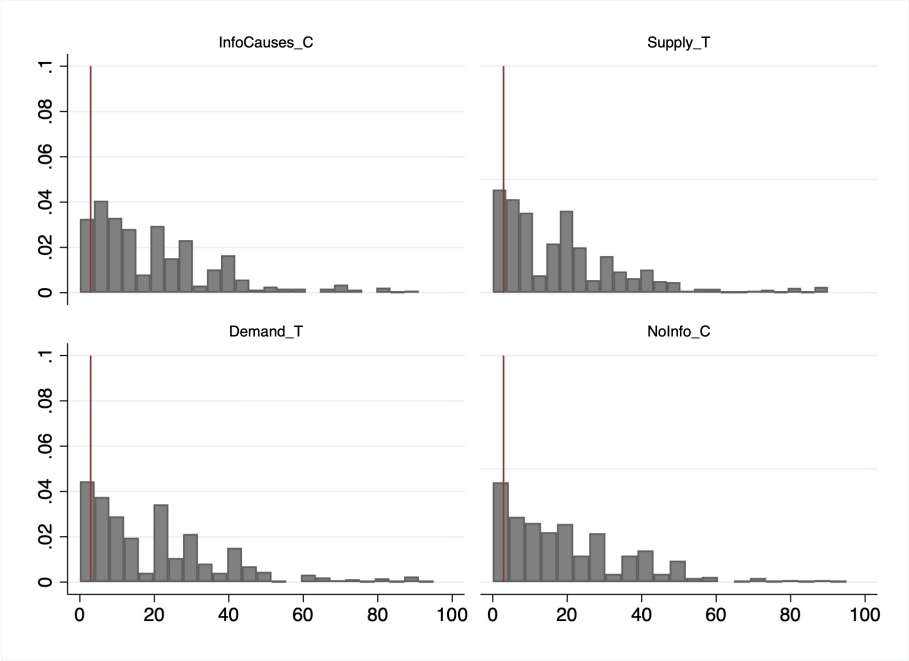

Figure C2: Distributions of prior beliefs about the % of CEO women. The red line identifies the right answer:2.9%
Question: *In your opinion, what percentage of CEOs of the 50 largest listed companies in Italy are women?*

N	Index GWG		t-test	p-value
	Supply	Demand		
	0.610	0.609	0.0065	0.995
1289				
	Supply	Info Causes		
	0.610	0.602	-0.401	0.688
1257				
	Supply	No Info		
	0.610	0.600	0.535	0.592
1178				
	Demand	Info Causes		
	0.609	0.602	-0.3932	0.694
1224				
	Demand	No Info		
	0.609	0.600	0.528	0.598
1145				

IndexGWG is the normalized absolute distance between the true value for the GWG question and the respondent's reported value.

Question: *Think about managers (full-time employees aged 30–49) working in Italian firms with more than 10 employees. In your opinion, on average, how much does a female manager with these characteristics earn for every €100 earned by a male manager with the same characteristics?*

Table C3: Mean and t-test for values of the IndexGWG (from 0 to 1) by treatment

N	Index CEO		t-test	p-value
	Supply	Demand		
	0.751	0.737	0.0065	0.995
1272				
	Supply	Info Causes		
	0.751	0.763	0.741	0.459
1244				
	Supply	No Info		
	0.751	0.751	0.055	0.956
1168				
	Demand	Info Causes		
	0.763	0.737	1.615	0.107
1210				
	Demand	No Info		
	0.737	0.751	-0.787	0.431
1134				

IndexCEO is the normalized absolute distance between the true value for the GWG question and the respondent's reported value.

Question: *In your opinion, what percentage of CEOs of the 50 largest listed companies in Italy are women?*

Table C4: Mean and t-test for values of the IndexCEO (from 0 to 1) by treatment

N	Underestimation Index		t-test	p-value
	Supply	Demand		
	0.681	0.673	0.621	0.535
1271				
	Supply	Info Causes		
	0.681	0.683	0.154	0.878
1244				
	Supply	No Info		
	0.681	0.676	0.389	0.697
1167				
	Demand	Info Causes		
	0.687	0.673	0.781	0.435
1209				
	Demand	No Info		
	0.673	0.676	-0.206	0.837
1132				

Underestimation Index equals the average of the CEO and GWG normalized-distance indices and is set to 0 for any respondent whose answer is at or below either true values (CEO $\leq$ 2.9 or GWG $\leq$ 55.64).

Table C5: Mean and t-test for values of the Underestimation Index (from 0 to 1) by treatment

	Employment rate	Part-time share
	(1)	(2)
Underestimation Index	0.171	-0.000
	(0.131)	(0.129)
Female education	1.281**	0.030
	(0.464)	(0.277)
Childacre uptake	0.656***	0.047
	(0.223)	(0.133)
Observations	60	60

Notes: Cluster-robust standard errors (region) in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
All specifications include year fixed effects and macro-geographical area fixed effects (omitted from the table).
Variables are scaled to unit intervals. Employment rate, part-time share, female education, and Childcare uptake are proportions in $[0, 1]$; the Underestimation Index is also bounded between 0 and 1.

Table C6: Underestimation Index and Regional Labor Market Outcomes

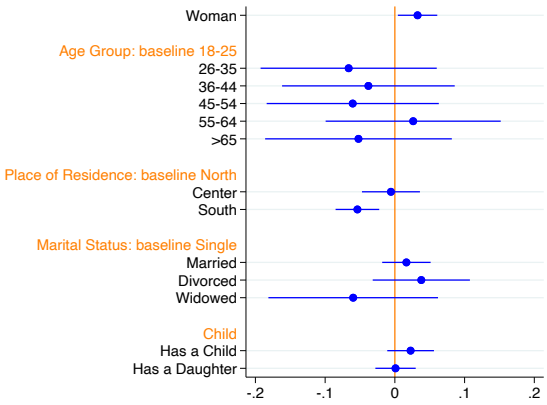

(a) Demographics

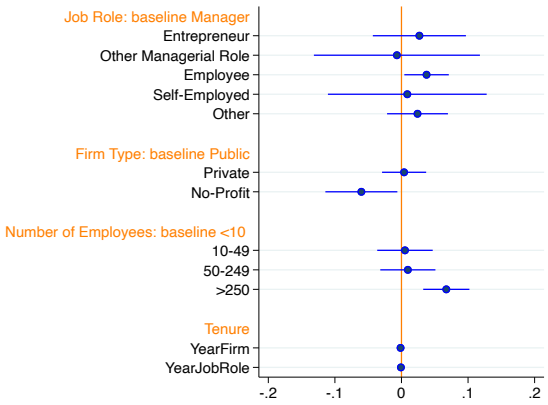

(b) Job Characteristics

Figure C3: Correlates of the *Underestimation Index*.

Notes: Points show OLS coefficients; bars denote 95% confidence intervals. Each point estimate comes from a *separate* regression (i.e., coefficients are not jointly estimated within a single model). The dependent variable is the *Underestimation Index*, ranging from 0 (accurate perception, no underestimation) to 1 (maximum underestimation of gender inequalities). Robust standard errors are used.

D OEQ: Causes of Gender Inequalities

This appendix complements Section 6.1.2 by detailing the construction of the causal-belief categories. We describe the preprocessing of open-ended responses, the dictionary-based classification procedure, and supporting validation and robustness checks.

D.1 Preprocessing of Open-Ended Responses

We follow the text-analysis approach proposed by Ferrario and Stantcheva 2022. Responses were first standardized by converting all text to lowercase and removing punctuation, extra spaces, numerical values, and common stopwords (e.g., "and," "the," "then"). Words were subsequently stemmed using the Snowball algorithm, which reduces morphological variants to a common root (e.g., "policy" and "policies" $\rightarrow$ "polic"). This preprocessing yields a cleaned textual representation for each response.

To ensure data quality, respondents were required to provide at least five characters. The median response length is 9 words (25th percentile: 5; 75th percentile: 18), indicating that most participants provided articulated explanations rather than minimal answers.

D.2 Dictionary Construction and Topic Assignment

Topic identification relies on custom dictionaries derived from the literature on the determinants of gender inequality. Each dictionary includes keywords and expressions associated with a specific conceptual category (Appendix Table D1 reports examples and keyword lists).

For each response, we construct a binary indicator equal to one if at least one keyword associated with a category appears in the cleaned text, and zero otherwise. Because responses may reference multiple mechanisms, category indicators are not mutually exclusive, except for the generic cultural category, which is coded as mutually exclusive from the specific mechanisms.

Topic selection was manual rather than algorithmic. Given the brevity and informal structure of many responses, manual validation ensures conceptual coherence and reduces the risk of context-insensitive keyword matches. All keywords were iteratively refined through inspection of potentially misclassified responses.

D.3 Handling of Ambiguous and Residual Categories

Responses that were not matched by the predefined dictionaries were manually reviewed to distinguish between (i) genuinely non-classifiable answers and (ii) explicit "no idea" statements, which were coded separately. In ambiguous cases, we used the ChatGPT API solely as a classification aid to assist in distinguishing between these two types of

responses.[12] This procedure affected a small subset of observations and did not alter the structure of the main categories. After this review, 5.62% of responses remain classified as "Not classifiable."

D.4 Distribution of Categories and Balance Across Treatments

Table 3 reports the overall distribution of categories. Work–life balance is the most frequently cited specific mechanism (33.47%), followed by women's ability (28.30%) and taste discrimination/reaction (10.37%). The *Women's ability* category warrants particular attention: this category is inherently heterogeneous and difficult to classify unambiguously. While in a few cases respondents clearly refer to external stereotypes or explicit discrimination against women's capabilities, in many others the wording blurs the distinction between perceived ability, social norms, and individual behavior. Typical responses refer to beliefs about women's competence, lack of confidence in women by employers or society, women's self-confidence, or stereotypical traits attributed to women, often combining these elements within the same sentence.[13] As a result, attributing these responses to either supply-side or demand-side mechanisms would require strong assumptions and substantial researcher interpretation. For this reason, we refrain from using the women's ability category for inference when distinguishing between supply-side and demand-side explanations.

Appendix Table D2 presents the frequency of each category by treatment group. Although some pairwise differences are statistically significant, no systematic patterns emerge. All subsequent regressions therefore control for pre-treatment causal beliefs.

Appendix Table D3 reports correlations across categories, confirming that respondents frequently invoke multiple explanations within the same response.

D.5 Heterogeneity by Demographics and Job Characteristics

Appendix Table D6 and D7 report regressions of category indicators on demographic and job-related characteristics. Women are significantly less likely to invoke generic explanations and more likely to cite work–life balance constraints. Married respondents disproportionately reference caregiving-related explanations. Older respondents are more likely to invoke discrimination-based accounts, while respondents with tertiary education

[12]The ChatGPT API was used exclusively to assist in classifying residual unmatched responses and did not generate categories or modify the predefined dictionaries. All final coding decisions were verified by the authors.

[13]Illustrative responses include: (i) "There is little trust in women; they are perceived as more incapable and emotional (which is not true)"; (ii) "Prejudices regarding women's managerial abilities, together with the objective difficulties women face in managing family responsibilities"; (iii) "Gender prejudice, a male-dominated managerial culture, and a lack of confidence both in women and among women themselves"; and (iv) "Women are more cautious and conservative traits in economic and financial decision-making." See Appendix Table E1 for additional examples and the full set of keywords used in the classification.

are less likely to attribute inequality to discrimination.

These patterns indicate that causal beliefs are systematically structured rather than randomly distributed across respondents.

D.6 Single-Mechanism Restriction

Because responses may contain multiple mechanisms, we also examine how many distinct causes are mentioned. Appendix Table D5 shows that 75% of respondents cite exactly one specific mechanism, 22% cite two, and fewer than 4% cite three or more.

To obtain a cleaner measure of interpretable priors, the main text restricts attention in selected analyses to respondents who (i) mention exactly one specific mechanism and (ii) do not invoke the generic cultural category (N = 1,077). Within this subset, 48% attribute inequality primarily to work–life balance constraints, whereas 11.8% cite taste discrimination or backlash. This asymmetry confirms that supply-side explanations are more prevalent than demand-side accounts in the baseline distribution of beliefs.

D.7 Beliefs and Misperceptions of Inequality

Appendix Table D8 reports regressions of category indicators on the Underestimation Index. Respondents who invoke generic cultural explanations exhibit significantly greater underestimation of gender gaps. Higher underestimation is also positively associated with citing taste discrimination/reaction as causes.

ID	Category	Definition	Examples	Main Keywords
Generic Causes				
1	Cultural causes	References to stereotypes/traditional social norms disadvantaging women, without specifying whether they concern work abilities or social roles.	• "Mentality and culture based on the masculine figure" • "Sexism and the outdated idea that men should dominate" • "Prejudices of society as a whole" • "Entirely cultural problem" • "Effects of a patriarchal society"	Mindset; patriarchy/patriarchal; culture/cultural legacy; male dominance; prejudices; society; biases; stereotypes; sexism; tradition; belief; misogyny; mentality.
Specific Causes				
2	Women's ability	Mentions women's (perceived) work capabilities/traits (often undervaluation).	• "Women are considered less at work" • "Belief that women are less perspicacious" • "Household/childcare duties impede performance" *(also category 3)* • "Women are more cautious and conservative traits in economic and financial decision-making." • "Prejudices regarding women's managerial abilities coexist with difficulties related to family responsibilities." *(also category 3)* • "Gender prejudice and a male-dominated managerial culture undermine confidence in women and women's self-confidence."	Ability; productivity; capacity; ambition; undervalue; underestimate; discrimination; trust; meritocracy.
3	Work–life balance	Attribution to caregiving roles/family responsibilities constraining careers.	• "Women are seen as homemakers, not career-oriented" • "Maternity and domestic chores hinder results" *(also category 2)*	Children; family; pregnancy; leave; caregiving; maternity; domestic duties; absences; time; manage; work–family.

Table D1: Description of categories for causes of gender inequalities

Table D1 (continued)

ID	Category	Definition	Examples	Main Keywords
4	Institutional problems	Insufficient institutional/organizational support (policies, services).	• "Little control from competent bodies" • "Lack of help/support for managing potential maternity" *(also 1/3)* • "Work does not guarantee adequate assistance" *(also 3)*	Assistance; policies; support; institutions; laws; state; welfare; unions.
5	Taste discrimination & reaction	Explicit discrimination by men in power and/or defensive reactions to women.	• "Male managers prevent women from entering" • "Fear of losing power by men" • "Fear because women are more skilled"	Leadership positions; disadvantage; obstruct; promote same gender; power; fear.
6	Denial	Denial of the existence/salience of inequalities.	• "In my experience, these inequalities do not exist" • "I don't think it's important" • "Maternity is simply an obstacle" *(also category 3)*	"I don't believe/think"; "there are none".
7	No idea	Respondent explicitly reports not knowing.	• "I really don't know" • "I wouldn't know what to say"	"I don't know"; "I wouldn't know".
8	Not classifiable	Responses that do not fit any category above.	—	—

	(1) Demad	(2) Supply	(3) Info Causes	(4) NoInfo	(5) Total	(6) p(D=S)	(7) p(D=IC)	(8) p(D=C)	(9) p(S=IC)	(10) p(S=C)	(11) p(IC=C)
Cultural Causes	0.275	0.313	0.264	0.280	682	0.134	0.673	0.854	0.057	0.216	0.560
Women's Ability	0.312	0.253	0.305	0.261	679	0.019	0.794	0.068	0.040	0.700	0.121
Work-Life Balance	0.342	0.287	0.345	0.375	804	0.036	0.903	0.250	0.028	0.001	0.308
Institutional Problem	0.062	0.044	0.069	0.060	140	0.146	0.620	0.879	0.053	0.216	0.535
Taste Discrimination Reaction	0.113	0.088	0.120	0.095	249	0.133	0.717	0.315	0.064	0.685	0.182
Denial	0.008	0.014	0.008	0.021	30	0.230	0.928	0.056	0.275	0.432	0.140
No Idea	0.017	0.042	0.032	0.031	74	0.009	0.102	0.137	0.334	0.303	0.334
N	629	661	594	518	2402						

Note that a single answer may be classified into more than one category; therefore, the number of observations (N) refers to the total number of answers. The shares are computed as the number of answers containing a given category over the total number of answers within each treatment group. *Cultural Causes* is the only mutually exclusive category.

Table D2: Causes of Inequalities: Share of observations and balance test, p-values from t-tests

	Women's Ability	Work-Life Balance	Institutional Problem	Taste Discrimination and Reaction
Women's Ability	1.000			
Work-Life Balance	-0.038*	1.000		
Institutional Problem	0.061*	0.095*	1.000	
Taste Discrimination and Reaction	0.008	-0.062*	0.050*	1.000

* $p < 0.05$

Table D3: Causes of Inequalities: correlation matrix

	Men	Women	t-test	p-value
Cultural Causes	0.321	0.253	3.680	0.000
Women's Ability	0.258	0.303	-2.397	0.017
Work–Life Balance	0.260	0.396	-7.066	0.000
Institutional Problem	0.046	0.068	-2.268	0.023
Taste Discrimination Reaction	0.093	0.113	-1.609	0.108
Denial	0.017	0.009	1.666	0.096
No Idea	0.041	0.023	2.549	0.011
N	2402			

Table D4: Categoriers Causes: means, t-tests and p-value by gender

Panel A: Responses mentioning exactly one category ($N = 1077$)	
Category	Share of responses
Women's Ability	0.371
Work–Life Balance	0.480
Taste Discrimination Reaction	0.118
Institutional Problem	0.031
Panel B: Responses mentioning more than one category ($N = 368$)	
Category	Share of responses
Women's Ability	0.758
Work–Life Balance	0.780
Taste Discrimination Reaction	0.332
Institutional Problem	0.291

Notes: Share of responses in which the category is present. Categories are not mutually exclusive, so shares do not sum to one.

Table D5: Share of responses mentioning each category

	Cultural causes	Work-life balance	Women's ability	Taste discrim. react.
	(1)	(2)	(3)	(4)
Demographics and background characteristics				
Female	-0.086***	0.107***	0.046**	0.014
	(0.021)	(0.021)	(0.020)	(0.014)
Has Children	-0.050*	0.036	0.030	-0.012
	(0.027)	(0.028)	(0.027)	(0.019)
Has a Daughter	0.028	-0.005	-0.041*	-0.004
	(0.022)	(0.023)	(0.022)	(0.015)
Born abroad	0.051	-0.139***	0.033	0.042
	(0.058)	(0.051)	(0.057)	(0.044)
Place of residence (ref. North)				
South	-0.023	-0.037	-0.004	0.003
	(0.022)	(0.023)	(0.022)	(0.015)
Center	0.009	-0.002	-0.074***	-0.017
	(0.028)	(0.029)	(0.026)	(0.018)
Bachelor's degree or more	0.040*	0.019	-0.030	-0.028**
	(0.020)	(0.021)	(0.020)	(0.013)
Age group (ref. $\leq$ 25)				
Age 26-35	0.082	-0.089	-0.112	0.061*
	(0.086)	(0.089)	(0.090)	(0.035)
Age 36-44	0.076	-0.135	-0.148*	0.054
	(0.085)	(0.088)	(0.089)	(0.035)
Age 45-54	0.073	-0.137	-0.167*	0.078**
	(0.085)	(0.088)	(0.089)	(0.036)
Age 55-64	0.078	-0.200**	-0.117	0.099***
	(0.087)	(0.090)	(0.091)	(0.038)
Age $\geq$ 65	0.088	-0.209**	-0.146	0.027
	(0.094)	(0.095)	(0.096)	(0.041)
Marital Status (ref. Single)				
Married	-0.041	0.066**	0.020	0.023
	(0.027)	(0.027)	(0.027)	(0.018)
Divorced	-0.042	0.038	0.022	0.028
	(0.051)	(0.051)	(0.050)	(0.035)
Widowed	-0.045	0.033	-0.029	-0.019
	(0.082)	(0.077)	(0.085)	(0.047)
Observations	2369	2369	2369	2369

Notes: Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
All specifications additionally control for work characteristics (job role, firm size, sector, firm type, and affirmative action at work) and belief measures (underestimation index and gender priority). These variables are omitted from the table for readability (see Table D7 - D8).

Table D6: Demographics and background characteristics correlates across beliefs about the causes of gender inequalities

	Cultural causes (1)	Work-life balance (2)	Women's ability (3)	Taste discrim. react. (4)
Work characteristics				
Job Role (ref. Manager)				
Entrepreneur	0.023	-0.043	0.051	0.015
	(0.048)	(0.046)	(0.050)	(0.034)
Other Managerial Role	0.023	0.144	0.146*	0.028
	(0.080)	(0.090)	(0.088)	(0.063)
Employee	-0.015	0.056**	0.035	0.006
	(0.024)	(0.026)	(0.024)	(0.017)
Self-Employed	-0.081	0.198**	0.038	-0.031
	(0.073)	(0.086)	(0.084)	(0.041)
Other	0.055*	0.034	0.037	-0.035*
	(0.032)	(0.033)	(0.032)	(0.020)
Number of employees (ref. <10)				
10-49	0.000	-0.048	-0.064**	0.018
	(0.029)	(0.031)	(0.030)	(0.021)
50-249	0.023	-0.017	-0.025	-0.011
	(0.031)	(0.032)	(0.032)	(0.021)
≥250	0.060**	0.033	-0.069**	-0.015
	(0.028)	(0.030)	(0.028)	(0.019)
Firm type (ref. Public)				
Private	0.044*	0.025	0.018	-0.008
	(0.027)	(0.027)	(0.026)	(0.019)
No profit	0.013	0.005	0.023	-0.008
	(0.042)	(0.041)	(0.043)	(0.028)
Affirmative Action at work	0.000	-0.013	0.004	0.003
	(0.021)	(0.022)	(0.021)	(0.015)
Observations	2369	2369	2369	2369

Notes: Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
All specifications also control for demographics, background characteristics (gender, parental status, country of residence, place of residence, educational level, age, marital status) and belief measures (underestimation index and gender priority). These variables are omitted from the table for readability; see Tables D6-D8).
Sector indicators are included in all specifications but omitted from the table for readability.

Table D7: Work characteristics correlates across beliefs about the causes of gender inequalities

	Cultural causes	Work-life balance	Women's ability	Taste discrim. react.
	(1)	(2)	(3)	(4)
Prior beliefs				
Underestimation index	0.079***	0.038	-0.013	0.042**
	(0.026)	(0.027)	(0.027)	(0.017)
Gender Priority	-0.005	-0.001	0.004	0.010**
	(0.006)	(0.006)	(0.006)	(0.004)
Observations	2369	2369	2369	2369

Notes: Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
All specifications also control for demographics, background characteristics ll specifications also control for demographics, background characteristics (gender, parental status, country of residence, place of residence, educational level, age, marital status) and work characteristics (job role, firm size, sector, firm type, and affirmative action at work). These variables are omitted from the table for readability; see Tables D6-D7).

Table D8: Prior beliefs correlates across beliefs about the causes of gender inequalities

D.8 Regional Beliefs and Labor Market Outcomes

To examine whether the distribution of causal beliefs reflects broader regional labor-market patterns, we aggregate responses by region and correlate them with administrative indicators of female labor-market outcomes provided by ISTAT. Because responses to the open-ended question may capture shared local narratives or prevailing social views, regional variation in belief categories can proxy for differences in the normative environment. This exercise follows the spirit of Fortin 2005 and Barigozzi and Montinari 2025.

For each region, we compute the share of respondents invoking a given belief category and estimate OLS regressions of regional female labor-market outcomes on these shares (Table D9). The dependent variables are: (i) the female employment rate[14] (Panel a) and (ii) the share of part-time employment among women[15] (Panel b).[16] Outcomes are observed for the years 2018–2020. All specifications control for regional female tertiary education and the share of children aged 0–2 enrolled in public daycare (ISTAT, 2022).

[14]ISTAT 2021 Labor Force Survey identifies employed women as women between the ages of 15 and 89 who fall into one of the following categories: (i) women who, during the reference week, worked for at least one hour for pay or profit, including unpaid family workers; (ii) women who, during the reference week, are temporarily absent from work because on vacation, on flexible hours (vertical part-time, recovery work hours, etc.), on sick leave, on mandatory maternity leave, or on employer-paid vocational training; (iii) women on parental leave, who receive and/or are entitled to work-related income or benefits, regardless of the duration of the absence; seasonal workers who, during the off-season, continue to regularly perform tasks and duties necessary for the continuation of the business, excluding the fulfillment of legal or administrative obligations; (iv) women temporarily absent for other reasons in all cases where the expected duration of absence is three months or less. The employment rate is, therefore, defined as the percentage of employed persons in relation to the corresponding total population.

[15]ISTAT 2021 defines part-time positions as jobs occupied by employed workers in the age range 15–64 whose regular working hours, certified by an individual contract, are lower than the standard hours established by the collective contract or company agreements. Data on part-time employment rates is obtained as a re-elaboration of data from ISTAT, downloaded in October 2023.

[16]The share of full-time employment is mechanically the complement of part-time employment; results are therefore identical with opposite sign and are available upon request.

Macro-area fixed effects are included based on joint Wald tests, and standard errors are clustered at the regional level.

The results indicate systematic associations between regional belief distributions and labor-market outcomes. Regions where generic cultural explanations are more prevalent display lower female part-time employment shares, whereas regions where work–life balance beliefs are more common exhibit greater reliance on part-time employment. By contrast, work–life balance beliefs are positively associated with part-time employment. Ability-related beliefs are also positively associated with part-time employment, although the coefficient is statistically significant only at the 10% level. Discrimination-based explanations are not significantly correlated with labor-market indicators.

A limitation of this exercise is the temporal mismatch between survey data (2022) and labor-market outcomes (2018–2020). While we cannot observe beliefs over time, the analysis exploits cross-regional variation under the assumption that regional differences in such views are relatively stable in the short run.

Category	Employment rate	Part-time share
Cultural causes	-0.019	-0.100***
	(0.046)	(0.034)
Female education	1.407**	-0.109
	(0.490)	(0.206)
Childcare uptake	0.705**	0.158
	(0.249)	(0.142)
Work-life balance	-0.002	0.138**
	(0.087)	(0.050)
Female education	1.436**	-0.268
	(0.568)	(0.287)
Childcare uptake	0.683**	0.146
	(0.257)	(0.143)
Women's ability	0.138	0.145*
	(0.100)	(0.079)
Female education	1.433***	0.031
	(0.480)	(0.207)
Childcare uptake	0.752***	0.119
	(0.196)	(0.150)
Taste discrimination reaction	0.120	0.246
	(0.280)	(0.206)
Female education	1.444***	0.054
	(0.485)	(0.246)
Childcare uptake	0.703***	0.086
	(0.223)	(0.142)
Observations	60	60

Notes: Cluster-robust standard errors (region) in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.
All specifications include year fixed effects and macro-geographical area fixed effects (omitted from the table).
Variables are scaled to unit intervals. Employment rate, part-time share, female education, and Childcare uptake are proportions in $[0, 1]$.

Table D9: Regional employment outcomes and beliefs about causes of gender inequalities

E OEQ: Beliefs Policies

Among the 2,161 respondents who supported greater state intervention, 3.47% (75/2,161) answered "no idea" and 15.3% (332/2,161) provided responses that were not classifiable. We apply the same text-processing pipeline and a binary coding rule (equal to one if a category is mentioned and zero otherwise). When gender quotas are explicitly referenced, we further classify the stance as pro- or anti-quota using the ChatGPT API.[17]

Appendix Table F2 reports category frequencies by treatment group and overall. The most frequently advocated policies target workplace *discrimination*, followed by measures aimed at *cultural change*. Overall, 10% of respondents (216 out of 2,161) express support for gender quotas, while 1.39% (30 out of 2,161) explicitly oppose them; we find no significant differences across treatments at the pre-treatment stage. Pairwise t-tests across treatments, reported in Appendix Table F2, reveal few statistically significant differences. Correlations across policy categories are shown in Appendix Table F3.

Finally, gender differences in policy preferences are reported in Appendix Table F4. Support for *family and motherhood* policies is markedly higher among women (27.5%) than among men (15.8%), with the difference being statistically significant ($p = 0.000$).

[17] Classification procedure as above; pro-/anti-quota tagging applied only when quotas are explicitly mentioned.

Table E1: Categories for beliefs about policies to reduce gender inequalities

ID	Category	Examples from answers	Main keywords
Policy categories			
1	Family support policy	• "Hire more women; provide free daycare; bonuses for babysitters." • "Promote flexibility for men and women." *(also fits Flexible work)* • "Provide workplace facilities like daycare centers." • "Social bonuses for newborns; enable smart working for pregnant women."	Daycare; babysitters; bonuses; flexibility; work–life balance; smart working; newborns; pregnancy; family; parental leave; paternity; maternity.
2	Pro-quota	• "Equal percentage between men and women, 50% in each company." • "Put 50% women in senior management positions." • "Require by law a minimum number of women in management."	Equality; 50/50; quotas; mandated share; legal requirement.
3	Anti-quota	• "Eliminate gender quotas; women have the same abilities—this should be understood."	Eliminate quotas; opposing quotas.

Table continues on the next page

ID	Category	Examples from answers	Main keywords
4	Anti-discrimination policy	• "Oversight of contracts by INPS and the Ministry of Labor." • "Equalize salaries for equal positions regardless of gender." • "Blind recruitment; remove gender from applications."	Equality; oversight; audits; equal pay; blind hiring; inform; law; enforce.
5	Incentives (non-quota)	• "Government should incentivize firms so women can access managerial roles." • "Signal that salaries must be equal; otherwise higher taxes for noncompliant firms."	Incentives; tax credits/penalties; bonuses.
6	Cultural change	• "Civic education in schools teaching gender equality (e.g., chores taught to all), removal of stereotypes." • "Structure education on equal opportunities; devalue sexist culture."	School; education; media; stereotypes; culture change.
7	Merit	• "Ensure fair salaries reflecting true abilities, regardless of gender." • "Pay and promotions based on skills, not gender." • "Evaluate competencies irrespective of gender."	Merit; skills; competence; performance; meritocracy.

Table continues on the next page

ID	Category	Examples from answers	Main keywords
8	No idea	• “I don’t know.” “No idea.”	I don’t know; no idea.
9	Not classifiable	—	—

To examine how perceived causes of gender inequality relate to the policy respondents propose, we estimate OLS regressions in which the dependent variable is an indicator for mentioning a given type of explanation and the key regressors are dummies for the policy categories mentioned in the same response (Figure E1).

Panel (a) focuses on the generic *cultural causes* category. Respondents who propose anti-discrimination policies and cultural-change interventions are more likely to reference cultural causes, whereas respondents who emphasize family-support policies are markedly less likely to invoke this generic explanation.

Panel (b) focuses on the specific beliefs. The strongest pattern concerns work–life balance: respondents advocating family-support policies are substantially more likely to frame inequality in work–life balance terms, and incentive-based proposals also align positively with this prior belief. By contrast, policy proposals centered on anti-discrimination and merit-based solutions are more consistently associated with explanations that stress misperceptions of women's ability as a main cause of inequality. Overall, the results indicate that respondents' policy preferences tend to co-occur with distinct prior beliefs about the causes of gender inequality.

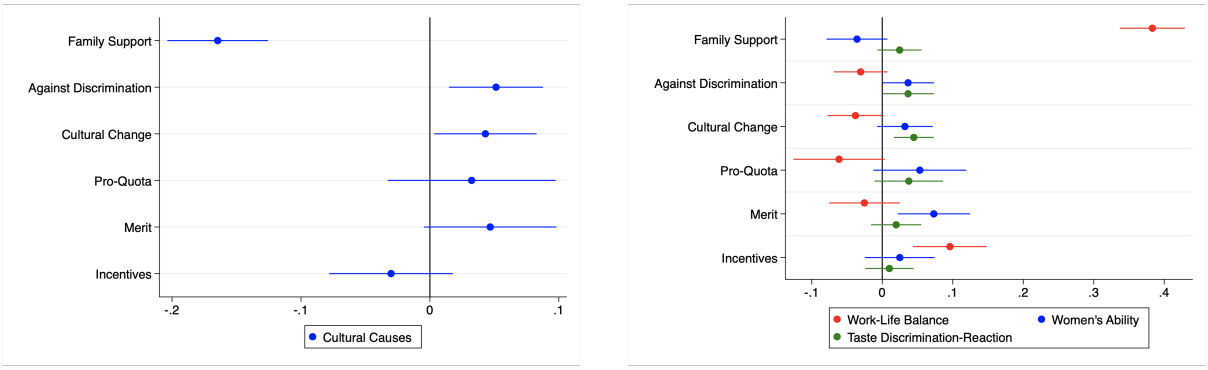

(a) Generic causes **(b)** Specific causes

Figure E1: Association between beliefs about causes and policy preferences

Notes: Points show OLS coefficients; bars denote 95% confidence intervals. Each point estimate comes from a *separate* regression (i.e., coefficients are not jointly estimated within a single model). Outcome is an indicator for invoking a given cause; regressors are dummies for policy categories mentioned. Models include demographic and job controls; standard errors are robust.

ID	Category	Definition	Examples	Main Keywords
Policy Categories				
1	Family support policy	Support for measures facilitating work–family reconciliation and caregiving.	• "Hire more women; provide free daycare; bonuses for babysitters." • "Promote flexibility for men and women." *(also fits Flexible work)* • "Provide workplace facilities like daycare centers." • "Social bonuses for newborns; enable smart working for pregnant women."	Daycare; babysitters; bonuses; flexibility; work–life balance; smart working; newborns; pregnancy; family; parental leave; paternity; maternity.
2	Pro-quota	Explicit support for mandatory gender quotas or equal gender representation.	• "Equal percentage between men and women, 50% in each company." • "Put 50% women in senior management positions." • "Require by law a minimum number of women in management."	Equality; 50/50; quotas; mandated share; legal requirement.
3	Anti-quota	Explicit opposition to gender quotas.	• "Eliminate gender quotas; women have the same abilities—this should be understood."	Eliminate quotas; opposing quotas.
4	Anti-discrimination policy	Policies aimed at preventing or sanctioning workplace discrimination.	• "Oversight of contracts by INPS and the Ministry of Labor." • "Equalize salaries for equal positions regardless of gender." • "Blind recruitment; remove gender from applications."	Equality; oversight; audits; equal pay; blind hiring; inform; law; enforce.

Table E2: Categories for beliefs about policies to reduce gender inequalities

Table E2 (continued)

ID	Category	Definition	Examples	Main Keywords
5	Incentives (non-quota)	Financial or regulatory incentives encouraging gender equality without quotas.	• "Government should incentivize firms so women can access managerial roles." • "Signal that salaries must be equal; otherwise higher taxes for noncompliant firms."	Incentives; tax credits/penalties; bonuses.
6	Cultural change	Interventions aimed at reshaping norms, stereotypes, and social attitudes.	• "Civic education in schools teaching gender equality (e.g., chores taught to all), removal of stereotypes." • "Structure education on equal opportunities; devalue sexist culture."	School; education; media; stereotypes; culture change.
7	Merit	Emphasis that pay and promotions should depend on merit and skills, not gender.	• "Ensure fair salaries reflecting true abilities, regardless of gender." • "Pay and promotions based on skills, not gender." • "Evaluate competencies irrespective of gender."	Merit; skills; competence; performance; meritocracy.
8	No idea	Respondent explicitly reports not knowing which policy to adopt.	• "I don't know." "No idea."	I don't know; no idea.
9	Not classifiable	Responses that do not fit any predefined policy category.	—	—

	(1) Demad	(2) Supply	(3) Info Causes	(4) NoInfo	(5) Total	(6) p(D=S)	(7) p(D=IC)	(8) p(D=C)	(9) p(S=IC)	(10) p(S=C)	(11) p(IC=C)
Should the government do more to decrease gender inequalities?											
No	52	91	48	52	243	0.002	0.892	0.299	0.001	0.052	0.248
Observations	629	661	596	518	2404						
Yes											
Family Support Policy	131	144	141	119	535	0.675	0.234	0.381	0.429	0.627	0.787
Pro-Quota	73	45	62	36	216	0.416	0.969	0.364	0.413	0.896	0.135
Anti-Quota	10	6	7	7	30	0.817	0.688	0.601	0.889	0.497	0.393
Against Discrimination Policy	398	341	351	276	1366	0.000	0.116	0.001	0.009	0.627	0.564
Incentives	103	99	116	89	407	490	0.159	0.716	0.035	0.305	0.327
Cultural Change	213	191	195	154	753	0.055	0.671	0.135	0.143	0.755	0.284
Merit	98	110	96	77	381	0.605	0.801	0.738	0.799	0.408	0.568
No Idea	19	45	26	30	120	0.002	0.212	0.021	0.061	0.478	0.276
Observations	629	661	595	518	2403						

Note that one answer can be classified with more than one category so the number of observation refer just to the total answers

Table E3: Policies to decrease inequalities: Number of observations and balance test, p-value from t-tests

	Family Support Policy	Pro-Quota	Anti-Quota	Against Discrimination Policy	Incentives	Cultural Change	Merit
Family Support Policy	1.000						
Pro-Quota	-0.238*	1.000					
Anti-Quota	0.093*	-0.467*	1.000				
Against Discrimination Policy	-0.075*	-0.104*	0.110*	1.000			
Incentives	0.140*	-0.146*	0.024	-0.036*	1.000		
Cultural Change	0.018*	-0.167*	0.176*	-0.102*	-0.066*	1.000	
Merit	-0.076*	-0.160*	0.174*	0.342*	-0.050*	0.142*	1.000

* $p < 0.05$

Table E4: Policies to decrease inequalities: Correlation matrix

	Men	Women	t-test	p-value
Family Support Policy	0.158	0.275	-6.945	0.000
Pro-Quota	0.661	0.688	-0.498	0.619
Anti-Quota	0.113	0.081	0.098	0.334
Against Discrimination Policy	0.544	0.588	-2.134	0.033
Incentives	0.140	0.194	-3.526	0.000
Cultural Change	0.302	0.322	-1.053	0.293
Merit	0.148	0.167	-1.289	0.198
No Idea	0.067	0.036	3.583	0.000
N	2402			

Table E5: Categoriers Policies to decrease inequalities: means, t-tests and p-value by gender

F OEQ: Motivation for Supporting/Opposing Quotas

This appendix documents the classification and empirical analysis of the open-ended justifications elicited after respondents stated whether gender quotas are a useful instrument to reduce inequalities. Immediately following the "Yes/No" question, participants were asked to explain their answer in their own words. These responses are central to the mechanism analysis in Section 6.2.1, where we examine how informational treatments reshape ex-post reasoning. Using the same preprocessing and dictionary-based procedure described in Appendix D, we coded responses into thematic categories capturing distinct motivations for supporting and opposing quotas. Categories are not mutually exclusive. Table F1 and Table F2 report the definitions, examples, and associated keywords used to classify motivations among supporters and opponents, respectively. Table F3 presents the distribution of motivations by treatment group together with pairwise tests. Table F4 reports regression estimates of the effect of the information treatments on the number of distinct motivations mentioned, while Table F5 examines heterogeneity by prior cultural beliefs. Finally, Table F6 and Table F7 analyze how the treatments affect the length (in words) of respondents' open-ended justifications.

ID	Category	Definition	Examples	Main Keywords
Motivations (Yes = quotas are useful)				
1	Fairness / equality	Quotas are justified as restoring equality or ensuring fairness between genders.	• "There shouldn't be gender differences." • "It would put men and women on equal footing." • "It's right for men and women to be equal."	Right; equality; equal footing; nondiscrimination; fairness; parity.
2	Signaling / information value	Quotas are viewed as revealing information about women's abilities and reducing prejudice.	• "Signals that commitment to equality is essential." • "Lets us assess a woman's qualifications." • "Helps dissolve prejudices about managers."	Information; signaling; skills; merit; quality; prejudice reduction.
3	For women (supply-side)	Quotas are seen as encouraging women's participation, confidence, or empowerment.	• "Women would feel less anxious." • "Some women feel guilty about prioritizing work." • "Quotas would encourage women to be more competitive."	Encourage; confidence; empowerment; competitiveness; role models.
4	No specific justification	Support expressed without a structured or policy-based rationale.	• "Personal reasons lead me to think it could work." • "No specific reason; it's just my view."	Personal view; unspecified reason; intuition.
5	Explicit mention of video	Support justified by reference to the informational treatment.	• "For the reasons outlined in the video I watched."	Video; treatment content.

Table F1: Categories for motivations supporting gender quotas

ID	Category	Examples	Main Keywords
Motivations (No = quotas are not useful)			
1	Not the right instrument / not effective	• "I don't believe it can work without incentives." • "If women aren't supported during key career stages, quotas alone won't solve it (e.g., more support for mothers)." • "It seems utopian: I like the concept, but it's hard to achieve."	Effectiveness; instrument; support policies; incentives; feasibility.
2	Reverse discrimination / meritocracy	• "This creates the opposite disparity; equality is equality only if applied correctly." • "It doesn't respect meritocracy." • "It doesn't reward merit but imposes presence."	Meritocracy; discrimination; disparity; fairness; equality.
3	Unneeded / backlash	• "Women are culturally equal; issues arise later (e.g., pregnancy leaves long absences and firms suffer)." • "Women don't need to prove anything." • "Women don't need quotas to express professionalism."	Unnecessary; unintended effects; backlash; professionalism.
4	No specific justification	• "It's how I feel based on my experience."	Personal view; unspecified reason.
5	Explicit mention of video	• "For the reasons outlined in the video I watched."	Video; treatment content.

Table F2: Categories for motivations opposing gender quotas

	(1) Demad	(2) Supply	(3) Info Causes	(4) NoInfo	(5) Total	(6) p(D=S)	(7) p(D=IC)	(8) p(D=C)	(9) p(S=IC)	(10) p(S=C)	(11) p(IC=C)
Do you think a quota system could be a useful tool for reducing inequalities?											
Yes											
Fairness	194	271	286	194	950	0.105	0.000	0.006	0.005	0.217	0.157
Signalling	214	165	143	78	600	0.000	0.000	0.000	0.690	0.000	0.000
For women's limits	49	49	105	58	261	0.000	0.536	0.890	0.000	0.001	0.467
Do not justify	21	34	29	20	104	0.108	0.177	0.635	0.822	0.296	0.415
Explicit mention to video	5	3	2	0	10	0.538	0.343	0.066	0.719	0.142	0.213
Observations	482	451	417	324	1674						
No											
Wrong Instrument	69	92	71	78	310	0.109	0.605	0.039	0.290	0.580	0.124
Reverse Discrimination	92	128	118	110	448	0.024	0.016	0.003	0.846	0.427	0.554
Unneeded and Backlash	26	25	40	18	109	0.746	0.046	0.564	0.019	0.780	0.015
Do not justify	14	23	11	25	73	0.178	0.639	0.016	0.075	0.246	0.005
Explicit mention to video	0	0	0	0	0	0					
Observations	147	210	179	194	730						

One answer can be classified with more than one category, the number of observation refer just to the total answers
The number of obs refers to the total number of respondents answering respectively YES or NO to the question:
Do you think a quota system could be a useful tool for reducing inequalities?

Table F3: Motivation for supporting/opposing gender quotas: Number of observations, p-value from t-tests

	Number of motivation mentioned		
	(1)	(2)	(3)
Info Causes	0.049	0.010	0.113
	(0.050)	(0.063)	(0.081)
Supply	0.184***	0.258***	0.057
	(0.048)	(0.062)	(0.076)
Demand	0.184***	0.212***	0.167*
	(0.053)	(0.065)	(0.094)
Cultural Causes	-0.045	-0.094**	0.026
	(0.038)	(0.047)	(0.064)
Female	0.028	0.089*	-0.052
	(0.038)	(0.046)	(0.061)
Observations	2,370	1,652	718
Sample	Full sample	QuotaUseful = 1	QuotaUseful = 0

Notes: The dependent variable is the number of distinct motivations mentioned in the open-ended justification for supporting/opposing quotas. All specifications include treatment dummies, demographics, job characteristics, and policy attitudes.
Justification to the question: *Do you think a quota system could be a useful tool for reducing inequalities?* — Why Yes/No.
$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$. Robust standard errors in parentheses.

Table F4: Effect of Information Treatments on the Number of Motivations Mentioned

	Number of motivation mentioned		
	(1)	(2)	(3)
Info Causes	0.043	-0.009	0.114
	(0.057)	(0.072)	(0.096)
Supply	0.193***	0.286***	0.035
	(0.057)	(0.074)	(0.093)
Demand	0.180***	0.203***	0.161
	(0.061)	(0.075)	(0.112)
Cultural Causes	-0.046	-0.096	0.005
	(0.073)	(0.100)	(0.108)
Info Causes × Cultural Causes	0.023	0.073	-0.010
	(0.106)	(0.135)	(0.169)
Supply × Cultural Causes	-0.028	-0.086	0.073
	(0.102)	(0.136)	(0.158)
Demand × Cultural Causes	0.014	0.031	0.021
	(0.105)	(0.133)	(0.188)
Observations	2,370	1,652	718
Sample	Full sample	QuotaUseful = 1	QuotaUseful = 0

Notes: OLS estimates with robust standard errors in parentheses. All models include the full set of demographic, job-related, and policy controls. Treatment indicators are interacted with the *Cultural Causes* indicator (generic belief).

The dependent variable is the number of distinct motivations mentioned in the open-ended justification for supporting/opposing quotas.

Justification to the question: *Do you think a quota system could be a useful tool for reducing inequalities?* — Why Yes/No.

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$. Robust standard errors in parentheses.

Table F5: Effects of Information Treatments on the Number of Motives Mentioned: Heterogeneity by Cultural Causes

	Number of words used in the open-ended response		
	(1)	(2)	(3)
Info Causes	0.049	0.263	0.298
	(0.701)	(0.776)	(1.640)
Supply	0.461	1.022	-0.478
	(0.688)	(0.814)	(1.473)
Demand	3.458***	3.008***	6.353**
	(0.908)	(0.839)	(2.735)
Cultural Causes	-2.853***	-3.332***	-2.921
	(0.591)	(0.508)	(1.832)
Female	0.859	1.285**	1.026
	(0.639)	(0.572)	(1.729)
Observations	2,370	1,652	718
Sample	Full sample	QuotaUseful = 1	QuotaUseful = 0

Notes: The dependent variable measures the length (in words) of the open-ended justification provided in response to the question "Do you think a quota system could be a useful tool for reducing inequalities?", conditional on whether the respondent answered yes or no. All specifications include demographics, job characteristics, and policy attitudes.
$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$. Robust standard errors in parentheses.

Table F6: Effects of Information Treatments on the Length of Open-Ended Justifications

	Number of words used in the open-ended response		
	(1)	(2)	(3)
Info Causes	0.301	0.577	0.082
	(0.911)	(0.949)	(2.209)
Supply	0.321	1.007	-0.678
	(0.861)	(1.032)	(1.732)
Demand	3.318***	3.128***	5.452*
	(1.086)	(1.041)	(2.982)
Cultural Causes	-2.875***	-2.853**	-3.864**
	(1.000)	(1.229)	(1.870)
Info Causes × Cultural Causes	-0.979	-1.262	0.728
	(1.547)	(1.606)	(3.297)
Supply × Cultural Causes	0.449	-0.075	0.535
	(1.326)	(1.537)	(2.780)
Demand × Cultural Causes	0.496	-0.512	3.174
	(1.685)	(1.619)	(5.009)
Observations	2,370	1,652	718
Sample	Full sample	QuotaUseful=1	QuotaUseful=0

Notes: The dependent variable measures the length (in words) of the open-ended justification provided in response to "Do you think a quota system could be a useful tool for reducing inequalities?", conditional on the Yes/No answer. All specifications include demographics, job characteristics, and policy attitudes. Robust standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table F7: Effects of Information Treatments on the Length of Open-Ended Justifications: Heterogeneity by Cultural Causes

G Complete Survey

Table G1: Survey sections, topics, and main questions

Section	Topics	Main Questions
1	Background; job & firm characteristics	• Year/country of birth; place of residence; marital status; # children; education. • Firm size (employees); sector; job role.
1	Importance of gender equality	*Rank* six challenges for the next five years (1 = most important): innovation; productivity; diversity & gender equality; technological progress; teamwork; environmental impact.
2	Incentivized beliefs on inequality (facts)	• For managers (age 30–49, FT, firms >10 employees): "For every €100 earned by a comparable man, how much does a woman earn?" • "What % of CEOs are women among the 50 largest listed firms in Italy?"
3	Beliefs on causes of inequality	• "What are the causes of gender inequalities in pay and managerial positions?"
3	Beliefs on policies	• "Should the government do more to promote gender equality in managerial positions? (Yes/No)" • *If Yes*: "Which policy would be effective?" • Brief description of gender quotas; "Were you already aware of quotas? (Yes/No)"
4	Information treatments	Assignment to *Demand*, *Supply*, or *Info Causes* (or *No Info*).
5	Self-reported policy demand	• "Some argue quotas are justified; others say they unfairly advantage women. Are quotas a useful tool to reduce inequalities? (Yes/No) Why?" • *If Yes*: "Would you agree to introduce quotas in your workplace? (Yes/No)"

Table continues on the next page

Section	Topics	Main Questions
5	Secondary outcomes (perceptions/mechanisms)	• If a quota (more women/fewer men) were introduced, to what extent would it: make people work less; apply for higher positions; stop teamwork? • Who is advantaged/disadvantaged: working women/men; non-working women/men? • "How important is it for men(women) to act against unfair dynamics?" (4-point scale) • Agreement (1 "Strongly agree"–5 "Strongly disagree"): – "If women wanted top positions, they should work harder." – "Past norms/biases hindered women's self-determination and access to top jobs." • Policy importance (4-point scale): gender-sensitive language; flexible hours; paternity leave; pay transparency.
6	Donation (incentivized behavior)	• Lottery: €500 voucher prize. Before knowing outcome, option to donate €0–€500 to the European Women's Lobby (50% match by researchers). • "Do you wish to donate part of potential winnings?" (Yes/No) → "Choose amount."
7	Final questions & feedback	• "Is gender important for career prospects in your workplace? (Yes/No)" • "Are there gender-equality programs in your organization? (Yes/No)" *If Yes & female*: "Have you benefited?" • UNI PdR 125:2022 certification: heard of it? (Yes/No) If Yes: has your org undertaken it? (Yes/No/Don't know) • Perceived political bias of the study: left-leaning / right-leaning / none. • *If treated video*: "Would you like more info/sources used in the video? (Yes/No)" • Incentive follow-up (beliefs task): if selected for €5/answer, choose receive voucher or donate to: Save the Children; Doctors Without Borders; Emergency; ActionAid. • Provide correct answers for the two incentivized belief questions (wage ratio; CEO share).

Table G1: Description of the survey

Alma Mater Studiorum - Università di Bologna

DEPARTMENT OF ECONOMICS

Strada Maggiore 45
40125 Bologna - Italy
Tel. +39 051 2092604
Fax +39 051 2092664
http://www.dse.unibo.it