



Designing Gender-Balanced Evaluation Committees with AI*

J. IGNACIO CONDE-RUIZ
MIGUEL DÍAZ SALAZAR
JUAN-JOSÉ GANUZA

Documento de Trabajo 2026/01
Febrero 2026

fedea

*Las opiniones recogidas en este documento son las de sus autores
y no coinciden necesariamente con las de Fedea.*

Designing Gender-Balanced Evaluation Committees with AI*

J. Ignacio Conde-Ruiz,^{a,c} Miguel Díaz Salazar,^a
Juan-José Ganuza,^{b†}

^aFedea

^bUniversitat Pompeu Fabra and Barcelona School of Economics

^cUniversidad Complutense de Madrid and ICAE

January 2026

Abstract

This paper combines artificial intelligence with economic modeling to design evaluation committees that are both efficient and fair in the presence of gender differences in economic research orientation. We develop a dynamic framework in which research evaluation depends on the thematic similarity between evaluators and researchers. The model shows that while topic balanced committees maximize welfare, this research-neutral-gender allocation is dynamically unstable, leading to the persistent dominance of the group initially overrepresented in evaluation committees. Guided by these predictions, we employ unsupervised machine learning to extract research profiles for male and female researchers from articles published in leading economics journals between 2000 and 2025. We characterize optimal balanced committees within this multidimensional latent topic space and introduce the Gender-Topic Alignment Index (GTAI) to measure the alignment between committee expertise and female-prevalent research areas. Our simulations demonstrate that AI-based committee designs closely approximate the welfare-maximizing benchmark. In contrast, traditional headcount-based quotas often fail to achieve balance and may even disadvantage the groups they intend to support. We conclude that AI-based tools can significantly optimize institutional design for editorial boards, tenure committees, and grant panels.

Keywords: Artificial Intelligence; Evaluation Committees; Committee Quotas; Research Orientation; Machine Learning; Topic Modeling; Institutional Design.

JEL Classification: D72; D82; J16; J78.

[†]Corresponding Author: Juan-Jose Ganuza, Universitat Pompeu Fabra, Ramon Trias Fargas 27, 08005, Spain; E-mail: juanjo.ganuza@gmail.com

*Thanks to Christian Zimmermann, Lorenzo Ductor, Nagore Iriberry, Manu Garcia, Luis A. Puch, Libertad Gonzalez and Judith Valls for helpful comments. We also thank participants at seminars at University of Siena, Universitat Pompeu Fabra, Universidad Carlos III and the SAEe Conference (Barcelona). José Ignacio Conde-Ruiz acknowledges the support of the Research Project of the Ministry of Science and Innovation, PID2023-148090NB-I00. Juan José Ganuza acknowledges the support of the Barcelona School of Economics and the Research Project of the Ministry of Science and Innovation PID2023-153318NB-I00 and from the Spanish Agencia Estatal de Investigación (AEI), through the Severo Ochoa Programme for Centres of Excellence in R&D (Barcelona School of Economics CEX2024-001476-S).

1 Introduction

Despite decades of institutional efforts to fight gender discrimination, women remain persistently underrepresented within the economic profession. While the share of female undergraduate majors has climbed to over 40%, this progress has failed to translate into senior academic roles. Recent assessments by (Lundberg and Stearns, 2019) and Chevalier (2021) reveal a discouraging stagnation: the fraction of female assistant professors has remained largely flat—slightly above 20%—since the mid-1990s. This gender gap is even more pronounced at elite levels; as documented by Siniscalchi and Veronesi (2020), the proportion of women in assistant professor positions at 'top 10' departments has actually declined to below 20%.

Publishing in leading journals is a key determinant of academic success in economics. Articles published in top journals strongly influence tenure and promotion decisions and contribute to defining the evolution of the discipline (Heckman and Moktan, 2020). Over time, competition for publication in these outlets has intensified markedly, as acceptance rates have declined sharply (Card and DellaVigna, 2013). In this context, the underrepresentation of women among authors publishing in top journals represents a potentially important bottleneck in academic career progression.

Figure 1 provides a first descriptive motivation by comparing female representation in two key dimensions of the economics profession in 2024: faculty positions in top-10 economics departments and authorship in leading journal publications. Women represent approximately 24% of faculty positions in top economics departments and about 25% of authors in our publication sample of the top 8 leading economic journals.

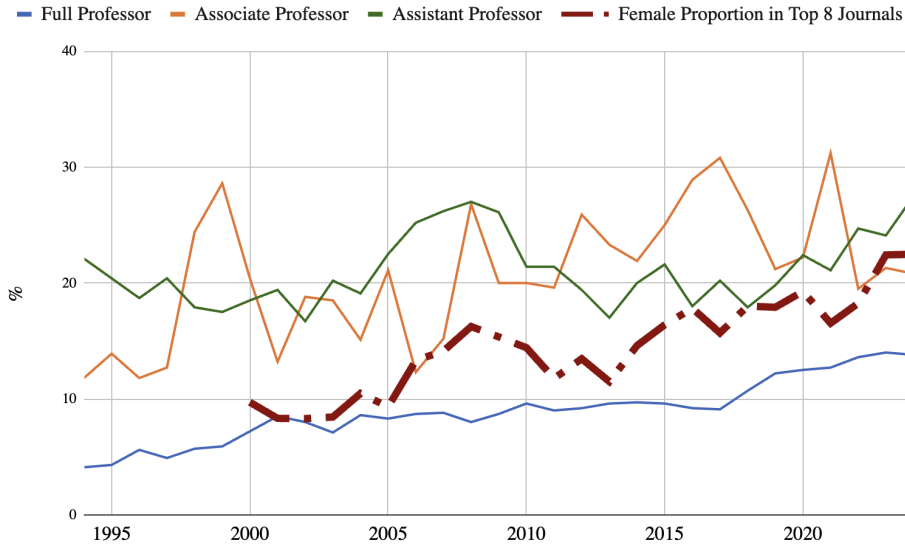


Figure 1 : Female faculty shares by academic rank in top-10 economics departments and on Top-8 Economic Journals, 1995–2024. Faculty data from CSWEP annual reports.

At the same time, existing evidence indicates that gender differences in academic publishing outcomes are not easily explained by direct discrimination in the evaluation process or by systematic differences in research quality. Once referee reports are taken into account, editorial and refereeing decisions in top economics journals appear largely gender-neutral (Card et al., 2020). Related evidence from editorial appointments points in the same direction: conditional on academic CVs, women are at least as likely as men to be selected into editorial roles at top economics journals, and in some periods even more likely. Moreover, female editors tend to handle and publish research in systematically different topics, without differences in ex post quality (Funk et al., 2025). The only controversial evidence relates to citations. (Card et al., 2020) also shows that, conditional on the review process, papers authored by women receive more—citations than papers authored by men. This citation gender gap has been documented by other papers and remains robust to alternative citation adjustments (Koffi, 2021; Hengel and Moon, 2023; Ductor et al., 2024), as well as to mechanisms of cumulative advantage such as the Matthew effect (Merton, 1968). If we take citations as a proxy for quality, the citation gap may raise doubts of whether or not female face tougher quality standards when they try to publish in leading economic journals. However, the literature shows that this gender citation gaps largely vanished when we control for research fields (JEL codes Koffi (2021) and Conde-Ruiz et al. (2025)) or research

latent topics (Conde-Ruiz et al. (2025)).¹

Taken together, these findings stand in tension with the persistent underrepresentation of women in senior academic positions and among authors publishing in the most prestigious journals (Lundberg and Stearns, 2019; Heckman and Moktan, 2020). If evaluation processes are largely gender-neutral and there is not signs of female direct discrimination, how can large and persistent gender gaps in academic careers be sustained? In this paper, we defend that a potential answer to this puzzle could be the gender horizontal differentiation in economic research.

An important feature of the economics profession is that male and female researchers are not evenly distributed across research fields. Empirical evidence shows that women tend to concentrate in specific subfields, while remaining underrepresented in others (Dolado et al., 2012; Bayer and Rouse, 2016; Beneito et al., 2021). These differences emerge early in academic careers and persist over time, shaping publication patterns and research trajectories. More recent work using text-based methods and latent topic models shows that such differences extend beyond coarse field classifications: women and men differ systematically in the research topics they pursue within journals and fields, even after controlling for journal outlets and publication cohorts (Conde-Ruiz et al., 2022b). ¿Through which mechanisms can persistent gender differences in research orientation translate into enduring disparities in academic careers?

A key institutional feature of academic careers is that evaluation and promotion decisions are typically made by small committees composed of senior researchers whose expertise is necessarily uneven across research areas. When evaluators assess work outside their own fields, they face greater informational frictions, making it harder to accurately infer candidates' underlying productivity. As a result, research that is closer to evaluators' own academic backgrounds tends to be assessed with greater precision and confidence. This mechanism ("homo-accuracy bias") has been formalized in models of evaluation with heterogeneous signal precision, where the accuracy of productivity signals depends on the match

¹For example, Conde-Ruiz et al. (2025) uses machine learning techniques over a database of all papers published in T5 economic journals in the last 25 years, and shows that there is an overall positive citation premium of 15 log points for articles authored by women controlling year fixed effects, and 9.4 log points when controlling by journal-year fixed effects. However, the gender citation gap becomes statistically insignificant after controlling for stimulated research latent topics.

between evaluators and candidates (Conde-Ruiz et al., 2022a).² In addition, evaluation may be shaped by self-image concerns, whereby senior researchers place disproportionate weight on research profiles similar to their own, even in the absence of explicit gender bias (Siniscalchi and Veronesi, 2020). When combined with systematic gender differences in research topics, these two mechanisms imply that groups that are initially underrepresented in evaluation committees face noisier assessments or tougher quality thresholds, which lead to weaker incentives to invest in human capital, and lower representation in future committees. Over time, these feedback effects can generate persistent gender disparities in representation and career outcomes, giving rise to what we refer to as a discrimination trap. This paper combines artificial intelligence-based measurement with economic modeling to study how evaluation committees should be designed in the context of gender differences in research orientation to maximize welfare and overcome gender research gaps.

We first develop a stylized dynamic model to analyze the interaction between horizontal research specialization and committee-based evaluation. In our framework, individuals invest in human capital and specialize—with gender-specific propensities—in either "theoretical" or "applied" research topics. Promotion probabilities depend not only on individual productivity but also on the thematic match between the researcher and the evaluation committee. We show that when a committee tilts toward a specific field, the group whose research orientation is underrepresented faces a lower expected return on research, which reduces their incentives to enter the academic labor market and results in fewer promoted researchers from that group. Because future committees inherit the thematic orientation of those currently promoted, small initial imbalances are self-reinforcing. This process generates a "statistical-discrimination trap" where persistent gaps in participation and career advancement emerge as an equilibrium phenomenon, even under gender-neutral evaluation rules. Finally, we demonstrate that while aggregate welfare is maximized by a balanced, research-neutral committee, this efficient allocation is dynamically unstable.

To bring the model to the data, we develop an AI-based measurement framework that uses large-scale text analysis and unsupervised machine learning to extract latent research

²Consistent with this mechanism, empirical evidence shows that evaluators' proximity to research—through coauthorship networks or thematic overlap—affects publication outcomes, even in the absence of explicit favoritism (Ductor and Visser, 2022).

topics from the universe of articles published in leading economics journals over the period 2000–2025³. Each article is represented as a distribution over topics, allowing us to compare research profiles along multiple dimensions simultaneously and to measure research similarity. Our empirical analysis shows persistent differences in thematic orientation between male and female researchers that are largely invisible to conventional field classifications such as JEL codes.

We characterize optimal balanced committees within this multidimensional latent topic space and introduce the Gender-Topic Alignment Index (GTAI) to measure the alignment between committee expertise and female-prevalent research areas. This index serves as a novel AI-based measurement tool to objectively assess whether an evaluation committee achieves gender research neutrality. A central contribution of this paper is to show how such an AI tool can implement “topic quotas”—balancing field representation directly—as a more effective alternative to prevailing demographic regulations that rely on traditional headcounting of demographic groups. Our simulation results confirm that AI-driven committee designs perform effectively in terms of the overall match between evaluation committees and researchers while preserving gender neutrality. In contrast, conventional headcount quotas often fail to resolve underlying informational frictions and may even disadvantage the groups they are intended to support.

The remainder of the paper is organized as follows. Section 2 develops a dynamic theoretical framework in which evaluation accuracy depends on similarity between evaluators and researchers. The model characterizes both static and dynamic implications, showing that although committees balanced in terms of research orientation maximize welfare, this efficient allocation is dynamically unstable. Section 3 brings the model to the data by developing an AI-based framework to measure horizontal differences in research orientation. Using a large corpus of articles published in leading economics journals and unsupervised machine learning methods, we estimate latent research topics and document systematic differences in thematic orientation across male and female researchers. Within this multidimensional topic framework, we characterize the optimally gender balanced committee and introduce

³The analysis covers all articles published between 2000 and 2025 in the Top-8 general-interest economics journals: *American Economic Review*, *Quarterly Journal of Economics*, *Journal of Political Economy*, *Econometrica*, *Review of Economic Studies*, *Economic Journal*, *Review of Economics and Statistics*, and *Journal of the European Economic Association*.

the Gender–Topic Alignment Index (GTAI) as a summary measure of thematic alignment. Finally, we propose designing committees by taking into account the research profiles of committee members in order to achieve perfectly balanced committees (topic based quota), and we compare this approach with the traditional regulation of counting-heads quotas. the Section 4 combines the theoretical model and the empirical measures in a simulation exercise to evaluate alternative committee formation rules, comparing unconstrained, quota-based, and research-balanced designs in terms of matching efficiency, evaluation accuracy, and committee composition. Section 5 concludes and discusses policy implications for the design of evaluation committees.

2 A Simple Theoretical Framework

2.1 Static Model

Consider two populations, men (M) and women (F), each of unit mass. Individuals draw a latent research ability $\theta \in [0, 1]$ from a uniform distribution and privately observe their type. Entering the economics research profession requires paying a fixed cost k , while opting out yields an outside option w . After entry, researchers specialize in either Theory (T) or Applied (A) research topics. Specialization is ex post and differs systematically across genders: with probability $\beta > \frac{1}{2}$, men specialize in Theory while women specialize in Applied research.

Researchers are evaluated by a committee and receive a payoff V if promoted. Let $\alpha \in [0, 1]$ denote the fraction of committee members specialized in Theory. Ex-ante match probabilities between the committee/evaluators and researchers differ across genders because topics differ across genders:

$$m_M(\alpha) = \beta\alpha + (1 - \beta)(1 - \alpha), \quad m_F(\alpha) = (1 - \beta)\alpha + \beta(1 - \alpha).$$

When $\alpha > \frac{1}{2}$, men enjoy better match conditions; the opposite holds when $\alpha < \frac{1}{2}$. A type- θ individual of gender g succeeds with probability $m_g(\alpha)\theta$. Then, entry is optimal whenever $Vm_g(\alpha)\theta - k \geq w$. Letting $\Delta = (w + k)/V$, the entry cutoff is

$$\theta_g^*(\alpha) = \frac{\Delta}{m_g(\alpha)}.$$

This entry cutoff determines the mass of researchers by gender

$$E_g(\alpha) = 1 - \theta_g^*(\alpha).$$

We assume that the promotion premium V is large enough compared with the entry and opportunity costs of research $k + w$, $\Delta < \frac{1}{2}$. This assumption guaranties that there is always a positive mass of researchers. The expected success is given by

$$S_g(\alpha) = \int_{\theta_g^*(\alpha)}^1 m_g(\alpha)\theta d\theta,$$

which is increasing in the match probability and decreasing in the entry threshold.

Proposition 1. *When the committee tilts toward Theory ($\alpha > \frac{1}{2}$), men enjoy higher match probabilities, lower entry thresholds, and hence higher entry mass and higher mass of promoted researchers than women.*

The underlying mechanism is as follows. Given that subject-matter experts on evaluation committees are skewed towards male-oriented research fields, women may revise their priors regarding promotion prospects downwards. This anticipation reduces the expected return on research for women, leading to a decline in both the number of female researchers and the aggregate number of successful female evaluations.

Proposition 2. *Whenever $\alpha > \frac{1}{2}$, conditional on entry, female researchers' average productivity exceeds that of male entrants.*

Since the expected return on research is lower for women, only the most productive females pursue this career path. It follows that, conditional on entry, the average productivity of women is higher than that of men.

Proposition 3. *Total entry and total success are maximized when the committee is balanced, $\alpha = \frac{1}{2}$.*

Mitigating the committee bias toward male topics facilitates a reallocation of talent. Female participation rises while male participation falls. Crucially, the marginal female entrant is more productive than the marginal male exiter. This substitution generates an efficiency surplus, which translates into higher aggregate entry and a greater number of favorable evaluations.

These results highlight that even in the absence of taste-based discrimination or intrinsic ability differences, horizontal specialization combined with committee imbalance is sufficient to generate gender gaps in participation and success, and also may lead to inefficiencies.

2.2 Dynamic Extension

We now extend the framework to a dynamic setting in which the composition of the evaluation committee evolves endogenously over time. Let α_t denote the fraction of Theory evaluators in period t . Researchers promoted in period t become committee members in period $t+1$, so the next-period committee inherits the topic composition of those promoted.

Promotions of theory researchers T occur with mass

$$S_T(\alpha_t) = \alpha_t \left[\int_{\theta_M^*(\alpha_t)}^1 \beta \theta d\theta + \int_{\theta_F^*(\alpha_t)}^1 (1 - \beta) \theta d\theta \right],$$

and promotions of applied researchers A with mass

$$S_A(\alpha_t) = (1 - \alpha_t) \left[\int_{\theta_M^*(\alpha_t)}^1 (1 - \beta) \theta d\theta + \int_{\theta_F^*(\alpha_t)}^1 \beta \theta d\theta \right].$$

The evolution of the committee is therefore governed by

$$\alpha_{t+1} = \mathcal{G}(\alpha_t) := \frac{S_T(\alpha_t)}{S_T(\alpha_t) + S_A(\alpha_t)}.$$

The explicit expression for this dynamic system is provided in Appendix A.

We summarize the qualitative properties of this dynamic system in the following proposition.

Proposition 4. *For the dynamic system $\alpha_{t+1} = \mathcal{G}(\alpha_t)$:*

1. *the fixed points are exactly $\{0, \frac{1}{2}, 1\}$;*
2. *the interior fixed point $\alpha^* = \frac{1}{2}$ is locally unstable, while the endpoints are locally stable.*

Figure 2 illustrates the implications of Proposition 4 by plotting the dynamic map $\alpha_{t+1} = \mathcal{G}(\alpha_t)$ together with the 45-degree line for different values of the topic-matching parameter




Figure 2 : Dynamic evolution of committee composition. The figure plots the dynamic map $\alpha_{t+1} = \mathcal{G}(\alpha_t)$ for two values of β together with the 45-degree line. Fixed points are given by intersections with the diagonal. The interior fixed point $\alpha^* = \frac{1}{2}$ is unstable, while $\alpha \in \{0, 1\}$ are stable attractors. Higher β strengthens the amplification mechanism and increases drift away from the balanced committee.

β . Fixed points correspond to intersections with the diagonal, yielding exactly $\alpha \in \{0, \frac{1}{2}, 1\}$. The figure also makes clear that the balanced committee $\alpha^* = \frac{1}{2}$ is dynamically unstable: for $\alpha_t > \frac{1}{2}$ the map lies above the 45-degree line and the system drifts toward $\alpha = 1$, while for $\alpha_t < \frac{1}{2}$ it lies below the diagonal and the system drifts toward $\alpha = 0$. As β increases, the curvature of \mathcal{G} relative to the diagonal becomes more pronounced, indicating a stronger amplification of small initial imbalances in committee composition.

The instability of $\alpha = \frac{1}{2}$ reflects a simple but powerful amplification mechanism. Let T_t and A_t denote the endogenous masses of Theory- and Applied-oriented candidates who reach evaluation in period t . A committee with share α_t of Theory evaluators promotes Theory candidates with probability α_t , so next period's composition satisfies

$$\alpha_{t+1} = \frac{\alpha_t T_t}{\alpha_t T_t + (1 - \alpha_t) A_t}.$$

If the pool of candidates were exogenous and symmetric ($T_t = A_t$), the dynamic map would be the identity and every committee would reproduce itself⁴. In our framework, however, the pools T_t and A_t are *endogenous*: they depend on α_t through gender-specific match probabilities, entry thresholds, and asymmetric topic choices. When $\alpha_t = \frac{1}{2}$, the evaluation environment is symmetric, leading to $T_t = A_t$ and making $\alpha = \frac{1}{2}$ a fixed point. But whenever α_t deviates even slightly from $\frac{1}{2}$, the induced asymmetries in match quality alter the relative success of Theory and Applied candidates. If $\alpha_t > \frac{1}{2}$, Theory candidates enjoy higher match probabilities and lower entry thresholds, implying $T_t > A_t$ and thus $\alpha_{t+1} > \alpha_t$; conversely, if $\alpha_t < \frac{1}{2}$, we obtain $T_t < A_t$ and $\alpha_{t+1} < \alpha_t$. Hence the dynamic map lies strictly above the 45-degree line for $\alpha_t > \frac{1}{2}$ and strictly below it for $\alpha_t < \frac{1}{2}$, making the efficient symmetric committee composition dynamically unstable. The combination of horizontal specialization and endogenous committee reproduction generates a self-reinforcing mechanism through which small initial imbalances expand over time, pushing committees toward increasingly Theory- or Applied-dominated compositions even in the absence of evaluative bias.

The dynamic instability of the efficient benchmark $\alpha = \frac{1}{2}$ has important implications for the long-run composition of the research populations. Horizontal specialization interacts with endogenous committee reproduction to create a statistical-discrimination trap: initial asymmetries in committee composition generate persistent, and potentially widening, gender gaps in participation and career advancement, even when evaluators themselves are unbiased. This perspective also clarifies the potential role of quota policies. A committee constrained to maintain $\alpha = \frac{1}{2}$ eliminates the dynamic amplification mechanism and restores both gender and topic balance in the long run. Moreover, because women face higher entry thresholds when $\alpha > \frac{1}{2}$, enforcing balance increases efficiency by allowing high-productivity women—who would otherwise be deterred—to enter and succeed, replacing lower-productivity male entrants. However, real-world gender quotas implemented through “counting heads” are an imperfect proxy for topic balance: even a numerically balanced committee may exhibit substantial variation in its topic composition due to sampling noise. This observation motivates our focus on “topic quotas”—that is, on ensuring balanced represen-

⁴Notice that $\alpha_{t+1} = \frac{\alpha_t T_t}{\alpha_t T_t + (1-\alpha_t) A_t} = \alpha_t f(T_t, A)$, where $f(T_t, A) = 1$ if $T_t = A_t$, $f(T_t, A) > 1$ if $T_t > A_t$ and $f(T_t, A) < 1$ if $T_t < A_t$.

tation of research fields rather than demographic groups—as a more direct and potentially more effective mechanism for preventing dynamic drift in committee composition.

Our simple dynamic framework provides a theoretical rationale for regulatory interventions, such as quotas, to mitigate imbalances in evaluation committees. However, conventional gender quotas—typically implemented through a “counting heads” approach—serve as an imperfect proxy for topic balance. Even a numerically balanced committee may exhibit substantial thematic variation due to sampling noise, failing to achieve the desired equilibrium in research expertise. This observation motivates our focus on “topic quotas”—prioritizing balanced representation across research topics rather than demographic groups—as a more direct and potentially more effective mechanism for escaping the statistical-discrimination trap. A potential explanation for the spread of demographic quotas is their ease of implementation. We acknowledge that to carry out a topic-based regulation is challenging, given that research is inherently multidimensional and difficult to categorize. The remainder of this paper demonstrates how unsupervised machine learning techniques can be leveraged to overcome these classification challenges and help us to implement a topic-balanced evaluation benchmark.

3 From the Model to the Data

We translate the fundamental mechanisms of our stylized model into an empirical framework, which allows us to discuss various policy instruments aimed at reducing gender gaps in economic research. Our roadmap is as follows:

1. We begin by testing the β -channel—the hypothesis that men and women systematically specialize in different research areas—using an unsupervised machine learning algorithm (a Structural Topic Model) to estimate gender-specific topic distributions across articles published in leading economics journals.
2. Building on these latent topic distributions, we generalize the matching technology by using cosine similarity to measure the thematic proximity between candidates and evaluators, a step that allows us to characterize the composition of optimal, research-neutral committees.

3. We introduce a new empirical tool, the Gender-Topic Alignment Index (GTAI), a continuous scalar measure that captures the degree to which a specific paper’s content, evaluator, or committee aligns with research topics that are relatively more prevalent among female researchers.
4. Finally, leveraging this index, we propose the implementation of “topic quotas” as a more direct and potentially more effective mechanism for preventing the statistical-discrimination trap than traditional demographic head-counting, which often serves as an imperfect proxy for thematic balance in evaluation committees.

3.1 Gender Horizontal Differences in Research Topics

3.1.1 Estimation of Latent Research Topics

The empirical analysis begins by documenting gender differences in economic research, a necessary first step to bring the theoretical framework to the data. Our analysis relies on a large corpus of 12,795 articles published between 2000 and 2025 in eight leading economics journals: *Econometrica*, *Journal of Political Economy*, *American Economic Review*, *Economic Journal*, *International Economic Review*, *Review of Economic Studies*, *Journal of the European Economic Association*, and the *Quarterly Journal of Economics*. Table 1 reports the number of articles published in each journal over the sample period.

Journal	AER	Econ	EJ	IER	JEEA	JPE	QJE	ReStud
Number of Articles	2343	1530	2326	1371	1355	1255	1113	1502

Table 1 : Number of articles published by journal.

We do not directly observe the gender of authors in our data. To analyze gender-related patterns, we classify authors by gender based on their first names.⁵ Table 2 summarizes the resulting distribution of authorship by gender, both at the article level and at the individual author level.

⁵We rely on three different databases: (i) the first-names database published by the U.S. Social Security Administration, created using data from Social Security card applications; (ii) the database constructed by Tang et al. (2011), which uses Facebook data on first names and self-reported gender; and (iii) the names database developed by Bagues and Campa (2017). We manually check any author who (a) falls within the [0.05, 0.95] probability interval of being male or female, or (b) cannot be found in any of the databases.

⁵Percentages are calculated over the total number of papers or authors.

Articles						Authors		
N articles	Male	Mostly Male	Neutral	Mostly Female	Female	Total	Male	Female
12795	69.80	11.33	10.69	2.51	5.62	12093	82.01	17.99

Table 2 : Distribution of authorship by gender.

Following, [Conde-Ruiz et al. \(2022b\)](#) and [Conde-Ruiz et al. \(2025\)](#), we use an unsupervised machine learning methodology to estimate latent research topics from article abstracts and to document gender differences in their distribution. This topic-based representation provides a multidimensional characterization of research content that is well-suited to the matching framework developed in the model and serves as the foundation for the similarity measures and alignment indices introduced in subsequent sections.

To uncover the latent thematic structure of economic research, we utilize the Structural Topic Model (STM), which allows for a probabilistic, low-dimensional representation of high-dimensional textual data while preserving essential informational content. Unlike foundational algorithms such as Latent Dirichlet Allocation (LDA), the STM is "structural" because it incorporates document-level metadata—specifically journal names and publication years—as covariates to inform the estimation of topic prevalence. This approach better captures the shifting relationships between words and latent themes across editorial lines and over time.

The implementation of this methodology begins with building an operative data base. To ensure semantically meaningful topics, we conduct a rigorous text-cleaning procedure on the corpus of our 12795 abstracts. During this stage, we convert all text to lowercase, remove stop-words based on the SMART list, apply linguistic stemming to consolidate related terms, and filter out infrequent words that appear only once or twice across the corpus. This procedure effectively reduces the initial vocabulary from 26093 words to a focused, high-information corpus of 5010 unique tokens. After this text-processing we represent our text data in a document-term matrix of D rows (12795 abstracts) and V columns (5010 unique words in our corpus) where the element (d, v) of the matrix is the number of times the v_{th} unique word appears in the d_{th} abstract. This document-term matrix that reduces the dimensionality of our original text variables is the input of the algorithm.

Upon this refined corpus, the STM algorithm jointly estimates the latent research themes—defined as probability distributions over the vocabulary (β_k)—and the proportional

allocation of each document across these themes (θ_d). The document-topic distribution, θ_d , captures the multidimensional nature of research by allowing individual abstracts to load on multiple topics simultaneously. Within this framework, each paper d is represented by a distribution vector $\theta_d = (\theta_{d1}, \dots, \theta_{dK})$, where θ_{dk} measures the share of the document's content associated with topic k , and $\sum_{k=1}^K \theta_{dk} = 1$. Finally, we determine the optimal dimensionality of the model by selecting the number of topics (K) that maximizes the model's likelihood, ensuring a robust balance between statistical fit with our data D and thematic interpretability. We estimated models ranging from $k = 15$ to $k = 65$, assessing them through *held-out likelihood*, *exclusivity*, and *semantic coherence*. We select $K = 50$ as the optimal benchmark, providing a parsimonious yet comprehensive mapping of the discipline's thematic landscape.

The estimated topics capture meaningful dimensions of economic research, including both substantive fields and methodological approaches. Figure 3 reports the prevalence of each topic in the corpus together with representative keywords, illustrating the semantic coherence of the estimated topic space. Importantly, topics are estimated in an unsupervised manner and independently of author gender, ensuring that any gender differences documented below reflect differences in research orientation rather than mechanical features of the estimation procedure.

The STM framework also allows us to visualize the structure of the topic space and the relationships between topics. Figure 4 displays the topic network for the full sample of articles.

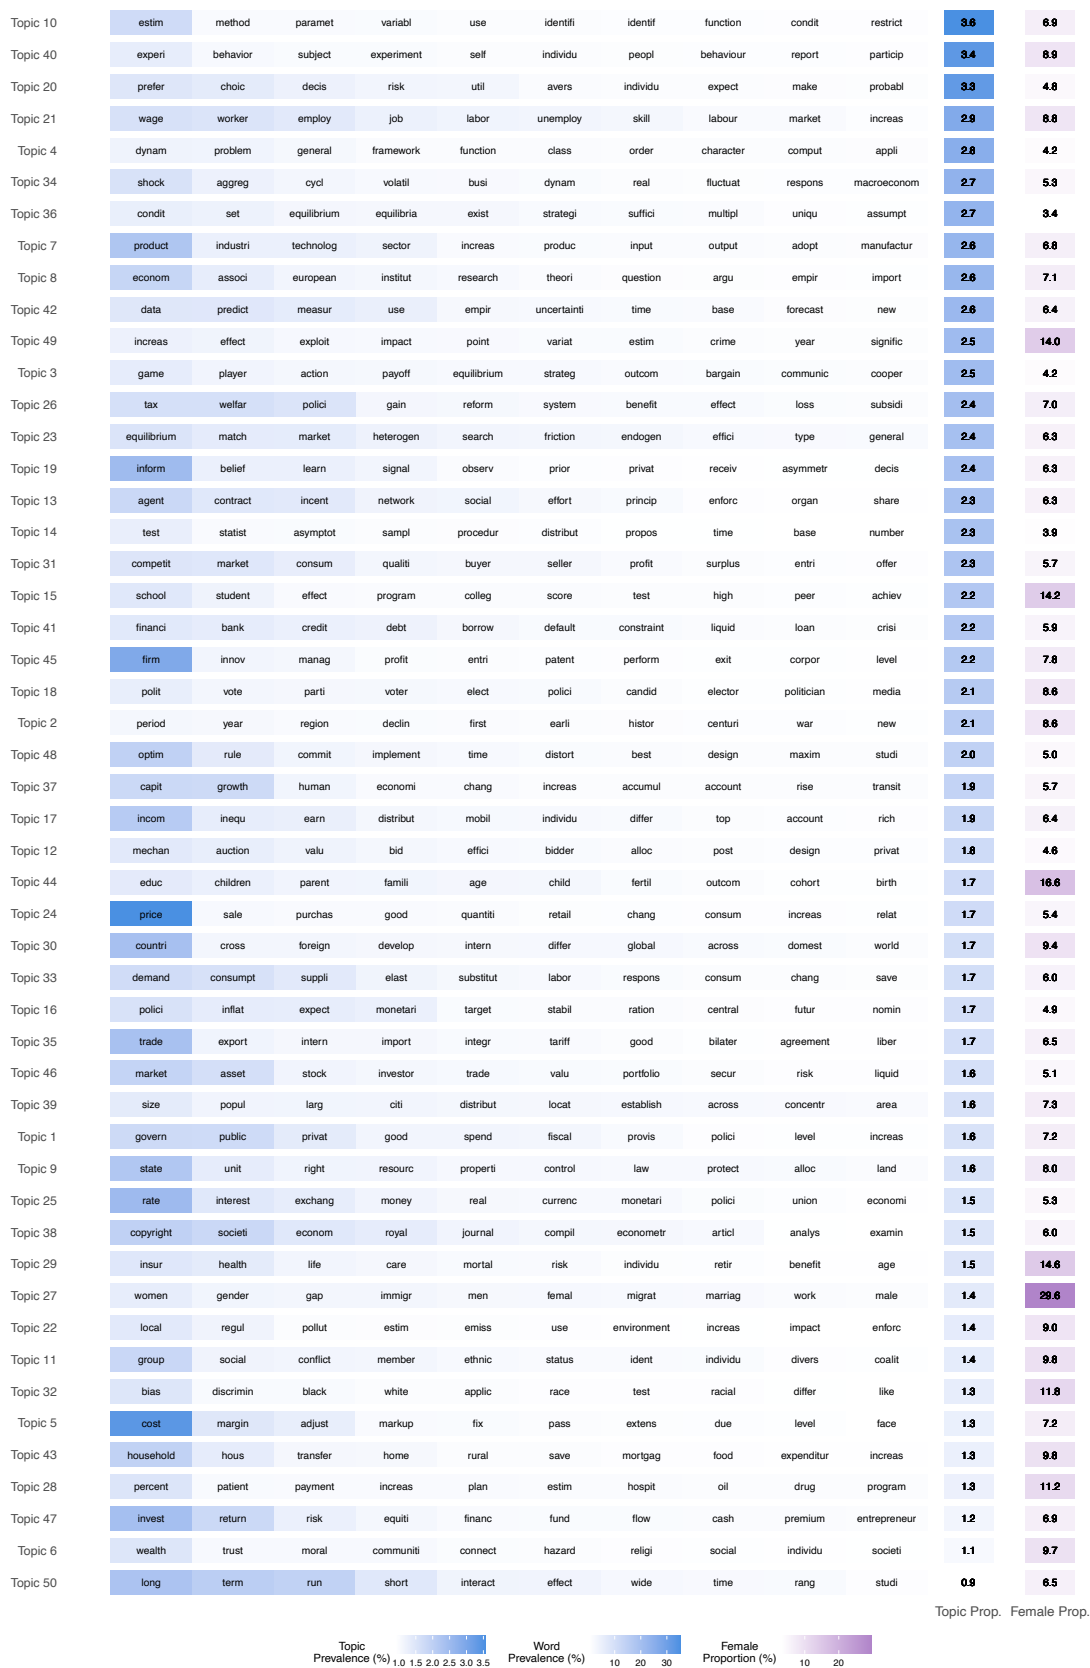


Figure 3 : Prevalence of latent topics, representative words per topic in the corpus and proportion of female papers in each topic.

Note: The first numerical column shows the topic's share in the overall corpus (i.e., its prevalence across all abstracts). The second column reports the proportion of female-authored papers associated with each topic. Topics are ordered by prevalence. Color shading reflects the share of female authorship: darker shades indicate higher female representation relative to the median across topics.




Figure 4 : Topic conectness and prevalence

This figure allows us to analyze the connectivity among research topics and how individual documents are distributed across the identified topics. This mapping is done using the document-topic distributions, θ_d . The connectivity between topics indicates semantic similarity, with applied areas such as health, education, and labor clustering together, while theoretical and econometric topics form distinct, more isolated nodes.

We can build up a similar figure with information about male and female authors. Using our classification of authors' names by gender and the allocation of documents to latent topics. Figure 5 shows latent topics where the sizes of circles are proportional to the percentage of male authors working in such topics. Notice the similarity between Figure 4 and Figure 5 because male are 80% of the authors.

However, Figure 6 is related to female economic research and provides initial evidence that economic research follows distinct patterns across genders, since it differs substantially

from the previous ones.




Figure 5 : Topic network for male-authored papers, reflecting prevalence across the overall corpus.

3.1.2 Research Diversification: The HHI Analysis

A natural concern is whether these horizontal differences are driven by variations in thematic concentration—that is, whether one gender tends to be more “specialized” or “diversified” in its research agenda than the other. To address this concern, we employ the Herfindahl–Hirschman Index (HHI) as a parsimonious measure of topic dispersion. In our context, we treat each author as a “market” and the latent research topics as “firms.” For any given author a , the share of their research dedicated to topic k , denoted by s_{ak} , is calculated by averaging the document–topic distributions (θ_d) across all of their published articles. The




Figure 6 : Topic network for female-authored papers, reflecting prevalence across the overall corpus.

HHI for author a is then defined as:

$$HHI_a = \sum_{k=1}^K s_{ak}^2. \quad (1)$$

A higher HHI indicates that an author’s research is highly concentrated in a narrow set of topics, while a lower value reflects a more diversified portfolio across the latent topic space. Our empirical results suggest that diversification patterns are remarkably consistent across genders. At the aggregate level, the HHI for female-authored papers is 0.023, compared to 0.022 for male-authored papers. As illustrated by the kernel density estimates in Figure 7, the two distributions largely overlap, suggesting that male and female researchers exhibit comparable levels of thematic concentration at both the article and author levels.




Figure 7 : Distribution of the Herfindahl–Hirschman Index (HHI) at the author level, by gender.

This reinforces our interpretation that gender differences in economics are primarily horizontal in nature—affecting the *direction* of research specialization rather than its concentration.

3.1.3 Gender Conditional Topic Distribution

As female authors are underrepresented, the most informative way to illustrate and analyze gender-based horizontal differences is to compute the conditional topic distribution (conditional on having published) for men and women. Let \mathcal{F} and \mathcal{M} denote the sets of papers authored exclusively by women and exclusively by men, respectively. We define the gender-specific average topic distributions as

$$\theta^f = \mathbb{E}[\theta_d \mid d \in \mathcal{F}], \quad \theta^m = \mathbb{E}[\theta_d \mid d \in \mathcal{M}],$$

These vectors summarize the conditional topic profiles by gender, that is, $\Pr(t \mid f)$ and $\Pr(t \mid m)$, the probability that a female or male author conducts research in topic t . For notational convenience, and to match the notation used in the topic-based matching framework below, we refer to these average topic profiles as

$$F_f \equiv \theta^f, \quad F_m \equiv \theta^m.$$

Figure [8](#) reports these conditional topic distributions. The figure reveals clear horizontal differences in research orientation: male and female research profiles are distributed differently across the topic space.




Figure 8 : Gender-conditional topic distributions in top-8 economics journals (2000–2025).

It is informative to describe the research areas in which male- and female-authored papers are most prevalent in absolute terms. Focusing on the highest-weight topics in the gender-specific profiles F_f and F_m , we find that female-authored papers are, on average, more concentrated in applied research areas such as health, education, labor, family economics, and policy-oriented and behavioral topics. By contrast, male-authored papers place relatively more weight on theoretical, quantitative, and methodological topics, including core economic theory, econometrics, and mathematical modeling. These patterns are consistent with earlier evidence on gender specialization across fields, but the topic-based representation reveals them in a continuous and multidimensional way. To further illustrate these gender differences, it is useful to focus on salient topics in which there are large differences in prevalence between male and female authors. To do so, we adapt to our context the concept of stereotypes.

3.1.4 Gender Stereotypes in Research

Since Phelps (1972) and Arrow (1973), we know that stereotypes can sustain inequality and hinder fair treatment and opportunities. Gender stereotypes (e.g., the underperformance of women in leadership or STEM fields) may undermine female self-confidence, affecting

performance, or lead to discrimination, as they often create inaccurate expectations about individuals based solely on gender. There is extensive evidence of stereotype biases in other contexts (see, for example, [Reuben et al. \(2014\)](#), [Bordalo et al. \(2019\)](#), and [Bohren et al. \(2019\)](#)). Here, we investigate whether horizontal gender differences in research topics may give rise to stereotypes.

To do so, we follow the model of stereotypes developed by [Bordalo et al. \(2016\)](#), which formalizes the representativeness heuristic introduced by [Kahneman and Tversky \(1972\)](#). Stereotypes are understood as simplified mental representations that emphasize the most distinctive traits of one group relative to another. While stereotypes often contain some truth, they can distort reality by exaggerating differences between groups. Importantly, these representations depend on context: the way a group is perceived depends on the reference group to which it is compared.

Kahneman and Tversky’s approach can be summarized as follows: stereotypes are formed by emphasizing the features that make one group stand out relative to another, rather than reflecting the full distribution of traits. [Bordalo et al. \(2016\)](#) develop a formal model based on probability distributions to characterize how beliefs about a group are shaped by its most distinctive characteristics, highlighting the role of comparative context in perception formation. In particular, they consider two populations, G and $-G$, characterized by their distributions over a set of types or features $T = \{t_1, t_2, \dots, t_K\}$. A type t^* is representative of group G relative to the reference group $-G$ if it maximizes the likelihood ratio:

$$t^* \in \arg \max_{t \in T} \frac{\Pr(t | G)}{\Pr(t | -G)}. \quad (2)$$

We apply this approach in our setting by identifying the latent topics that best represent female (male) authors relative to male (female) authors. Specifically, we select the topic that maximizes the likelihood ratio between the two groups. Using the conditional topic distributions $F(T | f)$ and $F(T | m)$, [Figure 9](#) reports the likelihood ratio by research topic,

$$\frac{\Pr(t_k | f)}{\Pr(t_k | m)}.$$




Figure 9 : Likelihood ratio by research topic.

Figure 10 illustrates the content of latent topics 27 and 44 using word clouds. Based on the keywords associated with each topic, we can infer their corresponding research areas. Topic 27 (Panel (a)) appears to relate to gender studies, racial discrimination, and social inequality, as indicated by words such as “women,” “gender,” “immigration,” and “gap.” This suggests a focus on labor economics, public policy, and research on diversity and inclusion. Topic 44 (Panel (b)) seems to correspond to family economics and child welfare, with words such as “children,” “family,” “parent,” “education,” and “birth,” pointing to research areas related to family dynamics, marriage, child development, and access to education, likely within development economics, social policy, and education economics. These topics highlight applied areas of research in which female authors are more strongly represented.

Similarly, Topics 36 and 14 play an analogous role for male authors. Figure 11a shows the word cloud for Topic 36. The most prominent terms, such as “condition,” “set,” “equilibrium,” “equilibria,” and “strategy,” suggest that this topic is centered on game theory and equilibrium analysis. Figure 11b displays the word cloud for Topic 14. The most salient terms, including “test,” “sample,” “asymptotic,” “distribution,” and “statistics,” indicate that this topic focuses on econometric and statistical inference.

Stereotypes may contain some truthful information about group characteristics but can also generate distorted beliefs and inaccurate perceptions. As illustrated by Bordalo et al. (2016), stereotypical associations—such as linking Florida with an elderly population or Ireland with red hair—rely on traits that are more prevalent in those groups than elsewhere,

them, with $\Pr(27 | f) = 0.0517$. Topic 44 is also relatively prevalent among female authors ($\Pr(44 | f) = 0.033$). However, even when a stereotypical topic coincides with the highest within-group prevalence, it does not imply numerical dominance: female authors account for only about 30% of total authors within these topics (see Figure 3). Similarly, for male researchers, Topics 36 and 14 are stereotypically associated with men but display much lower prevalence ($\Pr(36 | m) = 0.029$ and $\Pr(14 | m) = 0.025$) than Topic 10, which is the most prevalent topic among male authors ($\Pr(10 | m) = 0.36$).

3.2 Topic-Based Matching and the Design of Optimal Committees

To connect the stylized theoretical model with the empirical evidence on research topics, we generalize the matching technology to a multi-dimensional topic space. Researchers, papers, and evaluation committees are represented as probability distributions over topics, and match quality is measured by cosine similarity. This representation captures the idea that promotion and evaluation depend on thematic affinity between candidates and evaluators. We characterize the committee that maximizes ex-ante matching across genders and show that the optimal committee corresponds to a topic-balanced benchmark analogous to the efficient $\alpha = \frac{1}{2}$ outcome in the theoretical model.

3.2.1 Topic Representations and Cosine-Based Matching

To translate the theoretical notion of match quality into a multi-topic empirical setting, we represent both researchers and evaluation committees as probability vectors over a set of K research topics. The primary unit of analysis is the individual article d , which the Structural Topic Model (STM) characterizes as a latent topic distribution vector $\theta_d = (\theta_{d1}, \dots, \theta_{dK})$. To characterize an individual researcher a , we define their aggregate research profile, θ_a , as the arithmetic mean of the topic distributions across all their published articles in our sample:

$$\theta_a = \frac{1}{|D_a|} \sum_{d \in D_a} \theta_d \quad (3)$$

where D_a represents the set of papers authored by researcher a . The committee's expertise and orientation are measured directly through the observed research output of the




Figure 12 : Cosine-based topic matching between a candidate (p) and a committee (c). Match quality $m(p, c)$ increases as the angle θ between the topic vectors decreases.

evaluators. Following the previous logic of aggregation, we represent the thematic orientation of an evaluation committee C as the aggregate topic profile of its members. The committee’s distribution vector, c , is thus computed as the average of the individual topic profiles of its members:

$$c = \frac{1}{|C|} \sum_{a \in C} \theta_a \quad (4)$$

Summarizing a researcher or paper is characterized by a topic mixture $p = (p_1, \dots, p_K) \in \Delta^{K-1}$, while a committee is described by a vector $c = (c_1, \dots, c_K) \in \Delta^{K-1}$, where $\sum_k p_k = \sum_k c_k = 1$ and all entries are non-negative.

Match quality between a researcher and a committee is measured using cosine similarity, defined as

$$m(p, c) = \frac{p \cdot c}{\|p\| \|c\|} = \cos(p, c) \in [0, 1].$$

Cosine similarity captures the extent to which two topic distributions are aligned: it is high when p and c point in similar directions in topic space and low when they diverge. Figure [12](#) illustrates this geometry: the match $m(p, c)$ increases as the angle θ between the two vectors decreases.

Let F_m and F_f denote the empirical average topic distributions of male and female researchers in our sample—that is, the conditional topic profiles of those who have successfully published in top economics journals. Since these represent the observed pool of active researchers, we measure the expected match quality for each gender by the cosine between their average profile and the committee:

$$m_m(c) = \cos(F_m, c), \quad m_f(c) = \cos(F_f, c).$$

These expressions capture, in reduced form, how well a committee aligns with the research agendas of men and women currently present in the profession. Because promotion and evaluation are increasing in match quality (Section 2), these measures form the basis for the welfare objective analyzed in the next subsection.

3.2.2 Welfare Objective, Optimal Committees, and Equal Opportunity

Having defined gender-specific match functions, we now turn to the normative question of how evaluation committees should be designed. A key insight of the theoretical model is that committee composition affects promotion probabilities through match quality, and that imbalances can generate systematic gender differences even in the absence of evaluative bias. In the empirical setting, however, the observed population of researchers is already asymmetric, both in size and in topic composition. This raises the question of which notion of optimality should guide committee design.

One natural criterion is to maximize aggregate match quality across genders, while allowing for different weights for males and females. We model this objective as

$$W(\gamma; c) = \gamma m_m(c) + (1 - \gamma) m_f(c),$$

where $\gamma \in [0, 1]$ captures the relative welfare weight assigned to male researchers. This formulation encompasses several benchmarks. Setting $\gamma = \frac{1}{2}$ corresponds to equal normative weight across genders, consistent with the idea of the theoretical model of promoting future female participation despite their current underrepresentation.⁶ Choosing γ equal to the population share of men instead reflects a purely utilitarian objective that weights groups by size. We keep γ general to make explicit how different normative choices map into different committee designs.

⁶We could devise an alternative notion of optimality that focuses not on aggregate match quality but on equalizing evaluation conditions across genders. In this perspective, a committee is considered fair if a representative male and a representative female researcher face the same ex-ante probability of success. In our framework, this corresponds to the condition $m_m(c) = m_f(c)$. Among all committees that satisfy this equality-of-opportunity constraint, a natural choice is the one that maximizes the common match level. Our conjecture is that this *equal-opportunity committee*, should be very similar to the optimal committee with the normative weight of $\frac{1}{2}$.

Define the normalized topic profiles

$$u_m = \frac{F_m}{\|F_m\|}, \quad u_f = \frac{F_f}{\|F_f\|}.$$

Since cosine similarity depends only on directional alignment, the committee that maximizes $W(\gamma; c)$ must lie in the span of u_m and u_f .

Proposition 5. [*Utilitarian optimal committee*] For any $\gamma \in [0, 1]$, the committee that maximizes $W(\gamma; c)$ has topic vector

$$c^U(t; \gamma) = \frac{\gamma u_m(t) + (1 - \gamma) u_f(t)}{\sum_{j=1}^K [\gamma u_m(j) + (1 - \gamma) u_f(j)]}.$$

The proof, provided in Appendix A, shows that the welfare gradient points in the direction $\gamma u_m + (1 - \gamma) u_f$, which is then renormalized to lie in the simplex. The utilitarian optimal committee therefore interpolates between male and female topic profiles, with the degree of tilt governed by the welfare weights.

Two special cases of the utilitarian optimum are worth highlighting. First, when the norms of the empirical topic distributions are equal, $\|F_m\| = \|F_f\|$, normalization plays no role and the utilitarian optimal committee simplifies to a weighted average of the raw topic profiles,

$$c^U(t; \gamma) = \frac{\gamma F_m(t) + (1 - \gamma) F_f(t)}{\sum_{j=1}^K [\gamma F_m(j) + (1 - \gamma) F_f(j)]}.$$

This case is empirically relevant in our setting. F_f and F_m are probability vectors over topics, their HHIs satisfy $\text{HHI}(F_g) = \sum_{k=1}^K F_g(t_k)^2 = \|F_g\|_2^2$ for $g \in \{f, m\}$. We have shown that $\text{HHI}(F_f)$ and $\text{HHI}(F_m)$ are nearly identical which therefore implies that $\|F_f\|_2$ and $\|F_m\|_2$ are very similar.

Second, when both conditions hold—equal norms and equal welfare weights, $\gamma = \frac{1}{2}$ —the utilitarian optimal committee reduces to the simple midpoint of the two distributions,

$$c^U(t) = \frac{1}{2} F_m(t) + \frac{1}{2} F_f(t).$$

This expression mirrors exactly the efficient benchmark $\alpha = \frac{1}{2}$ in the theoretical model: a committee that balances topic representation across genders, maximizes aggregate match

quality when groups are treated symmetrically.

To illustrate the role of the vector norms and welfare weights over the design of the optimal committee, consider a simple two-topic environment. Suppose that the average topic profile of male researchers is $F_m = (1, 0)$, while that of female researchers is $F_f = (0.5, 0.5)$. The corresponding normalized profiles are $u_m = (1, 0)$ and $u_f = (1/\sqrt{2}, 1/\sqrt{2})$.

With equal welfare weights, $\gamma = \frac{1}{2}$, the utilitarian optimal committee is

$$c^U\left(\frac{1}{2}\right) \approx (0.71, 0.29).$$

Which generates a female and male matching values of $m_m(c^U(\frac{1}{2})) = 0.925$ and $m_f(c^U(\frac{1}{2})) = 0.922$.⁷ This optimal committee is not the average between F_m and F_f , which would be $\bar{c} = (0.75, 0.25)$ and would generate matching values of $m_m(\bar{c}) = 0.948$ and $m_f(\bar{c}) = 0.8944$ (and lower aggregate matching values $(m_m(c^U(\frac{1}{2})) + m_f(c^U(\frac{1}{2}))) = 1.847 > m_m(\bar{c}) + m_f(\bar{c}) = 1.8424$). While a standard average favors the distribution vector with larger norm, the optimal committee weights both distribution vectors equally, which in relative terms, means to get closer to the distribution vector of lower norm.

By contrast, if welfare weights reflect population shares, say $\gamma = 0.8$, the utilitarian optimal committee becomes

$$c^U(0.8) \approx (0.87, 0.13),$$

placing substantially more weight on the male-dominated topic. This committee maximizes aggregate match quality under the chosen weights, but it no longer equalizes evaluation conditions: male researchers enjoy substantially higher match quality than female researchers $m_m(c^U(0.8)) = 0.989$ and $m_f(c^U(0.8)) = 0.803$.

This example highlights a central trade-off. When group sizes differ, the committee that maximizes aggregate match quality may diverge sharply from the committee that guarantees equal opportunity. Weighting groups by their current representation tends to reinforce historical imbalances, especially when the majority group is thematically concentrated. By contrast, equal opportunity requires a committee that balances topic representation in a way that offsets these asymmetries.

⁷As we anticipated matching values are very close under equal normative weights.

3.3 The Gender–Topic Alignment Index (GTAI)

In this section, we introduce the Gender–Topic Alignment Index (GTAI), a summary measure designed to capture the relationship between a paper’s research content and gender-specific patterns of research specialization. The motivation for this index stems from the fact that research content is inherently multidimensional, as each paper loads on multiple latent topics, and the GTAI will provide a parsimonious scalar measure that summarizes the gender-related orientation of research content in a single, continuous statistic. For this reasons the GTAI may be a usefull tool for many empirical exercises and potentially for designing and implementing regulations.

As documented in Section 3, female- and male-authored papers exhibit systematically different conditional distributions over latent research topics. Let $F_f \in \mathbb{R}^K$ and $F_m \in \mathbb{R}^K$ denote the average topic distributions of female- and male-authored papers, respectively, where each vector lies in the probability simplex and summarizes the conditional topic profiles $P(t | f)$ and $P(t | m)$. A natural way to characterize gender-related differences in research orientation across the entire topic space is through the difference vector

$$F_f - F_m,$$

which captures, for each topic, the relative prevalence of that topic among female-authored papers compared to male-authored papers. This vector defines a direction in the latent topic space along which gender differences in research orientation are most pronounced.

We define the Gender–Topic Alignment Index of document d as the cosine similarity between its topic distribution $\theta_d = (\theta_{d1}, \dots, \theta_{dK})$ and the gender-difference vector $F_f - F_m$:

$$\text{GTAI}_d = \cos(\theta_d, F_f - F_m) = \frac{\sum_{k=1}^K \theta_{dk} (F_{f,k} - F_{m,k})}{\sqrt{\sum_{k=1}^K \theta_{dk}^2} \sqrt{\sum_{k=1}^K (F_{f,k} - F_{m,k})^2}} \in [-1, 1].$$

By construction, a positive GTAI indicates that the topic composition of a paper is more closely aligned with topics that are relatively more prevalent among female-authored papers, while a negative value indicates closer alignment with topics that are relatively more

prevalent among male-authored papers.⁸

After computing the GTAI for all papers in our sample, Table 3 reports illustrative examples of article titles with the highest and lowest values of the index. Papers with high GTAI values are predominantly concentrated in applied research areas such as family economics, health, education, and gender-related policy, whereas papers with low GTAI values are largely concentrated in theoretical and game-theoretic research. These examples provide an intuitive interpretation of the index and are fully consistent with the topic-level evidence on horizontal specialization documented in Section 3.

Table 3 : Illustrative titles with extreme values of the Gender–Topic Alignment Index (GTAI)

Panel A. Highest GTAI values	
Title	GTAI
-More Missing Women, Fewer Dying Girls: The Impact of Sex-Selective Abortion on Sex at Birth and Relative Female Mortality in Taiwan	0.750
-Social Interactions in High School: Lessons from an Earthquake	0.747
-Non-Native Speakers of English in the Classroom: What are the Effects on Pupil Performance?	0.746
Panel B. Lowest GTAI values	
Title	GTAI
-A General Formula for Valuing Defaultable Securities	-0.426
-On the Global Convergence of Stochastic Fictitious Play	-0.421
-Strategically Simple Mechanisms	-0.420

Notes: The table reports illustrative examples of papers with the highest and lowest values of the Gender–Topic Alignment Index (GTAI). High-GTAI papers tend to be concentrated in applied and policy-oriented research areas, while low-GTAI papers are predominantly theoretical and game-theoretic. The table is intended for illustrative purposes only.

Figure 13 shows the evolution of the average Gender–Topic Alignment Index over time, computed across all articles published in the sample period. For each year, the figure reports the mean GTAI across published articles, thereby capturing changes in the overall thematic orientation of published research with respect to topics that are relatively more prevalent among female versus male authors.

⁸The GTAI admits a simple geometric interpretation. Since cosine similarity depends only on the angle between two vectors, the index measures how closely the topic distribution of a document aligns with the direction $F_f - F_m$, independently of the overall dispersion or concentration of topics within the paper. Documents whose topic mixtures point more strongly in this direction receive higher GTAI values.

The figure reveals a clear upward trend in the average GTAI over time. In the early years of the sample, the average GTAI is negative, indicating that published research was, on average, more closely aligned with topics relatively more prevalent among male-authored papers. Over time, this pattern gradually reverses, with the average GTAI moving toward zero and becoming positive in more recent years. This evolution suggests a progressive shift in the thematic composition of published research toward topics that are relatively more prevalent among female authors, consistent with a gradual broadening of research focus within leading general-interest economics journals.




Figure 13 : Evolution of the average Gender–Topic Alignment Index (GTAI) of papers accepted in top-8 economics journals, by year.

Figure 14 reports the average Gender–Topic Alignment Index (GTAI) across leading general-interest economics journals, computed as the mean GTAI of all articles published in each outlet over the sample period. The figure reveals substantial cross-journal heterogeneity in thematic orientation: some outlets exhibit negative average GTAI values—indicating accepted output more closely aligned with topics that are relatively more prevalent among male-authored papers—whereas others display positive averages, reflecting greater alignment with topics relatively more prevalent among female-authored papers. Overall, these differences highlight that general-interest journals differ markedly in the thematic composition of published research along the gender-related topic dimension captured by the GTAI. Consistent with traditional editorial scope, outlets with a stronger emphasis on theoretical

and methodological contributions tend to display negative average GTAI values, whereas those publishing a larger share of applied and policy-oriented research exhibit positive averages.

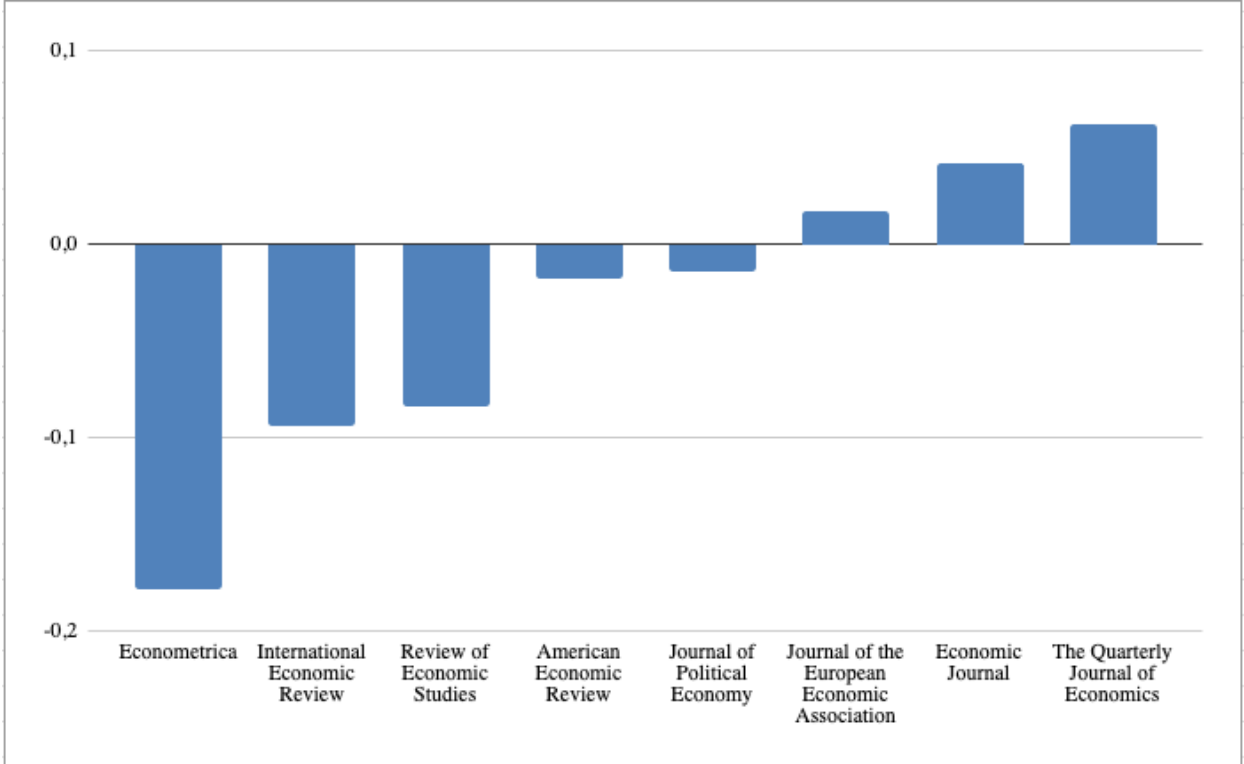


Figure 14 : Average Gender–Topic Alignment Index (GTAI) by journal. Bars report the mean GTAI across articles published in each top-8 journal over the sample period.

Our data base is concentrated in the period 2000 and 2025. However, we have been able to obtain data from articles published in the *Quarterly Journal of Economics* over more than a century. Using this long publication record, we can examine how gender-related topic alignment has evolved within a single leading outlet over a much longer horizon. Figure 15 reports the evolution of the Gender–Topic Alignment Index (GTAI) for articles published in the *Quarterly Journal of Economics* over more than a century.

Latent research topics are re-estimated using all publications available in each historical period. As coverage varies over time—particularly in earlier decades, when some journals did not yet exist—the underlying topic model is necessarily estimated on the set of available publications in each period. The figure is therefore intended as a descriptive illustration of long-run thematic change within a single journal, rather than as a direct quantitative comparison with the journal-level results reported above.

The figure reveals a pronounced long-run shift in thematic orientation. For much of the twentieth century, published research in the QJE was predominantly aligned with topics that are relatively more prevalent among male-authored papers, as reflected in persistently negative GTAI values. Beginning in the late twentieth century, this pattern gradually reverses, with the average GTAI moving toward zero and becoming positive in more recent decades.




Figure 15 : Evolution of the Gender–Topic Alignment Index (GTAI) in QJE publications. The figure reports the annual average GTAI for articles published in the *Quarterly Journal of Economics*. Latent research topics are re-estimated using the set of publications available in each historical period.

Finally, although the GTAI is defined at the document level, it can be naturally aggregated to characterize authors, journals, or editorial boards by averaging across their associated papers. This property will prove useful in the subsequent analysis of citation outcomes and in the discussion of research evaluation and committee composition developed in the following sections.

3.3.1 Gender-Topic Alignment and Citation Outcomes

By providing a continuous measure of an article’s thematic orientation along a gender-related topic dimension, the Gender–Topic Alignment Index (GTAI) offers a flexible tool for empirical analysis. In this section, we use the GTAI to examine gender differences

in citation outcomes in economics. As discussed in the introduction, papers authored by women in top-tier journals tend to receive more citations, but this pattern largely disappears once research area is taken into account. This suggests that topics more prevalent among female researchers may be associated with higher citation counts. The GTAI allows us to test this hypothesis directly by linking citation outcomes to thematic alignment.

In what follows, we relate citation outcomes to thematic alignment as measured by the GTAI. To assess whether topic alignment helps explain citation outcomes, we focus on articles published in Top Five economics journals. Citation data are obtained from RePEc and constructed using the same source and matching procedure as in [Conde-Ruiz et al. \(2025\)](#). As in that analysis, citation information is not available for all published articles, so the estimation sample consists of a well-defined subset of Top Five publications for which citation data can be reliably matched

We estimate the following regression model:

$$C_{p,t} = \beta_0 + \beta_1 \text{GTAI}_{p,t} + \gamma X_{p,t} + \alpha_t + \varepsilon_{p,t},$$

where $C_{p,t}$ denotes the inverse-hyperbolic-sine transformed citation count of paper p in year t , $\text{GTAI}_{p,t}$ is the Gender-Topic Alignment Index, $X_{p,t}$ is a vector of control variables, and α_t denotes year fixed effects.

Table 4 : Citations and GTAI

	Citations (asinh)	
	(1)	(2)
GTAI	73.171*** (3.894)	61.016*** (3.706)
Num.Obs.	7,214	7,214
R ²	0.387	0.413
R ² Adj.	0.385	0.410
FE: year	Yes	Yes
FE: journal		Yes

Notes: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. The dependent variable is the inverse hyperbolic sine of citations. Standard errors are reported in parentheses.

Table 4 reports the estimation results. Column (1) shows a strong positive association between GTAI and citation counts, indicating that papers more closely aligned with topics that are relatively more prevalent among female authors receive more citations. Column (2) adds journal fixed effects, absorbing time-invariant differences in citation practices and editorial scope across Top Five journals. Although the magnitude of the GTAI coefficient declines modestly, it remains precisely estimated, indicating that thematic alignment plays an important role in explaining citation outcomes within journals.

As an additional validity check, the positive association between GTAI and citations is fully consistent with the evidence in Conde-Ruiz et al. (2025), which shows that the apparent gender citation premium in Top Five journals largely disappears once horizontal differences in research topics are accounted for. The GTAI provides a continuous summary of this thematic dimension, and its strong relationship with citations reinforces the interpretation that citation outcomes are closely tied to research content.

Importantly, this result should not be interpreted as evidence that female authors themselves are cited more frequently conditional on content. Rather, it highlights that research topics that are more common among female researchers tend, on average, to attract higher citation counts within Top Five journals. Once thematic alignment is accounted for, gender differences in citations primarily reflect differences in topic orientation rather than differential treatment or recognition.

3.4 Topic versus Counting-Heads Quotas

The main insight of our analysis is that *balanced evaluation committees* may increase welfare. However, as our baseline model shows and the empirical evidence seems to suggest, balanced evaluation committees are unlikely to arise as an equilibrium outcome. This helps explain why, in most countries, regulations governing the composition of evaluation committees have been introduced. The most common regulatory instrument is a headcount quota that ensures a minimum percentage or number of women (or men) on the evaluation committee.

Alternatively, we propose designing committees by taking into account the research profiles of committee members in order to achieve perfectly balanced committees. This topic-based quota approach has two main advantages.

First, it reduces female researchers’ administrative burden. Suppose that the proportion of female full professors is 20% (which is close to the observed data). Requiring committees with an equal number of men and women would then imply that female professors participate in committees four times more often than their male counterparts. In practice, this burden can be reduced by replacing some female members with male researchers who exhibit a high GTAI.

Second, it improves accuracy. By construction, in expected terms, a committee with an equal number of men and women should be balanced, $\frac{1}{2}F_m + \frac{1}{2}F_f$. However, the law of large numbers does not apply to committees with a small number of members, and individual female or male researchers may have topic profiles that differ substantially from the corresponding population averages. Figure 16 illustrates the distribution of research profiles for male and female researchers in terms of their alignment with female-prevalent topics, as measured by the GTAI.




Figure 16 : Distribution of the Gender-Topic Alignment Index (GTAI) at the author level, by gender.

This figure shows that committees with an equal number of male and female members may nevertheless be strongly biased in topic space.

The main advantage of counting-heads quotas lies in their ease of implementation. Reg-

ulations based on minimum participation thresholds for a demographic group are straightforward to monitor and enforce. By contrast, designing a balanced committee based on the research profiles of its members is better understood as a normative principle rather than a directly enforceable regulatory rule. Nevertheless, tools such as the GTAI can be used to oversee whether editorial boards, tenure committees, or similar bodies are thematically balanced.

Proposition 6. *If $\|F_m\| = \|F_f\|$, a perfectly balanced committee $\frac{1}{2}F_m + \frac{1}{2}F_f$ has a GTAI index of 0.*

Proposition 6 supports the idea that a low aggregate GTAI index is a necessary condition for a research-neutral committee. A possible regulatory framework would grant organizations discretion in the design of their committees, provided that a topic-based metric such as the GTAI indicates that the resulting composition is not excessively biased. To reinforce this idea, we next explore the relationship between the thematic alignment of editorial boards—measured through the GTAI—and the characteristics of the papers they publish.

4 Designing Evaluation Committees: A Simulation Exercise

The theoretical analysis developed in the previous sections shows that evaluation committees that are balanced in terms of topic representation maximize aggregate matching and welfare. In practice, however, committees are composed of a small number of discrete evaluators drawn from a heterogeneous population, and institutional constraints limit the set of feasible committee compositions. In this section, we study through a simulation exercise how alternative rules for committee formation affect committee performance.

The simulation is conducted using a subset of our database, namely publication data from top-five economics journals over the period 2000–2025. The starting point of the exercise is that the population of researchers is given by the set of authors publishing in these journals, and that researchers are evaluated by committees drawn from this population according to a productivity-based eligibility rule, possibly combined with additional policy constraints. The goal of the simulation is to analyze how alternative committee selection rules shape both the composition of evaluation committees and their performance.

Each paper p is characterized by an empirical topic distribution vector θ_p , obtained

from the topic model described in Section 3. To avoid ambiguity in gender assignment, we restrict the simulation sample to papers authored exclusively by men and papers authored exclusively by women, excluding mixed-gender author teams. Let \mathcal{P} denote this restricted set of papers, and let \mathcal{P}_m and \mathcal{P}_f denote the subsets of all-male and all-female papers, respectively.

We define the pool of eligible evaluators \mathcal{E} to reflect the gender composition typically observed in editorial boards during our sample period. Rather than selecting the 100 most prolific authors overall—which would yield an extremely male-skewed pool—we construct \mathcal{E} by taking the 80 most prolific male authors and the 20 most prolific female authors, ranking each group by the number of publications in top-five journals. This procedure preserves a meritocratic notion of editorial eligibility while producing a pool with 20% female representation, consistent with the average share of women observed in editorial boards over 2000–2025.

Each evaluator $e \in \mathcal{E}$ is characterized by an aggregate research profile θ_e , constructed as the arithmetic mean of the topic distributions across all their published papers in the sample (as defined in Section 3.4):

$$\theta_e = \frac{1}{|D_e|} \sum_{p \in D_e} \theta_p,$$

where D_e denotes the set of papers authored by evaluator e .

Committees are formed by selecting $K = 4$ evaluators from the eligible pool. Let

$$\mathcal{C} = \{C \subset \mathcal{E} : |C| = 4\}, \quad |\mathcal{C}| = \binom{100}{4} = 3,921,225,$$

denote the space of all feasible committees. Because the committee size is fixed and the evaluator pool is finite, the committee selection problem is combinatorial but fully tractable. We therefore enumerate the entire set of feasible committees and evaluate each of them according to the matching criteria defined below.

Table 5 : Characteristics of the evaluator pool by gender (2000–2025)

	Male evaluators	Female evaluators
<i>Research productivity</i>		
Avg. number of top-5 papers	21.61	15.10
Median number of top-5 papers	19	14
<i>Thematic orientation (GTAI)</i>		
Mean evaluator GTAI	-0.047	0.053
Median evaluator GTAI	-0.072	0.007
SD evaluator GTAI	0.217	0.249

Notes: The table reports descriptive statistics for the pool of eligible evaluators used in the simulation exercise. The evaluator pool consists of the 80 most prolific male authors and the 20 most prolific female authors, ranked by the number of publications in top-five journals over the period 2000–2025. Evaluator GTAI is constructed from individual publication records using the same topic model and index as for published articles.

4.1 Academic promotion and hiring: committee evaluation at the author level

We first consider an academic promotion, hiring, or grant allocation setting in which committees evaluate researchers based on their overall research portfolios. Following Section 3.4, each researcher a is characterized by an aggregate research profile θ_a , defined as the arithmetic mean of topic distributions across all their published papers:

$$\theta_a = \frac{1}{|D_a|} \sum_{p \in D_a} \theta_p,$$

where D_a denotes the set of papers authored by researcher a . Let \mathcal{A} denote the resulting set of researchers, with \mathcal{A}_m and \mathcal{A}_f denoting male and female researchers, respectively.

Committee expertise and thematic orientation are summarized by the average research profile of its members. Consistent with the aggregation logic in Section 3.4, the committee topic profile is defined as

$$c(C) = \frac{1}{|C|} \sum_{e \in C} \theta_e.$$

We approximate evaluation quality using cosine similarity in topic space. The match between researcher a and committee C is defined as

$$m(a, C) = \cos(\theta_a, c(C)),$$

and the associated topic-based distance is

$$D(a, C) = 1 - m(a, C).$$

Lower values of $D(a, C)$ indicate closer proximity between the candidate’s research profile and the committee’s aggregate orientation, and therefore higher expected evaluation accuracy.

Committee formation rules are evaluated by aggregating distances across researchers. For any committee C , we define the average researcher–committee distance as

$$\widehat{D}(C) = \frac{1}{|\mathcal{A}|} \sum_{a \in \mathcal{A}} D(a, C),$$

with analogous definitions for male and female researchers.

Table 6 reports descriptive statistics for the author sample used in the simulation. Female researchers represent a minority of the population, but exhibit a markedly different thematic orientation. In particular, average GTAI is substantially higher for female authors than for male authors, indicating systematic horizontal differences in research focus. At the same time, levels of topic concentration, as measured by the Herfindahl–Hirschman Index (HHI), are very similar across genders. This pattern suggests that gender differences in research profiles primarily reflect differences in where researchers locate in topic space rather than differences in specialization intensity. These features of the data closely mirror the structure emphasized in the theoretical framework and provide a natural setting in which purely utilitarian committee formation may generate systematic disparities in evaluation accuracy across groups.

Table 6 : Descriptive statistics of the author sample

	Number of authors	Mean GTAI	SD GTAI	Mean HHI	SD HHI
All	12093	-0.01	0.21	0.13	0.09
Male	9817	-0.03	0.20	0.13	0.09
Female	2176	0.08	0.24	0.14	0.08

Figure 17 reports the distribution of $\widehat{D}(C)$ across the top 1,000 committees selected under each rule. The unconstrained rule yields the lowest average distances, reflecting maximal aggregate matching efficiency. Imposing gender quotas leads to a substantial deterioration in performance. By contrast, the equal-opportunity rule achieves much better performance than quotas and remains close to the unconstrained benchmark.




Figure 17 : Distribution of average researcher–committee distance across committee formation rules.

Figure 18 focuses on female researchers. Under the unconstrained rule, female candidates experience systematically weaker topic-based matching. Quota-constrained committees perform even worse. In contrast, equal-opportunity rules substantially improve matching for female researchers without incurring the large efficiency losses associated with quotas.




Figure 18 : Distribution of average distance to female researchers across committee formation rules.

Beyond matching efficiency, committee formation rules also affect the thematic orientation of selected committees. Figure 19 shows the distribution of committee-level GTAI across the top-performing committees. Quota constraints induce large shifts in committee thematic orientation, while equal-opportunity rules also reallocate committees in topic space, reflecting their emphasis on research neutrality.




Figure 19 : Distribution of committee-level GTAI across committee formation rules.

Table 9 summarizes the main performance metrics across rules. Consistent with the

graphical evidence, the equal-opportunity rule substantially improves evaluation accuracy for female researchers relative to both the unconstrained and quota-based rules, while preserving aggregate matching efficiency close to the utilitarian benchmark. By contrast, gender quotas generate efficiency losses without delivering comparable gains for female candidates.

Table 7 : Committee performance summary

	General	Quota	Equal opportunity
Mean distance	0.650	0.659	0.652
SD distance	0.002	0.002	0.003
Mean distance to females	0.658	0.659	0.651
SD distance to females	0.007	0.008	0.004
Mean GTAI	-0.030	0.047	0.065
SD GTAI	0.072	0.079	0.076

In Appendix [B](#), we consider an alternative institutional setting that more closely reflects editorial and referee decisions in academic journals. In that environment, the unit of evaluation is the individual manuscript rather than the researcher, and committee assessment follows a handling-editor rule in which each paper is evaluated by the committee member whose research profile is closest in topic space. Despite this different microfoundation, the qualitative patterns remain unchanged. In particular, equal-opportunity rules continue to dominate head-count quotas, improving evaluation accuracy for female-authored research while preserving substantially higher aggregate matching efficiency.

Taken together, the results highlight a general trade-off in committee design between aggregate matching efficiency, group-level evaluation accuracy, and thematic orientation. While reducing disparities in evaluation outcomes may require departures from purely utilitarian committee composition, equal-opportunity rules achieve this objective with substantially smaller efficiency costs than rigid quota-based constraints.

5 Conclusions

This paper investigates the optimal design of evaluation committees in an environment characterized by systematic differences in research orientation across genders. Our theoretical analysis identifies a fundamental "discrimination trap": while gender-topic balanced committees maximize welfare by minimizing informational frictions, they are dynamically unstable. Even under gender-neutral rules, small initial imbalances in committee topic composition are self-reinforcing, generating persistent gaps in participation and success without the need to invoke explicit bias. Our large-scale text analysis confirms the empirical relevance of this mechanism, documenting significant horizontal differentiation between male and female authors in leading economics journals that is often invisible to conventional field classifications.

We characterize optimal committee balance in this multidimensional topic setting and introduce the Gender–Topic Alignment Index (GTAI), an AI-based measurement tool that translates high-dimensional research profiles into a continuous measure of thematic alignment. We show that a low aggregate GTAI is a necessary condition for a gender-neutral research evaluation committee. Our simulation results demonstrate that gender-neutral research committees—formed using our AI-based methods—perform effectively in terms of welfare. In contrast, traditional headcount-based quotas, while easier to implement, often fail to resolve underlying informational frictions and may even disadvantage the groups they intend to support.

From a policy perspective, we propose replacing coarse demographic proxies with information-rich, topic-based quotas leveraging tools like the GTAI. Artificial intelligence makes this transition feasible by rendering research content observable and comparable at scale. Implementing such measures can significantly optimize the design of editorial boards, tenure committees, and grant panels, fostering evaluation institutions that are both more efficient and more gender neutral.

References

- Arrow, Kenneth J.**, “The Theory of Discrimination,” in Orley Ashenfelter and Albert Rees, eds., *Discrimination in Labor Markets*, Princeton University Press, 1973, pp. 3–33.
- Bayer, Amanda and Cecilia Elena Rouse**, “Diversity in the Economics Profession: A New Attack on an Old Problem,” *Journal of Economic Perspectives*, 2016, *30* (4), 221–242.
- Beneito, Pilar, José E. Boscá, Javier Ferri, and Manu García**, “Gender Imbalance across Subfields in Economics: When Does It Start?,” *Journal of Human Capital*, 2021, *15* (3), 469–511.
- Bohren, J. A., Alex Imas, and Michael Rosenberg**, “The Dynamics of Discrimination: Theory and Evidence,” *American Economic Review*, 2019, *109* (10), 3395–3436.
- Bordalo, Pedro, Katherine B. Coffman, Nicola Gennaioli, and Andrei Shleifer**, “Stereotypes,” *The Quarterly Journal of Economics*, 2016, *131* (4), 1753–1794.
- , **Katherine Coffman, Nicola Gennaioli, and Andrei Shleifer**, “Beliefs about Gender,” *American Economic Review*, 2019, *109* (3), 739–773.
- Card, David and Stefano DellaVigna**, “Nine Facts about Top Journals in Economics,” *Journal of Economic Literature*, 2013, *51* (1), 144–161.
- , —, **Patricia Funk, and Nagore Iriberry**, “Are Referees and Editors in Economics Gender Neutral?,” *The Quarterly Journal of Economics*, 2020, *135* (1), 269–327.
- Chevalier, Judy**, “Report: Committee on the Status of Women in the Economics Profession (CSWEP),” *AEA Papers and Proceedings*, 2021, *111*, 742–763. CSWEP annual report (dated Dec. 16, 2020).
- Conde-Ruiz, J. Ignacio, Juan José Ganuza, and Paola Profeta**, “Statistical Discrimination and Committees,” *European Economic Review*, 2022, *141*, 103994.
- Conde-Ruiz, José Ignacio, Juan-José Ganuza, Manu García, and Luis A. Puch**, “Gender Distribution across Topics in the Top Five Economics Journals: A Machine

- Learning Approach,” *SERIEs: Journal of the Spanish Economic Association*, 2022, 13 (1), 269–308.
- , **Miguel Díaz-Salazar, Juan-José Ganuza et al.**, “Citation Gender Gaps in Top Economics Journals,” *SERIEs*, 2025.
- Dolado, Juan, Florentino Felgueroso, and Miguel Almunia**, “Are Men and Women Economists Evenly Distributed across Research Fields? Some New Empirical Evidence,” *SERIEs: Journal of the Spanish Economic Association*, September 2012, 3 (3), 367–393.
- Ductor, Lorenzo and Bauke Visser**, “When a Coauthor Joins an Editorial Board,” *Journal of Economic Behavior & Organization*, 2022, 200, 576–595.
- , **Sergio Galletta, and Daniel Santamaría**, “Homophily and Gendered Citation Patterns in Economics,” *Journal of Economic Behavior & Organization*, 2024. Forthcoming.
- Funk, Patricia, Nagore Iriberry, and Nicole Venus**, “Women in Editorial Boards: An Investigation of Female Representation in Top Economic Journals,” July 2025. Unpublished manuscript (work in progress).
- Heckman, James J. and Sidharth Moktan**, “Publishing and Promotion in Economics: The Tyranny of the Top Five,” *Journal of Economic Literature*, 2020, 58 (2), 419–470.
- Hengel, Erin and Eunyoung Moon**, “Gender and Quality at Top Economics Journals,” Working Paper 202001, University of Liverpool, Department of Economics February 2023. mimeo / updated version of Working Paper 202001.
- Kahneman, Daniel and Amos Tversky**, “Subjective Probability: A Judgment of Representativeness,” *Cognitive Psychology*, 1972, 3 (3), 430–454.
- Koffi, Marlène**, “Gendered Citations at Top Economic Journals,” *AEA Papers and Proceedings*, 2021, 111, 60–64.
- Lundberg, Shelly and Jenna Stearns**, “Women in Economics: Stalled Progress,” *Journal of Economic Perspectives*, 2019, 33 (1), 3–22.
- Merton, Robert K.**, “The Matthew Effect in Science,” *Science*, 1968, 159 (3810), 56–63.

Phelps, Edmund S., “The Statistical Theory of Racism and Sexism,” *American Economic Review*, 1972, 62 (4), 659–661.

Reuben, Ernesto, Paola Sapienza, and Luigi Zingales, “How Stereotypes Impair Women’s Careers in Science,” *Proceedings of the National Academy of Sciences*, 2014, 111, 4403–4408.

Siniscalchi, Marciano and Pietro Veronesi, “Self-Image Bias and Lost Talent,” Working Paper 28308, National Bureau of Economic Research December 2020.

A Appendix A: Proofs

A.1 Explicit Expression for the Dynamic Map

The dynamic law of motion $\alpha_{t+1} = \mathcal{G}(\alpha_t)$ can be written explicitly as:

$$\alpha_{t+1} = \frac{\alpha_t a(\alpha_t)}{\alpha_t a(\alpha_t) + (1 - \alpha_t) b(\alpha_t)}, \quad (5)$$

where

$$a(\alpha) = (1 - \beta) \frac{m_F(\alpha)^2 - \Delta^2}{2 m_F(\alpha)^2} + \beta \frac{m_M(\alpha)^2 - \Delta^2}{2 m_M(\alpha)^2}, \quad b(\alpha) = \beta \frac{m_F(\alpha)^2 - \Delta^2}{2 m_F(\alpha)^2} + (1 - \beta) \frac{m_M(\alpha)^2 - \Delta^2}{2 m_M(\alpha)^2}.$$

The functions $m_M(\alpha)$ and $m_F(\alpha)$ are given by

$$m_M(\alpha) = \beta\alpha + (1 - \beta)(1 - \alpha), \quad m_F(\alpha) = (1 - \beta)\alpha + \beta(1 - \alpha).$$

A.2 Proof of Proposition 1

Proof. A direct computation shows

$$m_M(\alpha) - m_F(\alpha) = (2\beta - 1)(2\alpha - 1),$$

which is positive whenever $\beta > \frac{1}{2}$ and $\alpha > \frac{1}{2}$. Since $\theta_g^*(\alpha) = \Delta/m_g(\alpha)$, it follows that $\theta_M^*(\alpha) < \theta_F^*(\alpha)$, implying $E_M(\alpha) > E_F(\alpha)$. Expected success satisfies

$$S_g(\alpha) = \frac{1}{2} \left(m_g(\alpha) - \frac{\Delta^2}{m_g(\alpha)} \right),$$

which is strictly increasing in $m_g(\alpha)$. Therefore $S_M(\alpha) > S_F(\alpha)$. □ ■

A.3 Proof of Proposition 2

Proof. Since $\bar{\theta}_g(\alpha) = \frac{1 + \theta_g^*(\alpha)}{2}$ and $\theta_F^*(\alpha) > \theta_M^*(\alpha)$ for $\alpha > \frac{1}{2}$, the result follows immediately. ■

A.4 Proof of Proposition 3

Proof. Total entry is

$$E_M(\alpha) + E_F(\alpha) = 2 - \Delta \left(\frac{1}{m_M(\alpha)} + \frac{1}{m_F(\alpha)} \right),$$

and total success is

$$S_M(\alpha) + S_F(\alpha) = \frac{1}{2} - \frac{\Delta^2}{2} \left(\frac{1}{m_M(\alpha)} + \frac{1}{m_F(\alpha)} \right).$$

Then we want to find the $\alpha \in [0, 1]$ that maximizes $f(\alpha) = - \left(\frac{1}{m_M(\alpha)} + \frac{1}{m_F(\alpha)} \right)$

$$f'(\alpha^*) = (2\beta - 1) \left(\frac{m_F(\alpha^*) - m_M(\alpha^*)}{m_M^2(\alpha^*)m_F^2(\alpha^*)} \right) = 0 \rightarrow m_M(\alpha^*) = m_F(\alpha^*) \rightarrow \alpha^* = \frac{1}{2}$$

For the derivative we have used that $m'_M(\alpha) = -m'_F(\alpha) = 2\beta - 1$ and $m_M(\alpha) + m_F(\alpha) = 1$. $\alpha^* = \frac{1}{2}$ is a maximum because $f(\alpha)$ is concave

$$f''(\alpha^*) = -2(2\beta - 1)^2 \left(\frac{1}{m_M^3(\alpha)} + \frac{1}{m_F^3(\alpha)} \right) < 0.$$

■

A.5 Proof of Proposition 4

We first show the fixed points that satisfy $\mathcal{G}(\alpha) = \alpha$, which holds only at $\alpha = 0, \frac{1}{2}, 1$. Then, we will differentiate \mathcal{G} and evaluate at $\alpha = \frac{1}{2}$ which shows that $|\mathcal{G}'(\frac{1}{2})| > 1$, establishing instability, while the endpoints satisfy $|\mathcal{G}'(\alpha)| < 1$.

Fixed points of the dynamic map

Consider the dynamic law of motion $\alpha_{t+1} = \mathcal{G}(\alpha_t)$ given by (5). Define

$$a(\alpha) = (1-\beta) \frac{m_F(\alpha)^2 - \Delta^2}{2m_F(\alpha)^2} + \beta \frac{m_M(\alpha)^2 - \Delta^2}{2m_M(\alpha)^2}, \quad b(\alpha) = \beta \frac{m_F(\alpha)^2 - \Delta^2}{2m_F(\alpha)^2} + (1-\beta) \frac{m_M(\alpha)^2 - \Delta^2}{2m_M(\alpha)^2},$$

so that $\mathcal{G}(\alpha) = \frac{\alpha a(\alpha)}{\alpha a(\alpha) + (1-\alpha)b(\alpha)}$.

A fixed point α satisfies $\mathcal{G}(\alpha) = \alpha$, that is,

$$\alpha = \frac{\alpha a(\alpha)}{\alpha a(\alpha) + (1-\alpha)b(\alpha)}.$$

Multiplying both sides by the denominator yields

$$\alpha(1-\alpha)(b(\alpha) - a(\alpha)) = 0.$$

Hence, any fixed point must satisfy $\alpha = 0$, $\alpha = 1$, or $a(\alpha) = b(\alpha)$.

The first two conditions yield the endpoint fixed points $\alpha \in \{0, 1\}$. For an interior fixed point $\alpha \in (0, 1)$ we require $a(\alpha) = b(\alpha)$. By definition,

$$a(\alpha) - b(\alpha) = (2\beta - 1) \left[\frac{m_M(\alpha)^2 - \Delta^2}{2m_M(\alpha)^2} - \frac{m_F(\alpha)^2 - \Delta^2}{2m_F(\alpha)^2} \right].$$

For $\beta \neq \frac{1}{2}$, $a(\alpha) = b(\alpha)$ is therefore equivalent to

$$\frac{m_M(\alpha)^2 - \Delta^2}{2m_M(\alpha)^2} = \frac{m_F(\alpha)^2 - \Delta^2}{2m_F(\alpha)^2}.$$

Since the function $m \mapsto (m^2 - \Delta^2)/(2m^2)$ is strictly increasing for $m > 0$, this equality holds if and only if $m_M(\alpha) = m_F(\alpha)$. Using

$$m_M(\alpha) = \beta\alpha + (1-\beta)(1-\alpha), \quad m_F(\alpha) = (1-\beta)\alpha + \beta(1-\alpha),$$

we obtain

$$m_M(\alpha) - m_F(\alpha) = (2\beta - 1)(2\alpha - 1),$$

which equals zero if and only if $\alpha = \frac{1}{2}$. We conclude that the fixed points of \mathcal{G} are exactly $\{0, \frac{1}{2}, 1\}$.

Instability of the interior fixed point $\alpha = \frac{1}{2}$

We evaluate the derivative of \mathcal{G} at $\alpha = \frac{1}{2}$. At this point, $m_M(\frac{1}{2}) = m_F(\frac{1}{2}) = \frac{1}{2}$, implying $a(\frac{1}{2}) = b(\frac{1}{2})$ and hence

$$\mathcal{G}'(\frac{1}{2}) = 1 + \frac{1}{4} \frac{a'(\frac{1}{2}) - b'(\frac{1}{2})}{a(\frac{1}{2})}.$$

Using $m'_M(\alpha) = 2\beta - 1$ and $m'_F(\alpha) = -(2\beta - 1)$, together with

$$\frac{d}{d\alpha} \left(\frac{m(\alpha)^2 - \Delta^2}{2m(\alpha)^2} \right) = \frac{\Delta^2}{m(\alpha)^3} m'(\alpha),$$

a direct calculation yields

$$a'(\frac{1}{2}) - b'(\frac{1}{2}) = 16(2\beta - 1)^2 \Delta^2, \quad a(\frac{1}{2}) = \frac{1}{2}(1 - 4\Delta^2).$$

Substituting into the expression above gives

$$\mathcal{G}'(\frac{1}{2}) = 1 + \frac{8(2\beta - 1)^2 \Delta^2}{1 - 4\Delta^2}.$$

For $\Delta < \frac{1}{2}$ and $\beta \neq \frac{1}{2}$, this derivative is strictly greater than one. Therefore, the interior fixed point $\alpha = \frac{1}{2}$ is locally unstable.

Stability of the extreme fixed points $\alpha = 1$ and $\alpha = 0$

We evaluate the derivative of \mathcal{G} at $\alpha = 1$. At this point, $\alpha a(\alpha) + (1 - \alpha)b(\alpha) = a(1)$ and $\alpha a(\alpha) = a(1)$, implying

$$\mathcal{G}'(1) = \frac{b(1)}{a(1)}.$$

a direct calculation yields

$$a(1) = (1-\beta) \frac{m_F(1)^2 - \Delta^2}{2m_F(1)^2} + \beta \frac{m_M(1)^2 - \Delta^2}{2m_M(1)^2}, \quad b(1) = \beta \frac{m_F(1)^2 - \Delta^2}{2m_F(1)^2} + (1-\beta) \frac{m_M(1)^2 - \Delta^2}{2m_M(1)^2}.$$

For $\beta > \frac{1}{2}$, we have $m_M(1) = \beta > 1 - \beta = m_F(1)$, implying

$$\frac{m_M(1)^2 - \Delta^2}{2m_M(1)^2} > \frac{m_F(1)^2 - \Delta^2}{2m_F(1)^2}.$$

Substituting into the expressions above gives

$$a(1) - b(1) = (2\beta - 1) \left[\frac{m_M(1)^2 - \Delta^2}{2m_M(1)^2} - \frac{m_F(1)^2 - \Delta^2}{2m_F(1)^2} \right] > 0.$$

Therefore,

$$0 < \mathcal{G}'(1) = \frac{b(1)}{a(1)} < 1.$$

Hence, the endpoint $\alpha = 1$ is locally stable. By symmetry, the same arguments applies to $\alpha = 0$ ■

A.6 Proof of Proposition 5

Proof. Recall that the utilitarian welfare objective is

$$W(\gamma; c) = \gamma \cos(F_m, c) + (1 - \gamma) \cos(F_f, c),$$

with $c \in \Delta^{K-1}$. Using the definition of cosine similarity, we can write

$$W(\gamma; c) = \frac{1}{\|c\|} \left[\gamma \frac{F_m}{\|F_m\|} + (1 - \gamma) \frac{F_f}{\|F_f\|} \right] \cdot c.$$

Define the normalized topic profiles $u_m = F_m/\|F_m\|$ and $u_f = F_f/\|F_f\|$. Since $\|c\| > 0$ for any $c \in \Delta^{K-1}$, maximizing $W(\gamma; c)$ over the simplex is equivalent to maximizing the linear functional

$$(\gamma u_m + (1 - \gamma) u_f) \cdot c$$

subject to $c \in \Delta^{K-1}$.

The maximizer must therefore lie in the direction of the vector

$$v(\gamma) = \gamma u_m + (1 - \gamma) u_f.$$

Imposing the simplex constraint $\sum_k c_k = 1$ yields the unique solution

$$c^U(t; \gamma) = \frac{\gamma u_m(t) + (1 - \gamma) u_f(t)}{\sum_{j=1}^K [\gamma u_m(j) + (1 - \gamma) u_f(j)]}.$$

This committee maximizes $W(\gamma; c)$ and is unique. ■

A.7 Proof of Proposition 6.

Recall that for any topic-profile vector x , the Gender–Topic Alignment Index (GTAI) is defined as

$$\text{GTAI}(x) = \cos(x, F_f - F_m) = \frac{x \cdot (F_f - F_m)}{\|x\| \|F_f - F_m\|}.$$

Consider a perfectly balanced committee with topic profile

$$c^* = \frac{1}{2}F_m + \frac{1}{2}F_f.$$

Its inner product with the gender-difference direction satisfies

$$\begin{aligned} c^* \cdot (F_f - F_m) &= \frac{1}{2}F_m \cdot (F_f - F_m) + \frac{1}{2}F_f \cdot (F_f - F_m) \\ &= \frac{1}{2}(F_m \cdot F_f - \|F_m\|^2) + \frac{1}{2}(\|F_f\|^2 - F_f \cdot F_m) \\ &= \frac{1}{2}(\|F_f\|^2 - \|F_m\|^2). \end{aligned}$$

If $\|F_m\| = \|F_f\|$, then $c^* \cdot (F_f - F_m) = 0$. Consequently,

$$\text{GTAI}(c^*) = \frac{0}{\|c^*\| \|F_f - F_m\|} = 0.$$

■

B Appendix B: Editorial Committees: Evaluation at the Paper Level

In the main text, committee performance is evaluated using the average committee profile, a structure that naturally captures academic hiring, promotion, or grant allocation decisions, where committee members jointly assess candidates based on their overall research portfolios. In this appendix, we consider an alternative institutional environment that more closely reflects editorial handling and referee assignment in academic journals.

In editorial processes, submitted manuscripts are typically managed by the editor whose research expertise is closest to the paper. To capture this feature, we model evaluation as being carried out by the committee member with the highest topic proximity to each manuscript. This appendix therefore shifts the unit of analysis from researchers to papers and replaces collective deliberation with a handling-editor rule.

Each paper p is characterized by an empirical topic distribution vector θ_p , obtained from the topic model described in Section 3. Committees are formed exactly as in the main text by selecting $K = 4$ evaluators from the eligible pool \mathcal{E} . Each committee member $e \in C$ is characterized by an aggregate research profile θ_e , defined as the average topic distribution across their published papers.

Committee evaluation follows a closest-member (handling-editor) rule. Formally, the match between a paper p and committee C is defined as

$$\hat{m}(p, C) = \max_{e \in C} \cos(\theta_p, \theta_e),$$

and the associated topic-based distance is

$$D(p, C) = 1 - \hat{m}(p, C).$$

Lower values of $D(p, C)$ indicate closer proximity between the manuscript and the most relevant committee member, and therefore higher expected evaluation accuracy.

As in the main analysis, we restrict attention to papers authored exclusively by men or exclusively by women, excluding mixed-gender author teams to avoid ambiguity in gender

assignment. Let \mathcal{P} denote the resulting set of papers, with \mathcal{P}_m and \mathcal{P}_f representing male-authored and female-authored papers, respectively.

Table 8 reports descriptive statistics for the article sample used in this appendix. Female-authored papers represent a small fraction of total output, but differ systematically in their thematic orientation, as captured by the Gender–Topic Alignment Index (GTAI). Differences in GTAI coexist with similar levels of topic concentration, measured by the Herfindahl–Hirschman Index (HHI), indicating horizontal rather than vertical differences in research focus.

Table 8 : Descriptive statistics of the article sample

	Number of articles	Mean GTAI	SD GTAI	Mean HHI	SD HHI
All	9,650	-0.05	0.20	0.12	0.10
Male-authored	8,931	-0.06	0.19	0.16	0.10
Female-authored	719	0.11	0.26	0.16	0.09

Committee formation follows the same three rules analyzed in the main text: a general (unconstrained) rule, a quota-constrained rule imposing gender parity on committee membership, and an equal-opportunity rule that assigns equal aggregate weight to male-authored and female-authored papers in the evaluation objective. As in the main text, Committees are ranked according to the objective implied by each rule, and for each rule we retain the top 1,000 committees.

Figure 20 reports the distribution of $\widehat{D}(C)$ across the top committees. As in the main text, the unconstrained rule yields the lowest average distances, while the quota rule induces a substantial efficiency loss. The equal-opportunity rule performs markedly better than quotas and lies much closer to the unconstrained benchmark.




Figure 20 : Distribution of average paper–committee distance under the closest-member rule.

Figure 21 focuses on female-authored papers. Under the unconstrained rule, papers authored by women experience systematically weaker topic-based matching. Both the quota and equal-opportunity rules improve proximity for this group. Importantly, the equal-opportunity rule achieves these gains with substantially smaller efficiency losses than head-count quotas.




Figure 21 : Distribution of average distance to female-authored papers under the closest-member rule.

Beyond matching efficiency, committee formation rules also affect the thematic orien-

tation of selected committees. Figure 22 reports the distribution of committee-level GTAI across the top committees under each rule. As in the main analysis, quota constraints induce a large shift in committee thematic orientation, while the equal-opportunity rule leads to a more moderate reallocation in topic space.




Figure 22 : Distribution of committee-level GTAI under the closest-member rule.

Finally, Table 9 summarizes the main performance metrics across rules. The table highlights that the equal-opportunity rule substantially improves evaluation accuracy for female-authored papers relative to both the general and quota rules, while preserving aggregate matching efficiency close to the unconstrained benchmark. By contrast, head-count quotas generate large efficiency losses without delivering comparable gains.

Table 9 : Committee performance summary

	General	Quota	Equal opportunity
Mean distance	0.525	0.536	0.533
SD distance	0.002	0.003	0.004
Mean distance (female-authored)	0.559	0.560	0.519
SD distance (female-authored)	0.021	0.021	0.004
Mean GTAI	-0.115	-0.099	0.008
SD GTAI	0.056	0.059	0.040

Taken together, the appendix confirms that the main qualitative findings of the paper are robust to an alternative microfoundation of the evaluation process. Whether evaluation is modeled as a deliberative committee assessing researchers or as a decentralized editorial process assigning manuscripts to the closest expert, equal-opportunity rules dominate rigid quota-based constraints by delivering more equitable outcomes at a substantially lower efficiency cost.