Journal of Economic Behavior and Organization 240 (2025) 107334

Contents lists available at ScienceDirect

Journal of Economic Behavior and Organization

journal homepage: www.elsevier.com/locate/jebo

Dynamic effects of affirmative action: Experimental evidence ☆,☆☆

Gary Charness [a], Ramón Cobo-Reyes [b,*], Lamis Kattan [b], Simone Meraglia [c], Ángela Sánchez [d,e]

[a] University of California Santa Barbara
[b] Georgetown University Qatar
[c] University of Exeter
[d] Comillas Pontifical University – ICADE, Spain
[e] Institute for Research in Technology, ICAI School of Engineering, Comillas Pontifical University

ARTICLE INFO

JEL numbers:
A14
C91
J16
J70

Keywords:
Gender stereotypes
Discrimination
Affirmative action

ABSTRACT

This paper studies the impact of Affirmative Action (AA) policies on workforce sorting in competitive payment schemes, workers' productivity, and firms' performance. Using a laboratory experiment, we analyse an environment where participants repeatedly choose between two firms with tournament-based pay or a firm offering piece-rate compensation. Experimental treatments vary based on whether participants can switch firms each period and whether one of the tournament-based firms implements AA policies. Our findings reveal that AA significantly boosts female participation in competitive settings without reducing male engagement or productivity, effectively closing the gender gap in tournament selection. However, AA firms face mild talent retention challenges. Importantly, AA induces strong sorting by gender and productivity, with highly productive women preferring AA firms and highly productive men opting for competitive environments in which AA is not in place. These findings suggest that while AA policies successfully attract more women to competitive workplaces, they may also contribute to gender-based clustering – where men and women self-select into different firms based on the presence of AA – while maintaining overall organizational efficiency.

1. Introduction

Despite the significant progress made in recent decades, substantial gender disparities persist across labor markets in both the private and public sectors. These disparities range from wage gaps to limited career advancement opportunities for women (e.g., Blau et al., 2013; Blau and Kahn, 2000; Weichselbaumer and Winter-Ebmer, 2007). To address these persistent imbalances, Affirmative Action programs (AA, hereafter) have emerged as a policy tool aimed at promoting equal opportunities and rectifying potential gender-based discrimination, ultimately striving for equitable representation of women in upper-tier positions. Advocates for the introduction of gender quotas argue that AA strategies offer a rapid and tangible means of enhancing female representation within

☆ We are thankful to the Georgetown University in Qatar FRG for financially supporting this project. This project has received IRB approval from Georgetown University in Qatar, application no STUDY00005803.☆☆ In Memoriam: This paper is dedicated to the memory of Gary Charness, who passed away as this work was nearing completion. Gary was a brilliant scholar and a remarkable coauthor, mentor, and friend whose influence on the field is beyond measure. As we share this work, we do so with deep gratitude for the lasting impact he made on us all.

* Corresponding author.
E-mail address: rc1557@georgetown.edu (R. Cobo-Reyes).

https://doi.org/10.1016/j.jebo.2025.107334
Received 22 March 2025; Received in revised form 22 September 2025; Accepted 7 November 2025
Available online 21 November 2025
0167-2681/

specific roles or domains. However, these programs do not come without controversy, and the use of AA has been largely debated.

Experimental evidence examines Affirmative Action (AA) policies from multiple perspectives. One key dimension is whether AA policies discourage the non-favored group, potentially leading to a decline in performance. Research generally suggests that AA does not significantly reduce performance for the non-favored group (Fallucchi and Quercia, 2018; Balafoutas and Sutter, 2012; Calsamiglia et al., 2013). However, these results are not always conclusive and may depend on individual abilities (Bracha et al., 2019; Cotton et al., 2022).

Beyond their direct impact on performance, AA policies also influence workforce composition by acting as a sorting mechanism, attracting or deterring workers from joining a firm. This is particularly relevant because workforce composition plays a crucial role in shaping organizational outcomes. Extensive evidence shows that workers often select firms based on the incentives offered (Dohmen and Falk, 2011; Eriksson and Villeval, 2008; Larkin and Leider, 2012; Flory et al., 2015; Barbulescu and Bidwell, 2013). Thus, AA policies may impact company performance not only through direct incentive effects but also by shaping the firm's talent pool, influencing long-term organizational success.

Evidence on the effect of Affirmative Action (AA) as a sorting mechanism is relatively scarce, with much of the existing research concentrating on the immediate impacts of AA policies in attracting or deterring certain types of candidates in competitive payment systems (see Ibanez and Riener, 2018; Fallucchi and Quercia, 2018; Niederle et al., 2013; Balafoutas and Sutter, 2012; Sutter et al., 2016). To the best of our knowledge, all the previous studies limit participants to a binary choice between AA and piece rates, which may not fully capture the preferences of individuals who want a competitive environment but face limited options. Moreover, none of the previous papers account for the impact of AA policies on talent retention within firms in repeated settings when learning plays a role**.** Consequently, there is a need for further empirical research that explores the sorting effects of AA within more diverse labor market settings, offering a broader range of alternatives for a more comprehensive understanding of its impact on workforce dynamics.

We take this question as a starting point for our experimental design. Upon arriving to the lab, subjects are allocated to a group and assigned the role as a firm or a worker.[1] In our *Affirmative Action No Moving Treatment*, workers are asked to perform a real effort task consisting on adding five two-digit numbers (Niederle and Vesterlund, 2007) for three minutes. There are three types of firms, and workers' payment depends on the firm they work for. Workers working for Firm A are paid 10 tokens per correct answer. Workers working for Firm B are paid on a tournament basis: the winner gets 15 tokens per correct answer; the loser earns 5 tokens per answer. Workers in Firm C, are paid the same as in Firm B, with the only difference that two extra correct answers will be added to the total performance of female participants. The *Affirmative Action Moving Treatment* mirrors *Affirmative Action No Moving Treatment*, with key distinction that, in this case, workers can decide, at the end of each period, whether they want to continue working for the same firm or if they prefer to move to another firm. Finally, the *No-Affirmative Action Moving Treatment* is the same as *Affirmative Action Moving Treatment* with the only difference that no AA policy is implemented in any firm.

Our results suggest that AA policies significantly influence workforce sorting by gender: Among women who opt for competitive environments, the vast majority prefer firms that implement AA policies, whereas men predominantly select firms without such initiatives. Furthermore, women are more likely to choose a competitive environment over a piece-rate scheme when AA policies are in place. This mechanism effectively eliminates the gender gap in tournament selection. We observe no sorting effect of AA when analysing worker selection based on productivity levels: While high-productivity workers generally prefer tournament-based compensation schemes over piece-rate systems, their selection into tournaments remains unaffected by the presence of AA. This indicates that while AA policies alter gender-based sorting, they do not distort the self-selection of workers based on ability.

In terms of productivity, our results show that payment systems do not affect workers' performance when groups are fixed, but allowing self-selection leads to significantly higher productivity in firms using tournaments, irrespective of AA implementation. This suggests that self-selection leads to efficiency gains, as higher-performing individuals naturally gravitate toward competitive environments. Interestingly, AA policies do not reduce productivity within tournament firms, implying that their implementation does not come at an efficiency cost.

We also find that, despite the productivity advantage of tournaments, firm profitability remains balanced across payment schemes. This is due to the larger number of workers attracted to piece-rate firms, which compensates for the lower individual productivity in those firms. However, in the long run, firms implementing AA experience a slight decline in the retention of high-productivity workers, suggesting that while AA effectively eliminates the gender gap in tournament selection, it may create mild disincentives for top performers over time.

The contribution of this paper is threefold. First, it provides a deeper analysis of the sorting effect of Affirmative Action (AA) policies by introducing a market where workers have multiple alternatives. Unlike previous studies, where the only alternative to AA was a piece rate without competition, this paper proposes a more complex and realistic market scenario, allowing workers to choose from a broader set of options. This richer market context is expected to yield more robust results. Second, the paper not only examines whether AA policies are effective in attracting workers but also investigates their impact in a repeated setting, including retention rates. Third, it sheds light on the role of AA policies on productivity, offering new insights into how these policies influence not only the composition of the workforce, but also its performance.

The rest of the paper is organized as follows. Section 2 outlines the experimental design and introduces a simple theoretical model to derive our main hypotheses. Section 3 presents our main empirical results, while Section 4 concludes.

[1] Measuring the sorting effects of Affirmative Action (AA) in the field poses significant challenges, primarily due to the complexities of implementing the necessary conditions and controlling for individual worker productivity. In contrast, laboratory-based economic experiments provide a more precise and insightful perspective on the impact of AA programs.

2. Experimental design and procedures

2.1. Experimental design

The experimental design consists of three treatments: the *Affirmative Action No Moving Treatment*, the *Affirmative Action Moving Treatment*, and the *No Affirmative Action Moving Treatment.* All three treatments involve a real-effort task. The task consists of adding five two-digit numbers for three minutes.[2]

Affirmative Action No Moving Treatment (AANMT). Upon arrival to the lab, subjects are randomly allocated to a group and randomly assigned the role as a firm or a worker. Each group is composed of three firms and twelve workers (six male workers and six female workers). Workers will be randomly paired with a firm and will perform a real effort task. We make use of a standard task in the literature: adding five two-digit numbers (Niederle and Vesterlund, 2007) for three minutes. Workers' payment depends on the firm they work for. Workers working for Firm *A* are paid on a piece rate basis, depending on their performance. In particular, they are paid 10 tokens per correct answer. Workers working for Firm *B* are paid on a tournament basis: workers are randomly matched within the firm, the worker who performs better earns 15 tokens per correct answer; the worker who performs worse earns 5 tokens. Workers in Firm *C* are paid according to the same rules as in Firm *B* with the only difference that two extra correct answers are added to the total performance of female participants when they do the adding task.[3] No extra correct answer would be added for male participants. This rule is common knowledge to all participants.[4]

All firms are paid depending on workers' performance. The firm makes 4.5 tokens per correct answer provided by all the workers in that firm. The firm has to pay 8 tokens per worker.[5]

The task is repeated for 15 periods of three minutes each. Each firm is matched with the same four workers (two male and two female workers) for the duration of the experiment. At the end of each period, workers receive information about: i) their own payoffs, ii) their performance, iii) whether they are first (winner) or second (loser) in the tournament (only for workers in Firms *B* and *C*), and iv) partial information about the performance of the competitor of the worker in the tournament. This information will be provided to the workers in the form of intervals like [0–10] [11–20].[6] All workers are also informed about the payment system of all firms.

Affirmative Action Moving Treatment (AAMT): This treatment is exactly the same as *AANMT* with the only difference that, at the end of each period, after the workers have completed the real effort task, they participate in a *moving stage*. This moving stage works as follows. After the real effort task is finished and workers have received the corresponding information (same as in *AANMT*), they decide whether they want to stay working for the same firm or if they prefer to move to another firm.

Firms have no choice over the workers that work for them. They cannot accept or reject workers that apply to work under their payment system. That is, workers can move freely within the different payment systems available and there is no unemployment.

No-Affirmative Action Moving Treatment (NAAMT): This treatment is exactly the same as *AAMT* with the only difference that no AA policy is implemented in any firm: Workers working for Firm *A* are paid on a piece rate basis, and workers working for Firms *B* and *C* are paid on a tournament basis without any extra correct answers added for female workers.

2.2. Procedures

The experiment was conducted at the University of Castellon with 750 participants, who were recruited using the online recruitment system ORSEE (Greiner, 2004). The data collection took place from March to October 2024. No subject participated in more than one treatment. We conducted 15 sessions of 45 subjects each. On average, each person received 14 Euros for (up to) a 90-minute session. Subjects were paid for three randomly selected periods. The conversion rate is 10 tokens equals 0.50 Euros.

Before participating in the first period of the experiment, regardless of the treatment, participants played under a piece rate system in order to collect individual skills of the workers.

2.3. Theoretical model

In this section, we introduce a simple model that captures the main features of our experimental design, as outlined in Section 2.1. This theoretical framework aims to inform our hypotheses regarding the impact of Affirmative Action on workers' behavior in terms of productivity and location choices. For simplicity, we focus squarely on workers' decisions and disregard the role of firms.

The game consists of 12 workers, each indexed by $i \in \{1,...,12\}$. Each worker is characterized by *a)* their group membership $g_i \in \{0,$

[2] See Appendix A for the instructions for the *Affirmative Action Moving Treatment.*

[3] Note that the two extra correct answers are only relevant for determining the tournament winner. Once the piece rate has been set, the payment for that period will depend solely on the workers' performance, without considering any additional correct answers.

[4] In the first period of the experiment all participants performed the task and were paid on a standard piece rate basis. This was done to collect information regarding subject's baseline performance.

[5] Note that, the 8 tokens paid by the firm per worker will not be received by the workers. This will be a cost that firms will pay but will not be part of the salary of the worker. The goal of this payment is to make the profit of the firm depend not only on the number of workers, but also on workers' productivity.

[6] Workers in Firm *A* will receive partial information about the relative performance of one of the members in the group who will be randomly chosen. This is to keep the information set the same across firms.

1}, *b)* their productivity $\theta_i \in \{\underline{\theta}, \overline{\theta}\}$, with $\overline{\theta} > \underline{\theta}$. In what follows, we interpret the group membership g_i as representing worker i's gender, where $g_i = 0$ corresponds to males, and $g_i = 1$ corresponds to females. We assume an equal number of males and females in the game. Additionally, for both genders, half of the workers are characterized by low productivity ($\underline{\theta}$), and the other half by high productivity ($\overline{\theta}$). In summary, the game includes 3 low-productivity males, 3 high-productivity males, 3 low-productivity females, and 3 high-productivity females. We assume that θ_i is common knowledge, for $i \in \{1, ..., 12\}$.

Each worker i chooses one of three firms in which to locate, with $f_i \in \{A, B, C\}$ denoting the firm f chosen by worker i, with $f \in \{A, B, C\}$.[7] Let Ω denote the set of all possible allocations of workers to firms. An allocation $\omega \in \Omega$ specifies the location decisions of all 12 workers:

$$\omega = (f_1, f_2, ..., f_{12}).$$

Given $\omega \in \Omega$, we obtain the firm size vector $\boldsymbol{n}(\omega) = (n_A(\omega), n_B(\omega), n_C(\omega))$, where $n_f(\omega)$ denotes the number of workers in firm f. Moreover, ω determines the composition of each firm, $N_f(\omega)$, defined as the set of worker types (g_i, θ_i) located in firm f.[8]

Each worker i also chooses an effort $e_i \in [0, E(\theta)]$ in solving a number of (identical) tasks for a firm f, where $E(\underline{\theta}) = \underline{e} < E(\overline{\theta}) = \overline{e}$. Let $c(e_i) = \frac{e_i^2}{2}$ denote the cost of effort. Moreover, we denote by $y_i(e_i, \theta_i) = e_i(\theta_i)$ the number of tasks worker i solves correctly. Worker i receives a compensation x_i^f per-task solved, where the piece-rate varies depending on the firm f for which the task is performed. In particular, we have:

Firm A: All workers located in firm A receive a piece-rate $x_i^A = x$.

Firm B: Let N_B denote the set of workers who choose firm B, and let $n_B = |N_B|$ denote its size. Worker i located in firm B is matched uniformly at random with one of the other $n_B - 1$ workers in the same firm. The two matched workers compete in a tournament: The worker who solves the highest (respectively, lowest) number of tasks receive a piece-rate $x_i^B = \overline{x}$ (respectively, $x_i^B = \underline{x}$), with $\overline{x} > \underline{x}$ and $\frac{\overline{x}+\underline{x}}{2} = x$. Formally, for any pair (i,j), for $i,j \in N_B$ and $i \neq j$, we define the function:

$$\rho_{ij}^B(e_i, \theta_i, e_j, \theta_j) = \begin{cases} 0 \text{ if } y_i < y_j \\ \frac{1}{2} \text{ if } y_i = y_j \\ 1 \text{ if } y_i > y_j \end{cases} \tag{1}$$

According to (1), the function $\rho_{ij}^B(\cdot)$ takes value 0 (respectively, 1) if worker i solves a strictly lower (respectively, greater) number of tasks than worker j, whereas the function takes value $\frac{1}{2}$ if the two workers solve the same number of tasks. Given an allocation ω, since opponents are drawn uniformly at random from $N_B(\omega) \setminus \{i\}$, the ex-ante probability that worker i wins the tournament in firm B is:

$$\overline{\rho}_i^B(e_i, \theta_i | \omega) = \begin{cases} \frac{1}{n_B(\omega) - 1} \sum_{j \in N_{B(\omega)} \setminus \{i\}} \rho_{ij}^B(e_i, \theta_i, e_j, \theta_j) \text{ if } n_B(\omega) \geq 2 \\ 1 \text{ if } n_B(\omega) = 1, \end{cases} \tag{2}$$

where, from (2), worker i receives $x_i^B = \overline{x}$ with probability 1 when he/she is the only worker located in firm B.[9]

Firm C: Let N_C denote the set of workers who choose firm C, and let $n_C = |N_C|$ denote its size. Like in firm B, worker i located in firm C enters a tournament with one of the other workers located in the same firm, where the opponent is randomly selected. The rule according to which a worker receives the high ($x_i^C = \overline{x}$) or low ($x_i^C = \underline{x}$) piece-rate differs from that in place in firm B – as shown in (1) – in that a bonus $b_i(g_i) \geq 0$ is added to the number of tasks solved (y_i) to determine the winner of the tournament. Importantly, the bonus varies with worker i's gender g_i. Specifically:

$$b_i(g_i) = \begin{cases} 0 \text{ if } g_i = 0 \\ b \text{ if } g_i = 1 \end{cases}, \tag{3}$$

[7] Figure 2. (a) Cell voltage of the redox compartment as a function of current density (1–5 mA cm^{-2}) and Fe-DTPA concentration (10, 25, 50, 100 mM), which was measured using a two-compartment flow cell divided by a cation exchange membrane at a flow rate of 10 mL min^{-1}. Error bars indicate the standard deviation from duplicate experiments as well as positive/negative current densities. (b) Voltage distribution between the redox and feed compartments of the RFD system for freshwater production (C_d = 0.5 g L^{-1}). The overall cell voltage for freshwater production was interpolated from the data presented in Figure 1(d) between 4 and 4.5 mA cm^{-2}. The corresponding current densities were used to calculate the cell voltage of the redox compartments from the data presented in Figure 2(a) between 4 and 5 mA cm^{-2} by linear interpolation. The difference between these two yielded the voltage contributed from the feed compartments.

[8] More formally, we have $n_f(\omega) = |\{i : f_i = f\}|$ and $N_f(\omega) = \{(g_i, \theta_i) : f_i = f\}$.

[9] Figure 3. (a) Cyclic voltammograms of Fe-DTPA as a function of concentration (25, 50, and 100 mM) at a scan rate of 10 mV s^{-1}. (b) Cyclic voltammograms of Fe^{3+}-DTPA (10 mM ferric DTPA in 0.03 M Na_2CO_3 and 1.4 M NaCl) and Fe^{3+}-CN (10 mM ferricyanide in 1.5 M NaCl) at a scan rate of 100 mV s^{-1}.

where $b \geq 0$.[10] As a consequence, for any pair (i,j), for $i,j \in N_C$ and $i \neq j$, we define the function:

$$\rho^c_{ij}\left(e_i, \theta_i, g_i, e_j, \theta_j, g_j\right) = \begin{cases} 0 \text{ if } y_i + b_i < y_j + b_j \\ \frac{1}{2} \text{ if } y_i + b_i = y_j + b_j. \\ 1 \text{ if } y_i + b_i > y_j + b_j \end{cases} \tag{4}$$

Given an allocation ω, since opponents are drawn uniformly at random from $N_C(\omega) \setminus \{i\}$, the ex-ante probability that worker i wins the tournament in firm C is:

$$\overline{\rho}^C_i(e_i, \theta_i, g_i|\omega) = \begin{cases} \frac{1}{n_C(\omega) - 1} \sum_{j \in N_C(\omega)\setminus\{i\}} \rho^C_{ij}\left(e_i, \theta_i, g_i, e_j, \theta_j, g_j\right) \text{ if } n_C(\omega) \geq 2 \\ 1 \text{ if } n_C(\omega) = 1. \end{cases} \tag{5}$$

Workers are risk-neutral. The payoff of worker i located in firm A is deterministic and is given by:

$$U^A_i = xy_i(e_i, \theta_i) - c(e_i) = x\theta_i e_i - \frac{e_i^2}{2}. \tag{6}$$

Conditional on allocation $\omega \in \Omega$, the expected payoff of worker i located in firm B is:

$$U^B_i(e_i, \theta_i|\omega) = \left[\overline{\rho}^B_i(e_i, \theta_i|\omega)\overline{x} + \left(1 - \overline{\rho}^B_i(e_i, \theta_i|\omega)\right)\underline{x}\right]y_i - \frac{e_i^2}{2}, \tag{7}$$

where $\overline{\rho}^B_i(\cdot|\omega)$ is given by (2).

Likewise, the expected payoff of worker i located in firm C conditional on ω is:

$$U^C_i(e_i, \theta_i, g_i|\omega) = \left[\overline{\rho}^C_i(e_i, \theta_i, g_i|\omega)\overline{x} + \left(1 - \overline{\rho}^C_i(e_i, \theta_i, g_i|\omega)\right)\underline{x}\right]y_i - \frac{e_i^2}{2}, \tag{8}$$

where $\overline{\rho}^C_i(\cdot|\omega)$ is given by (5).

Workers interact for one period. The timing of the game is as follows:

1. Workers simultaneously choose the firm where they locate f_i, for $i \in \{1, 2, 3, \ldots, 12\}$;
2. Workers simultaneously choose their effort levels e_i, for $i \in \{1, 2, 3, \ldots, 12\}$;
3. Workers receive their payoffs in accordance with the payment rule implemented by the firm they locate in.

Importantly, mirroring our experimental design, workers do not observe the location choices of others before selecting their effort. Our solution concept is Nash Equilibrium.

In line with the experimental design presented in Section 2.1, we analyse three versions of this game. In the first (benchmark) version, referred to as the *No-Moving Game*, workers are exogenously assigned to a firm (f) – either *A, B*, or *C* – and cannot move from their assigned firm. This version parallels the experimental treatment *AANMT*, with *Stage 0* excluded from the game and Affirmative Action in place in firm C ($b > 0$). Specifically, in the *No-Moving Game* we assume that each firm consists of four workers: one high-productivity male, one low-productivity male, one high-productivity female, and one low-productivity female. In the second version, referred to as the *Moving Game* (paralleling the experimental treatment *AAMT*), workers can choose which firm to join, meaning that Stage 0 is included in this version of the game. Affirmative Action is also in place in firm C ($b > 0$) in this version of the game. Finally, in the third version, referred to as the *No-Bonus Moving Game* (paralleling the experimental treatment *NAAMT*), workers can choose which firm to locate in However, unlike the previous version, we set $b = 0$, meaning no Affirmative Action is implemented in firm C.

Before turning to the equilibrium analysis, we clarify how our stylized theoretical model relates to the laboratory experiment described in Section 2.1. The model is framed as a one-period game with common knowledge of workers' productivity, whereas the laboratory setting involves repeated interactions under incomplete information about others' types. This difference is important: in the laboratory, workers can use the repeated nature of the game to update beliefs about others' productivity, influencing their decisions to move across firms. While the model abstracts from this dynamic learning process, it is designed to capture the core structure and strategic trade-offs of the experimental environment – namely, how workers select firms based on their own productivity, their information about others, and the incentive schemes offered by different firms.

2.3.1. *No-moving game (Benchmark)*

We begin by analysing the benchmark version of the game where workers are exogenously assigned to one of the three firms, with b

[10] Setting $b > 0$ corresponds to the experimental treatment *AAMT*, where Affirmative Action is in place in firm C. Conversely, setting $b = 0$ corresponds to the experimental treatment *NAAMT*, where no Affirmative Action is implemented in firm C.

> 0. To determine the equilibrium, we make the following assumption:

$\boldsymbol{A}1: \ \underline{x} > \overline{e}.$

Assumption **A1** simplifies the analysis by stating that the marginal benefit of effort for a worker exceeds its marginal cost when they anticipate being compensated with a piece-rate of at least $\underline{x}$. In practice, this means that workers always exert maximum effort, regardless of *i)* which firm they are located in and *ii)* the number or type of co-workers within the same firm. This assumption is plausible in our experimental laboratory setting, where participants primarily focus on the task at hand.

Using **A1**, the following remark states workers' equilibrium behavior in the *No-Moving Game*.[11]

Remark 1 (Benchmark: No-Moving Game)

*Under assumption **A1**, in the No-Moving Game, all workers exert maximum effort* $E(\theta_i)$*, for* $\theta \in \{\underline{\theta}, \overline{\theta}\}$ *and* $i \in \{1, 2, 3, \ldots, 12\}$.

From Remark 1, it follows that, in all firms, high-productivity workers solve more tasks correctly than low-productivity workers. Moreover, there should be no performance differences among high-productivity workers across different firms, nor among low-productivity workers across different firms.

2.3.2. Moving game

We now examine the version of the game where, in Stage 0, workers can select the firm they wish to join. As in the *No-Moving Game*, we set $b > 0$. Under assumption **A1**, as in the benchmark case (Remark 1), workers continue to exert maximum effort, irrespective of *(i)* the firm they choose to join, and *(ii)* the composition of workers – both in terms of number and type – within the same firm.

The following proposition states the implications of this result and firms' characteristics on workers' optimal location choices. We focus on pure-strategy equilibria.

Proposition 1 (Moving Game)

*Under assumption **A1**, in the Moving Game, at equilibrium:*

i. *All workers exert maximum effort* $E(\theta_i)$*, for* $\theta \in \{\underline{\theta}, \overline{\theta}\}$ *and* $i \in \{1, 2, 3, \ldots, 12\}$*;*
ii. *At least one high-productivity worker – either male or female – joins firm B;*
iii. *At least one high-productivity female worker joins firm C;*
iv. *All low-productivity workers – both males and females – join firm A;*
v. *No male worker – whether high or low productivity – joins firm C.*

Proof. See Appendix B.■

When workers have the option to select which firm to join, the differing incentives offered by the three firms result in two distinct types of sorting: by productivity and by gender. Specifically, the structure of piece-rate compensation in firms *B* and *C*, which depends on the outcome of a tournament, leads to *all* low-productivity ($\underline{\theta}$) workers to join firm *A*. In firm *A*, the piece-rate is uniform and does not depend on the performance of other workers, making it an attractive option for low-productivity workers. Moreover, the Affirmative Action policy implemented by firm *C* implies that no male (high-productivity) worker opts to join this firm. As a consequence, firm *C* is *exclusively* populated by female workers.[12]

2.3.3. No-bonus moving game

We conclude by examining the version of the game where, in Stage 0, workers can select the firm they wish to join. Unlike the *Moving Game* analysed in Section 3.3.2, we set $b = 0$.

The following proposition states workers' optimal effort and location choices. We focus on pure-strategy equilibria.

Proposition 2 (No-Bonus Moving Game)

*Under assumption **A1**, in the No-Bonus Moving Game, at equilibrium:*

i. *All workers exert maximum effort* $E(\theta_i)$*, for* $\theta \in \{\underline{\theta}, \overline{\theta}\}$ *and* $i \in \{1, 2, 3, \ldots, 12\}$*;*
ii. *At least one high-productivity worker – either male or female – joins firm B;*
iii. *At least one high-productivity worker – either male or female – joins firm C;*
iv. *All low-productivity workers – both males and females – join firm A.*

Proof. See Appendix B.■

When workers have the option to select which firm to join, but no Affirmative Action policy is in place, sorting occurs only by

[11] We relegate a formal proof for Remark 1 to Appendix B.

[12] More formally, Proposition 1 allows for the existence of multiple equilibria. Specifically, there are equilibria in which one or more high-productivity male workers choose to join firm *B*. Similarly, there are equilibria in which one or more high-productivity female workers decide to join either firm *B* or firm *C*.

productivity. As in Proposition 1, the structure of piece-rate compensation in firms *B* and *C* implies that *all* low-productivity ($\underline{\theta}$) workers join firm *A*. However, unlike Proposition 1, the absence of an Affirmative Action policy in firm *C* implies that both male and female (high-productivity) worker can now opt to join either firm *B* or firm *C*.[13]

2.3.4. Testable hypotheses

Building on the results established in Remark 1, and Propositions 1 and 2, we formulate the following hypothesis:

Hypothesis 1 (sorting by gender). We observe:

a. In the *AAMT* treatment, no male worker joins Firm *C*.
b. Compared to *AAMT*, a greater proportion of male workers join Firm *C* in the *NAAMT* treatment.

Hypothesis 2 (sorting by productivity). In both the *AAMT and NAAMT* treatments, low-productivity workers join Firm *A*, avoiding Firms *B* and *C*.

Hypothesis 3. There is no significant difference in performance among high-productivity workers, either across treatments (*AANMT, AAMT*, and *NAAMT*) or across firms within the same treatment. Similarly, the performance of low-productivity workers remains consistent across treatments and firms.

The next section tests these hypotheses by presenting our experimental findings.

3. Results

This section is structured as follows. First, we analyze group formation dynamics and the main determinants influencing workers' firm selection. Next, we compare worker productivity and firm profits across the various treatments and payment systems.

3.1. Group formation

To explore whether AA policies influence the sorting of the workforce across different firms, we examine mobility choices within *AAMT* and *NAAMT*. Our first analysis focuses on gender. Then, we focus on productivity as a source of heterogeneity.

3.1.1. Group formation by gender

We start by examining whether AA policies influence the sorting of workers based on their gender. Fig. 1 displays the average number of male and female workers across different firm types, for *AAMT* in Fig. 1(a) and *NAAMT* in Fig. 1(b).

Fig. 1(a) shows that both genders initially prefer tournaments, but gradually shift to piece rate systems over time. Specifically, the average share of females (66 %) and males (64 %) opting for the tournament system in the first five periods is nearly double compared to the number of females (34 %) and males (36 %) selecting piece-rate.[14] However, in the final five periods, this pattern reverses: 44 % of males and 37 % of females select the tournament system, while 56 % of males and 62 % of females choose piece-rate.[15]

Furthermore, Fig. 1(a) highlights the significant impact of AA policies on workers' sorting. Among those who choose tournament-style environments, AA policies tend to attract female workers while deterring male workers. Specifically, on average, 87 % of female workers who chose the tournament system selected the option with AA, compared to just 13 % who opted for the tournament without AA.[16] Conversely, for male workers, this pattern is reversed, with an average of 11 % choosing the tournament with AA and 89 % preferring the tournament without AA policies.[17] This result supports Hypothesis 1a.

Fig. 1(b) illustrates that in the absence of AA policies, the market dynamics change drastically. While the proportion of males choosing a tournament payment system in the first five periods under *NAAMT* (64 %) closely resembles that under *AAMT* (64 %), the percentage of females opting for a competitive environment drops under NAAMT (46 %) compared to *AAMT* (66 %).[18] Consequently, the overall proportion of males and females workers selecting the tournament scheme is quite similar in *AAMT* ($z = -0.430$, $p = 0.676$, two-tailed Mann-Whitney test), but the gender differences in competitive choices become statistically significant in *NAAMT* ($z = 2.373$, $p = 0.015$, two-tailed Mann-Whitney test).

[13] Like Proposition 1, Proposition 2 also allows for the existence of multiple equilibria. Specifically, there are equilibria in which one or more high-productivity male workers choose to join firm *B* or *C*. Similarly, there are equilibria in which one or more high-productivity female workers join either firm *B* or firm *C*.

[14] Differences between the average number of females and males choosing tournaments and those choosing piece rate is statistically significant ($z = 2.420$, $p = 0.014$, and $z = 2.267$, $p = 0.022$ two-tailed Mann Whitney test for females and males, respectively)

[15] Differences between the average number of females and males choosing tournaments and those choosing piece rate is statistically significant ($z = 4.132$, $p < 0.001$, and $z = 3.279$, $p = 0.001$ two-tailed Mann Whitney test for females and males, respectively)

[16] Differences are statistically significant ($z = 3.580$, $p < 0.001$, two-tailed Mann Whitney test).

[17] Differences are statistically significant ($z = 3.587$, $p < 0.001$, two-tailed Mann Whitney test).

[18] Differences are not statistically significant for males ($z = 0.402$, $p = 0.708$, two-tailed Mann Whitney test), but they are significant for females ($z = 3.199$, $p < 0.001$, two-tailed Mann Whitney test).

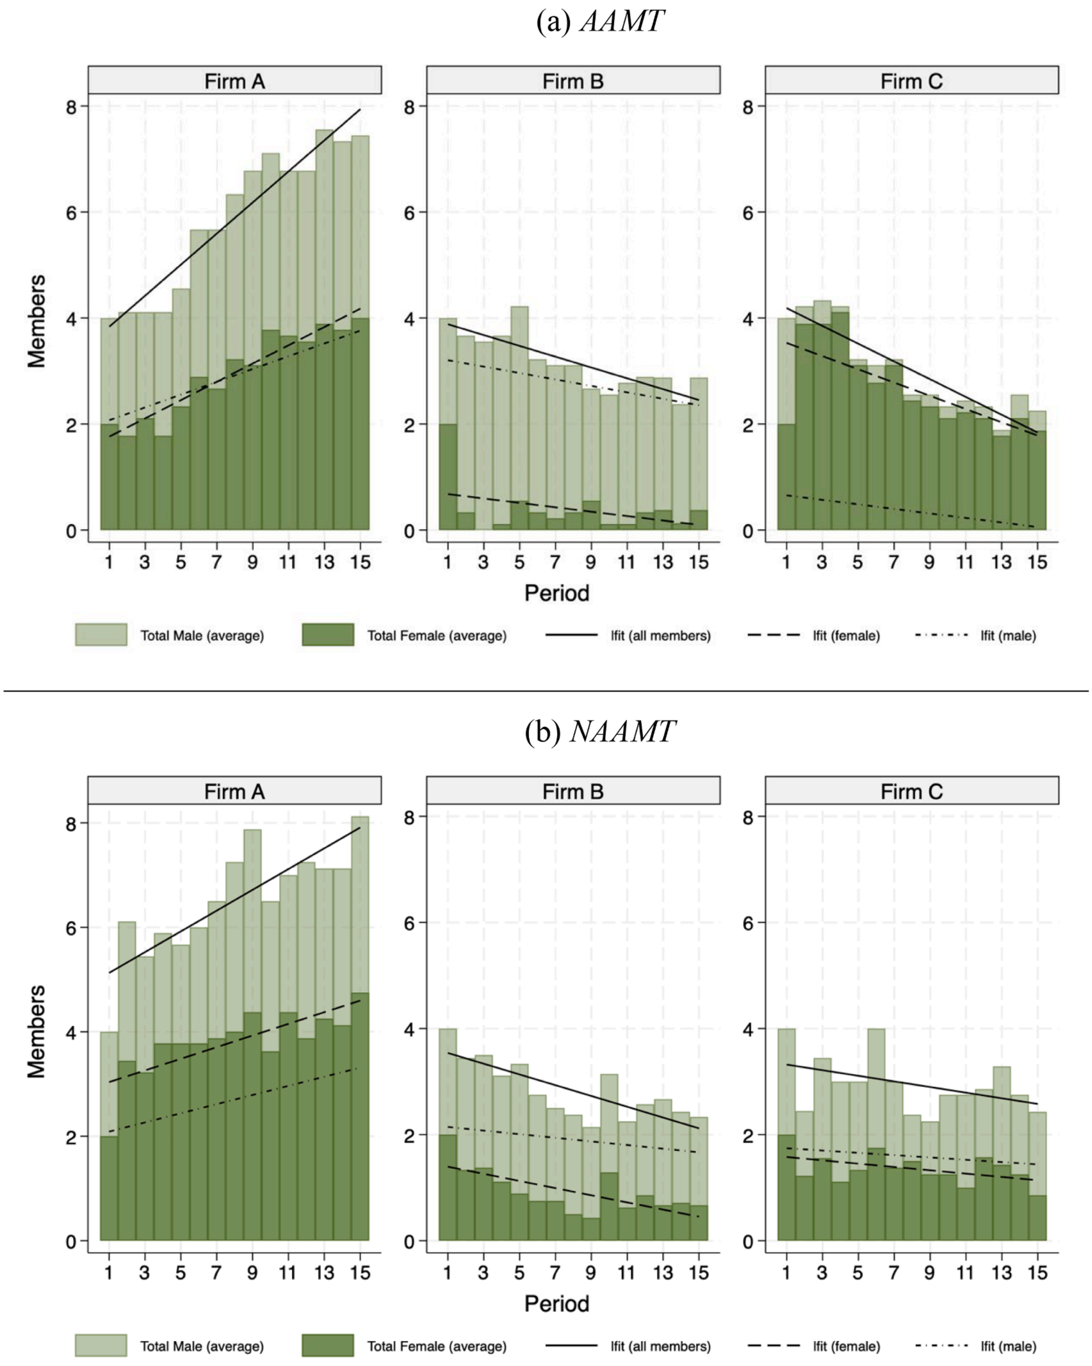

Fig. 1. Group formation in *AAMT* and *NAAMT*, by gender.

Regarding sorting, Fig. 1(b) confirms that, as expected, there is no gender-based sorting in the absence of AA policies. Male and female workers who choose the tournament system are evenly distributed between Firm B and Firm C.[19] On average, 42 % of female participants who chose a tournament system are located in Firm B, while 58 % are in Firm C.[20] Male workers exhibit a similar distribution, with 54 % selecting tournaments in Firm B and 46 % in Firm C.[21] As expected, the proportion of male workers opting for Firm C is significantly larger in *NAAMT* compared to *AAMT* ($z = 3.497$, $p < 0.001$, two-tailed Mann-Whitney test). This finding provides support for Hypothesis 1b

Result 1. a: *Affirmative Action policies have a strong impact on how workers sort themselves into firms (by gender). Among those who opt for a competitive environment, females tend to choose companies that implement AA, while males are more likely to select companies without such systems.*

Result 1. b: *Affirmative Action policies are an effective way of increasing willingness to compete for female workers, and closing the gender gap in tournament selection entirely.*

3.1.2. Group formation by productivity

We present an analysis of group formation, but now focused on productivity rather than gender. Fig. 2 mirrors Fig. 1 but considers workers' initial productivity levels instead. Specifically, workers are categorized into two groups: high, and low productivity levels. High-productivity workers are those whose performance in the first period (where all workers participate in a baseline treatment) is above the 75th percentile, while low-productivity workers are defined as those whose performance in the first period falls below the 25th percentile.[22]

Fig. 2(a) and (b) show that the piece-rate system consistently attracts a higher proportion of low-productivity workers in both *AAMT* and *NAAMT*. In *AAMT*, 61 % of low-productivity workers opt for the piece-rate system, compared to just 32 % of high-productivity workers. A similar pattern is observed in *NAAMT*, where a larger proportion of low-productivity (60 %) workers choose the piece-rate system, as opposed to 37 % of high-productivity workers.[23] Overall, these results support Hypothesis 2.

To lend further support to Hypothesis 2, we analyse the choices of low-productivity workers regarding which firm to join across periods. The average number of low-productivity workers shows a statistically significant decline for those companies implementing tournaments in both *AAMT* and *NAAMT*.[24] By contrast, when tournament is not in place (firm *A*), there is a significant positive trend in the average number of low-productivity workers ($t = 7.17$, $p < 0.001$ for *AAMT*, $t = 3.03$, $p = 0.003$ for *NAAMT*).

Fig. 2(a) further suggests that the sorting observed at the gender level in *AAMT* due to AA does not occur at the productivity level. Specifically, tournaments without AA policies in place (Firm *B*) attract an average of 36 % of high-productivity workers, a proportion nearly identical to that observed in tournaments with AA policies (Firm *C*, with 32 %). This difference is not statistically significant ($z = 0.442$, $p = 0.681$, two-tailed Mann-Whitney test).[25] Fig. 2(a) also demonstrates that, when firms with and without AA policies coexist (*AAMT*), tournament systems without AA policies are more effective at retaining talent. The average number of high-productivity workers shows no statistically significant decline ($t = 0.72$, $p = 0.472$) for those companies not implementing AA policies (Firms B). However, when AA policies are in place (Firms C), there is a significant negative trend in the average number of high-productivity workers ($t = 2.44$, $p = 0.016$).[26]

Result 2. a: *Tournament systems are effective in attracting high-productive workers compared to piece rates. By contrast, the majority of low-productive workers chooses piece-rate.*

Result 2. b: *In AAMT, the same proportion of high-productive workers choose tournament systems with and without AA. However, tournament systems without AA policies do better in retaining high-productive workers.*

3.1.2. Econometric analysis

To better understand the mechanisms underlying the group-level patterns reported above, this section presents individual-level analyses of group choice. Table 1 reports the marginal effect from a Probit model fitted to our data in the moving treatments

[19] In *NAAMT*, we observe some switching between the two firms implementing a tournament system: About 17% of subjects in Firm B move to Firm C in a given period, while about 12% of subjects in Firm C move to Firm B. These flows are balanced, as the difference between B to C and C to B transitions is not statistically significant (McNemar's $\chi^2(1)=1.75$, $p = 0.186$).

[20] Figure 5. (a) Fe^{2+}-CN and (b) Fe^{2+}-DTPA optimized structures using DFT. Dimensions indicate the nitrogen-nitrogen distance plus the van der Waals radius of nitrogen (in (a)) or twice the farthest iron-carbon distance plus the van der Waals diameter of oxygen (in (b)).

[21] Differences are not statistically significant ($z = 1.015$, $p = 0.3401$, two-tailed Mann Whitney test).

[22] Appendix C replicates the analysis and split the sample based on the initial median-level productivity. Our results remain consistent.

[23] Differences are statistically significant ($z = 2.608$, $p = 0.007$, two-tailed Mann Whitney test)

[24] $t = 3.30$, $p = 0.001$ for Firm *B* in *AAMT*; $t = 4.67$, $p < 0.001$ for Firm C in *AAMT*; $t = 2.45$, $p = 0.016$ for Firm *B* in *NAAMT*; $t = 1.45$, $p = 0.150$ for Firm C in *NAAMT*.

[25] Similar results are observed in Figure 2(b) for *NAAMT*, where, as expected, the tournament systems in Firm *B* and *C* attract comparable proportions of high-productivity workers (32% and 31%, respectively). Differences are not statistically significant ($z = 0.133$, $p = 0.931$, two-tailed Mann Whitney test)

[26] One potential explanation for this finding is that high-productive females are more prone to leave a competitive setting than high productive males. We explore this possibility in Appendix D.

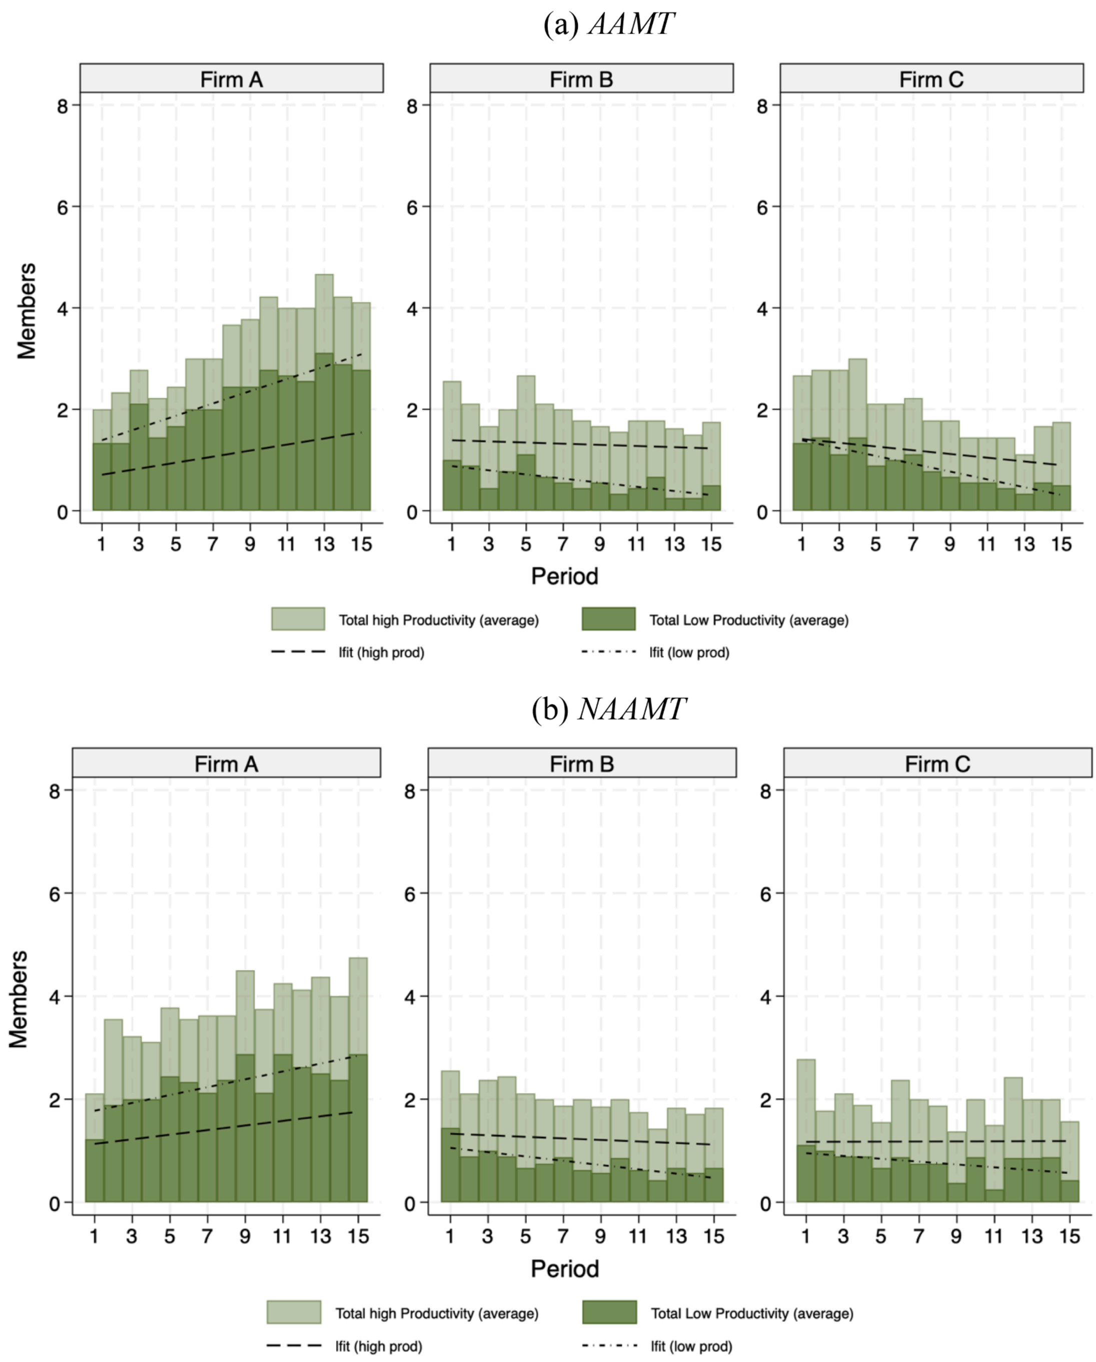

Fig. 2. Group formation in *AAMT* and *NAAMT*, by productivity.

Table 1
Probability of choosing a tournament.

	AAMT			NAAMT		
	(1)	(2)	(3)	(4)	(5)	(6)
Female	0.033	0.008	−0.027	−0.140**	−0.163***	−0.144**
	(0.058)	(0.071)	(0.054)	(0.057)	(0.056)	(0.064)
Initial Productivity	0.151***			0.078**		
	(0.038)			(0.036)		
Lagged Performance		0.003			0.021**	
		(0.011)			(0.010)	
Lagged Win			0.107*			0.072**
			(0.058)			(0.031)
Age	✓	✓	✓	✓	✓	✓
Degree	✓	✓	✓	✓	✓	✓
N	1605	1498	776	1440	1344	623

Notes: The coefficients represent average marginal effects from a probit model estimated using the full sample of participants in a panel format over 15 periods. The outcome variable is a dummy equal to 1 if player *i* chooses a tournament-type game (i.e., Firm B or C) and 0 otherwise (Firm A). Initial productivity is a categorical variable based on percentiles: it takes the value 0 for low productivity (below the 25th percentile), 1 for medium productivity (between the 25th and 75th percentiles), and 2 for high productivity (above the 75th percentile). Lagged performance is the number of correct answers in period $t - 1$. Lagged win is a dummy variable equal to 1 if player *i* had a higher score than their opponent in period $t - 1$, conditional on having participated in a tournament in that period. The first period is omitted in columns (2), (3), (5), and (6). Standard errors in parentheses clustered at the group level: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

(*AAMT* and *NAAMT)*. The outcome variable "*probability of choosing a tournament*" is a dummy variable equals to 1 if player *i* chooses a tournament-type system (i.e., Firm B or C), and 0 otherwise (Firm A). As explanatory variables, we include: "*Female*", a bivariate variable that takes value 1 if the worker is a female, and 0 otherwise; "*Initial productivity*" is a categorical variable based on percentiles, where 0 corresponds to low productivity (below the 25th percentile), 1 to medium productivity (between the 25th and 75th percentiles), and 2 to high productivity (above the 75th percentile). Additionally, we include "*Lagged performance*", a proxy for the worker's performance, which captures the number of correct answers provided by worker *i* in period t-1; and "*Lagged win*", a dummy variable equals to 1 if player *i* has a higher score that their opponent in period $t - 1$, and 0 otherwise.[27]

Results in Table 1 indicate that females are less likely to participate in a tournament when no AA policies are in place. However, this pattern changes when AA policies are introduced, effectively eliminating the gender gap in tournament participation.[28] Additionally high-productivity workers are more inclined to select a tournament in both AAMT and NAAMT treatments. These findings are in line with the patterns observed in Fig. 2a and b Furthermore, workers who had a successful experience in the tournament during the previous period are more likely to choose a tournament again. However, this effect is only marginally significant in societies where AA policies are implemented.

Next, we explore the selection into AA systems. Table 2 reports the marginal effect from a Probit model fitted to our data in *AAMT*. The analysis is conducted only for those workers choosing a tournament system (i.e., Firms B or C). The dependent variable "*probability of choosing AA*" is a dummy variable equals to 1 if player i chooses a tournament-type system with AA (i.e., Firm C), and 0 for a tournament with no AA policies in place (Firm B). We use the same set of explanatory variables as in Table 1.

Results in Table 2 further support the conclusions drawn from Figs. 1(a) and 2(a), indicating that women are significantly more likely to select into Affirmative Action systems. In contrast, initial productivity levels do not appear to predict the selection into these systems.

3.2. Productivity

We start by comparing workers' productivity across the three payment systems in the three treatments. Specifically, Table 3 shows the average number of correct sums for workers working for Firms *A, B*, and *C*.

The first row of Table 3 (*All Workers*) reports the average productivity of workers across different treatments. The difference in productivity is not significant between workers in *AANMT* and *NAAMT* ($z = 0.221$, $p = 0.8633$, two-tailed Mann-Whitney test) and between workers in *AAMT* and *NAAMT* ($z = 0.927$, $p = 0.3865$, two-tailed Mann-Whitney test). When comparing *AAMT* and *AANMT* the difference in productivity is only marginally significant ($z = 1.796$, $p = 0.0789$, two-tailed Mann-Whitney test). Overall, these

[27] Note that "*Lagged win*" is only generated for those workers who were in a tournament in period *t* – *1*.

[28] Table E1 in Appendix E shows that the lower willingness of females to participate in a tournament when no Affirmative Action policies are in place is driven by high-productivity females. In the NAAMT treatment, high-productivity females are less likely to choose a tournament compared to their male counterparts, whereas in the AAMT treatment, their behavior is indistinguishable from that of males.

Table 2
Probability of choosing affirmative action (AA).

	AAMT		
	(1)	(2)	(3)
Female	0.487***	0.427***	0.426***
	(0.008)	(0.024)	(0.025)
Initial Productivity	0.030		
	(0.028)		
Lagged Performance		−0.001	
		(0.004)	
Lagged Win			−0.012
			(0.038)
Age	✓	✓	✓
Degree	✓	✓	✓
N	816	745	595

Notes: The coefficients represent average marginal effects from a probit model estimated using the full sample of participants in a panel format over 15 periods. The outcome variable is a dummy equal to 1 if player *i* chooses a tournament type with Affirmative Action (Firm C), and 0 for a tournament with no Affirmative Action (Firm B). Initial productivity is a categorical variable based on percentiles: it takes the value 0 for low productivity (below the 25th percentile), 1 for medium productivity (between the 25th and 75th percentiles), and 2 for high productivity (above the 75th percentile). Lagged performance is the number of correct answers in period $t-1$. Lagged win is a dummy variable equal to 1 if player *i* had a higher score than their opponent in period $t-1$, conditional on having participated in a tournament in that period. The first period is omitted in columns (2) and (3). Standard errors in parentheses clustered at the group level: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3
Workers' productivity.

	AANMT	*AAMT*	*NAAMT*
All Workers	5.3 (0.563)	6.182 (1.272)	5.500 (1.936)
Workers in Firm A	5.32 (0.783)	4.529 (0.760)	4.118 (0.832)
Workers in Firm B	5.086 (1.345)	7.698 (2.107)	6.259 (1.647)
Workers in Firm C	5.493 (1.337)	6.319 (1.634)	6.124 (2.355)
Workers in Firm A (high)	7.536(1.945)	6.594(1.910)	6.515(1.414)
Workers in Firm B (high)	7.787(3.090)	9.514(2.589)	8.602(2.009)
Workers in Firm C (high)	7.158(1.879)	8.106(2.640)	7.798(2.751)
Workers in Firm A (low)	3.417(1.330)	3.377(1.094)	3.054(1.219)
Workers in Firm B (low)	3.665(1.731)	4.048(2.612)	3.333(1.592)
Workers in Firm C (low)	4.361(2.830)	2.510(1.190)	3.683(1.873)

findings support the model's assumption (**A1**) according to which workers exert maximum effort under any of the piece-rate payment schemes implemented in the different firms.

Table 3 shows that the difference in productivity is not statistically significant among low-productivity workers when comparing Firm *A* to Firm *B*, Firm *A* to Firm *C*, and Firm *B* to Firm *C*. This result holds for all treatments.[29] Similarly, virtually no significant differences are observed among high-productivity workers when comparing Firm *A* to Firm *B*, Firm *A* to Firm *C*, and Firm *B* to Firm *C*.[30] Overall, these results are in line with Hypothesis 3.

We now analyze the effect of the mobility process in *AAMT* and *NAAMT* on average productivity. In both cases, workers are allowed to self-select their preferred payment scheme. As illustrated in Figures 2(a) and 2(b), high-productivity workers are more likely to choose tournament-based systems, while low-productivity workers predominantly opt for piece rate systems. We explore whether this self-selection behavior leads to differences in productivity across payment schemes.

The results in Table 3 for *AANMT* indicate that the three different payment systems—piece rate, tournament without AA, and tournament with AA—yield very similar productivity outcomes when participants cannot move across firms. The average productivity

[29] Firm *A* vs. Firm *B*: $z = 0.530$, $p = 0.6262$; $z = 0.398$, $p = 0.7299$; $z = 0.927$, $p = 0.3865$, two-tailed Mann Whitney test for *AANMT, AAMT* and *NAAMT*, respectively. Firm *A* vs. Firm *C*: $z = 0.476$, $p = 0.6806$; $z = 1.280$, $p = 0.2224$; $z = 0.574$, $p = 0.6048$, two-tailed Mann Whitney test for *AANMT, AAMT* and *NAAMT*, respectively. Firm *B* vs. Firm *C*: $z = 0.309$, $p = 0.7962$; $z = 1.193$, $p = 0.2482$; $z = 0.044$, $p = 1.000$, two-tailed Mann Whitney test for *AANMT, AAMT* and *NAAMT*, respectively.

[30] Firm *A* vs. Firm *B*: $z = 0.210$, $p = 0.8785$; $z = 2.213$, $p = 0.0274$; $z = 1.987$, $p = 0.0503$, two-tailed Mann Whitney test for *AANMT, AAMT* and *NAAMT*, respectively. Firm *A* vs. Firm *C*: $z = 0.526$, $p = 0.6289$; $z = 0.866$, $p = 0.4234$; $z = 0.839$, $p = 0.4363$, two-tailed Mann Whitney test for *AANMT, AAMT* and *NAAMT*, respectively. Firm *B* vs. Firm *C*: $z = 0.105$, $p = 0.9591$; $z = 1.280$, $p = 0.2224$; $z = 0.618$, $p = 0.5613$, two-tailed Mann Whitney test for *AANMT, AAMT* and *NAAMT*, respectively.

under each system is 5.320, 5.086, and 5.493, respectively, with no statistically significant differences between them.[31] This suggests that the type of payment system does not substantially affect worker productivity under *AANMT*. The findings remain consistent when disaggregated by gender, showing that the policy has a uniform effect across both male and female workers.[32]

When the workers have the opportunity to choose the payment system they prefer, results change dramatically. In *AAMT*, the average productivities under piece rate, tournament without AA, and tournament with AA are 4.529, 7.698, and 6.319, respectively. Differences are statistically significant when we compare piece rate to tournament without AA ($z = 3.135$, $p < 0.001$, two-tailed Mann-Whitney test) piece rate to tournament with AA ($z = 2.428$, $p = 0.014$, two-tailed Mann-Whitney test), but not when we compare tournament with to that without AA ($z = 1.280$, $p = 0.222$, two-tailed Mann-Whitney test). Thus, it appears that both tournament schemes result in higher average performance compared to the piece-rate system, in line with participants sorting by productivity (Result 2a). Moreover, the presence of AA policies does not impact worker performance under a tournament system. When disaggregating by gender, results remain similar.[33]

A similar result to that in *AAMT* is found in *NAAMT*: The average productivities under piece rate, tournament without AA, and tournament with AA (firms A, *B,* and *C*) are 4.118, 6.259, and 6.124, respectively. Differences are statistically significant when we compare piece rate to tournament ($z = 2.958$, $p = 0.002$ and $z = 1.987$, $p = 0.050$, two-tailed Mann-Whitney test for the comparison with Firm *B* and Firm *C*, respectively), but not when we compare the two tournaments ($z = 0.050$, $p = 0.931$, two-tailed Mann-Whitney test for comparing Firm *B* to Firm *C*). So, it seems that self-selection dynamic, enabled by the flexibility in *AAMT* and *NAAMT*, results in observed differences in productivity across payment systems.[34]

Result 3. *While the payment system does not impact workers' productivity when groups are fixed, the sorting process leads to tournaments (with and without AA policies) significantly improving average worker performance in treatments where participants are allowed to choose their preferred payment scheme.*

3.2.2. *Econometric analysis*

To formally examine the potential effects of the different payment systems on workers' productivity, Table 4 presents the results from an OLS model where the dependent variable is the number of correct answers provided by individual i at time t, for treatments AANMT (column 1), AAMT (column 2), and NAAMT (column 3). The explanatory variables include: "Female", a binary variable equal to 1 if the worker is female and 0 otherwise; "Initial productivity", a categorical variable based on percentiles, where 0 represents low productivity (below the 25th percentile), 1 indicates medium productivity (between the 25th and 75th percentiles), and 2 corresponds to high productivity (above the 75th percentile). Additionally, we include two dummy variables, Firm B and Firm C, with Firm A as the omitted reference category.

Notes: The coefficients represent OLS estimates using the full sample of participants in a panel format over 15 periods. The outcome variable is the number of correct answers of individual *i* in time t. Initial productivity is a categorical variable based on percentiles: it takes the value 0 for low productivity (below the 25th percentile), 1 for medium productivity (between the 25th and 75th percentiles), and 2 for high productivity (above the 75th percentile). Coefficients on Firm B and Firm C are reported relative to Firm A (omitted). Standard errors in parentheses clustered at the group level: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4 shows that the gender of the workers does not affect their performance, while the initial level of productivity has a positive and significant effect on outcomes. This relationship holds across all treatments. Additionally, the results in column (1) indicate that when workers cannot move between firms, the payment system has no impact on productivity. However, this changes substantially when group formation is endogenous and workers are allowed to choose their preferred payment system. In this context, as shown in columns (2) and (3), firms using a tournament-based system outperform those offering piece rates, regardless of whether the tournament includes Affirmative Action policies.

3.3. *Profits*

Although highly correlated with workers' performance, assessing firms' profits is not straightforward, as they also depend on both the size and composition of the firm. To further investigate how the payment system and AA policies affect firm profits, Fig. 3 illustrates the dynamic of firms' average profit across different payment system and treatment condition.

Fig. 3(a) shows that firms' profits in *AANMT* are very similar across the three different payment systems ($z = 0.756$, $p = 0.481$, $z = 0.151$, $p = 0.912$, and $z = 0.530$, $p = 0.617$; two-tailed Mann-Whitney test when we compare Firm *A* vs. Firm *B*, Firm *A* vs. Firm *C*, and

[31] Two-tailed Mann-Whitney test: $z = 0.756$, $p = 0.481$; $z = 0.151$, $p = 0.912$; $z = 0.530$, $p = 0.617$; for the comparison Firm *A* vs Firm *B*; Firm *A* vs Firm *C*; and Firm *B* vs Firm *C*, respectively.

[32] Two-tailed Mann-Whitney test for the females (males) sample: $z = 0.397$, $p = 0.730$ ($z = 0.530$, $p = 0.621$); $z = 0.397$, $p = 0.730$ ($z = 0.442$, $p = 0.683$) ; $z = 0.927$, $p = 0.387$ ($z = 0.751$, $p = 0.489$); for the comparison Firm *A* vs Firm *B*; Firm *A* vs Firm *C*; and Firm *B* vs Firm *C*, respectively.

[33] Two-tailed Mann-Whitney test for the females' sample: $z = 3.135$, $p < 0.001$, for the comparison Firm *A* vs Firm *C* in *AAMT*. Two-tailed Mann-Whitney test for the males' sample: $z = 2.517$, $p = 0.011$, for the comparison Firm *A* vs Firm *B* in *AAMT*. Note that, due to sorting by gender, we do not have enough observations (i.e., females in Firm *B* and males in Firm *C*) to provide a statistical analysis of the remaining comparisons.

[34] These findings hold when disaggregating by gender under NAAMT: Two-tailed Mann-Whitney test for the females (males) sample: $z = 1.193$, $p = 0.2480$ ($z = 2.428$, $p = 0.0142$) ; $z = 1.810$, $p = 0.0770$ ($z = 1.457$, $p = 0.1615$) ; $z = 0.751$, $p = 0.4748$ ($z = 0.927$, $p = 0.3865$); for the comparison Firm A vs Firm B; firm A vs Firm C; and firm B vs Firm C, respectively.

Table 4
Level of productivity.

	AANMT (1)	AAMT (2)	NAAMT (3)
Female	0.478	−0.427	−0.055
	(0.400)	(0.540)	(0.401)
Initial Productivity	1.837***	2.131***	1.867***
	(0.320)	(0.278)	(0.292)
Firm B	−0.188	1.316**	1.458**
	(0.506)	(0.458)	(0.447)
Firm C	0.168	0.864***	0.965**
	(0.463)	(0.257)	(0.285)
Age	✓	✓	✓
Degree	✓	✓	✓
χ^2: $\beta_{\text{Firm B}} = \beta_{\text{FirmC}}$	0.31	1.27	2.89
Prob > χ^2	0.589	0.292	0.133
N	1800	1605	1440

Firm *B* vs. Firm *C*, respectively). This result is not surprising since productivity levels are comparable across the three payment systems in *AANMT* (as shown in Table 1), and group sizes remain constant throughout the experiment.

For *AAMT*, the Mann-Whitney test results show no significant differences in profits across the three payment systems: Firm *A* vs. Firm *B* ($z = 0.132$, $p = 0.931$), Firm *A* vs. Firm *C* ($z = 0.751$, $p = 0.489$), and Firm *B* vs. Firm *C* ($z = 1.104$, $p = 0.297$). This indicates that, although productivity levels are higher under tournament systems compared to the piece-rate system, the larger number of workers attracted to the piece-rate system balances profits across firms.

However, when focusing on the last five periods, a different pattern emerges. While differences in profits remain statistically insignificant between Firm *A* and Firm *B* ($z = 0.751$, $p = 0.4894$, two-tailed Mann-Whitney test) and between Firm *B* and Firm *C* ($z = 1.280$, $p = 0.2224$, two-tailed Mann-Whitney test), significant differences are observed when comparing Firm *A* and Firm *C* ($z = 1.987$, $p = 0.0503$, two-tailed Mann-Whitney test). This suggests that, over time, the larger number of workers attracted to Firm *A* compared to Firm *C* not only offsets the higher productivity observed under the tournament system but eventually reverses its effect on profits.

Similarly, for *NAAMT*, Mann-Whitney test results indicate no significant differences in profits across the three payment systems: Firm *A* vs. Firm *B* ($z = 0.662$, $p = 0.5457$), Firm *A* vs. Firm *C* ($z = 1.369$, $p = 0.1903$), and Firm *B* vs. Firm *C* ($z = 0.839$, $p = 0.4363$). For the last five periods, the Mann-Whitney tests gives: $z = 1.575$, $p = 0.1304$; $z = 2.310$, $p = 0.0207$; and $z = 1.050$, $p = 0.3282$; for the comparison between Firm *A* and Firm *B*, Firm *A* and Firm *C*, and Firm *B* and Firm *C*, respectively.

Result 4. *While workers' productivity is higher under tournament systems when workers are free to move between groups, the larger number of workers attracted by piece rate systems balances out the profits across firms.*

4. Conclusions

This paper introduces a novel experimental design that provides a comprehensive approach to analysing the effects of Affirmative Action (AA) on workforce sorting, worker's productivity, and firm's performance. By allowing workers to repeatedly choose among multiple firms with different payment structures, including competitive environments with and without AA, we capture a broader and more nuanced view of how AA influences labor market dynamics.

Our findings reveal several key insights. First, AA policies significantly increase female participation in competitive workplaces without reducing male engagement or productivity, effectively closing the gender gap in tournament selection. This rise in female participation aligns with previous results (see Niederle et al., 2013; Ibáñez and Riener, 2018, or Sutter et al., 2016). However, unlike earlier studies, we do not find evidence that men reduce their participation under AA. Second, we find clear evidence of sorting by productivity: low-productivity workers predominantly choose firms with piece-rate payment schemes, whereas high-productivity workers are more likely to opt for competitive compensation structures. This result aligns with prior evidence (Falucchi and Quercia, 2018; Niederle et al., 2013; Sutter et al., 2016). Third, we identify a strong sorting effect by gender: highly productive women are drawn to AA firms, while highly productive men choose to join firms without AA. This gender-based sorting extends the existing literature by highlighting how AA policies influence not only the overall workforce composition but also its internal distribution across firms. Finally, our results indicate that AA policies do not reduce productivity within tournament firms, suggesting that their implementation carries no efficiency cost, consistent with prior findings (Sutter et al., 2016; Balafoutas and Sutter, 2012).

To avoid overgeneralizing these findings, it is important to highlight certain characteristics of our study that may limit their external validity. First, the observed dynamics stem from a repeated setting run in the laboratory, which cannot fully capture the long-term effects of AA policies in real-world organizational contexts. Second, our design does not incorporate the potential for backlash against women benefiting from AA – a phenomenon shown to negatively influence women's willingness to enter environments where

(a) *AANMT*

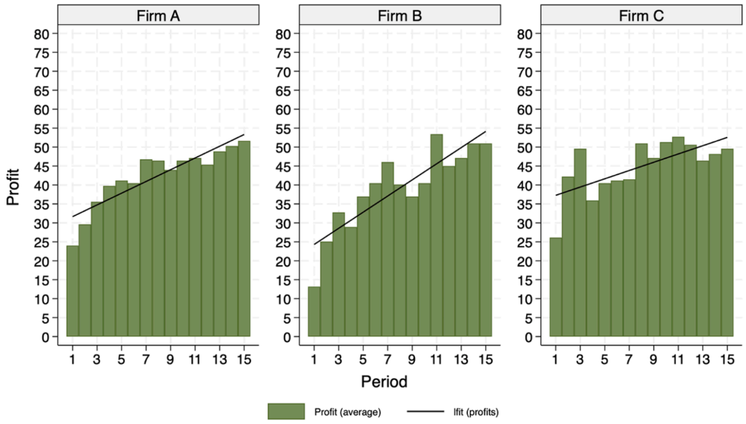

(b) *AAMT*

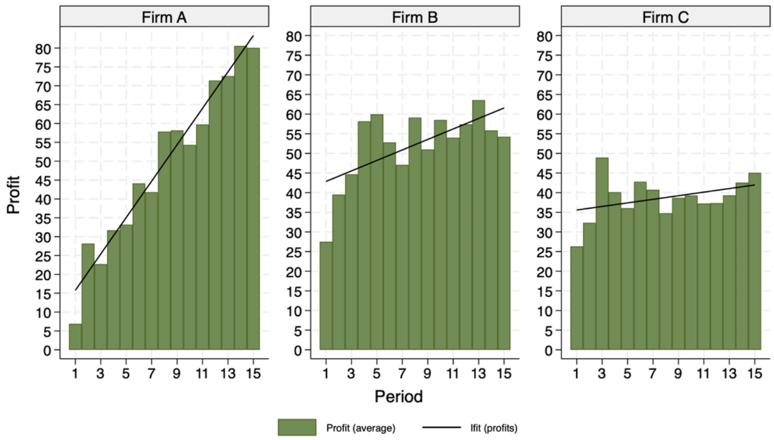

(c) N*AAMT*

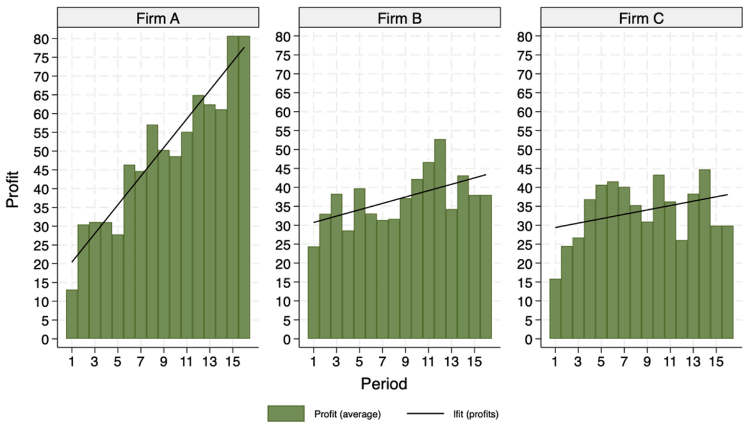

Fig. 3. Firm's average profits.

they receive preferential treatment (Leibbrandt et al., 2018). This omission limits the extent to which our results can be generalized beyond the laboratory setting.

Our findings raise important avenues for future research. A first question concerns whether a firm should adopt AA given its competitive environment. Specifically, if a firm adopts AA while a competitor does not, it faces a trade-off: On the one hand, AA attracts highly productive women who might otherwise opt for piece-rate firms; on the other, it risks losing highly productive men to competitors that do not implement AA. A second question is whether the gains from AA would persist once the policy is removed. While our results show that AA closes gender gaps in tournament entry in the short-term, its long-term impact is uncertain. Repeated exposure may build confidence and normalize women's participation, yet participation could revert if structural barriers remain. Finally, our findings also raise questions about the implications of gender-based sorting for team dynamics and collaboration. When teamwork plays a central role in productivity, strong gender clustering within firms could introduce challenges related to diversity, innovation, and performance (Bear and Woolley, 2011; Azmat, 2019). More research is needed to assess the strategic trade-offs of AA, its persistence once policies are removed, and its broader implications for team composition and organizational performance.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests:

Ramon Cobo Reyes reports financial support was provided by Georgetown University. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Instructions *AAMT*

First of all, thank you for participating in an experimental study on the labor market. The purpose of this experiment is to study how individuals make decisions in this context. The instructions are simple, and if you follow them carefully, you will receive a cash payment in a confidential manner, as no one will know how much the other participants are paid. Please keep in mind that your decisions will affect the amount of money you earn in the experiment. You may ask questions at any time by raising your hand. Apart from these questions, any kind of communication between you is not allowed. This experiment is funded by Georgetown University.

1. To ensure anonymity and confidentiality, you have been assigned a random code at the beginning of the experiment.
2. There are two types of roles: employer and worker.
3. Within the employer role, there are three types: Employer Type A, Employer Type B, and Employer Type C.
4. In the first period, each employer is randomly matched with four workers.
5. The computer will randomly determine whether you are an employer (Type A, B, or C) or a worker. You will be informed of this at the beginning of the experiment. The identity of other participants, as well as your own, will remain completely anonymous before, during, and after the experiment.
6. The experiment consists of 15 rounds. In each round, workers will earn money based on their performance in the following task. Workers must add a series of five two-digit numbers during a three-minute period.

The type of sum the worker will have to solve is: 22+34+16+05+75.

7. The amount of money the worker earns depends on the employer they are working for.
7.1 If the worker is working for Employer Type A, the worker earns 10 points per correct answer.
7.2 If the worker is working for Employer Type B, their earnings depend on a tournament. Each worker will be paired with another worker from the same company, and the worker who solves more sums in each pair will earn 15 points per correct answer. The worker who solves fewer sums will earn 5 points per correct answer.
7.3 If the worker is working for Employer Type C, earnings also depend on the result of a tournament. As before, the worker who solves more sums earns 15 points per correct answer, and the one who solves fewer earns 5 points per correct answer. The main difference is that for female workers, 2 correct answers will be added to the total number of sums they solved. Note that these 2 correct answers are used only to decide who wins the tournament. When calculating payments for the employer, female workers, and male workers, the actual number of correct answers (without any additions) will be used.
8. The employer's earnings (Type A, B, or C) depend on the performance and number of workers in their group. In each round, employer earnings are calculated as follows:

Employer's earnings = 4.5 points × number of correct sums by the four workers – 8 points per worker employed

9. At the end of each round, each worker will receive information on:
9.1 i) their performance, ii) their payment, iii) whether they won or lost the tournament (for workers in companies B and C), iv) information on the opponent's performance in the tournament. This information will be given as a range. For example, whether the opponent's performance was in the range [0–10 correct answers], [11–20 correct answers], etc., depending on the opponent's actual performance. Workers in Company A will receive the same performance range information about another randomly selected member of their company.
9.2 Each worker will also receive information on the payment system used in each of the other companies.

10. At the end of each round, each worker will have to decide whether to continue working for the same employer or switch companies. Each worker may choose which type of employer they want to work for in the next round.

11. At the end of each round, each employer will receive information on: i) the money generated by their workers, ii) the number of workers employed by them, iii) the number of workers employed by the other employers, iv) the payment system used by the other employers, v) total output of the other employers. This information will be provided as a range, e.g., total correct answers by the workers in other companies in the range [0–10], [11–20], etc.
12. All participants, regardless of their role as employer or worker, will take part in a practice round at the start of the experiment to become familiar with the summing task. In this practice round, all subjects will earn 10 points per correct answer.
13. At the end of the experiment, your payment will be based on three randomly selected rounds out of the 15. Payment will be made privately and confidentially. Points will be converted to euros at a rate of 10 points = 1 euro.

Appendix B. Proofs of Remark and Propositions

Proof of Remark. From (6) and **A1**, worker i located in firm A at equilibrium chooses:

$$e_i^A(\theta_i) = E(\theta_i) \text{ for } \theta_i \in \left\{\underline{\theta}, \overline{\theta}\right\}. \tag{A.1}$$

From (A.1), it follows that, at equilibrium, both types of workers exert maximum effort. Likewise, from (1), (2), (7) and **A1**, we have:

$$e_i^B(\overline{\theta}) = E(\theta_i) \text{ for } \theta_i \in \left\{\underline{\theta}, \overline{\theta}\right\} . \tag{A.2}$$

Finally, from (3), (4), (5), (8) and **A1**, we have:

$$e_i^C(\overline{\theta}) = E(\theta_i) \text{ for } \theta_i \in \left\{\underline{\theta}, \overline{\theta}\right\} . \tag{A.3}$$

■

Proof of Proposition 1. We begin the proof by establishing the existence of two equilibria that satisfy points *i)–v)* of the proposition. The first equilibrium illustrates how gender segregation (between firms *B* and *C*) and productivity segregation (with high-productivity workers in firms *B* and *C* and low-productivity workers in firm *A*) can arise in our game. The second equilibrium shows how both high- and low-productivity workers may choose firm *A* (where tournaments are absent), as long as at least one high-productivity worker joins firm *B* and firm *C*. Building on these results and the associated intuitions, we then demonstrate how points *i)–v)* of the proposition hold more generally at equilibrium.

Putative Equilibrium 1: Consider the following putative equilibrium:

1.a) All workers exert maximum effort $\mathrm{E}(\theta_\mathrm{i})$, for $\theta \in \{\underline{\theta}, \overline{\theta}\}$ and $\mathrm{i} \in \{1, 2, 3, \ldots, 12\}$;
1.b) All high-productivity male workers join firm *B*, i.e., $f_i^* = B$, for all $i : (g_i, \theta_i) = (0, \overline{\theta})$;
1.c) All high-productivity female workers join firm *C*, i.e., $f_i^* = C$, for all $i : (g_i, \theta_i) = (1, \overline{\theta})$;
1.d) All low-productivity workers – both males and females – join firm *A*, i.e., $f_i^* = A$, for all $i : \theta_i = \underline{\theta}$, $\forall g_i \in \{0, 1\}$.

It is straightforward to verify that points *1.a)–1.d)* satisfy conditions *i)–v)* of the proposition. To show that this profile ω^* constitutes an equilibrium of the *Moving Game*, we proceed as follows:

Part 1.a): This directly follows from assumption **A1**.

Part 1.b): Consider a high-productivity male worker. Given the choices of all other workers (both effort and firm selection ω^*), and using (1)-(8), $\frac{\overline{x}+\underline{x}}{2} = x$, and $b > 0$, such a worker is indifferent between joining firm *B* and firm *A*, but strictly prefers firm *B* over firm *C*. Therefore, joining firm *B* is a best response for a high-productivity male worker.
Part 1.c): Consider a high-productivity female worker. Under the same conditions, she is indifferent across all three firms. Therefore, joining firm *C* is a best response for a high-productivity female worker.
Part 1.d): Finally, consider a low-productivity worker, male or female. Given the choices of all other workers (both effort and firm selection), and using (1)-(8), and $\frac{\overline{x}+\underline{x}}{2} = x$, such a worker strictly prefers joining firm *A* over either firm *B* or firm *C* (where they would receive a piece-rate equal to $\underline{x}$).

From this reasoning, we conclude that points *1.a)*–1.d*)* constitute an equilibrium of the *Moving Game*.
Putative Equilibrium 2: Consider the following putative equilibrium:

2.a) All workers exert maximum effort $\mathrm{E}(\theta_\mathrm{i})$, for $\theta \in \{\underline{\theta}, \overline{\theta}\}$ and $\mathrm{i} \in \{1, 2, 3, \ldots, 12\}$;
2.b) One high-productivity male worker joins firm *B*, while all the remaining high-productivity male workers join firm *A*;
2.c) One high-productivity female worker joins firm *C*, while all the remaining high-productivity female workers join firm *A*;
2.d) All low-productivity workers – both males and females – join firm *A*.

It is straightforward to verify that points *2.a)*–*2.*d*)* satisfy conditions *i)*–*v)* of the proposition. To show that this profile constitutes an equilibrium of the *Moving Game*, we proceed as follows:

Part 2.a): This directly follows from assumption **A1**.
Part 2.b): Consider the high-productivity male worker in firm *B*. Given the choices of all other workers (both effort and firm selection), and using (1)-(8), $\frac{\overline{x}+\underline{x}}{2} = x$, and $b > 0$, such a worker is strictly better-off in firm *B* than in either of the other two firms. Now consider a high-productivity male worker in firm *A*: Under the same conditions, he is indifferent between firm *A* and *B*, and strictly prefers firm *A* to firm *C*.
Part 2.c): Consider the high-productivity female worker in firm *C*. Under the same conditions, she is strictly better-off in firm *C* than in either of the other two firms. A high-productivity female worker in firm *A* is indifferent across all three firms.
Part 2.d): Finally, consider a low-productivity worker, male or female. Given the choices of all other workers (both effort and firm selection), and using (1)-(8), and $\frac{\overline{x}+\underline{x}}{2} = x$, such a worker strictly prefers joining firm *A* over either firm *B* or firm *C* (where they would receive a piece-rate equal to $\underline{x}$).

From this reasoning, we conclude that points *2.a)*–2.d*)* constitute an equilibrium of the *Moving Game*.

Building on these results, we now formally establish points *i)-v)* in the proposititon. We proceed by backward induction. In Stage 1, for any composition of workers in each firm, assumption **A1** guarantees that all workers exert maximum effort. This proves part *i* in the proposition.

We now prove parts *ii* and *iii* in the proposition. Given (1)-(8), $\frac{\overline{x}+\underline{x}}{2} = x$, $b > 0$, workers' effort choices, and given the presence of at least one female high-productivity worker in firm *C*, a high-productivity worker – either male or female – is weakly better-off joining firm *B* than either firm *A* or *C*. Specifically, regardless of the composition of workers in firm *B*, a high-productivity worker can *expect* a piece-rate not lower than *x* by joining firm *B*, where a high-productivity worker expects a strictly higher piece-rate by joining firm *B* rather than firm *A* or firm *C* when no other worker joins firm *B*. This latter observation proves that, at equilibrium, there must be at least one high-productivity worker located in firm *B*, proving part *ii* in the proposition.

Given the results established in parts *i* and *ii*, a similar reasoning to that used to prove part *ii* can be implemented to prove part *iii* in the proposition. Unlike part *ii*, only high-productivity female workers join firm *C*. This occurs because $b > 0$, implying that a high-productivity male worker can profitable deviate by moving from firm *C* to firm *B*. This result proves both part *iii* and *v*.

To prove part *iv* in the proposition, first note that, from (6) and part *i* in the proposition, a low-productivity worker's payoff when located in firm *A* is:

$$U_i^A\left(\underline{\theta}\right) = x\underline{e} - \frac{\underline{e}^2}{2}. \tag{A.4}$$

Given (1)-(8), $\frac{\overline{x}+\underline{x}}{2} = x$ and part *i* in the proposition, a low-productivity worker can expect a payoff higher than (A.4) by joining either firm *B* or *C* if and only if no high-productivity worker joins those firms. However, by parts *ii* and *iii* in the proposition, at equilibrium there is at least one high-productivity worker located in both firm *B* and *C*, implying that all low-productivity workers join firm *A* (proving part *iv* in the proposition). ■

Proof of Proposition 2. The proofs for all parts in the proposition follow directly from the proof of Proposition 1. Unlike Proposition 1, in the *No-Bonus Moving Game*, because $b = 0$, there is no difference between firm *B* and firm *C*. As a consequence, the same reasoning implemented for firm *B* in Proposition 1 applies to both firm *B* and firm *C* in Proposition 2. This proves part *iii* in Proposition 2. ■

Appendix C. Productivity-Based Sorting (Median Split)

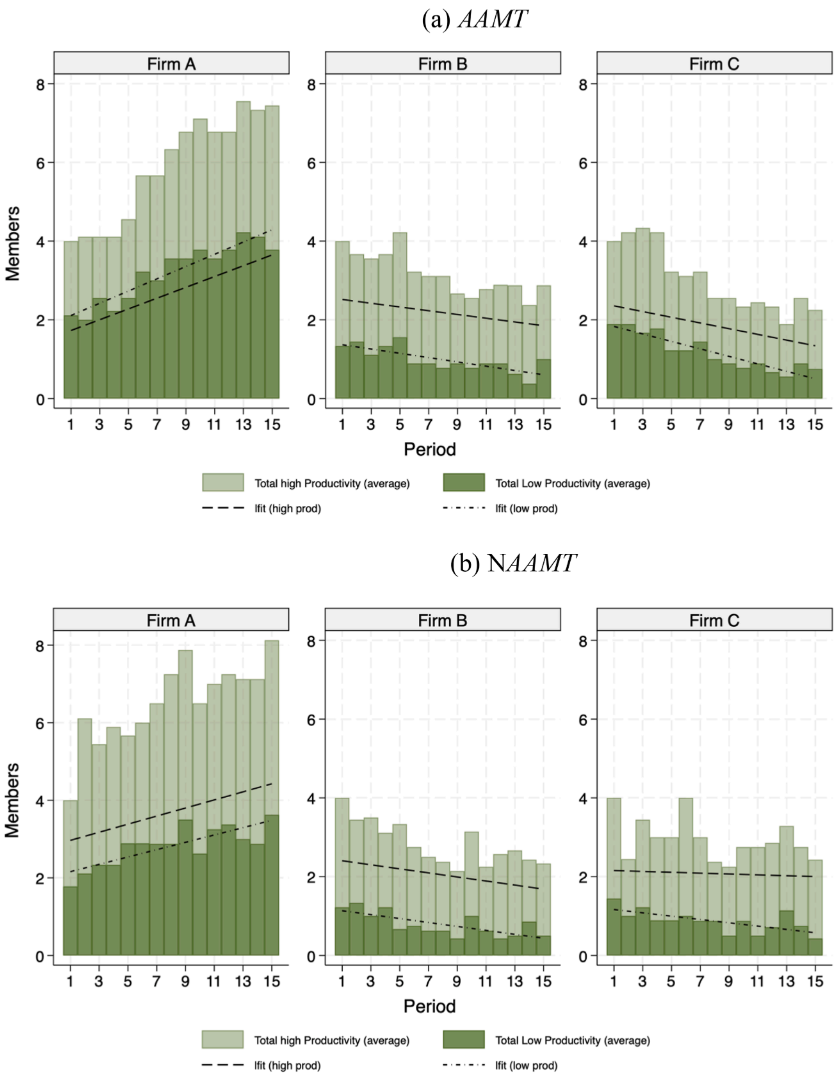

Fig. C1. Group formation in *AAMT* and *NAAMT*, by median productivity.

Similar to Figs. 2(a) and 2(b), Fig. C1(a) and (b) show that 60 % of low-productivity workers compared to just 40 % of high-productivity opt for the piece-rate system, in AAMT.[35] Similarly, the proportions in NAAMT are 62 % and 46 %.[36]

The proportion of high productive workers in Firms B and C are 32 % and 28 %, respectively. This difference is not statistically significant (z = 0.972, p = 0.352, two-tailed Mann-Whitney test).

[35] Differences are statistically significant (z = 2.605, p = 0.008, two-tailed Mann Whitney test)

[36] Differences are statistically significant (z = 3.182, p < 0.001, two-tailed Mann Whitney test)

Appendix D. Sorting by productivity and gender

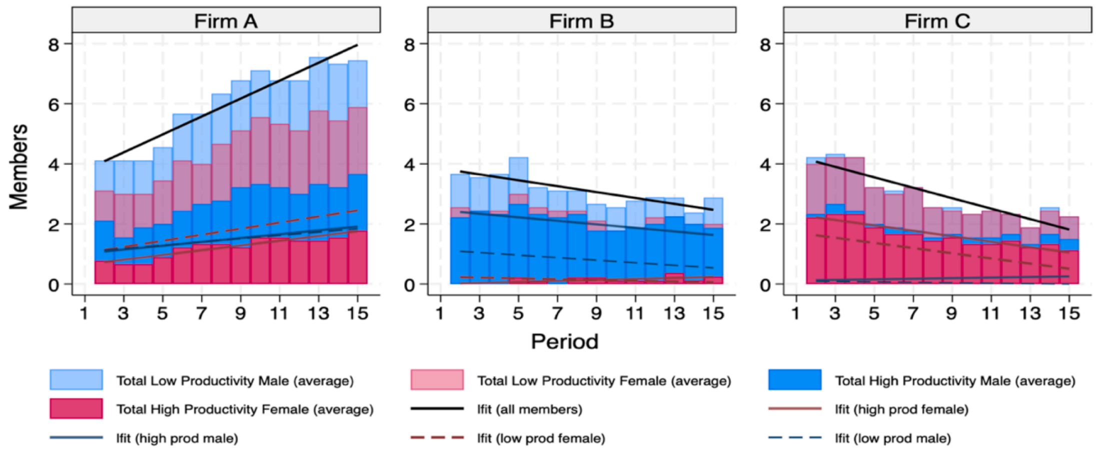

Fig. D1. Group formation in *AAMT*, by productivity and gender.

Fig. 2(a) in the main text illustrates that tournament systems without Affirmative Action (AA) policies are more effective at retaining talent compared to those with AA policies. A potential explanation for this observation is that high-productive females may be more likely to leave competitive environments than their male counterparts.

In order to explore this possibility, Fig. D1 extends the analysis in Fig. 2 by combining two factors: gender and initial productivity. The results in Fig. D1 show that, in Firm C, the decrease in the number of high-productive females is statistically significant ($t = 3.22$, $p = 0.002$). In contrast, the decrease in the number of high-productive males in Firm B is not statistically significant ($t = 1.29$, $p = 0.198$). These findings support the notion that behavioral differences between high-productivity male and female workers are driving the patterns observed in Fig. 2(a).

Appendix E. Analysis of tournament choice by productivity level

Table E1
Probability of choosing tournament by level of productivity.

	AAMT			NAAMT		
	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A: High productivity					
Female	0.001	0.018	−0.048	−0.185*	−0.198**	−0.120*
	(0.111)	(0.116)	(0.087)	(0.097)	(0.093)	(0.065)
Lagged Performance		0.007			0.031***	
		(0.012)			(0.006)	
Lagged Win			0.037			0.045
			(0.081)			(0.050)
N	480	448	177	450	420	166
	Panel B: Low productivity					
Female	0.089	0.066	0.132	−0.070	−0.091	−0.071
	(0.121)	(0.102)	(0.124)	(0.099)	(0.094)	(0.096)
Lagged Performance		−0.017			0.001	
		(0.019)			(0.023)	
Lagged Win			0.252*			−0.001
			(0.137)			(0.182)
N	480	448	308	450	420	256
Age	✓	✓	✓	✓	✓	✓
Degree	✓	✓	✓	✓	✓	✓

Notes: The coefficients represent average marginal effects from a probit model estimated using the sample of high productive (Panel A) and low productive (Panel B) participants in a panel format over 15 periods. The outcome variable is a dummy equal to 1 if player *i* chooses a tournament-type game (i.e., Firm B or C) and 0 otherwise (Firm A). High productivity is a dummy variable equal to 1 for productivity above the 75th percentile and 0 for below the 25th percentile (i.e., low productivity). Lagged performance is the number of correct answers in period $t-1$. Lagged win is a dummy variable equal to 1 if player *i* had a higher score than their opponent in period $t-1$, conditional on having participated in a tournament in that period. The first period is omitted in columns (2), (3), (5), and (6). Standard errors in parentheses clustered at the group level: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Data availability

Data will be made available on request.

References

Azmat, G., 2019. Gender diversity in teams. IZa World Labor. 29. https://doi.org/10.15185/izawol.29.v2, 2019.
Balafoutas, L., Sutter, M., 2012. Affirmative action policies promote women and do not harm efficiency in the laboratory. Science 335, 579–582.
Barbulescu, R., Bidwell, M., 2013. Do women choose different jobs from men? Mechanisms of application segregation in the market for managerial workers. Organ. Sci. 24 (3), 737–756.
Bear, J.B., Woolley, A.W., 2011. The role of gender in team collaboration and performance. Interdiscip. Sci. Rev. 36 (2), 146–153. https://doi.org/10.1179/030801811X13013181961473.
Blau, F., Kahn, L., 2000. Gender differences in pay. J. Econ. Perspect. 14 (4), 75–99.
Blau, F., Kahn, L., 2013. Female labor supply: why is the United States falling behind? Am. Econ. Rev. 103 (3), 251–256.
Bracha, A., Cohen, A., Conell-Price, L., 2019. The heterogeneous effect of affirmative action on performance. J. Econ. Behav. Organ. 158, 173–218.
Calsamiglia, C., Franke, J., Rey-Biel, P., 2013. The incentive effects of affirmative action in a real-effort tournament. J. Public Econ. 98, 15–31.
Cotton, C., Hickman, B., Price, J., 2022. Affirmative action and Human capital investment: evidence from a randomized field experiment. J. Labor. Econ. 40, 157–185.
Dohmen, T., Falk, A., 2011. Performance pay and multidimensional sorting: productivity, preferences, and gender. Am. Econ. Rev. 101 (2), 556–590.
Eriksson, T., Villeval, M., 2008. Performance-pay, sorting and social motivation. J. Econ. Behav. Organ. 68 (2), 412–421.
Fallucchi, F., Quercia, S., 2018. Affirmative action and retaliation in experimental contests. J. Econ. Behav. Organ. 156, 23–40.
Flory, J.A., Leibrandt, A., List, J.A., 2015. Do competitive workplaces deter female workers? A large-scale natural field experiment on job entry decisions. Rev. Econ. Stud. 82, 122–155, 1 (290).
Greiner, B., 2004. An online recruitment system for economic experiments. Published in: Forschung und wissenschaftliches Rechnen 2003 GWDG Ber. 63, 79–93.
Ibáñez, M., Riener, G., 2018. Sorting through affirmative action: three field experiments in Colombia. J. Labor. Econ. 36 (2), 437–478.
Larkin, I., Leider, S., 2012. Incentive schemes, sorting, and behavioral biases of employees: experimental evidence. Am. Econ. J.: Microecon. 4 (2), 184–214.
Niederle, M., Vesterlund, L., 2007. Do women shy away from competition? Do men compete too much? Q. J. Econ. 122 (3), 1067–1101.
Leibbrandt, A., Wang, L., Foo, C., 2018. Gender quotas, competitions, and peer review: experimental evidence on the backlash against women. Manage Sci. 64 (8), 3501–3516.
Niederle, M., Segal, C., Vesterlund, L., 2013. How costly is diversity? Affirmative action in light of gender differences in competitiveness. Manage Sci. 59 (1), 1–16.
Sutter, M., Glätzle-Rützler, D., Balafoutas, L., Czermak, S., 2016. Cancelling out early age gender differences in competition: an analysis of policy interventions. Exp. Econ. 19, 412–432.
Weichselbaumer, D., Winter-Ebmer, R., 2007. The effects of competition and equal treatment laws on gender wage differentials. Econ. Policy. 22 (50), 236–287.