

Breaking the Glass Ceiling with a Gavel? Gender Promotion Gap in the Brazilian Judiciary

Helena Laneuville

Monday 13th May, 2024

Abstract

I identify a gender promotion gap among judges in the state of São Paulo, Brazil, and characterize its labor-demand and labor-supply-related drivers. To do so, I construct a novel dataset that measures characteristics, skills and observable performance on the job and the process of career advancement. Women take significantly longer to be promoted within the Trial Court and, consequently, acquire less experience with larger court districts. At the end of their career, they are four times less likely than men to reach the Appeals Court. On the labor supply side, women are 14% less likely to seek career advancement within the Trial Court, which accounts for 40% of the gender promotion gap at this stage. On the labor demand side, another 41% of this gap can be explained by a differential probability of receiving a promotion when applying. Among applicants for merit-based promotions, female candidates are 21% less likely to be successful than their male counterparts. This effect is not explained by an observable difference in skills at the time of hiring or productivity on the job.

1 Introduction

Having more women in positions of power changes how decisions are made both because women make different choices¹ and because they can bring new information and perspectives (Boyd et al., 2010). Despite the social benefits of gender parity in the workplace, women’s professional progression can be hindered for reasons related to both labor supply and labor demand. On the labor-supply side,

¹See Hessami and da Fonseca (2020) for a review of the literature on how female politicians make different decisions from their male counterparts. Adams-Prassl et al. (2022) show that, in the case of workplace male-on-female violence, the boss’s gender impacts the outcome of the case. Papers such as Cai et al. (2022) and Bielen and Grajzl (2020) provide evidence of different decisions made by judges of different genders.

a large body of literature describes gender gaps motivated by different choices (Ngai and Petrongolo, 2017; Säve-Söderbergh, 2019; Azmat and Ferrer, 2017; Bertrand et al., 2010). On the labor-demand side, women can be less recognized at work even when they make choices similar to men's (Benson et al., 2023). Most papers, however, focus on either labor supply *or* labor demand.

The main contributions of this paper are to discuss both the supply- and demand-side drivers of the gender promotion gap and measure their relative importance in a high-profile occupation. Specifically, I focus on state judges in São Paulo, Brazil, for three reasons. First, job titles and performance measures are comparable in the long run. Since this occupation has very little turnover,² men and women who start their judicial careers can be followed for a long time. After all, the entry costs are high, and the benefits are excellent. The only way to become a judge in any Brazilian state is to pass a competitive admissions exam with a below-2% acceptance rate. Securing judicial employment brings with it a salary among the highest in the country and guaranteed job security after two years. Second, I can clearly characterize reasons for promotion and career progression because the criteria are transparent. Half of open positions are chosen by a mechanical seniority-based criterion and the other half by a more discretionary merit-based criterion. Third, there is a wealth of useful publicly available information. The São Paulo Court system publicizes judicial admission exam grades, an annual list of all active judges, several measures of on-the-job productivity, and the entire process of promotion, including applicants, vacancies, selection criteria, and final decisions. I collect and digitalize this information to create a novel dataset.

I have four main findings. First, there is a gender promotion gap that widens at later stages of the judicial career despite the fact that men and women are similar at the time they enter the judiciary and rarely exit before retirement. Although female and male judges are similar in age, ability, and job attachment, women take longer to be promoted within the Trial Court. Furthermore, eligible women are four times less likely to be promoted to the Appeals Court. Only 0.53% of eligible women are promoted each year, compared to 2.10% of eligible men.

Second, supply-side factors can account for some, but not all, of the gap. Women apply more selectively and are less recognized by the merit criterion than male applicants competing for the same promotions. Eligible female Trial Court judges are 14% (3.89 p.p.) less likely to apply for promotion than their male counterparts and are more selective in the positions to which they apply. I find supporting evidence that part of this latter effect is because women face a higher cost of moving away

²Less than 2% of men and 5% of women who started between 2007 and 2012 had left the justice system within 10 years.

from their current job location than do men.

Third, on the labor-demand side, female judges who apply for promotion within the Trial Court are 20.7% (3.22 p.p.) less likely to obtain a merit-based promotion and 10.3% (2.53 p.p.) less likely to be promoted than equally senior male judges. This disparity in merit promotions is not explained by differences in productivity or the number of options to which judges apply to. As a result, women eligible for promotion to the Appeals Court spend 1.34 fewer years than their male counterparts in the last stage of Trial Court.

The fourth finding comes from a long-run simulation and decomposition of the gap in a simulated scenario based on the various estimates discussed. In this setting, the promotion gender gap widens when men start getting their third and fourth promotions in the career ladder. The first widening of the gap is due to women taking longer to reach the final stage of Trial Court and persists for 10 years. Although women eventually catch up, the gap widens again around year 23 as men start getting promoted for the fourth time, now to the Appeals Court. By the time judges reach retirement, a large gap remains, with men being 12.08 p.p. more likely to have received all four promotions by year 30.

The promotion gender gap reaches its first peak at 11 years of tenure. At this point, 40.0% of the gender disparity is explained by different propensities to seek career advancement, while 40.6% is attributed to different probabilities of being promoted conditional on applying. The remaining 19.4% is explained by differential attrition. The second spike in the gender gap is largely due to women not spending enough time in the preceding step of the career ladder. If women were to seek promotions as frequently as their male counterparts throughout their careers, there would be smaller and more short-lived promotion gender gaps throughout the years. As a result, women would have enough time to catch up before retirement and 94% of the gap would have been closed by year 30.

Finally, I discuss other factors that could potentially impact judges' endogenous labor supply choices — such as productivity and applying for promotion — and, consequently, the gender promotion gap. I find no evidence that children or intra-couple bargaining differentially affects female productivity. Even during child-bearing years, women are more productive than men in the same age group. Furthermore, I find evidence that a judge's likelihood of applying for promotion is not independent of their expected probability of success; each judge becomes more likely to apply when they become more competitive by crossing a relevant seniority threshold. This result is robust to controlling for characteristics of their current occupation and the vacancies offered at each promotion opportunity. This suggests that there are sunk costs to promotion application, as otherwise judges would be expected to apply to all

vacancies that dominate their current position regardless of their chances of success.³ Consequently, a change in the promotion mechanism can impact the gender gap beyond the mechanical effect on the male and female candidates' probability of success into judges' application decisions.

This paper contributes directly to the literature on gender pay gaps by combining two of its branches. The first group of authors focuses on supply-side factors. There are several papers discussing how women choose different occupations (Ngai and Petrongolo, 2017), negotiate starting salaries and benefits less fiercely (Säve-Söderbergh, 2019), work fewer hours, and/or take more career pauses following the birth of children (Azmat and Ferrer, 2017; Bertrand et al., 2010). Other papers are connected more broadly to the discrimination literature (Becker, 1971; Phelps, 1972; Arrow, 1974) and find that women are less recognized even when they make similar choices. For instance, Benson et al. (2023) show a case where women are more productive but not perceived as promising employees by their bosses.

Furthermore, I add to the body of evidence on the interaction between screening mechanisms and the professional recognition of minority workers by showing that a subjective criterion can have disparate impacts on women in the judiciary system. Several papers have already shown that standardized tests or objective mechanisms either don't hurt (Autor and Scarborough, 2008) or are significantly less discriminatory than subjective assessments, such as interviews or discretionary performance evaluations (Mocanu, 2023; Hoffman et al., 2018; Shukla, 2022).

The paper is structured as follows. Section 2 describes the institutional setting and the promotion criteria. Section 3 describes the data and presents basic summary statistics. Section 4 documents that there is a promotion gender gap despite the fact that male and female judges have similar ability and almost equally high job attachment. Section 5 shows the contribution of labor supply- and labor demand-related decisions to the promotion gender gap. Section 6 discusses how judges' endogenous choices are constrained by exogenous factors and how these restrictions affect the gender gap. Section 7 concludes.

2 Institutional Setting

This section explains the entry and career advancement process for judges in the São Paulo Justice System.

Judges in State Courts deal with all cases that do not concern the federal administration, elections,

³This result holds even as one incorporates the costs associated with actually taking the promotion, such as moving costs or having new colleagues.

labor disputes, or the military. Examples of cases they consider include common crimes, divorce, and child custody. These judgeships are therefore extraordinarily high-responsibility and high-status positions. After two years in office, judge's positions become permanent, guaranteeing job stability and security. That means they cannot be fired unless it has been proven that they have become mentally unfit or committed a felony.⁴

In Brazil, trial judges are selected by a public admissions exam to integrate the State Courts. In addition to fulfilling minimum eligibility requirements, judges must rank highly enough among competitors to secure a slot. Minimum requirements include a minimum score on each test, a bachelor's degree in law,⁵ and a minimum of three years of professional experience in the area. Examinations are fiercely competitive, with approval rates ranging from 0.37% to 2.54% between 2006 and 2020 (see Table 21 in Appendix A.1). This competition is likely due to the high salary, prestige, and job stability inherent in judicial positions.

After passing the exam, state judges start their careers as substitutes across multiple court districts. They fill in for other magistrates as needed and occasionally work alongside local judges to learn from more experienced magistrates. Within the Trial Court, judges can be promoted up to three times. The first promotion consists of becoming a presiding judge in a small, or initial-stage, court district. Subsequent promotions are shifts towards districts classified as intermediate and final stage based on adult population size.⁶ In São Paulo, initial-stage districts have up to 50,000 adults, intermediate districts range from 50,000 to 100,000 adults, and final districts have over 100,000 adults. As shown in Table 1, promotion leads to a small pay increase, made slightly more substantial by bonuses and benefits that make up the final annual salary. However, career advancement also means moving to a larger and often more prosperous city. For instance, the median municipal GDP per capita of a final-stage court district is double that of an initial-stage court district. Furthermore, because most final-stage jobs are in the capital, being promoted to a final-stage judgeship means moving from a relatively small city, with little more than 200 thousand people, to Brazil's largest city.

⁴This job security is a right granted by Article 95 of the Brazilian Constitution to shield judges from persecution and intimidation.

⁵In contrast to the US, undergraduate students in Brazil must choose a major before entering university. Law is one of the available options.

⁶The population measure used is the number of registered voters as voting is mandatory from ages 18 to 65 across Brazil. Hence, nearly every Brazilian in that age group is registered and assigned a voting place.

Table 1: Characteristics of Judges' Jobs by City of Residence and Career Stage

Stage	Median Job 31st Dec 2009		Pay Increase from Immediately Preceding Stage	
	Per Capita GDP	Population	Base Salary	Final Pay Before Taxes
	(BRL 2010)	Census 2010	% 2010 Annual Salaries	% 2010 Annual Salaries
Initial	16,249	28,371	4.8%	6.5%
Intermediate	30,383	210,082	6.5%	7.7%
Final	39,451	11,216,559	5.3%	6.7%
Appeals Court	39,451	11,216,559	4.2%	14.1%

Notes: To construct columns 1 and 2, I link each judge to a career stage and a municipality and take the median GDP per capita. The data reflect the place where the median judge in each category works. Column 3 is the percentage difference between the annual base salary of a judge in each career stage and the annual base salary in the immediately preceding stage. Column 4 shows the percentage increase in total salary (before taxes) when a judge moves from the immediately preceding stage to the one in x .

Every year, a small group of judges is promoted from final-stage court districts to the Appeals Court. While 80% of vacancies are filled this way, 20% of Appeals Court judges are elected from a list of candidates proposed by the Prosecutor's Office or the state branch of the Brazilian Bar Association.⁷ All promotions are decided by a commission of Appeals Court judges and sanctioned by the governor, meaning that the decision is increasingly politicized at higher promotion stages. However, the law sets salaries and benefits, preventing individual negotiation.

Promotion slots must be filled according to either a merit-based or a seniority-based criterion.⁸ If one slot is filled by seniority, the next available slot must be filled by merit, and vice-versa. The seniority criterion is mechanical. A judge is considered more senior than another in the same career stage if they have spent more time in their current career stage. If there is a tie, seniority is determined by their total experience as magistrates in the São Paulo State Court. The merit criterion is more nuanced. In theory, judges are only eligible for merit promotions if they spend at least two years in their current career stage and are among the 20% most senior judges in that category. However, other applicants may be considered for a slot if no applicant satisfies this requirement.⁹

If multiple qualified judges compete for merit promotion, evaluators must grade and compare them on five dimensions.¹⁰ The first is qualitative performance, defined as being known to write clear and well-justified sentences. The second is productivity, measured by the structure and volume of a judge's

⁷Brazilian Constitution article 94.

⁸Article 93 of the Brazilian Constitution. Originally Article 80 of Supplementary Law no 35/1979, later incorporated into the current Brazilian Constitution, enacted in 1988.

⁹The two-year cutoff is the most binding restriction.

¹⁰Resolution 106/2010 from the National Justice Council

workload. The third is promptness, as in being punctual, not missing work, and efficiently performing administrative work. The fourth is technical betterment, measured by frequency in a selection of courses offered by the court to the judges. The fifth criterion is compliance with the Ethical Code of the Magistrates. Failing to meet this criterion is immediately disqualifying for merit promotion.

Promotions in the justice system are well-documented and must always follow the same procedure. First, vacancies must be publicly announced in the São Paulo Justice Diary along with the promotion criteria and application instructions. After the application deadline, the Court publishes tables with information on the vacancies each judge is competing for. Judges are allowed to submit two separate lists of slots they are interested in being promoted to: one for vacancies filled by seniority and one for those filled by merit. In promotions by seniority, the most senior applicant gets promoted unless 2/3 of the Appeals Court hiring committee judges vote against. In promotions by merit, the candidate with the highest number of votes is promoted. The list of promoted judges at the end of the process is also publicized in the São Paulo Justice Diary.

3 Data and Summary Statistics

This section describes the dataset constructed for this paper and provides basic summary statistics.

3.1 Data

The dataset has two main sources: the online records of the São Paulo Justice Diary and the São Paulo State Court of Justice website. The São Paulo Justice Diary has been publicly available online since October of 2007. It contains all public records of promotions and admission exams, as well as annual lists of active judges (ranked by seniority) within the Trial and the Appeals Courts. The São Paulo State Court of Justice website has detailed data on judges' monthly productivity outcomes in each court district from January 2011 to the present. The dataset constructed for this paper is summarized in Table 2.

Table 2: Summary of All Pieces of Data and Coverage

Data	Frequency	Oct/07	Dec/07	Jan/11	...	Dec/22
Promotions	Each Opportunity					
Exam Grades	Each Exam					
Seniority List	Yearly					
Productivity	Monthly					

Next, I describe each of these components and the resulting datasets constructed.

3.1.1 Seniority Lists and Career Trajectory Datasets

Every January, the State Court of Justice publicizes in the São Paulo Justice Diary two lists (one for the Trial Court and one for the Appeals Court) of all active judges as of December 31st of the previous year, ordered by seniority.

The Trial Court annual seniority lists contain the judge’s name, job title as of December 31st, career stage, starting date in the profession, and starting date in their current career stage. Job titles are informative of a judge’s court district and, if applicable, specialization in an area of law. In some years, the sub-sample of substitute judges also contained their date of birth. I collected all annual Trial Court seniority lists available in the São Paulo Justice Diary, covering 2007 to 2022.

Each Appeals Court seniority list contains the magistrate’s name and a letter indicating whether the judge was promoted from the final-stage Trial Court by promotion or was elected from a list of candidates appointed by the public prosecutor’s office or the bar association.

I constructed the following datasets from these lists. The first is a panel of all active judges and their respective career stages from 2007 to 2022. The second is a dataset containing the summary of the career of each trial judge ever recorded as being active during the sample period.¹¹ If a judge was not promoted to at least the intermediate stage before December 31st, 2007, the pooled seniority lists from 2007 to 2022 contain the date of every promotion they received until December 31st, 2022. Otherwise, the information on the lists included the date when they started as magistrates, the date when they started the stage they were in as of December 31st, 2007, and the date of every subsequent promotion they received until December 2022.

¹¹The exact dates of promotions to the Appeals Court were recovered by linking the judges’ name and year of promotion to the promotion records later in this part.

Pooled seniority lists are difficult to combine because the same person may be recorded with slightly different names in different years.¹² This may yield gender-based bias in results as women are more likely to change their names throughout their lives in response to marriage and divorce. To minimize this concern, I create a unique identifier for all records of the same judge in the following way. For every year t from 2008 to 2022, a judge should have a new entry if they began their judicial career in year t . If they were an active judge prior to January 1st of the year t , they should have a record in year $t - 1$. I therefore searched the data for matching records in the prior year. If there was no perfect name match, I searched for records of judges with the closest name and the same starting date (conditional on having over 90% similarity using the Jaro-Wrinkler distance). When such a match did not exist,¹³ I matched similar names to the same starting date in the current career stage. With this approach, I got a 100% match rate for all records of the set of judges who appeared in any seniority list in year $t \in \{2008, \dots, 2022\}$. Finally, I added to the dataset all records from judges who were active on December 31st, 2007, but not on December 31st, 2008.

Finally, I created female dummies for all datasets using the package `genderBR` in R, which matches first names to Brazilian census records to attribute gender. For the names that could not be classified in this way, I manually searched for the judges in Google to identify their gender.

3.1.2 Promotion Data

To track down the different phases of the promotion process, I used records from the São Paulo Justice Diary from October 2007 to December 2022, yielding two datasets:

- i) *Judges' stated preferences*: I collected public announcements of judges' promotion preferences for every promotion opportunity since October 2007. That is, for each career stage and promotion criterion (merit or seniority), I collect a table such that each entry is the ranking judge j attributes to vacancy v . If judge j is not interested in vacancy v , the entry remains empty. For promotions to the Appeals Court, judges mark positions in which they are interested with an S rather than in ranked order.
- ii) *Promotion Records*: On the official date promotions go into effect, the São Paulo Justice Diary lists all promoted judges, the promotion criterion, and the judge's new job title. When the promotion results from a competitive process, the list includes the judge's old position. Oth-

¹²For instance, if a judge changed their last name or if it has many words or a slightly more complicated spelling.

¹³E.g., because there was a one or two-day mismatch in the judges' starting date relative to the previous year.

erwise, the judge is promoted mechanically because their court district increased in size (and is therefore attributed to a higher career stage). In this case, the announcement mentions the legislation establishing that judges from court districts that grew in size and changed categories are prioritized for promotion.

I first linked the dataset of promotion records to the judges' career dataset in the following way: for each promotion date, I matched the names of each promoted judge to someone in the career summary dataset with a similar name¹⁴ who received a corresponding promotion within a 14-day bandwidth.¹⁵ All promotions were matched, and each judge was associated with the same identifier as in the other datasets.

I linked the promotions and the preferences datasets by matching the date the preferences were stated to the date the promotions were announced. To constitute a match, the sub-sample of preferences reported on a given date had to fulfill two requirements: 1) all job titles of the promoted judges must have been listed as available choices; 2) every promoted judge had to be listed as an applicant for promotion at the career stage and criterion reported in the promotion dataset.

3.1.3 Exam Grades

The result of each admissions exam between October 2007 and December 2022 was publicized in the São Paulo Justice Diary, including the final grades of all candidates who passed. From 2014 to 2022, the exam records also include the breakdown of grades in each component of the exam. There are five different components: two tests of general legal knowledge, one multiple choice and one open-ended; two tests of the candidate's ability to sentence life-like cases, one focused on civil cases and the other on criminal cases; and a score for the candidate's academic titles.

Using that information, I constructed a dataset of exam grades for all selected candidates. I later linked this dataset to the panel of active judges and the record of judges' career trajectories.

3.1.4 Productivity

I use each judge's productivity measures by year, month, court district, and job description found on the website of the São Paulo Court system.¹⁶ The outcomes include the number of hearings, final

¹⁴Over 95% similarity to at least one of the variations of the person's name in the seniority lists according to the Jaro-Wrinkler distance.

¹⁵In some cases, the entries are posted in the São Paulo Justice Diary on one date refer promotions that went into effect on a slightly different date. This precaution was added to account for these discrepancies.

¹⁶<https://www.tjsp.jus.br/produktividadesweb/> Access on June 1st, 2023.

decisions, intermediate decisions, orders, and case records waiting for a judge’s action by the end of each month. Because this dataset contained job descriptions, I also constructed dummies for whether the judge was working in a position that specialized in a given area of law each month. This is to account for the fact that a judge in a specialized courtroom could have a distinct workload because of the nature of the cases they are assigned. Consequently, their performance outcomes might not be directly comparable to their colleagues in different specialties.

It is important to note that judges are randomly assigned to cases within a court district and specialization (if applicable). Hence, these outcomes are likely to be informative of a judge’s effort and efficiency in dealing with a given workload. Furthermore, because the São Paulo State Court itself provides this data, it is the most likely source of productivity statistics that Appeals Court judges use to decide on promotions.

Productivity Measures

Finding the right productivity measure is challenging because Resolution 106/2010 of the National Council of Justice (which regulates judicial promotions in state courts) does not specify: 1) how many months of productivity data should be considered for promotion decisions and 2) the relative importance of each measure of productivity. It is said, however, that productivity should be evaluated for at least two years. Judges who spent less time in their current career stage will only be considered for a promotion in the absence of candidates who fulfill this minimum requirement.

I consider two different measures of productivity. First, I follow Dahis et al. (2020) in using the number of cases disposed of each month as a benchmark. Second, I calculate a linear combination of the productivity-related outcome variables that best predict merit promotions of male candidates.

To construct the second measure, I aggregate both productivity measurements and specialization dummies that might change how productivity is perceived. The productivity measurements I consider are the number of hearings, final decisions, intermediate decisions, orders, and case records waiting for a judge’s action by the end of each month. The specialization dummies indicate whether a judge worked in a courtroom specialized in criminal law, family law, civil law, public treasury affairs, or infancy and youth rights. I aggregate these measures by semester and rescale them in the following way: if the outcome i is a dummy for specialization in a given area of law, then for each of the four semesters $\tau \in \{t-1, t-2, t-3, t-4\}$ preceding a promotion opportunity t , I assign the value $\hat{X}_{ji\tau} = 1$ if judge j has worked in such a specialized courtroom during semester τ and $\hat{X}_{ji\tau} = 0$ otherwise. If the

outcome is a productivity-related measure,¹⁷ I aggregate the measure by semester to obtain an input $X_{ji\tau}$ for judge j , outcome i , and semester τ . This outcome is then rescaled using

$$\hat{X}_{ji\tau} = \frac{X_{ji\tau} - m_{jt}}{M_{jt} - m_{jt}},$$

where

$$m_{jt} = \min\{X_{ki\tau} : k \text{ eligible for the same career stage as } j \text{ on date } t\}, \text{ and}$$

$$M_{jt} = \max\{X_{ki\tau} : k \text{ eligible for the same career stage as } j \text{ on date } t\}.$$

I then estimate a random forest model on promotion using the sample of male judges who applied for a merit promotion and did not get a seniority promotion. For this set of judges, it is reasonable to assume that the better-ranked candidates in terms of merit were promoted. I later use the estimated model to predict judges' productivity. The model considers that judges who spent less than two years in their current position may be evaluated differently:

$$\begin{aligned} \text{Promoted Merit}_{jt} = & \alpha + \gamma 1_{\{\text{tenure} \geq 2y\}} + \sum_{\tau=t-4}^{t-1} \sum_i \beta_{i\tau} \hat{X}_{ji\tau} \\ & + \sum_{\tau=t-4}^{t-1} \sum_i \eta_{i\tau} \hat{X}_{ji\tau} 1_{\{\text{tenure} \geq 2y\}} + \varepsilon_i. \end{aligned}$$

Volume of work measures

Some productivity-related outcomes (the number of hearings, the number of decisions (except sentences), and the number of sentences) can be used to proxy for the volume of work a judge puts in every month. I use these measures to investigate the possibility that women are being passed over for promotion for being less willing than their male counterparts to put in the effort necessary for career progression, for instance, as a consequence of having children and more responsibilities at home.

3.2 Summary Statistics

The São Paulo State Court is a huge sub-national court. At the end of each year (2007-2022), the Court had, on average, 2,081.7 judges in the Trial Court and 354.9 magistrates in the Appeals Court, as shown in Table 3. Although the gender composition is relatively egalitarian among substitute, initial-stage, and intermediate-stage judges, women become rarer in the biggest court districts of the

¹⁷E.g., number of hearings, final decisions, intermediate decisions, orders, or case records waiting for a judge's action by the end of each month.

Trial Court and practically disappear from the Appeals Court. By law, one-fifth of the members of the Appeals Court must be either prosecutors, public defenders, or lawyers of known legal expertise, as indicated by the Prosecutor’s Office and the Brazilian Bar Association.¹⁸ Among Appeals Court judges promoted from the Trial Court, the proportion of women is even lower, averaging just under 5%.

Table 3: Summary Statistic for Judges in the São Paulo Justice System

<i>Trial Court (Annual Average of Active Judges 2007-2022)</i>		
<i>Career Stage</i>	<i>Number of judges</i>	<i>% Fem</i>
Substitute	104.69	43.77%
Initial (court districts \leq 50k voters)	201.19	41.50%
Intermediate (court districts between 50k and 100k voters)	684.88	45.56%
Final (court districts above 100k voters)	1090.94	33.51%
Total	2081.69	38.76%
<i>Appeals Court (Annual Average of Active Judges 2007-2022)</i>		
<i>Career Stage</i>	<i>Number of judges</i>	<i>% Fem</i>
Appeals Court Judges	354.94	6.45%
Appeals Court Judges Promoted from Trial Court	286.25	4.67%

Notes: Summary of career stage and gender composition of São Paulo’s State Justice Courts from 2007-2022, constructed from the annual seniority list released in the São Paulo Justice Diary. Each list contains the name, career stage, starting date on the job, and starting date in the current career stage of all active judges as of December 31 of the previous year.

This paper shows that the pattern of gender composition in Table 3 is not in the process of converging to an egalitarian distribution because career advancement is still an unequal process. In order to show this, I analyze what happens each time a promotion opportunity arises. A typical promotion opportunity involves the public announcement of a batch of vacancies in the subsequent career stage, instructions on how eligible judges should apply, and the selection criterion for each posting. I investigate if and why there is a gender difference in the application patterns for posted vacancies and in the likelihood of receiving a promotion. Table 4 summarizes the total number of promotion opportunities and judges promoted in each career stage.

¹⁸Brazilian Constitution, article 94.

Table 4: Summary of Promotion Opportunities by Career Stage

	Promotion Opportunities	Judges Promoted
Final to Appeals Court	102	259
Intermediate to Final	37	925
Initial to Intermediate	43	811
Substitute to initial	41	857

Notes: Each promotion opportunity consists of a list of vacancies to be filled. The first column captures the number of promotion opportunities in the sample period (Oct 2007-Dec 2022), and the second column contains the total number of judges promoted to these opportunities.

4 Gender Promotion Gap

In this section, I provide evidence that female judges take substantially longer to get promoted within the Trial Court and rarely reach the Appeals Court. This promotion gap exists even though judges are very similar at the beginning of their careers and rarely leave their jobs before retirement.

There are no gender differences in age or ability among rising judges who passed their examinations between October 2007 and December 2022. Table 5 shows the difference in exam scores and age¹⁹ by gender. Male and female judges are roughly the same age on their first day at the job and have similar final grades in admission exams. I then verify if similar grades actually mean similar ability, which would imply that the rankings of top candidates are not affected by gender bias in evaluations. To do so, I compare the gender difference in grades awarded on the blinded and non-blinded components of the test, which are available from 2014 onward. The blinded components are the multiple choice test, the written exam about general knowledge of the law, the exams simulating civil and criminal cases to be sentenced by candidates, and the pre-determined scoring system for the degrees and titles each candidate has. The oral exam must be ‘non-blinded’ by construction. The gender difference is very similar and null in all components, which indicates that differential treatment by gender does not affect these judges’ final grades, and their abilities are indeed similar.

¹⁹Measured as time elapsed between their date of birth and their first day on the job as a judge.

Table 5: Judges' Age and Ability at Entry

	Age (years)	Final Score (z-scored)	Admission Exam Components					Oral Exam (z-scored)	Academic Titles (z-scored)
			Multiple Choice (z-scored)	Written (z-scored)	Civil Sentence (z-scored)	Crim. Sentence (z-scored)			
Fem	-0.38 (0.28)	-0.03 (0.07)	-0.06 (0.10)	-0.04 (0.10)	0.13 (0.10)	-0.10 (0.10)	-0.05 (0.10)	-0.11 (0.10)	
Num. obs.	634	885	453	453	453	453	453	453	
Exam. FE	✓	✓	✓	✓	✓	✓	✓	✓	

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

Notes: Z-scores of judges' admission exam grades from October 2007 to December 2022. The sample size is different here because, for the final grade, the sample spans the years of 2007 to 2022. However, the final grade breakdown is only available from 2014 to 2022.

Once a judge is selected in the admissions exam and accepts the job offer, they rarely quit their job before they reach retirement age, as shown in Table 6. Among the judges who started their careers between 2007 and 2012, over 95% of women and 98% of men are still active after ten years. Among judges who stopped being active between 2007 and 2022, the average number of years of experience at exit is 26 years for men and 21 for women (see Table 6).²⁰ Hence, the sample constructed in this paper is an accurate picture of a large portion of most of these judges' careers. Importantly, there is little evidence of judges taking long pauses or changing to less demanding or part-time jobs after giving birth, one of the main channels of the maternity penalty.

Table 6: Job Attachment and Exit Patterns

	Male judges	Female judges	All judges
Prob. of exiting with less than 10 years of experience (% of judges starting in the 2007-2012 period)	1.72%	4.27%	2.79%
% of active judges in year t exiting in year t (Annual average 2007-2021)	1.06%	0.68%	0.92%
Average number of years of experience conditional on exit (Annual average 2007-2021)	26.47	21.11	24.97
Median number of years of experience conditional on exit (Annual average 2007-2021)	27.13	23.44	26.47

Notes: This table shows information from judges who exit the sample from one year to the next. The exit pattern is compatible with retirement, and very few judges exit the sample without working for at least ten years.

²⁰The gender difference in experience at exit is consistent with the minimum retirement age in Brazil. Note that women were eligible for their retirement pension five years earlier than men during the sample period.

Nevertheless, men’s and women’s careers do not progress at the same pace. I provide two forms of evidence of this. First, among the judges who received a promotion, women waited longer than the men. Table 7 shows that the first two promotions (to the initial and intermediate stage) are relatively fast and egalitarian²¹; among the judges who received their second promotion between 2008 and 2022, the median wait time between the first day on the job and the second promotion was 3.24 years for men and 3.51 for women. However, the third promotion (from the intermediate to final stage) takes significantly longer for women. While the median wait time for promoted female judges was 8.82 years, their median male counterparts waited only 7.5 years. Women’s promotions to the final stage therefore take 17.65% more time than men’s. When including controls for the year of promotion, women still take 12.27% (0.92 years) longer than the men.

Table 7: Median Time Until Promotion for Judges Promoted from 2008 to 2022

	Sample: Promotion 2 in 2008-2022		Sample: Promotion 3 in 2008-2022			
	Entry - Promotion 2		Promotion 2 - Promotion 3		Entry - Promotion 3	
	Median Wait	Median Wait	Median Wait	Median Wait	Median Wait	Median Wait
Female	0.27	0.29	1.32***	0.92***	1.46***	1.24***
	(0.29)	(0.18)	(0.30)	(0.15)	(0.33)	(0.26)
Year Dummies		✓		✓		✓
Median Men	3.24	3.24	7.50	7.50	11.87	11.87
Num. obs.	772	772	870	870	870	870

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$

Notes: The regressions in this table concern the gender difference in the time judges waited for each promotion. Columns 1 and 2 concern the sample of judges who reached intermediate stage (Promotion #2) between 2008 and 2022. The outcome in these columns is the median wait time from the judges’ starting day until their second promotion. The sample in Columns 3-6 is the sample of judges who reached the final stage (Promotion #3) between 2008 and 2022. The outcome in Columns 3-4 is the median time judges waited from their second to their third promotion. The outcome in Columns 5-6 is the median time judges waited from their first day as judges until their third promotion. Column 2 includes year dummies for the year judges were promoted for the second time. Columns 4 and 6 include year dummies for the year judges were promoted for the third time.

The second way I verify a gender promotion gap is by showing that female judges are less likely to receive a promotion to the next career stage in a given amount of time. Here, I focus on the set of judges who reached the immediately preceding career stage from 2007 to 2012. Table 8 shows the percentage of men and women who received the promotion listed in Column 2 in the time window indicated in Column 3. The time windows were selected to depict the first semester where more than

²¹I omit the wait time from entry to the first promotion for conciseness. Most judges receive their first promotions quickly, so gender differences in wait time are less consequential.

50% of the men received the career advancement in Column 2. Although the first two promotions in a judge’s career take a similar amount of time for both genders, the third promotion is not as egalitarian. Among the judges who reached the intermediate stage between 2007 and 2012, there is a large gender gap in the likelihood of being promoted for the third time within 10 years. While 53.5% of men waited less than a decade between their second and their third promotions, the same was true for only 34.8% of women.

Table 8: Gender and Timing of Promotions in the Trial Court

Cohort	Advancement	Time Window (years)	Men Promoted %	Female-Male (p,p)		
				Coefficient	Min (95% CI)	Max (95% CI)
Start 2007-2012	Promotions 1-2	4.5	55.2	-5.4	-14.3	3.4
Intermediate 2007-2012	Promotion 3	10.0	53.5	-18.7	-28.4	-9.1

Notes: The time window was chosen as the first semester when more than 50% of men from the cohorts described in Column 1 completed the career advancement stage in Column 2. All regressions include year fixed effects for the year when judges started the immediately preceding career stage (i.e., the beginning of their judicial careers in row 1 and the day they were promoted to the intermediate stage in row 2).

Few women are promoted to Appeals Court. Part of this is due to female judges not reaching the stage of their careers when they become eligible for this promotion. However, Table 9. shows that eligible women are still four times less likely than their male counterparts to reach the Appeals Court. While an average of 2.1% of eligible male judges reach the Appeals Court every year, only 0.5% of eligible women advance to that stage. This means that the Appeals Court will continue to be male-dominated for the foreseeable future unless there are considerable changes to promotion criteria.

Table 9: Promotion to Appeals Court by Gender (Annual Average 2007-2022)

	Average Promoted (#)	Average Eligible (#)	Promoted/Eligible
Women	1.86	351.14	0.53%
Men	14.6	695.34	2.10%

Notes: The first column of this table shows the average number of judges of each gender promoted to the Appeals Court every year in the 2007-2022 period. The second column shows the number of judges eligible for promotion to the Appeals Court (i.e., judges in the final stage of the Trial Court at the end of every year from 2007 to 2022). The third column is the ratio of the first and the second columns.

The remainder of the paper investigates the reasons for this promotion gender gap, simulates the

impact of each channel at each point in the judges' career (Section 5), and discusses the constraints behind judges' choices (Section 6).

5 Main Results: Explaining the Gender Promotion Gap

The main results of this paper are divided into two parts: identifying the drivers of the gap and quantifying their importance.

In section 5.1, I analyze the labor-supply and labor-demand channels that drive the gender promotion gap and find that both are important: women are more selective in seeking career advancement, partially because of costs associated with accepting a promotion, and they are less likely to be successful than equally senior men when they do apply for promotion. The latter is entirely driven by merit promotions and cannot be explained by observable on-the-job productivity and application patterns.

In section 5.2, I simulate judges' 30-year career trajectory to describe the long-run impacts of each driver of the gap. The probability of applying conditional on eligibility and the probability of being promoted conditional on applying contribute almost equally to the gender gap. Together, these two drivers explain over 80% of the difference between the number of promotions given to men and women by year 11, the year when a substantial fraction of men start getting their third promotion. Although women are eventually promoted for the third time, they are still subject to a virtually impenetrable glass ceiling at the Appeals Court. By their thirtieth year in the judiciary, the gender gap reaches a second peak, which could almost entirely disappear if women applied for promotion similarly to men throughout their entire careers.

5.1 Supply and Demand Drivers of the Gap

In this section, I analyze the supply- and demand-side drivers of the promotion gender gap based on what happens at each promotion opportunity. One advantage of this setting is that the same promotion opportunity typically has sets of vacancies selected by seniority and merit criteria. The promotion criterion for each slot is quasi-random and, therefore, independent of any specific vacancy's characteristics. Article 93 Part II of the Brazilian Constitution states that the courts should promote judges to new vacancies using seniority and merit criteria alternately. That is, if one vacancy is awarded by seniority, the next one must be chosen by merit, and so on. Table 10 shows that the choice of seniority as a selection mechanism is indeed uncorrelated with vacancy characteristics. The fact that the set of choices in seniority and merit is not systematically different has an important implication:

if a male and a female candidate apply for the same promotion opportunity and are equally likely to be promoted by seniority, but the woman is less likely to be awarded on the basis of merit, the gender gap in merit can be attributed to the differences between these selection mechanisms and the economic incentives they generate. In other words, the gap is not a result of reactions to consistently different sets of choices. Likewise, systematic gender differences in the probability of eligible judges seeking career advancement to vacancies selected by seniority and merit can be attributed to the criteria themselves.

Table 10: Independence of the Selection Criterion from Vacancy Specialization

	Chosen by seniority Criterion	
	Mean	Standard Deviation
Criminal	-0.04	0.03
Infancy and Youth	-0.06	0.06
Civil	-0.01	0.03
Family matters	-0.01	0.05
Business	-0.19	0.22
Public Treasury	0.03	0.05
Num. obs.	5,062	
F-statistic(proj model)	0.8715	
Career Phase FE	✓	
Date FE	✓	

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$

Notes: This table shows the correlation between the probability of a vacancy being filled by seniority and the different characteristics of the vacancy. The low value of the F-statistic of the projected model (with position and date fixed effects) indicates that one cannot reject the null hypothesis that all the vacancy characteristic coefficients are zero.

5.1.1 Labor-Demand Drivers

Table 11 shows the gender differences in the probability of an applicant receiving a promotion by gender and criterion. For promotions inside the Trial Court, female candidates are 20.68% (3.22 p.p.) less likely to receive a merit promotion and equally likely to receive a seniority promotion as male applicants of similar experience and tenure in their current career stage. Overall, female applicants are 10.31% (2.53 p.p.) less likely to succeed in applications for promotion. In the Appeals Court, there is no significant difference between the likelihood of equally senior men and women being promoted.

However, because there are fewer promotions at that stage, the regressions have limited statistical power to identify statistically significant effects.

Table 11: Probability of Promotion Conditional on Candidacy and Selection Criterion

	Trial Court			Appeals Court		
	Promotion Merit	Promotion Seniority	Promoted	Promotion Merit	Promotion Seniority	Promoted
Fem	-3.22*** (0.79)	-0.07 (0.81)	-2.53*** (0.84)	-1.26 (2.10)	-0.03 (2.03)	-1.27 (2.40)
mean men	15.57	13.90	24.53	7.86	7.83	13.97
Sample	Cand Merit	Cand Seniority	Cand	Cand Merit	Cand Seniority	Cand
Experience/Tenure Ctrl	✓	✓	✓	✓	✓	✓
Date FE	✓	✓	✓	✓	✓	✓
Career Stage FE	✓	✓	✓			
Num. obs.	8339	7757	9872	1662	1699	1896

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

Notes: The first three columns concern promotions within the trial court (substitute to initial, initial to intermediate and intermediate to final stage), while the last three columns concern promotions from the final stage of the trial court to the appeals court. The outcome of the first and fourth columns is whether a judge received a merit promotion while the sample consists of judges who applied for a merit promotion. The outcome of the second and fifth column is whether a judge received a seniority promotions and the sample consists of judges who applied for seniority promotions. The outcome of the third and the sixth column is whether a judge received any promotion (seniority or merit) and the sample consists of all the judges who applied for a promotion. All regressions include fixed effects for the date of the promotion opportunity and experience and tenure controls. Regressions in columns 1-3 include fixed effects for the judges' current position.

Results on merit promotions are robust to introducing controls for other factors that may make candidates less likely to be successful, such as applying for a lower number of positions and on-the-job productivity. The sample for the regressions in Table 12 is the set of candidates in each promotion opportunity from January 2013 to December 2022.²² Regressions include controls for productivity, the number of options selected by each candidate for merit promotion, and the number of options selected for seniority promotion. As shown in Table 12, the gender difference is exacerbated by the inclusion of these controls both in the Trial and in the Appeals Court.

²²This is a subset of the data in Table 11 for which it is possible to calculate the productivity measure that maximizes the likelihood of promotion for men.

Table 12: Probability of Promotion by Merit Conditional on Candidacy with Productivity and Vacancy Choice Controls (2013-2022)

	Trial Court		Appeals Court	
	Promotion Merit	Promotion Merit	Promotion Merit	Promotion Merit
Fem	-3.60*** (0.72)	-2.13** (0.95)	-4.82** (2.33)	-1.97 (2.44)
Prod RF	18.49*** (0.31)		16.05*** (1.59)	
# Choices Merit	0.28*** (0.07)		-5.04 (26.27)	
# Choices Seniority	-0.49*** (0.13)		3.89 (26.00)	
mean men	11.32	11.32	7.49	7.49
Sample	Cand Merit	Cand Merit	Cand Merit	Cand Merit
Ten/Exper Ctrls	✓	✓	✓	✓
Date FE	✓	✓	✓	✓
Career Stage FE	✓	✓		
Num. obs.	4585	4585	949	949

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

Notes: Regressions of merit promotions on observable variables. Each observation is a candidate to merit promotions taking place from January 2013 until 2022, the time span for which I can calculate the productivity measure that best explains male promotions (Prod RF). Controls also include the number of vacancies each candidate applied to in merit and seniority promotions. The first and second columns contemplate all merit promotions within the trial court from January 2013 to December 2022. The third and fourth columns contemplate all merit promotions from the final stage of trial court to the appeals court from January 2013 to December 2022.

In Appendix A.1.2, I discuss how the addition of control variables impacts the probability of promotion by seniority and of promotion overall. I show that women are still less likely to be promoted overall, which is entirely explained by the fact that they are not chosen as often through the merit criterion. Further, seniority promotions are not positively impacted by productivity. This is consistent with the structure of seniority promotions which by construction do not take performance into account.

5.1.2 Labor-Supply Drivers

Table 13 displays the gender difference in the probability of applying for a promotion when eligible. Among judges eligible for promotion in the Trial Court, women are 14% (3.89 p.p.) less likely to apply than men of similar experience and tenure at their current position. There is no substantial difference between merit and seniority criteria: women are 17.3% (4.14 p.p.) less likely to apply for at least one merit vacancy and 17.8% (3.95 p.p.) less likely to apply for at least one seniority promotion. There is no significant difference in the probability of applying to the Appeals Court conditional on spending a similar amount of time in the final stage and in the justice system overall.

Table 13: Probability of an Eligible Judge Becoming a Candidate

	Trial Court			Appeals Court		
	Apply Merit	Apply Seniority	Apply	Apply Merit	Apply Seniority	Apply
Female	-4.14*** (0.42)	-3.95*** (0.41)	-3.89*** (0.44)	-0.04 (0.08)	-0.05 (0.08)	-0.03 (0.08)
Mean Men	23.88	22.22	27.75	1.96	2.00	2.22
Sample	Eligible	Eligible	Eligible	Eligible	Eligible	Eligible
Experience/Tenure Ctrls	✓	✓	✓	✓	✓	✓
Date FE	✓	✓	✓	✓	✓	✓
Career Stage FE	✓	✓	✓	✓	✓	✓
Num. obs.	37935	37935	37935	111197	111197	111197

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

Notes: The first three columns are the subsample of all judges eligible for promotion within the trial court at each given promotion opportunity. The outcome of the first column is whether the judge offered themselves as a candidate for at least one vacancy chosen by merit promotion, the second column outcome is whether the judge was a candidate for seniority promotion and the third column outcome is whether the judge was a candidate for any promotion. The last three columns depict the same outcomes, but for the subsample of judges eligible for promotion to the appeals court. Regressions include date fixed effects. The regressions in the first three columns also include career stage fixed effects.

Nevertheless, men and women rarely compete for the same slots in the Appeals Court. From October 2007 to December 2022, women applied for the Appeals Court 250 times, while men applied 1646 times.²³ This difference is explained by the fact that judges typically only seek advancement to

²³See Appendix A.1.3 for a breakdown of candidacies and promotions of men and women to the Appeals Court by

the Appeals Court after spending many years in the final stage, which is a rarer occurrence for female judges. The average tenure of candidates is remarkably similar across genders at slightly above 12 years.²⁴ However, eligible men have spent, on average, 1.34 more years in the highest stage of the Trial Court than their female counterparts, as shown in Table 14. This pattern is consistent with women taking longer in between each promotion within the Trial Court and with women being able to retire earlier with a full pension. I show in Appendix A.1.3 that, when excluding the control variable for tenure in the final stage, women are significantly less likely to apply for promotion to the Appeals Court than men with the same number of years in the justice system.

Table 14: Years in Current Career Stage at Each Promotion Opportunity

	All Final	Candidates for Promotion		
		All Candidates	Candidates Merit	Candidates Seniority
Female	-1.34*** (0.02)	0.17 (0.13)	0.15 (0.14)	0.15 (0.14)
Mean men	7.30	12.59	12.65	12.56
Date FE	✓	✓	✓	✓
Num. obs.	111,197	1,896	1,662	1,699

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

Notes: This table shows the gender difference in tenure in the final stage, measured as the number of years between the day each judge got promoted to the final stage and the date of each opportunity for promotion to the Appeals Court. Each column includes a different sample of judges included in the analysis. Column 1 features all judges in the final stage, column 2 features all candidates, column 3 features all candidates for merit promotion, and column 4 shows all candidates for seniority promotion.

While all promotion vacancies to the Appeals Court are the same, this is not the case in the Trial Court. Men and women may, therefore, vary in how they rank their desired slots. The most visible and consequential way a judge can compare two vacancies a and b is by applying to a and not b — that is, ranking a as an acceptable alternative, and not b . This is indicative of which opportunities a judge found worth pursuing based on their outside option, the value they place on the vacancy, and the probability of being promoted to each vacancy offered. I investigate which vacancy features encourage (or discourage) men and women to apply given their career trajectories and previous job postings.

selection criterion.

²⁴Given that the average amount of time from the first day of work until promotion to the final stage is, on average, nine years, this entails a considerable amount of experience in the justice system as a whole.

Appendix A.2 discusses details of the mechanism for allocating judges to vacancies and sufficient conditions to map judges' application choices to their actual preferences for slots.

The vacancy features considered in this analysis are specialization in a given area of law, the position's proximity to the capital, and whether it is close to each candidate's workplace as of December 31st of the previous year.²⁵ To capture the idea of closeness between courthouses (e.g., the absence of significant moving costs if the position is accepted), I link each courthouse to its corresponding municipality. I then define two cities as "close" if they are located in the same immediate region according to the Brazilian Institute of Geography and Statistics (IBGE). Immediate regions are municipalities within an urban network structured around a city that satisfies the immediate needs of the population, such as the purchase of goods, services such as health and education, and jobs. I define judges who live in the same immediate region as the state capital as "Capital" and the other judges as "Out-of-Capital".²⁶

Table 15 is constructed using a sample of candidate-vacancy pairs. It measures the probability of a judge applying for a given promotion slot conditional on having applied to at least one vacancy offered in the same promotion opportunity. I split the sample into judge-vacancy pairs in which the judge's current position is outside the immediate region of the capital and judge-vacancy pairs in which the judge's current position is within the immediate region of the capital. I do this because a judge whose current position is near the capital would be able to accept the promotion without incurring moving costs and while maintaining geographical proximity to Appeals Court Judges (as opposed to a promotion that would bring them further from the people who will decide on their future career advancement). If the judge was not posted in the capital, I can separate the effects of having a preference for the capital vicinity and facing a moving cost. All regressions include career-stage-by-date fixed effects to capture the variation within each promotion opportunity where the set of available options is fixed.

The first two columns in Table 15 suggest that moving costs are an important driver of the gender difference. For female candidates outside the capital, the likelihood of applying for a vacancy in the same immediate region is similar to a man's. In other words, the sum of the female coefficient with the interaction of female and same region is close to zero. However, women are less likely to apply to

²⁵I use the judges' posting on December 31st instead of the court districts where they worked each month as substitute judges can work in multiple districts, but their official posting dictates where they have to reside.

²⁶Immediate and Intermediate regions are updated versions of micro-regions and meso-regions of Brazil, territorial divisions made in 1989. Sources used to construct this division include data from the late 2000s and earlier 2010s. See IBGE (2017) for a detailed explanation.

vacancies that involve moving outside their current posting, represented by the statistically significant and negative coefficient on the female indicator variable. Although male judges are more interested in the capital vacancies, the same is not true for women. However, the difference in interest amounts to access to jobs that specialize in given areas of law. From the second to the third columns of the table, I add controls for the interaction of specialization dummies with the indicator for female to capture judges of different genders reacting to specialization with distinct levels of enthusiasm. The interaction of female and capital in the third column captures the gender difference in interest in capital vacancies when the vacancy is unspecialized. In that case, women are more interested in going to the capital than are men. For judges in the capital region, women are less interested than their male counterparts in promotions both within the immediate region and outside the region, regardless of the vacancy's specialization. It is important to note that male candidates posted in the capital are more likely to apply aggressively; they apply to 51.1% of vacancies, whereas men outside the capital only apply to 37.92% of vacancies. One possible explanation for this pattern is that men are more likely to move to the capital to pursue a job they like, especially if it has a specialization they prefer. On the other hand, women are more interested in the capital as a location, possibly as a reflection of better networking opportunities and opportunities for spouses and children.

Table 15: Likelihood of a Candidate Applying for a Vacancy

	Apply	Apply	Apply	Apply	Apply	Apply
Fem	-5.54*** (0.34)	-5.54*** (0.34)	-1.65*** (0.40)	-9.61*** (0.47)	-9.59*** (0.46)	-6.22*** (0.81)
Capital	2.84*** (0.35)	2.70*** (0.35)	1.44*** (0.35)			
Same Immediate region	41.23*** (1.05)	41.23*** (1.05)	40.99*** (1.05)	33.57*** (0.42)	33.20*** (0.42)	32.82*** (0.42)
Fem x Same Immediate region	5.17*** (1.66)	5.09*** (1.66)	5.99*** (1.66)	-6.89*** (0.59)	-6.93*** (0.59)	-6.15*** (0.61)
Fem x Capital	-2.13*** (0.53)	-2.12*** (0.53)	1.66*** (0.57)			
mean men	37.92	37.92	37.92	51.11	51.11	51.11
Sample	Candidates Out-of-Capital	Candidates Out-of-Capital	Candidates Out-of-Capital	Candidates Capital	Candidates Capital	Candidates Capital
Tenure/Experience ctrls	✓	✓	✓	✓	✓	✓
Specialization ctrls		✓	✓		✓	✓
Specialization X Fem ctrls			✓			✓
Current Career Stage x date FE	✓	✓	✓	✓	✓	✓
Num. obs.	116490	116490	116490	85472	85472	85472

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

Notes: The sample for this table are judge-vacancy pairs in which each judge applied to at least one vacancy. In Columns 1-3, regressions are done with the sub-sample of judge-vacancy pairs in which the judge’s position on December 31st of the previous year is outside the immediate region of the state capital. In Columns 4-6, regressions are done with the sub-sample of judge-vacancy pairs in which the judge’s position on December 31st of the previous year is within the immediate region of the state capital. All regressions include career-stage-by-date fixed effects to capture judges’ choices within the same promotion opportunity.

5.2 Long-run Impacts of the Promotion Gender Gap and its Drivers

I show in section 5.1 that the promotion gender gap is due to both limited supply of female judges seeking promotions and lower demand for them in higher positions. I now estimate how these differences impact the evolution of the promotion gender gap over 30 years, approximating the duration of a judge’s career from initial hiring to retirement. I abstract away from any pre-existing differences and cohort effects to answer the following questions: how many more promotions would men accumulate over time under the current application and promotion regime? How much of this difference can be explained by judges’ application, attrition, and likelihood of being promoted?

I simulate judicial career trajectories based on the following idea: suppose 1,000 identical judges of each gender start their careers at the same time at the end of year 0. These identical judges

are compared monthly to their real counterparts who are similar to them in terms of gender, career stage, and number of years of experience in both the judiciary and in their current career stage. Opportunities for career advancement appear randomly every month over the course of 30 years.²⁷ At each opportunity, the simulated judges face the same probability of applying for a promotion and getting chosen for the next career stage as their real counterpart judges. They also face the same probability of exiting the judicial system at the end of every year. Exit is permanent and represents the end of their career. Once they leave, they can no longer apply for (or receive) a promotion, and the total number of promotions they have accumulated is frozen. I then compute the number of times each judge has been promoted from year 0 until the end of each year $t \in \{1, \dots, 30\}$. See Appendix A.4 for details on the simulation algorithm.

Figure 1 shows the promotion gender gap in two important ways. On the left-hand side, it displays the average number of promotions accrued by men and women at the end of every year. The right-hand side displays the difference in the number of male and female judges who have reached each career stage (by accumulating a minimum number of promotions) by the end of each year.

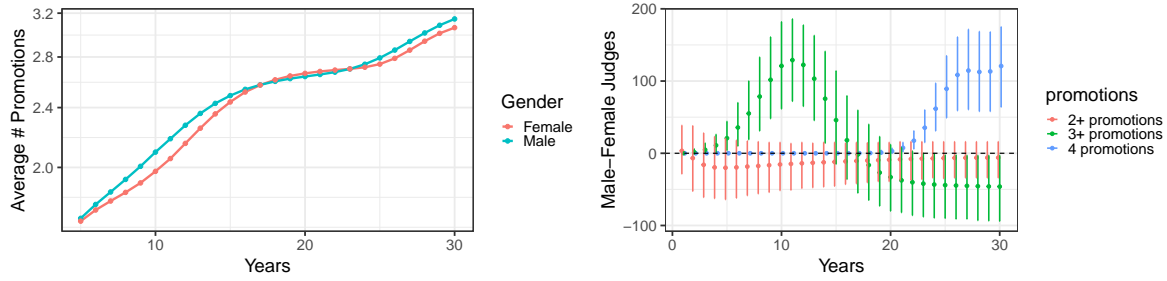
The gap in the average number of promotions for men and women is relatively modest,²⁸ with both series showing strong time trends. This is a result of the use of the seniority criterion for 50% of promotions.²⁹ The graphs show that the promotion gender gap widens from years 5-15, closes, and then widens again after year 23. The first widening is due to women taking longer to attain their third promotion (from the intermediate to final stage), while the second divergence is due to gender disparity in promotions to the Appeals Court. The gender gap therefore peaks in year 11 of the simulation, where there are 129 more men than women who have reached at least the final stage of Trial Court promotions (3 or more promotions), and in year 30, where there are 120.8 more men than women who have been promoted to the Appeals Court (4 promotions).

²⁷The likelihood of one of these opportunities happening every month is designed to match the annual probability of a judge moving from one career stage to the next.

²⁸Note that a 0.1 difference is equivalent to 10% of the judges receiving 1 fewer promotion.

²⁹See Section 5.1 for a longer discussion on how the seniority criterion is an equalizing factor.

Figure 1: Evolution of the Promotion Gender Gap



Notes: The graphs in this figure are outputs for a 30-year simulation of the career trajectories of 1000 male and 1000 female judges who enter the justice system at the same time (year 0) and face probabilities of exit, candidacy, and promotion compatible with real judges of their gender and career stage. The graph on the left shows the average number of promotions by year and by gender. The graph on the right shows the difference in the number of male and female judges in each career stage in any given number of years with 95% confidence intervals.

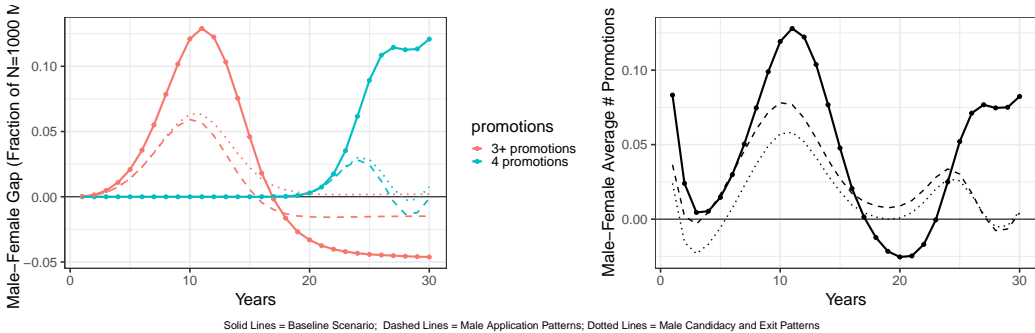
Figure 2 decomposes the gap into three sources of gender disparity: 1) probability of promotion conditional on application, 2) likelihood of seeking career advancement, and 3) exit patterns (e.g., retirement or career change). This is attained by contrasting the baseline scenario (solid lines) and the two counterfactual exercises. In the first exercise, shown in dashed lines, all probabilities are equal to baseline with the exception that female judges are assigned the same probability of applying for promotions as their male counterparts with similar career trajectories. The difference between the gender gap verified in this scenario (dashed) and the baseline (solid) allows us to distinguish the effect of application probabilities. In the second exercise, shown in dotted lines, female judges have the same probability of applying for promotion and exiting the justice system as their male counterparts conditional on having the same career path. Hence, this exercise only allows men and women to differ in the likelihood of receiving a promotion conditional on applying, thereby capturing the importance of this driver to the gender gap.³⁰ The role of differential job attachments can be seen by the difference between the dashed and the dotted curves. Here, the only difference between the two counterfactuals is that the dashed curves allow women to decide differently on whether to leave their jobs.

By the end of year 11, when the difference between the number of men and women receiving their third promotion peaks, 40.60% of the gender gap is explained by different probabilities of promotion, 39.95% by different probabilities applying for the promotion, and the remaining 19.45% by differential

³⁰If women were equally likely to be promoted conditional on applying, to apply for promotion conditional on eligibility, and to be eligible, there would be no reason for the simulation to yield any gender gap at all because all judges are ex ante equal.

job attachment. However, when it comes to getting the promotion to the Appeals Court, application decisions can explain the lion’s share of the effect. In year 30, 94.95% of the gap is explained by women’s decision to (not) apply for posted vacancies. This is intuitive — if the gender promotion gap at the final stage (the third promotion) is smaller and shorter-lived (as in the exercises represented by the dotted and dashed lines), then female judges are expected to catch up and get to the Appeals Court by year 30. If the world works as in the baseline scenario (solid lines), however, the gap remains high until their expected retirement age. Women are therefore “timed out” of their fourth promotion.

Figure 2: Decomposition of the Promotion Gender Gap



Notes: This figure shows outputs relating to three scenarios of a career trajectory simulation where 2000 identical judges, 1000 men and 1000 women, enter the justice system at the end of year zero. The solid lines represent the baseline scenario, the dashed lines represent a scenario in which women make decisions to apply for a promotion in the same way men do throughout their careers, and the dotted lines represent the counterfactual in which women make both exit and candidacy decisions in the same way as men do. In this second counterfactual, women only differ from men in the likelihood of being promoted conditional on applying. The graph on the left shows the gender difference in the average number of promotions by the end of each year. The graph on the right shows the difference in the number of male and female judges who accumulate each given number of promotions by the end of the year in the x-axis.

6 Discussion: Constraints, Exogenous Shifters, and the Gender Gap

In the previous section, I discuss how and to what extent the gender gap is explained by attrition, supply, and demand for female judges at higher career stages. However, all of these components are the results of endogenous choices: seeking career advancement, leaving their current jobs, and being more productive to make themselves better candidates for merit promotion. In this section, I discuss if and how two exogenous factors – responsibilities at home and the sunk costs of application — constrain these choices.

The increase in women’s responsibilities at home following the birth of a child is a major driver of

gender disparities in career advancement in multiple settings.³¹ Even in the absence of significantly low job attachment (as in the setting studied in this paper), this shift in priorities could still lead to lower productivity (Azmat and Ferrer, 2017). I discuss this possibility in Section 6.1 and find no evidence that women are less productive overall or specifically during their child-bearing years. Hence, application choices are the remaining way by which spousal and parenting responsibilities at home can affect women differently from men.

Sunk costs of application can influence how a change in the likelihood of an applicant receiving a promotion (e.g., because of the imposition of a quota) can affect the gender promotion gap. If judges face a cost when applying regardless of the outcome (such as a fear of exposure to evaluation and scrutiny by Appeals Court judges), they may be reluctant to apply when they perceive their chances are low. Alternatively, they may choose to wait for the next opportunity if their probability of success is expected to increase in the near future (such as by crossing some tenure threshold). In Section 6.2, I show suggestive evidence that judges indeed consider their present and future odds when deciding whether to apply for promotion in addition to comparing the quality of their current job to posted vacancies. In that case, the effect of a change in promotion mechanisms can go beyond a mechanical impact on the demand for male and female judges in higher positions to affect their likelihood of seeking career advancement.

6.1 Additional Evidence of Gender Differences in Productivity

Next, I investigate if 1) women are different in their productivity; 2) their output is consistent with them putting in fewer hours or less effort into their work than their male counterparts; 3) they are less competitive candidates for merit-based promotions based on their productivity; and 4) the results change substantially during childbearing years. The answer for all four inquiries is no.

Table 16 shows gender differences in the two main measures of productivity in the paper: the best predictors of merit promotions for male judges and the number of cases disposed of by each judge in a given month (Dahis et al., 2020). The women eligible for promotions within the Trial Court (i.e., judges who are currently substitutes or in initial- and intermediate-stage court districts) are slightly more productive than their male counterparts. On the other hand, female judges eligible for Appeals Court promotion were somewhat less productive in terms of cases disposed.

³¹See Cukrowska-Torzewska and Matysiak (2020) for a meta-analysis on the topic.

Table 16: Productivity by Gender

Sample:	Outcome: Productivity Random Forest (z-scored)		Outcome: Cases Disposed (Cases/Month)	
	Eligible Trial	Eligible Appeals	Eligible Trial	Eligible Appeals
Fem	0.02*** (0.01)	0.00 (0.00)	1.23** (0.62)	-5.42*** (0.96)
mean_men	0.49	-0.41	11.01	53.20
Area of Law Controls			✓	✓
Month-Year FE			✓	✓
Career Stage FE	✓	✓	✓	✓
Preference Dates FE	✓	✓		
Court District FE	✓	✓		
Num. obs.	85565	105118	463629	261558

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

Notes: The table was constructed using productivity measures from judges eligible for promotion within the Trial Court and for promotion to the Appeals Court. The outcome of the first and second columns is the two-year productivity measure constructed with the random forest algorithm to predict men's promotions. The sample consists of every eligible judge at every promotion opportunity and includes fixed effects for the date preferences were stated and the eligible positions. The third and fourth columns' outcome is the number of disposed cases, the main productivity measure in Dahis et al. (2020). I use month-year, court district, and career-stage fixed effects, as well as controls for whether the judge's legal specialization (e.g., criminal cases, family matters, civil cases).

I use three monthly measures of effort to capture the volume of work done by a judge: the number of hearings, the number of other decisions (except for sentences), and the number of sentences (final or not). Table 17 shows the gender difference in work volume. Notably, this difference goes in the same direction as the productivity difference.

Table 17: Volume of Work By Gender

	Trial Court			Appeals Court		
	Hearings	Decisions (ex. sent)	Sentences	Hearings	Decisions (ex. sent)	Sentences
	Per Month	Per Month	Per Month	Per Month	Per Month	Per Month
Fem	1.33*** (0.09)	10.05*** (2.06)	3.52*** (1.31)	-0.27** (0.12)	-22.62*** (2.86)	-8.80*** (1.31)
Mean Men	11.01	178.65	47.79	15.76	365.67	78.02
Sample	Eligible	Eligible	Eligible	Eligible	Eligible	Eligible
Area of Law Controls	✓	✓	✓	✓	✓	✓
Month-Year FE	✓	✓	✓	✓	✓	✓
Career Stage FE	✓	✓	✓			
Court District FE	✓	✓	✓	✓	✓	✓
Num. obs.	463629	463629	463629	261558	261558	261558

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

Notes: This table includes measures of intensity of work for Trial Court judges eligible for Trial Court promotion (columns 1-3) and promotion to the Appeals Court (columns 4-6). The monthly outcomes are the number of hearings, decisions other than sentences, and the total number of sentences (final or not). Regressions include month-year, court district, and (for judges eligible for promotion within the Trial Court) current position fixed effects. I include controls for specializations in a given area of law.

Taken together, the results on productivity and effort show minimal differences between male and female judges doing the same job. Furthermore, because the Prod RF measure captures a judge’s competitiveness for promotion, it is clear that women are *not* self-selecting into jobs less likely to be acknowledged.

Another concern is that women may have a decline in productivity during child-bearing years, generating an irreversible gender gap. Table 18 shows that, although being of child-bearing age affects women more negatively than it does men, the effect is smaller in magnitude than the female coefficient. Hence, women of child-bearing age are not worse productivity-wise or effort-wise than their male counterparts.

Table 18: Monthly Measures of Productivity and Volume of Work by Gender and Age

	Cases Disp	Hearings	Decisions (ex sentences)	Sentences
Fem	4.22*** (1.01)	-0.01 (0.24)	23.96*** (6.24)	7.99*** (1.45)
Less 40 y.o.	0.09 (0.79)	0.34* (0.18)	-10.82** (4.88)	1.25 (1.13)
Fem x Less 40 y.o.	-2.34** (1.15)	0.05 (0.27)	-13.03* (7.12)	-5.58*** (1.65)
mean_men	31.65	11.01	178.65	47.79
Sample	Eligible	Eligible	Eligible	Eligible
Admission Exam	✓	✓	✓	✓
Court District FE	✓	✓	✓	✓
Month-Year FE	✓	✓	✓	✓
Num. obs.	271,923	271,923	271,923	271,923

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

Notes: This table is constructed using data from all judges whose date of birth is known. (Data availability depends only on the judge's entry year.) Regressions include admission exam, month-year, and court district fixed effects. The first column's outcome is the number of disposed cases, the main productivity measure in Dahis et al. (2020). The outcome in columns 2-4 are measurements of the volume of work: the number of hearings conducted, the sum of other decisions made (such as intermediate decisions and dispatches a judge signs throughout cases), and the sum of final and non-final sentences written

Hence, the only way responsibilities with children and partners can impact women more intensely than men is through candidacy decisions. Indeed, results in subsection 5.2 provide supporting evidence that women face a higher cost of moving when accepting promotions. This is consistent with the findings in Jayachandran et al. (2023). They find that heterosexual couples' re-location decisions in Germany and Sweden are more likely to be motivated by career opportunities for the husband. This can be thought of as the couple placing a larger weight on a dollar earned by the man than on a dollar earned by the woman.

6.2 Sunk Costs of Application

Apart from valuing vacancies differently and having distinct opinions about outside options, judges could face individual costs for participating in a promotion opportunity – even if they are unsuccessful.

Finding and applying for vacancies is not hard given that it is an online process widely disclosed by the court system. However, publicly applying for promotion and having that request evaluated by a prominent group of Appeals Court judges entails a certain degree of exposure. Suppose a judge of any gender faces this kind of sunk cost when applying. That means, when evaluating potential promotion opportunities, they will consider their likelihood of getting promoted in the current opportunity and the value of waiting for the next one, possibly when seniority promotion is more likely.

The promotion mechanism has the following features I explore: 1) a judge with less than two years' tenure in the current career stage can only be considered for promotion (seniority or merit) if no other judge wants that promotion; 2) promotions to the Appeals Court only occur in the data after judges have spent at least 11.71 years in the final stage; and 3) after a given number of years of tenure in the current career stage, waiting for the next opportunity will not substantially improve the likelihood of promotion because judges are already very senior.

In a model shown in Appendix A.3, I show that:

- i) If a judge does not face any application sunk costs, candidacy should be entirely driven by finding a sufficiently good vacancy, where “good enough” can be described by characteristics of the posted vacancies (dummies for specialization in a given area of law, immediate region where it is located, and whether it is in the same immediate region as the judge's current job).
- ii) If application costs matter, then the relationship between productivity, tenure, and candidacy should be as follows. Before a threshold number of years of experience (I use two years for Trial Court promotions and 11.71 years for Appeals Court promotions), waiting has a high return, especially during the months leading up to these thresholds. The return to waiting is especially high for high-productivity judges, making lower-productivity judges more likely to apply during these periods. After a given number of years (I use the medium tenure of judges promoted by seniority within each career stage), the returns to waiting are low, and candidates focus more on their likelihood of current promotion now. Higher-productivity judges are therefore more likely to apply at that point than before.

Table 19 supports the existence of a sunk cost to applying. This table captures the likelihood of an eligible judge applying to each promotion opportunity given the presence of jobs with each characteristic in the application process, the characteristics of the judge's previous job(s), and the judge's tenure and productivity at each opportunity. While columns 1 and 2 are constructed using

a sample of judges within six months of the minimum tenure thresholds (2 years for judges eligible for promotion within the Trial Court and 11.71 years for judges eligible for promotion to the Appeals Court), columns 3 and 4 include all eligible judges in each corresponding stage. All regressions include judge and career-stage fixed effects.

Table 19: Relationship Between Application Decisions and Characteristics that Affect Likelihood of Promotion - Test for the Existence of Application Sunk Cost

	Candidate Trial Court	Candidate Appeals Court	Candidate Trial Court	Candidate Appeals Court
Tenure	4.44 (6.53)	8.72*** (2.02)	1.31*** (0.16)	-0.10*** (0.02)
Tenure > Min Tenure	6.93* (3.55)	4.02*** (1.12)	1.82** (0.87)	4.32*** (0.16)
Productivity RF	-5.36*** (1.63)	-3.68** (1.46)	-2.12*** (0.48)	-0.68*** (0.08)
Productivity RF × (Tenure > Min Tenure)	1.47 (1.59)	6.54*** (0.88)	-4.23*** (0.67)	2.50*** (0.18)
Productivity RF × (Tenure > Median Tenure Seniority Promotion)			5.90*** (0.65)	-0.75*** (0.21)
Sample	Eligible Tenure ∈ [1.5y, 2.5y]	Eligible Tenure ∈ [11.25y, 12.25y]	All Eligible	All Eligible
Judge FE	✓	✓	✓	✓
Controls Job Last 2y	✓	✓	✓	✓
Controls Vacancies offered	✓	✓	✓	✓
Career Stage FE	✓	✓	✓	✓
Num. obs.	2235	4004	22255	76015

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

Notes: This table is constructed using the sub-sample of judges eligible for promotion at each promotion opportunity. Column 1 is constructed using all judges eligible for promotion in the Trial Court with between 1.5 and 2.5 years of tenure in the current career stage. Column 2 is constructed using all judges eligible for promotion to the Appeals Court with 11.21 to 12.21 years of tenure in the final stage of Trial Court. Column 3 is constructed using all judges eligible for a promotion within the Trial Court. Column 4 is constructed using all judges eligible for promotion to Appeals Court.

Table 20 shows that this pattern is overall consistent across genders, indicating that both men and women face a sunk cost of application.

Table 20: Test for Cost of Application by Gender

	Candidate	Candidate	Candidate	Candidate
	Trial M	Trial F	Appeals M	Appeals F
tenure>minT	6.68*** (1.03)	2.43** (1.24)	4.46*** (0.19)	2.38*** (0.18)
tenure>median	0.46 (1.36)	7.78*** (1.45)	0.22 (0.22)	0.21 (0.21)
Prod RF	-1.69*** (0.59)	-1.68** (0.83)	-0.91*** (0.12)	-0.26*** (0.09)
Prod RF*(tenure \geq minT)	-6.50*** (0.81)	-0.11 (1.15)	2.91*** (0.24)	1.55*** (0.23)
Prod RF*(tenure \geq median)	8.88*** (0.79)	-0.82 (1.14)	-0.93*** (0.28)	-0.29 (0.27)
Judge FE	✓	✓	✓	✓
Controls Job Last 2y	✓	✓	✓	✓
Controls vacancies offered	✓	✓	✓	✓
Career Stage FE	✓	✓	✓	✓
Num. obs.	12510	9745	48729	27286

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

Notes: Column 1 is constructed using all male judges eligible for promotion within the Trial Court. Column 2 is constructed using all female judges eligible for promotion within the Trial Court. Column 3 is constructed using all male judges eligible for promotion to the Appeals Court. Column 4 is constructed using all female judges eligible for promotion to the Appeals Court. All regressions include judge and career-stage fixed effects.

This finding has an important policy implication. The effect of a gender quota on the gender promotion gap is not limited to the mechanical impact of changing the probability of candidate promotion. It can also differentially affect male and female judges' decision to apply.

7 Conclusion

I identify a promotion gender gap in the Brazilian judiciary and two major drivers of this gap. Together, these drivers account for more than 80% of the observed gap. Women are less likely to

apply for promotion *and* less likely to receive it conditional on application. This result is robust to the inclusion of controls for the number of vacancy applications and their observable productivity. The evidence provided in this paper shows that the criterion used for merit promotions has a disparate impact. Still, women have a lower probability of applying for promotion despite having productivity and commitment to their jobs comparable to that of their male counterparts. Although part of this effect is due to women facing higher moving costs when accepting promotions, there is suggestive evidence that judges of both genders face an individual sunk cost when applying for promotion, leading them to weigh the probability of a successful candidacy in the present and in the future when deciding to seek career advancement.

Taken together, the results indicate that a well-designed gender promotion quota might achieve the goal of promoting more women without impacting judicial productivity. Women are not less productive than their male counterparts despite the lower probabilities of promotion. Furthermore, a higher likelihood of promotion may incentivize more women to apply, allowing the justice system to promote higher-performing judges from a larger pool in which women are better represented. Investigating the potential impacts of such a policy on gender parity and productivity is an important avenue for further work with a practical application. The issue is a live one for policy. In September 2023, the Brazilian National Council of Justice established a 50% gender quota for seniority and merit promotions from the final stage of the Trial Court to the Appeals Court until parity is reached. This norm is valid for all state, federal military, and labor courts.

References

- Adams-Prassl, A., Huttunen, K., Nix, E., and Zhang, N. (2022). Violence against women at work.
- Arrow, K. J. (1974). THE THEORY OF DISCRIMINATION. In *Discrimination in Labor Markets*, pages 1–33. Princeton University Press.
- Autor, D. H. and Scarborough, D. (2008). Does job testing harm minority workers? Evidence from retail establishments*. *The Quarterly Journal of Economics*, 123(1):219–277.
- Azmat, G. and Ferrer, R. (2017). Gender Gaps in Performance: Evidence from Young Lawyers. *Journal of Political Economy*, 125(5):1306–1355. Publisher: The University of Chicago Press.
- Becker, G. S. (1971). *The Economics of Discrimination*. Economic Research Studies. University of Chicago Press, Chicago, IL.

- Benson, A. M., Li, D., and Shue, K. (2023). Potential and the gender promotion gap. *Academy of Management Proceedings*, 2023(1):19580.
- Bertrand, M., Goldin, C., and Katz, L. F. (2010). Dynamics of the Gender Gap for Young Professionals in the Financial and Corporate Sectors. *American Economic Journal: Applied Economics*, 2(3):228–255.
- Bielen, S. and Grajzl, P. (2020). Gender-based judicial ingroup bias in sex crime sentencing: Evidence from Belgium. *International Journal of Law, Crime and Justice*, 62:100394. tex.lccn: Gender.
- Boyd, C. L., Epstein, L., and Martin, A. D. (2010). Untangling the Causal Effects of Sex on Judging. *American Journal of Political Science*, 54(2):389–411.
- Cai, X., Li, P., Lu, Y., and Song, H. (2022). Do judges exhibit gender bias? Evidence from the universe of divorce cases in china. Issue: 4206653 Place: Rochester, NY.
- Cukrowska-Torzewska, E. and Matysiak, A. (2020). The motherhood wage penalty: A meta-analysis. *Social Science Research*, 88-89:102416.
- Dahis, R., Schiavon, L., and Scot, T. (2020). Selecting Top Bureaucrats: Admission Exams and Performance in Brazil. Issue: 3584725 Place: Rochester, NY Type: SSRN Scholarly Paper.
- Hessami, Z. and da Fonseca, M. L. (2020). Female political representation and substantive effects on policies: A literature review. *European Journal of Political Economy*, 63:101896.
- Hoffman, M., Kahn, L. B., and Li, D. (2018). Discretion in Hiring*. *The Quarterly Journal of Economics*, 133(2):765–800.
- IBGE (2017). *Divisão regional do Brasil em regiões geográficas imediatas e regiões geográficas intermediárias: 2017*. IBGE.
- Jayachandran, S., Nassal, L., Notowidigdo, M., Paul, M., Sarsons, H., and Sundberg, E. (2023). Moving to Opportunity, Together.
- Mocanu, T. (2023). Designing Gender Equity: Evidence from Hiring Practices and Committees.
- Ngai, L. R. and Petrongolo, B. (2017). Gender Gaps and the Rise of the Service Economy. *American Economic Journal: Macroeconomics*, 9(4):1–44.

Phelps, E. S. (1972). The Statistical Theory of Racism and Sexism. *The American Economic Review*, 62(4):659–661. Publisher: American Economic Association.

Shukla, S. (2022). Making the Elite: Top Jobs, Disparities, and Solutions.

Säve-Söderbergh, J. (2019). Gender gaps in salary negotiations: Salary requests and starting salaries in the field. *Journal of Economic Behavior & Organization*, 161:35–51.

A Appendix

A.1 Extra Tables

A.1.1 Competitiveness of Admissions Exams

Table 21: Competitiveness of Admission Exams in the São Paulo Justice System

Examination #	Year	Applicants	Selected	Selected (%)
178 ^o	2006	4,130	105	2.54%
179 ^o	2006	4,710	86	1.83%
180 ^o	2008	5,459	76	1.39%
181 ^o	2009	7,625	98	1.29%
182 ^o	2010	11,452	90	0.79%
183 ^o	2012/13	13,716	94	0.69%
184 ^o	2014	13,391	108	0.81%
185 ^o	2014/15	11,640	81	0.70%
186 ^o	2015/16	12,648	77	0.61%
187 ^o	2017/18	20,385	81	0.40%
188 ^o	2018/20	23,122	86	0.37%

Source: São Paulo Justice Diary from January 29th, 2020, page 1.

A.1.2 Application and Productivity Controls with All Promotion Criteria

In this part, I show that adding the productivity and application controls has the expected effect.

In promotions within the Trial Court, productivity is positively correlated only with merit promotions — which is consistent with seniority promotions being decided solely based on experience and

tenure at the current career stage. Furthermore, the probability of a judge being promoted by merit decreases as the judge adds more seniority options (increasing their likelihood of being promoted by seniority) and increases with the number of merit options they choose. The probability of a seniority promotion is unaffected by the number of merit choices because of a particular characteristic of the allocation of applicants to vacancies – seniority promotions are decided first (See appendix A.3 for details).

Notably, the addition of controls does not significantly affect the gender difference in seniority promotions, only the merit ones.

Table 22: Probability of Promotion within Trial Court Conditional on Candidacy with Productivity and Application Controls (2013-2022)

	Promoted Merit	Promoted Merit	Promoted Seniority	Promoted Seniority	Promoted	Promoted
Fem	-3.60*** (0.72)	-2.13** (0.95)	0.29 (0.99)	-0.20 (0.98)	-2.21** (0.89)	-1.63 (1.03)
Prod RF	18.49*** (0.31)		0.45 (0.50)		17.29*** (0.41)	
# Choices Merit	0.28*** (0.07)		-0.03 (0.10)		0.26*** (0.09)	
# Choices Seniority	-0.49*** (0.13)		0.68*** (0.18)		0.52*** (0.16)	
mean_men	11.32	11.32	10.62	10.62	18.27	18.27
Sample	Cand Merit	Cand Merit	Cand Seniority	Cand Seniority	Cand	Cand
Tenure/Experience Controls	✓	✓	✓	✓	✓	✓
Date FE	✓	✓	✓	✓	✓	✓
Career Stage FE	✓	✓	✓	✓	✓	✓
Num. obs.	4585	4585	4134	4134	5362	5362

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

Differences in the likelihood of promotions conditional on candidacy to promotions within the Trial Court (substitute to initial, initial to intermediate, and intermediate to final). Columns 1-2 show the likelihood of promotion by merit for merit candidates. Columns 3-4 show the likelihood of a seniority promotion for seniority promotion candidates. Columns 5-6 show the likelihood of a promotion for those who applied for at least one position. All samples refer to candidates who have been in the justice system for at least four semesters. T span from January 2013 to December 2022. This restriction is driven by the availability of productivity data.

A similar pattern is observed for promotions to the Appeals Court — although with more noise and a difference: now seniority promotions are negatively correlated with productivity. This effect could be due to the fact that promotions to the Appeals Court take a long time to happen — particularly so in the seniority criterion. This could mean judges who waited enough to be promoted by this criterion envision a retirement very soon or overall less intense.

Table 23: Probability of Promotion to Appeals Court Conditional on Candidacy with Productivity and Application Controls (2013-2022)

	Promoted Merit	Promoted Merit	Promoted Seniority	Promoted Seniority	Promoted	Promoted
Fem	-4.82** (2.33)	-1.97 (2.44)	1.00 (2.37)	0.07 (2.37)	-3.32 (2.74)	-1.67 (2.75)
Prod RF	16.05*** (1.59)		-5.80*** (1.68)		9.41*** (1.97)	
# Choices Merit	-5.04 (26.27)		10.16 (27.15)		5.69 (33.59)	
# Choices Seniority	3.89 (26.00)		10.93 (26.87)		15.25 (33.25)	
mean_men	7.49	7.49	7.48	7.48	13.09	13.09
Tenure/Experience Controls	✓	✓	✓	✓	✓	✓
Date FE	✓	✓	✓	✓	✓	✓
Career Stage FE	✓	✓	✓	✓	✓	✓
Num. obs.	949	949	983	983	1110	1110

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

Differences in the likelihood of promotions conditional on candidacy to the Appeals Court (final to Appeals). Columns 1-2 show the likelihood of promotion by merit for merit candidates. Columns 3-4 show the likelihood of a seniority promotion for seniority promotion candidates. Columns 5-6 show the likelihood of a promotion for those who applied for at least one position. All samples refer to candidates who have been in the justice system for at least four semesters. The sample spans from January 2013 to December 2022. This restriction is driven by the availability of productivity data.

A.1.3 Gender Gap in Promotions to the Appeals Court

Table 24 shows the breakdown of the number of applications and promotions from the final stage to the Appeals Court. Women are much less likely to seek career advancement. Conditional on trying, their likelihood of success is similar.

Table 24: Application and Promotion to Appeals Court by Gender and Selection Criterion

	Merit		Seniority		All	
	Candidates	Promoted	Candidates	Promoted	Candidates	Promoted
Women	212	15	218	14	250	29
Men	1450	114	1481	116	1646	230

Notes: Sum of all records of promotions and candidacies by gender and selection criterion from October 2007 to December 2022.

Table 25 shows the probability of an eligible judge applying for a promotion to the Appeals Court with and without controlling for tenure in the final stage. Although the application patterns of men and women are very similar conditional on both tenure at the final stage and experience, women with similar experience are roughly 20% less likely to apply. Because this analysis still conditions judges reaching the final stage and being experienced, it still does not account for cohort effects — in the past, the number of women entering the justice system was much lower — and for selection into the final stage.

Table 25: Probability of an Eligible Judge Applying for Promotion to the Appeals Court

	Apply Merit	Apply Merit	Apply Seniority	Apply Seniority	Apply	Apply
Female	-0.39*** (0.08)	-0.04 (0.08)	-0.40*** (0.08)	-0.05 (0.08)	-0.41*** (0.08)	-0.03 (0.08)
Experience	-0.67*** (0.04)	-1.08*** (0.04)	-0.66*** (0.04)	-1.07*** (0.04)	-0.78*** (0.04)	-1.23*** (0.04)
Experience ²	0.03*** (0.00)	0.03*** (0.00)	0.03*** (0.00)	0.03*** (0.00)	0.03*** (0.00)	0.03*** (0.00)
Tenure in Final		0.59*** (0.01)		0.58*** (0.01)		0.64*** (0.02)
mean_men	1.96	1.96	2.00	2.00	2.22	2.22
Date FE	YES	YES	YES	YES	YES	YES
Num. obs.	111197	111197	111197	111197	111197	111197

*** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$

Notes: This table is constructed with the sample of judges who were eligible for a promotion to the Appeals Court at each promotion opportunity. Columns 1-2 show the likelihood of an eligible judge applying for an Appeals Court vacancy selected by the merit criterion. Columns 3-4 show the probability of an eligible judge applying for a seniority vacancy. Columns 5-6 show the likelihood of an eligible judge applying for any Appeals Court vacancy. All regressions include fixed effects for the date of the promotion opportunity.

A.2 Vacancy-Candidate Allocation Algorithm and Revealed Preferences

In this section, I explain the details of the allocation mechanism that assigns judges who want to be promoted to posted vacancies in each promotion opportunity. Then, I use a theoretical framework to determine how much about the court merit ranking of judges and the judges' preferences for vacancies can be extracted from the data.

After judges submit their preferences, vacancies are filled in the order previously assigned in the promotion opportunity announcement. The best judge for each vacancy will be determined by the ranking imposed by the criterion previously established. Judges' merit and seniority rankings are independent of the characteristics of the vacancy the candidate is applying to. Furthermore, each ranking should be a complete and transitive preference order within the set of judges competing for the same career stage on the same date.

Assignment happens as in a sequential dictator game: the best judge who accepts vacancy 1 (that is, is a candidate for promotion and lists vacancy 1 as one of the options they find acceptable) gets promoted to vacancy 1. The best judge who accepted vacancy 2 and was not promoted to vacancy 1 gets promoted to vacancy 2, and so on. In the sample analyzed, all the seniority vacancies were announced before the merit ones.

That being said, it is possible to establish which judges are being directly compared — and, therefore, who was declared the best.

Definition 1 *A judge j is revealed preferred to a judge j' according to the merit criterion if:*

1. *There is n such that:*
 - *The n^{th} vacancy is selected by the merit criterion;*
 - *Judges j and j' applied for the n^{th} vacancy and were not promoted to vacancies $\{1, \dots, n-1\}$ posted before;*
 - *Either judge j gets promoted to vacancy n while judge j' does not or judge j gets shortlisted for promotion to vacancy n and judge j' does not.*
2. *There is a sequence of judges j_1, \dots, j_k such that $j \succsim j_1 \succsim j_2 \succsim \dots \succsim j_k \succsim j'$ where the order \succsim is defined by item 1.*

Regarding the judges' preferences for vacancies, the amount of information rankings can give depends on assumptions about judge preferences - and, more generally, on the game they are playing.

Let a promotion opportunity be one single posting in which vacancies $V \equiv \{v_1^s, \dots, v_N^s\}$ of a given career stage s are announced. The set of players J is comprised of all judges in the immediately preceding career stage, who can either apply or not apply for each vacancy v_1^s, \dots, v_N^s . Hence, their pure strategy space is $A = \{0, 1\}^N$. Given each strategy profile $a = (a_1, \dots, a_J) \in A^J$, judges perceive a probability $p_{jv}(a)$ of judge j being promoted to vacancy v when a is played. The uncertainty in that case comes from judges not knowing exactly how they rank in merit promotions.

Each judge's preferences are represented by expected utility $U_j : (\Delta A)^J \rightarrow \mathbb{R}$. If $\sigma \in (\Delta A)^J$ is a mixed strategy profile, then:

$$U_j(\sigma) = \sum_{\tilde{a} \in A^J} \sum_{v \in V} \sigma(\tilde{a}) p_{jv}(\tilde{a}) u_j(v) + \left(1 - \sum_{\tilde{a} \in A^J} \sum_{v \in V} \sigma(\tilde{a}) p_{jv}(\tilde{a}) \right) o_j$$

Where the function $u_j : V \rightarrow \mathbb{R}$ attributes a value to each vacancy being offered; $p_{jv}(\sigma)$ is the probability of the court promoting j to v , and o_j is j 's outside option value - which is the utility j gets

from remaining in their current job and waiting for the next promotion opportunity. Because vacancies are filled in following a serial dictator game, judge j is promoted to a vacancy v_i^s if and only if the court prefers judge j over every other judge that applied to v_i^s and was not promoted to v_1, \dots, v_{i-1}^s , which depends on application behaviour of every judge.

Under a very ambitious assumption, one can extract a lot of information from rankings.

Assumption 1 *At every promotion opportunity, each judge j faces an outside option lottery that yields expected utility $o_j > 0$ and evaluates each vacancy v' being offered with $u_j(v') \in \{0, 1\}$.*

One way to interpret the outside option lottery is the value of staying where one is and waiting for future promotion opportunities.

Proposition 1 *Suppose σ is an equilibrium strategy profile, Assumption 1 holds and j does not face a probability one of promotion, that is, $\sum_{\tilde{a} \in A^j} \sum_{m=1}^N \sigma(\tilde{a}) p_{jv_m}(\tilde{a}) < 1$. Then, the following implication holds:*

$$\text{Judge } j \text{ applies to } v \in V \text{ and not to } v' \in V \implies u_j(v) > u_j(v').$$

Proof of Proposition 1. If a judge applied to vacancy v_i with positive probability, it must be that $u_j(v_i) \geq o_j$. To see this, consider the one-shot deviation where judge j decides to apply to vacancy v_i with probability 0 (and make the same application decisions concerning every other vacancy).

If $i = N$ and j was not promoted to $\{v_1, \dots\}$ then this one-shot deviation would mean getting utility o_j with probability 1 instead of being promoted to v_N with some positive probability p (note that if the other players play all-mixed strategies, there is at least some small probability that j is the only judge competing for v_N) and staying on o_j with probability $1 - p$. That would be a profitable deviation if and only if $u_j(v_N) = 0$. Note that The only way it is optimal for j to apply to v_N is if $u_j(v_N) = 1$.

Suppose, to obtain an induction, that for every $m > i$, it is true that if $a_{jm} = 1$ implies $u_j(v_m) = 1$. Then, if j is not promoted to v_1, \dots, v_{i-1} , not applying to v_i means facing a lottery that can only give utility 1 (in case j gets promoted to any vacancy v_m for $m > i$) or utility o_j . This would be a profitable deviation if $u_j(v_i) = 0$. Hence, if $a_{ji} = 1$, it must be that $u_j(v_i) = 1$. By induction, for all $n \in \{1, \dots, N\}$ the following implication holds:

$$a_{jn} = 1 \implies u_j(v_n) = 1.$$

I will show now that, if j applied to some vacancy v , but not to v' , then $u_j(v') = 0$.

Suppose a judge j applies to a non-empty set of vacancies $V_j = \{v \in V : u_j(v) = 1\}$. Hence, j is facing a lottery in which they get either utility 1 (with positive probability) or utility o_j . Because the judge had the choice not to apply to anything (and get at least o_j) and still chose to apply, it must be that $o_j < 1$. Hence, if $u_j(v') = 1$, applying to v' would only increase j 's probability of getting utility 1, their preferred outcome (if other players play all-mixed strategies, it would strictly increase this probability). Hence, if $u_j(v') = 1$, j would have applied to v' . By Assumption 1, it must be that $u_j(v') = 0$

Consequently, if j applied to v and not to v' $u_j(v) = 1 > u_j(v') = 0$. ■

If one focuses on trembling-hand perfect equilibrium strategies, every judge should use this type of strategy for any criterion, except for the most senior judge in each career stage. This judge will know precisely how they rank relative to other candidates in terms of seniority. They know they can get a promotion to the first vacancy they apply to, regardless of what others do. Otherwise, this proposition should apply to applications chosen by merit criterion for those who know they will not be promoted by seniority.

A.3 Cost of Application Model

In $t = 0$, the judge can choose a level of effort, which can be $e \in \{h, l\}$. Each effort level e happens at a lump-sum cost ψ_e , where $\psi_h = \psi > 0$ and $\psi_l = 0$. However, exerting high effort means facing higher promotion probabilities conditional on candidacy in periods 1 and 2.

In $t = 1$, the court offers a good vacancy with probability α and a bad vacancy with probability $1 - \alpha$. After seeing the vacancy, the judge chooses whether to apply on period 1 ($a = 1$) at a lump-sum cost of γ or not ($a = 0$). If the judge decides to apply, the court promotes the judge with probability p_1^e . If promoted, the judge gets payoff $U_1 - \gamma - \psi_e$, where $U_1 = U_g > 0$ if the vacancy is good, and $U_1 = U_B < 0$ if the vacancy is bad and the game ends. If the judge does not get promoted, either because they applied and got rejected or because they didn't apply, the game goes to period $t = 2$.

At the beginning of $t = 2$, the court offers a good vacancy with probability α , and a bad vacancy with probability $1 - \alpha$. Once again, the judge can apply, choosing $a_2 = 1$ at a lump-sum cost of γ , or not ($a_2 = 0$). If the judge applies ($a_2 = 1$), then they get promoted with probability p_2^e and receive a payoff $U_2 - \gamma - \psi_e - \gamma 1_{a_1=1}$, where $U_2 = U_G$ if the vacancy is good and $U_2 = U_B$ otherwise. If the judge doesn't apply or doesn't get promoted, the judge's payoff is $-\gamma 1_{a_1=1} - \gamma 1_{a_2=1} - \psi_e$ and the

game ends.

Assumption 2 *I assume the following with respect to the probabilities of promotion:*

- i) *High effort always brings a weakly higher probability of promotion when applying: $(p_1^h, p_2^h) \geq (p_1^l, p_2^l) \geq 0$.*
- ii) *Conditional on effort choice, waiting never reduces the likelihood of promotion: $p_2^e \geq p_1^e$.*

I consider the parameters (α, γ, ψ) as preference parameters, where α represents how picky a judge is about the vacancies they like, γ is the personal cost of putting their name forward for promotion and ψ is their dislike for effort. These are primitives of the model. The remainder of this section will relate the observable outcomes with these primitives. On the other hand, the promotion probabilities p_1^e, p_2^e for each given level of effort can change over time, for instance, because of the possibility of seniority promotion.

Proposition 2 *The judge never applies to the bad vacancy in period 1. Conditional on effort choice, a fellow chooses to apply to a good vacancy in period 1 if and only if:*

$$\gamma \leq \frac{1 - \alpha p_2^e}{1 - \alpha p_1^e} p_1^e U_G \quad (1)$$

Proof of Proposition 2. Let's proceed by backward induction. In $t = 2$, the judge will never apply to a bad vacancy because $U_B < 0$, but will apply to a good vacancy if and only if $p_2^e U_G > \gamma$. After all, applying in the second period means getting a payoff of:

$$p_2^e U_2 - \gamma - \gamma 1_{a_1=1} + \psi_e$$

Not applying means:

$$\gamma 1_{a_1=1} + \psi_e.$$

Hence, it is only worthwhile to apply if $p_2^e U_2 - \gamma \geq 0$.

Hence, the expected ex-antes payoff from the second period is:

$$\alpha \max\{p_2^e U_G - \gamma, 0\} - \psi_e - \gamma 1_{a_1=1}.$$

Hence, applying to a vacancy in the first period yields an expected payoff of:

$$E(\pi|a_1 = 1, e) = p_1^e U_1 + (1 - p_1^e) \alpha \cdot \max\{p_2^e U_G - \gamma, 0\} - \gamma - \psi_e$$

Not applying yields a payoff of:

$$E(\pi|a_1 = 0, e) = \alpha \max\{p_2^e U_G - \gamma, 0\} - \psi_e$$

That means the judge will choose to apply if and only if:

$$E(\pi|a_1 = 1, e) - E(\pi|a_1 = 0, e) = p_1^e \left(U_1 - \alpha \max\{p_2^e U_G - \gamma, 0\} \right) - \gamma \geq 0 \quad \iff \quad (2)$$

$$p_1^e \left(U_1 - \alpha \max\{p_2^e U_G - \gamma, 0\} \right) \geq \gamma \quad (3)$$

Using equation (3), we get that, because $\gamma \geq 0$ and $\max\{p_2^e U_G - \gamma, 0\} \geq 0$, the judge will never find it a good idea to apply to the bad vacancy. Furthermore, because $p_2^e \geq p_1^e$, if the judge finds it optimal to apply in period 1, they will certainly do so in period 2. As a consequence, if the vacancy is good, I can rewrite equation (3) as:

$$p_1^e \left(U_G - \alpha(p_2^e U_G - \gamma) \right) \geq \gamma$$

$$p_1^e (1 - \alpha p_2^e) U_G + p_1^e \alpha \gamma \geq \gamma$$

$$p_1^e (1 - \alpha p_2^e) U_G \geq (1 - \alpha p_1^e) \gamma$$

$$\frac{(1 - \alpha p_2^e) p_1^e}{1 - \alpha p_1^e} U_G \geq \gamma.$$

■

If the costs of application are important, then the likelihood of a judge becoming a candidate is:

$$Pr(candidate) = \begin{cases} \alpha, & \text{if } \gamma < \frac{(1 - \alpha p_2^e)}{1 - \alpha p_1^e} p_1^e U_G \\ 0, & \text{otherwise.} \end{cases}$$

If every judge has a zero cost of application, the probability of any given judge applying for a

vacancy should depend only on the profile of the vacancies offered (which is directly related to whether there is a good vacancy being offered, from the point of view of the judge), and not on productivity and seniority (i.e., factors that change over time and alter their likelihood of being promoted).

Note also that if a judge thinks their efforts will be more rewarded in the future than in the present, that is, $p_1^H \sim p_1^L$ but $p_2^H > p_1^L$, then they will be more likely to apply after choosing low effort because:

$$\frac{(1 - \alpha p_2^H)}{1 - \alpha p_1^H} p_1^H U_G < \frac{(1 - \alpha p_2^L)}{1 - \alpha p_1^L} p_1^L U_G.$$

Extension to Multiple Vacancies

The promotion from the final stage to appeals court is a particular case of this model where α is either zero or one, and only one type of promotion is ever offered: appeals court judge. However, there are multiple vacancies at every promotion opportunity within the Trial Court. If Assumption 1 holds, the main message remains: if there are no costs of application, then every judge should apply whenever there is at least one good vacancy being offered because there would be better than staying in their current job and waiting for another opportunity to present itself in the future. Suppose vacancies can have different utility values for judges. In that case, it might be that a judge who sees themselves as less likely to get a good promotion in the future (e.g., for not being senior enough) is more willing to apply to each vacancy in the present because their outside option is less attractive. In that case, they should become more selective as they become more senior - and therefore more likely to be promoted by seniority, not only in the decision to apply but also in the profile of vacancies they choose.

A.4 Simulation Algorithm

At the end of year 0, there are 2000 identical judges with zero experience and zero tenure starting off as substitute judges: 1000 male and 1000 female. Each judge has four relevant state variables associated with them: 1) whether the judge has exited the system, 2) the number of promotions the judge has ever gotten (which corresponds to the higher career stage they have reached); 3) the number of years of tenure in the same career stage they have accrued; 4) the number of years of experience they have. The simulation algorithm changes the state variables in the following way:

1. At the beginning of every year (day 1 of month 1), each judge faces an individual probability of exiting the justice system conditional on having been active the previous year. In other words, each judge's exit decision is a Bernoulli with probability $p_{j,s,g}^e = Pr_j(\text{Exit year } t | \text{Active year } t - 1, \text{Career Stage} = s, \text{Gender} = g)$. If a judge exits, they are not eligible for any promotion

opportunity in the future. If they stay, they are eligible for every promotion opportunity for their current career stage that year.

2. By the end of every month, each judge gets one extra month of experience and one extra month of tenure in the current position. There is a toss of a coin for each career stage within the Trial Court, defining whether or not there will be an opportunity for promotion — which applies to judges of all genders. The monthly probability of there being a promotion opportunity for career stage s is p^s .
3. For any career stage such that there is a promotion opportunity, each eligible judge (i.e. a judge who did not exit the sample and is currently in that stage) faces an individual probability of becoming a candidate given by $p_{s,g,j}^c = \Pr(\text{Candidate} \mid \text{Eligible}, \text{Gender}=g, \text{Career Stage}=s)$ and an individual probability of being promoted conditional on being a candidate $p_{s,g,j}^p = \Pr(\text{Promoted} \mid \text{Candidate}, \text{Gender}=g, \text{Career Stage}=s)$.
4. If a judge gets promoted, then they instantly move to the next career stage and receive zero tenure at that stage. Otherwise, their state variables stay the same at the end of the month.

I estimate the individual probability of exit conditional on activity in the previous year is estimated separately by gender x career stage combination using equation (4).

$$\ln\left(\frac{p_{j,s,g}^e}{1 - p_{j,s,g}^e}\right) = \beta_{0,s,g}^e + \beta_{1,s,g}^e \text{Tenure Stage}_{j,t} + \beta_{2,s,g}^e \text{Experience}_{j,t} + \beta_{3,s,g}^e \text{Experience}_{j,t}^2 + \varepsilon_{j,t}, \quad (4)$$

The probability of an eligible judge in career stage s and gender g applying for promotion in each promotion opportunity is determined by equation (5).

$$\ln\left(\frac{p_{j,s,g}^c}{1 - p_{j,s,g}^c}\right) = \beta_{0,s,g}^c + \beta_{1,s,g}^c \text{Tenure Stage}_{j,t} + \beta_{2,s,g}^c \text{Experience}_{j,t} + \beta_{3,s,g}^c \text{Experience}_{j,t}^2 + \varepsilon_{j,t}. \quad (5)$$

The probability of an eligible judge in career stage s and gender g getting a promotion in each promotion opportunity conditional on wanting one is determined by equation (6).

$$\ln\left(\frac{p_{j,s,g}^p}{1 - p_{j,s,g}^p}\right) = \beta_{0,s,g}^p + \beta_{1,s,g}^p \text{Tenure Stage}_{j,t} + \beta_{2,s,g}^p \text{Experience}_{j,t} + \beta_{3,s,g}^p \text{Experience}_{j,t}^2 + \varepsilon_{j,t}. \quad (6)$$

To obtain the monthly probability of a promotion opportunity p^s for each career stage s (equally available to both genders), I used the following identity to determine the average annual number of promotions for career stage s , N_s .

$$\begin{aligned} Pr(\text{Same Career stage } s \text{ in Dec 31st of year } t \text{ and Dec 31st year } t - 1) = \\ Pr(\text{Eligible but unpromoted at every promotion opportunity for } s \text{ in year } t) = \\ (1 - Pr(\text{Promoted at a promotion opportunity}|\text{Eligible}))^{N_s} \end{aligned}$$

Using the number of promotions N_s taking place in a year, I estimated that an average of $n_s = N_s/12$ promotions had to take place in a month, where n_s was between 0 and 1 for all career stages. I then established $p^s = n_s$.