

The Leaky Pipeline: Retention, Performance, and Promotion Policies*

Kate Reinmuth and Brian E. Higgins[†]
(*Job Market Paper*)

Last Updated: October 10, 2026

[\[Click here for latest version.\]](#)

Abstract

Despite substantial progress toward labor market parity, women remain less likely than men to be promoted in many professions. This paper studies how firm promotion policies shape these gaps in large law firms, where women make up half of junior associates but only a third of new equity partners. Using new administrative panel data from a large U.S. law firm, we find that differences in associate retention to the promotion stage account for 31% of the gap, while differences in performance among retained associates, measured using total billable hours, account for the remaining 69%. The hours gap itself reflects gender differences in part-time work – and particularly the persistence of part-time spells. Although individual promotions are not observed, we infer the promotion policy from aggregate promotion rates and find that a gender-neutral promotion policy can rationalize the data. We then estimate a dynamic labor supply model to match these empirical patterns and evaluate how alternative promotion policies would affect worker behavior, promotion gaps, and firm productivity. Several counterfactual promotion policies can substantially reduce the gap, but with different effects on firm productivity. Accounting for part-time schedules when measuring performance yields the most favorable tradeoff among the policies we consider, cutting the promotion gap in half with weakly positive productivity effects.

*We are deeply grateful to our data provider for making this project possible. The views and conclusions stated herein are those of the authors and do not necessarily reflect the views of the data provider, individuals/organizations associated with the After the JD study, or any other entity. We thank Nick Bloom, Liran Einav, Neale Mahoney, and Alessandra Voena for their continual guidance and support on this project, as well as Matt Gentzkow, Claudia Goldin, Katja Hofmann, Caroline Hoxby, Larry Katz, Paul Oyer, Elena Pastorino, Petra Persson, Torsten Persson, Luigi Pistaferri, Martin Schneider, Tess Snyder, David Wilkins, Shoshana Vasserman, Ali Yurukoglu, and participants in Stanford’s IO, applied, and labor/public workshops and Middlebury College and IIES seminars, as well as the 2026 SOLE Annual Meeting, BSE Summer Forum, and All-California Labor Economics Conference, among others, for valuable feedback. Kate Reinmuth gratefully acknowledges financial support from Knight-Hennessy Scholars and the Clayman Institute at Stanford University as well as the John M. Olin Program at Stanford Law School.

[†]Reinmuth: Stanford University, reinmuth@stanford.edu; Higgins: IIES, Stockholm University, brian.higgins@iies.su.se.

1 Introduction

Despite substantial progress toward labor market parity in recent decades (Goldin 2006), women are still less likely than men to be promoted in many professions (e.g., Blau and Devaro 2007; Benson, Li, and Shue 2026). Even when firms achieve gender balance at entry-level positions, women tend to fall behind at later career stages – a pattern often described as the “leaky pipeline.” These promotion gaps also contribute to broader gender gaps in earnings and organizational leadership (e.g., Okuyama, Murooka, and Yamaguchi 2025; Tannenbaum and Timpe 2025). Understanding the mechanisms behind promotion gaps is therefore central to understanding gender differences in contemporary labor markets.

In this paper, we analyze the firm’s role in shaping gender promotion gaps. Specifically, we ask how promotion policies affect promotion outcomes and firm productivity, both directly through the promotion decisions and indirectly through the incentives they create for workers over the course of their careers. Existing work has provided important insights into promotion gaps (e.g., Bertrand, Goldin, and Katz 2010; Azmat and Ferrer 2017; Benson, Li, and Shue 2026), but we still know relatively little about whether counterfactual promotion policies could reduce observed gaps, how workers would respond, and what tradeoffs firms would face.

We study these gaps in the setting of large law firms, where women make up half of junior associates but only about 24% of existing equity partners and 33% of new equity partners (Choi and Violante 2020; NALP 2024).¹ Large law firms generally place considerable weight on total billable hours in bonus and promotion decisions (Landers, Rebitzer, and Taylor 1996; Azmat and Ferrer 2017), providing an unusually direct measure of promotion-relevant performance.² Accordingly, our analysis utilizes new administrative panel data on individual associates’ billable hours and retention over time, which we supplement with information on compensation, promotion rates, firm financials, and industry trends. Individual promotions are not observed in the administrative data; instead, this combination of data allows us to infer an implied promotion policy and study how the aggregate promotion gap arises.

We begin by providing descriptive evidence on the gender promotion gap in this setting.³ We organize the analysis around three mechanisms that jointly determine promotion

¹ Throughout the paper, we focus on *equity* partners rather than non-equity partners. Equity partnership is the relevant high-stakes promotion margin in our setting because it confers ownership, profit-sharing rights, and governance authority. Non-equity partnership generally involves different compensation, responsibilities, and rights, and it is often closer to an extended senior associate track than to partnership.

² Other common measures of performance like performance reviews rely on managerial perceptions.

³ This industry’s promotion gap is similar in magnitude to the gaps observed in other elite careers. E.g., women make up 23% of full professors in top economics departments (Antecol, Bedard, and Stearns 2018; Chen, Liu, and Kim 2022; Tesar 2025), 32% of senior managers in central banks (Hospido, Laeven, and Lamo 2022; Comunale, Bruxelles, Kochhar, Raskauskas, and Unsal 2023), 6% of executives in publicly traded firms

outcomes: retention, or who remains at the firm long enough to be considered for promotion; performance among retained associates; and the firm’s promotion policy, which maps retained associates’ performance into promotions. These mechanisms are sequential components of observed promotion outcomes, rather than independent causal forces – e.g., worker retention and performance may themselves respond to the promotion policy. However, this framework allows us to understand where and when the status quo promotion gap arises at a high level.

Starting with retention, we find that women account for 53% of first-year associates but only 44% of eighth-year associates. Performance differences further widen the gap, as eighth-year women are significantly less likely than men to appear in the right tail of total billable hours. Finally, estimating the implied mapping from hours to promotion, we find no evidence of different promotion standards being applied to men and women.

We combine these descriptive patterns in a static decomposition to quantify how much each mechanism contributes to the observed promotion gap. We find that differences in performance, measured using total billable hours, account for 69% of the overall gap, while differences in retention account for the remaining 31%. We also show that the retention and performance gaps develop over the associate life cycle, with women being slightly *over*-represented and having slightly *higher* hours in the early years, but falling behind over time. Finally, we show that the performance gap itself is driven by differences in part-time work. Part-time work is closely linked to parenthood, and part-time spells are substantially more persistent for women than for men, contributing to persistent gender gaps in hours.

This static evidence provides a simple accounting of the margins through which the observed promotion gap arises, but it cannot speak to the underlying economic forces or evaluate how alternative promotion policies would affect outcomes once workers adjust their behavior. For example, a change in promotion policy may alter not only which associates are selected for partnership, but also associates’ incentives to remain at the firm and supply billable hours. These responses matter for both the gender promotion gap and firm productivity.

To study these responses, we develop and estimate a dynamic labor supply model in which associates choose whether to remain at the firm, whether to work full-time or part-time, and how many hours to supply each year. These choices determine current pay, firing risk, and future promotion prospects, allowing the model to capture the key tradeoffs associates face during the pre-promotion years. For example, working more hours raises bonus eligibility and promotion probabilities and lowers firing risk, but it is costly because it requires associates

(Keller, Molina, and Olney 2023), 14% of district managers in large retail chains (Benson, Li, and Shue 2026), and less than 30% of full professors in academic medicine (Marhoffer et al. 2024).

to give up more non-work time.

In the model, we allow key primitives to differ by gender and include two sources of individual-level heterogeneity. The first is a permanent type, which captures persistent differences in willingness or ability to supply long hours. The second is a persistent latent state that evolves between a low-cost state and a high-cost state over time. The high-cost state captures life events outside work, such as parenthood, that increase non-work demands on an associate’s time. The type distributions and latent state transition processes are also allowed to differ by gender.

We estimate the model by Simulated Method of Moments (SMM), matching key gender-specific patterns in the data on hours and contract choices over time, as well as promotion outcomes. The model fits the targeted moments well and suggests that the main difference across genders appears in the latent state process rather than baseline labor supply. Specifically, spells in the high-cost state are estimated to be more than three times as persistent for women as for men, helping generate the greater time in part-time work and higher exit rates among women that we see in the data.

We then use the estimated model to evaluate three alternative promotion policies that mirror real-world proposals: accounting for part-time schedules, setting gender-based representation targets, and limiting long hours. All three reduce the promotion gap, but through different margins and with different tradeoffs for firms. Accounting for part-time schedules can cut the promotion gap in half with little effect on the productivity margins captured by the model. Representation targets can close more of the gap by construction, but their productivity effects depend on the target: moderate targets preserve productivity, while more stringent targets, such as parity, introduce costs by weakening selection into partnership. Hours caps can also substantially increase the female share of promotions by limiting long hours, but they generate the largest productivity costs among the policies we consider.

Related Literature. This paper relates to three distinct parts of the existing literature: research on gender gaps in pay and promotion, models of labor supply under non-linear and long-term incentives, and work on law firms and the professional services industry.

First, a large literature documents persistent gender gaps in earnings and promotions, even as gender differences in education and entry-level hiring have narrowed substantially. An early and influential strand of this literature focused on decomposing pay disparities within and across job types and levels (Lazear and Rosen 1990; Wood, Corcoran, and Courant 1993; Bertrand and Hallock 2001; Bertrand, Goldin, and Katz 2010; Okuyama, Murooka, and Yamaguchi 2025), bringing attention to the “leaky pipeline” and to promotions as a key margin where pay disparities emerge. Since then, gender promotion gaps have been docu-

mented across a wide range of occupations, including academia, business, government, law, medicine, and retail (Antecol, Bedard, and Stearns 2018; Chen, Liu, and Kim 2022; Hospido, Laeven, and Lamo 2022; Keller, Molina, and Olney 2023; Lassen and Ivandić 2024; Benson, Li, and Shue 2026; Cairo, Ivandic, Lassen, and Tartari 2026). Prior work has identified several mechanisms that may contribute to these gaps, including the role of children (Goldin 2014; Goldin and Katz 2016; Azmat and Ferrer 2017; Bertrand 2018; Kleven, Landais, and Egholt-Søgaard 2019), aspirations (Azmat and Ferrer 2017), harassment (Azmat, Cuñat, and Henry 2025), reactions to competition (Niederle and Vesterlund 2007), gender-specific or otherwise biased standards (Lazear and Rosen 1990; Huang, Mayer, and Miller 2024; Benson, Li, and Shue 2026), and non-wage amenities (Wiswall and Zafar 2018; Gupta, Gallen, Hedeager, and Holzheu 2026). Relative to this literature, our contribution is to study the role of firm promotion policies in shaping these gaps.

Second, this paper relates to the literature on labor supply responses to non-linear incentives and long-term payoffs. A substantial body of work shows how individuals respond to non-linearities, and especially discontinuities, in incentive schedules (e.g., Saez 2010; Chetty, Friedman, Olsen, and Pistaferri 2011; Kleven and Waseem 2013; Bastani and Selin 2014; Blomquist, Newey, Kumar, and Liang 2021). A separate strand of the literature considers dynamic labor supply decisions, including intertemporal substitution (e.g., MaCurdy 1981; Altonji 1986; Hotz, Kydland, and Sedlacek 1988; Farber 2005; Martínez, Saez, and Siegenthaler 2021) and long-term incentives such as promotions (Landers, Rebitzer, and Taylor 1996; Holmström 1999; Bell and Freeman 2001). We use a dynamic labor supply model to estimate how salaried associates’ labor supply decisions respond to both current compensation and future promotion prospects. This allows us to evaluate counterfactual promotion policies in a setting where marginal incentives operate primarily through bonuses and career concerns rather than hourly wages.⁴

Third, this paper relates to the literature on law firms and professional service firms. Early work in this literature modeled promotions in law firms as an option to be invested in by associates while they acquire valuable training (Leibowitz and Tollison 1978), or as a tournament to be won by the highest performing associates that also governs the size of the firm (Galanter and Palay 1991). Landers, Rebitzer, and Taylor (1996) then examined how promotion policies affect hours by inducing a “rat race” among associates, with more recent work focusing, like us, on gender differences in law firm careers (Goldin 2014; Azmat and Ferrer 2017; Azmat, Cuñat, and Henry 2025). Separate but related strands of the literature

⁴In doing so, this paper also connects to the literature on careers in organizations and internal labor markets (e.g., Doeringer and Piore 1971; Cappelli and Cascio 1991; Baker, Gibbs, and Holmstrom 1994; Gibbons and Waldman 1999).

focus on specialization and supervision (Garicano and Hubbard 2007), how lawyers sort into firms (Oyer and Schaefer 2016), the returns to attending an elite law school (Oyer and Schaefer 2019), the broader market for lawyers and legal services (Rosen 1992), and litigation as a signal of lawyer skill (Ferrer and Okulicz 2025). Law firms are also an example of professionalized, knowledge-intensive firms in the professional services industry more broadly (Greenwood and Empson 2003; Nordenflycht 2010). This literature suggests that up-or-out promotion systems are closely tied to the partnership structure common in professional service firms (Levin and Tadelis 2002; Rebitzer and Taylor 2007; Barlevy and Neal 2019). We add to this literature by using within-firm administrative data to study how associates’ hours, part-time arrangements, and retention relate to gender gaps in promotion, and by evaluating how alternative promotion policies would affect outcomes.

Roadmap. The paper proceeds as follows. Section 2 describes the institutional setting and data. Section 3 measures the observed promotion gap and three key mechanisms – retention, performance, and the firm’s promotion policy – before decomposing the observed promotion gap into the relative contribution of each mechanism. Section 4 endogenizes retention and performance by estimating a dynamic labor supply model in which associates respond to both short-term incentives and future promotion prospects. Finally, Section 5 uses the estimated model to evaluate counterfactual promotion policies and their effects on promotion gaps and firm productivity.

2 Setting and Data

Setting. We study large U.S. law firms. Lawyers enter these firms as associates and are considered for promotion to partner after several (usually eight) years at the firm (Fatima 2025). Associates earn a salary and may also earn a bonus if they meet the firm’s billable hour threshold.⁵ Promotion to equity partner gives lawyers the right to buy in to the firm and receive profit shares that can be up to 25 times first-year associate salaries (Simons 2018). Associates who do not make partner usually move to smaller firms, government jobs, or in-house roles (Harrington and Hsi 2007).⁶ Thus, law firms are a setting with highly non-linear returns to performance (Goldin 2021).

Billable hours are a central performance metric in the industry and measure the time a lawyer spends working on matters that are billed to clients at a set hourly rate. These

⁵Associate salaries and bonuses are standardized both across associates and across firms through the Cravath scale, an industry-wide lockstep compensation schedule. See Boone (2021).

⁶Appendix Table 3 shows the declining share of law firm work over time for a given cohort.

hours are directly tied to firm revenue, determine bonus eligibility, and are used in promotion decisions.⁷ Prior work argues that law firms rely on billable hours because they provide a signal of an associate’s propensity to supply effort, a relevant input into partner suitability when profit sharing may otherwise generate shirking or adverse selection (Landers, Rebitzer, and Taylor 1996). As in Azmat and Ferrer (2017), this feature of the setting allows us to observe an objective administrative measure of performance that does not rely on subjective managerial assessments, which are themselves the object of study in related settings (e.g., Benson, Li, and Shue 2026; Minni, Nguyen, Sarsons, and Srebot 2026).

These law firms also have sizable gender gaps in promotion: although women make up 53% of first-year associates in our sample, they account for only 24% of existing equity partners and 33% of newly promoted equity partners in recent years (Choi and Violante 2020; NALP 2024). One industry response to these gaps has been the use of part-time arrangements (Above the Law 2015).⁸ Associates in these arrangements generally work fewer days and receive proportionally lower pay: for example, an associate working four days per week would earn 80% of full-time salary and 80% of any bonus, subject to 80% of the full-time bonus threshold.⁹

Data. Our primary data come from a large U.S. law firm and cover approximately 1,000 unique associates observed in an unbalanced panel over five years. Because associates enter and leave the firm during this period, the data contain about 2,300 associate-year observations. The data are fully anonymized but include a unique identifier that allows us to follow associates over time.

For each associate-year observation, we see annual billable hours, gender, year in the associate track (e.g., second-year versus fifth-year associates), and contracted share of full-time work. Annual billable hours are recorded by the firm in 6-minute increments for billing purposes and then aggregated to the associate-year level. The data also provide FTE billable hours, which are calculated by dividing total billable hours by the associate’s contracted share of full-time work. These are the same records the firm uses for internal operations, including bonus eligibility. We measure retention using the unique associate identifier. For associates below the eighth year, we define retention from year t to year $t + 1$ as appearing in the firm

⁷See Appendix Figure 2, which shows a strong relationship between billable hours and promotion – even in data collected across firms and with self-reported hours. As one lawyer described: “billable hours are the way to live: one 0.1 [6-minute increment] at a time.” As illustrated in Appendix Figures 3 and 4, billable hours are also positively correlated with other potential dimensions of performance, such as new client revenue generation and number of informal mentors.

⁸Firms have also experimented with targeted professional development, mentoring, and alternative senior positions (McLymore 2021; Rummel 2024; Lateral Link 2026).

⁹With weekend work, these part-time shares may actually over-state the share of days off. Williams and Calvert (2002) explain how some associates adopt a 60% schedule to keep hours “in the 80% range.”

in the following year.

The main limitation of the administrative data is that we do not observe individual-level promotion outcomes. We therefore use supplementary sources, described below, combined with the gender composition of the promotion-eligible cohort (i.e., associates retained through their eighth year) in the administrative data to construct promotion rates by gender.

We supplement the law firm data with several external sources. First, we use publicly available firm-level promotion announcements, which report the size of new partner classes, to calculate the average promotion rate among eighth-year associates in our sample. We then combine this firm-level promotion rate with evidence on the gender composition of new equity partner cohorts from Choi and Violante (2020) and the gender composition of eighth-year associates in our sample to construct gender-specific promotion rates. We also use BigLaw Investor data to measure historical associate salaries and bonuses and Am Law 100 data to measure firm revenue, headcount, and profits. We use Revelio Labs employment histories to characterize lawyers’ outside options after leaving large law firms. We also use the After the JD longitudinal survey to provide additional evidence on career dynamics, beliefs, and expectations. Appendix A provides more details on these supplementary sources.

3 Descriptive Evidence

This section documents the observed promotion gap and examines the mechanisms that generate it. We begin by describing female representation across career stages. We then study three mechanisms through which promotion gaps may arise: retention, performance, and the firm’s promotion policy. For each mechanism, we present descriptive evidence from the data. We then combine these objects in a static decomposition that quantifies each mechanism’s contribution to the observed promotion gap. This exercise provides a transparent accounting of the margins through which the observed promotion gap arises in the data. In doing so, it highlights the empirical patterns that motivate the dynamic model in Section 4 and clarifies the policy margins most relevant for the counterfactuals in Section 5.

3.1 Promotion Gap


Promotion gaps describe differences in representation across seniority levels. In law, female representation is high among first-year law students, summer associates, and associates, but falls substantially at the equity partner stage, as shown in Figure 1.¹⁰ The figure reports newly promoted and existing equity partners separately because the stock of existing

¹⁰Benson, Li, and Shue (2026) conducts a similar job levels exercise for retail operations workers.

partners reflects a longer history of promotion decisions. By either measure, however, female representation is much lower than at junior levels.

To remain agnostic about the exact starting point of the partner pipeline, we define the “promotion gap” as a deviation from gender parity. This definition yields a symmetric measure: positive values indicate female under-representation, while negative values indicate male under-representation. Under this definition, the female share of newly promoted partners implies a promotion gap of 17 percentage points.

Figure 1: Female Share of Lawyers by Job Level



Notes: Figure reports female representation at each stage of the large law firm pipeline. Bars show the female share of 1L students, summer associates, associates, newly promoted equity partners, and existing equity partners. The horizontal line marks gender parity at 50%. Newly promoted partners refer to recent new equity partner cohorts; existing partners refer to the stock of equity partners. **Data:** Student share from Russian (2024); new partner share from Choi and Violante (2020); other shares from NALP (2024).

3.2 Mechanisms

In general, promotions (and therefore promotion gaps) can be understood as the result of three sequential mechanisms. First, workers are either retained or exit the firm, whether voluntarily or involuntarily.¹¹ Second, conditional on remaining at the firm, workers perform at some level.¹² Third, they face the firm’s promotion policy, which determines how perfor-

¹¹Accounting for this mechanism allows us to assess whether a gap would persist even if representation in the eighth-year cohort were equalized. Because the data do not distinguish voluntary departures from involuntary departures, we cannot decompose the retention margin by reason for exit. Some departures might alternatively be viewed as part of the performance margin if they were related to low performance.


¹²Performance reflects both workers’ effort choices and factors such as innate ability.

mance translates into advancement. These three mechanisms – retention, performance, and promotion policy – guide our analysis below.

Retention. Retention determines who remains at the firm long enough to become eligible for promotion. Rather than examining associates as a whole, as in Figure 1, we focus here on representation within the associate ranks. We compare the female share of first-year associates to the female share of eighth-year associates, where we consider eighth-year associates as constituting the promotion-eligible pool in our setting. Differential attrition over these years can change the composition of this pool. And, in some settings, retention is the primary driver of gender gaps (e.g., Zhao, Thoma, Hertwig, and Wulff 2025).

In our setting, retention does differ by gender over the associate years. On average, women make up 53% of the first-year associate class, but only 44% of eighth-year associates, as shown in Figure 2. This decline implies that retention contributes to the gender promotion gap by reducing women’s representation in the pool from which partners are selected. However, retention is not the only relevant mechanism, as women account for 44% of the promotion-eligible cohort but only 33% of newly promoted partners. This difference suggests that performance and promotion policies may also play a role.

Figure 2: Female Share of Associates by Year



Notes: Figure plots the female share of associates by associate year. Blue markers show the empirical female share among associates observed in each associate year in the Law Firm Dataset. The gray line shows the best-fit linear relationship between female share and associate year, with the shaded region showing the 95% confidence interval. The horizontal line marks gender parity at 50%. Labels report the female share in the first and eighth associate years. **Data:** Law Firm Dataset.

Performance. Performance captures how well workers do their jobs and is a key input into rewards like bonuses and promotions. For promotions, the top of the performance distribution is likely to be the most relevant, since advancement is generally reserved for top performers. Therefore, observing the full within-firm performance distribution, rather than only group averages, is useful for understanding whether men and women differ along the part of the distribution most relevant for promotion.

As discussed above, large law firms often measure performance using total billable hours (Landers, Rebitzer, and Taylor 1996; Azmat and Ferrer 2017). We interpret billable hours as an output that associates produce at heterogeneous and increasing marginal cost.¹³ With this in mind, Figure 3A plots the counter-CDF of eighth-year associates’ billable hours separately for men and women. Hours are normalized by the maximum observed value to preserve confidentiality. At each value z on the horizontal axis, the counter-CDF reports the share of associates with billable hours above z . Across most of the distribution, the male curve lies above the female curve, indicating that men are more likely to appear in the right tail of annual billable hours.

Although total billable hours are a key promotion-relevant performance metric in the industry, firms also track an alternative metric – FTE billable hours – for annual bonuses. Figure 3B shows that the gender gap in hours disappears when measured in FTE terms, suggesting that the performance gap itself is driven by differences in part-time arrangements.¹⁴


Promotion Policies. Promotion policies determine how a firm maps performance into advancement.¹⁵ Unlike retention and performance, the promotion policy is not directly observed in our data. We therefore recover a descriptive measure of the policy from the objects we do observe: gender-specific promotion rates among eighth-year associates and the performance distributions of those same workers. This exercise allows us to describe the implied mapping from performance to promotion and to ask whether gender-specific promotion standards are needed to rationalize the data.

¹³This cost may reflect standard labor supply forces, such as the opportunity cost of non-work time, but it may also reflect differences in associate quality. For example, if partners preferentially assign work to higher quality associates, those associates can generate additional hours at lower effective marginal cost.

¹⁴The median female eighth-year associate works 85% of full-time, compared to 100% for the median male eighth-year associate.

¹⁵These policies specify not only the performance metrics used to evaluate workers, as discussed above, but also features such as the timing and frequency of promotions. For now, we take the status quo promotion structure described in Section 2 as given. In Section 5, we vary these policy dimensions.

Figure 3: Billable Hours Distributions



Notes: Each panel plots the counter-CDF of eighth-year associates' billable hours separately by gender. At each value on the horizontal axis, the figure reports the share of eighth-year associates with hours above that value. Panel A uses total annual billable hours. Panel B uses full-time-equivalent (FTE) billable hours, defined as total billable hours divided by the associate's contracted share of full-time work. Hours are normalized by the empirical maximum within each measure to preserve confidentiality. Men are shown in orange and women in gray. **Data:** Law Firm Dataset.

To implement this exercise, we adopt a flexible logistic mapping from performance to promotion probability:¹⁶

$$p(z; \gamma, c_z^p) = [1 + \exp(-\gamma(z - c_z^p))]^{-1}, \quad (1)$$

where z is performance, $\gamma > 0$ governs the steepness of the promotion policy, and c_z^p is the point at which the promotion probability equals 50%. Under a common policy, $p_f(z) = p_m(z) = p(z)$, the promotion rate for gender g can be written as:

$$\psi_g(\gamma, c_z^p) = \int p(z; \gamma, c_z^p) dF_g(z),$$

where $F_g(z)$ is gender g 's performance distribution among eighth-year associates.

We estimate the promotion policy parameters (γ, c_z^p) using generalized method of moments (GMM), as described in Appendix Section B.1. We find that a gender-neutral promotion policy can rationalize the gender-specific promotion rates. Specifically, both groups' promotion rates are well approximated by a threshold rule where promotion is concentrated among associates with normalized hours above 0.57. Figure 4 illustrates this policy by shading the estimated promoted region of the performance distributions. Although an hours-based threshold may seem unintuitive, it is consistent with systems like rank-based promotion in which advancement depends on a worker's position in the performance distribution (Lazear and Rosen 1997). In our setting, the estimated threshold corresponds to approximately the 40th percentile of the eighth-year hours distribution. And if we assume that the promotion policy depends *only* on billable hours, this finding rules out gender-specific promotion standards.¹⁷


3.3 Decomposing the Observed Promotion Gap

We now combine the three mechanisms discussed above – retention, performance, and the promotion policy – to decompose the observed promotion gap. Let r_g denote gender g 's share of eighth-year associates, and, following the descriptive evidence above, suppose promotion is governed by a gender-neutral policy, $p(z)$. Gender g 's share of promotions can

¹⁶This function is weakly increasing in performance (z), allowing for both deterministic thresholds and more gradual, probabilistic promotion policies that reflect other unobserved dimensions of performance.

¹⁷However, if promotion also depends on other factors, the exercise cannot rule out gender differences in how those factors are evaluated.

Figure 4: Total Billable Hours Distributions with Estimated Promotion Policy



Notes: Figure plots the counter-CDF of total annual billable hours for eighth-year associates, separately by gender. Hours are normalized by the empirical maximum to preserve confidentiality. The shaded region shows the promoted region implied by the gender-neutral promotion policy estimated in Appendix Table 2. The estimated policy is approximately a threshold rule, with promotion probabilities rising sharply once normalized hours exceed $\hat{c}_z^p = 0.57$. Horizontal markers report the empirical promotion rates used to discipline the policy estimation: 45% for women and 72% for men. Men are shown in orange and women in gray. **Data:** Law Firm Dataset and promotion rates constructed in Appendix Table 1.

then be written as:

$$\pi_g = r_g \cdot \frac{\int p(z) dF_g(z)}{\int p(z) dF(z)} = r_g \cdot \frac{\psi_g}{\psi}, \quad (2)$$

where $\psi_g = \int p(z) dF_g(z)$ is gender g 's promotion rate among eighth-year associates, $F_g(z)$ is gender g 's eighth-year performance distribution, $\psi = \int p(z) dF(z)$ is the overall eighth-year promotion rate, and $F(z)$ is the pooled eighth-year performance distribution.

Taking logs of (2) yields a convenient additive decomposition:

$$\ln \pi_f = \ln r_f + \ln \psi_f - \ln \psi.$$

This representation allows us to define three gap measures as deviations from gender parity, as described in Section 3.1:

$$\begin{aligned} \text{Promotion Gap: } \Delta &:= \ln 0.5 - \ln \pi_f, \\ \text{Retention Gap: } \Delta^R &:= \ln 0.5 - \ln r_f, \text{ and} \\ \text{Performance Gap: } \Delta^P &:= \ln \psi - \ln \psi_f. \end{aligned}$$

The retention gap captures female under-representation among eighth-year associates, while the performance gap captures how gender differences in eighth-year performance translate into promotion outcomes under the firm’s promotion policy. The decomposition takes the promotion policy estimated in Section 3.2 as fixed. Because this policy is gender-neutral, there is no separate promotion policy gap in the decomposition. Instead, the decomposition focuses on differences in representation and performance given the promotion policy.¹⁸

Table 1 reports the empirical gaps implied by the decomposition and the data described in Section 3.2. Both retention and performance contribute to the promotion gap, but performance is the larger component: differences in total billable hours account for roughly two-thirds of the overall gap.¹⁹ The central role of total billable hours is consistent with industry-wide survey evidence from Azmat and Ferrer (2017).²⁰

Table 1: Decomposition of Promotion Gap

	Promotion Gap		Retention Gap		Performance Gap
	Δ	=	Δ^R	+	Δ^P
Log Points:	0.391		0.122		0.269
Share of Δ :			31%		69%

Notes: Table decomposes the observed promotion gap for eighth-year associates into retention and performance components. The promotion gap is defined as $\Delta = \log(0.5) - \log(\pi_f)$, where π_f is the female share of newly promoted partners. The retention gap is $\Delta^R = \log(0.5) - \log(r_f)$, where r_f is the female share of eighth-year associates. The performance gap is $\Delta^P = \log(\psi) - \log(\psi_f)$, where ψ_f is the female promotion rate among eighth-year associates and ψ is the pooled promotion rate. Gaps are reported in log points. The percent contribution of each component is its share of the total log-point gap. **Data:** Hours and retention shares are from the Law Firm Dataset. Promotion rates are constructed in Appendix Table 1.

The same decomposition can be applied over the associate life cycle to show how and when these gaps emerge. Figure 5A applies the decomposition separately by associate year, treating each year’s cohort as if it were the relevant promotion cohort.²¹ As above, these are descriptive accounting components rather than independent causal mechanisms. This exercise shows how retention and performance evolve *before* associates reach the standard eighth-year promotion stage. In the early years, women make up a larger share of associates

¹⁸Importantly, though, the decomposition describes where the observed promotion gap arises in the data rather than its deeper economic causes. For example, interpreting Δ_R as the share of the gap that would disappear if retention were equalized requires one to assume that the women who would be newly retained have the same eighth-year performance distribution as the women who are retained at baseline.

¹⁹As in Section 3.2, we focus here on associates’ performance in their final year before promotion. In Section 4, we allow the firm to also consider performance in other associate years.

²⁰Bertrand, Goldin, and Katz (2010) also find a similar result for MBAs.

²¹To calculate the performance gap, we hold the overall promotion rate fixed.

and slightly outperform men, so both gaps are negative. Over time, however, both gaps become positive and grow in magnitude, illustrating the “leaky pipeline.”

We also repeat the exercise using the FTE hours metric discussed in Section 3.2. Figure 5B shows that the performance gap disappears when performance is measured in FTE terms, with $\Delta^{P'} = -0.06$. This result suggests that higher rates of part-time work among women are a primary driver of the observed promotion gap.

Part-Time Dynamics. The higher prevalence of part-time work among women primarily reflects differences in the *persistence* of part-time spells. Men and women transition into part-time work at similar rates, but men return to full-time substantially faster. Appendix Figure 7 shows that 45% of mid-career male associates working part-time return to full-time the following year, compared with only 27% of female associates. For example, among associates who first go part-time in their fourth year, women who remain at the firm are four times as likely as men to still be working part-time in their eighth year.


These part-time spells are closely linked to parenthood. Although we do not observe children or parenthood status in the administrative data, the After the JD survey shows that, among lawyers working part-time in large law firms, 86% cite children as the primary reason for doing so.²² Thus, persistent child-related demands for flexibility seem to play an important role in the gender differences in part-time work documented above.²³

Summary. Together, these descriptive results characterize the margins through which the observed promotion gap arises. However, they cannot speak to how workers would respond to alternative promotion policies or to the resulting net effects of such changes. Retention outcomes, part-time work, and hours are jointly determined and may all respond to promotion incentives. In the following section, we therefore turn to a model that endogenizes workers’ labor supply decisions, allowing us both to evaluate counterfactuals and to study the economic forces underlying the descriptive patterns outlined above.


²²Among the 102 respondents at large law firms who report working part-time, 88 cite children, 1 cites disability, 1 cites family, and 12 cite another reason or do not provide one.

²³Associate age is also consistent with this interpretation: in the fourth through eighth years, associates are, on average, 32 to 36 years old, as shown in Appendix Figure 6. Goldin and Katz (2008) document that, among women in an elite college-graduate sample, the median age at first birth was 31 to 32, conditional on having a birth. This timing suggests that further delay may become increasingly costly for associates who plan to have children.

Figure 5: Promotion Gap Decompositions



(A) Across Associate Years



(B) Across Hours Metrics

Notes: Figure decomposes the promotion gap into retention and performance components. Positive values indicate female under-representation relative to parity, while negative values indicate female over-representation. The blue component is the retention gap, which captures female under-representation among associates in the relevant cohort. The red component is the performance gap, which captures how gender differences in hours translate into promotion outcomes under the estimated gender-neutral promotion policy. Panel A applies the decomposition separately by associate year, treating each year's cohort as if it were the promotion-eligible cohort. Panel B applies the decomposition to eighth-year associates using either total billable hours or FTE billable hours as the performance measure. The horizontal line marks parity at zero. **Data:** Law Firm Dataset and promotion rates constructed in Appendix Table 1.

4 Endogenizing Labor Supply

Our primary objective is to evaluate how promotion policies affect gender promotion gaps and firm productivity. Doing so requires understanding not only how these policies select workers at the promotion stage, but also how workers adjust their labor supply decisions during the pre-promotion years. Therefore, whereas Section 3 documented descriptive evidence on the observed promotion gap holding worker behavior fixed, this section develops a dynamic model of labor supply to capture how workers respond to promotion incentives.

4.1 Model

Time is discrete and finite, with periods $t = 1, \dots, T$. Workers are indexed by i and enter the firm in period $t = 1$. Periods $t < T$ are pre-promotion periods, while period T is the promotion period. In each pre-promotion period, workers choose a contract and number of hours to work under that contract. On the firm's side, we maintain the structure of the gender-neutral, hours-based promotion policy estimated in Section 3.2 and use the model to quantify the incentives workers face under this policy. To simplify notation, we suppress the i subscript on workers' choices and individual state variables below.

Contracts. We model the contract choice set as having a full-time and part-time option within the firm and an outside option:

$$\mathcal{J} = \{\mathcal{F} \text{ (full-time)}, \mathcal{P} \text{ (part-time)}, \mathcal{L} \text{ (leave)}\}.$$

Each contract is associated with a share of days worked, ϕ_j , and a compensation schedule, $w_{jt}(z)$, made up of a fixed salary and a bonus that depends on hours. We assume that workers who leave do not re-enter the firm.

In the model, in addition to choosing to leave, workers can also be fired by the firm. We model the probability of being fired in period $t+1$ as decreasing in period- t FTE hours, such that firing risk is adjusted for whether the worker is on a full-time or part-time contract:

$$q_{t+1}(z_t, j_t) = 1 - [1 + \exp \{-\gamma_q(z_t/\phi_{j_t} - c_z^q)\}]^{-1} \quad (3)$$

where γ_q governs the steepness of the firing rule and c_z^q is the FTE performance level associated with a 50% firing probability.²⁴

²⁴Both firing risk, $q_{t+1}(\cdot)$, and promotion probability, $p_T(\cdot)$, are common across workers and evaluated at individuals' choices and states.

Firm Capital. To aggregate hours-based performance across years, we define a “firm capital” term that captures factors like reputation, firm-specific human capital, and signaling accumulated over time at the firm. We assume this firm capital evolves with total (not FTE) hours according to:

$$h_{t+1}(h_t, z_t) = (1 - \delta)h_t + z_t, \quad (4)$$

where $\delta \in [0, 1]$ governs the depreciation of past hours.

We then adapt the gender-neutral promotion policy estimated in Section 3.2 to make promotion probabilities a function of firm capital instead of eighth-year hours:

$$p_T(h_T) = [1 + \exp \{-\gamma_p(h_T - c_h^p)\}]^{-1}, \quad (5)$$

where γ_p governs the perceived steepness of the promotion policy and c_h^p is the firm capital level associated with a 50% promotion probability.²⁵

Timing and State Variables. The timing within each pre-promotion period $t < T$ runs as follows. First, at the beginning of each period $t \geq 2$, the firm fires workers based on last period’s performance according to the rule in (3). Second, workers still employed at the firm observe their period- t state, given by firm capital h_t from (4) and a time-varying, persistent latent state η_t that captures life events outside of work, such as parenthood. They also observe the taste shocks ε_{jt} . Third, workers employed at the firm choose either an inside-firm contract $j_t \in \{\mathcal{F}, \mathcal{P}\}$ and hours z_t , or the outside option $j_t = \mathcal{L}$. Workers in the outside option work a fixed number of hours z_o .²⁶ Finally, for workers who remain at the firm, period- t choices determine next period’s firm capital, $h_{t+1}(h_t, z_t)$, and firing risk, $q_{t+1}(z_t, j_t)$.

Worker’s Problem. Workers differ by a permanent type n , which is drawn at entry and captures heterogeneity in workers’ willingness or ability to supply long hours.²⁷ This type enters the disutility of work, $v_t(z_t, j_t; n, \eta_t)$, as specified below. Given type n and state

²⁵Appendix Section B.2 shows that workers can *perceive* a threshold rule promotion policy to be smoother than it is under incomplete information about others’ performance. Therefore, we estimate γ_p below and do not simply adopt the threshold rule estimated in Section 3.2.

²⁶We set outside option hours to a 40-hour workweek-equivalent, motivated by After the JD evidence that most exiting lawyers move to non-firm jobs, as shown in Appendix Table 3.

²⁷This parameter captures underlying heterogeneity in things like ability (Kleven 2016), the opportunity cost of non-work time (Cook, Diamond, Hall, List, and Oyer 2020), preferences for competition (Niederle and Vesterlund 2007), and also nests a world in which higher quality workers are able to obtain “cheaper” marginal hours because partners prefer to work with them.

variables (h_t, η_t) , flow utility from choosing contract $j_t \in \mathcal{J}$ and hours z_t is:

$$u_t(j_t, z_t; n, \eta_t) = w_{j_t}(z_t) + \psi_{j_t}(\eta_t) - v_t(z_t, j_t; n, \eta_t) + \varepsilon_{j_t},$$

where $w_{j_t}(z_t)$ is compensation, $\psi_{j_t}(\eta_t)$ is the non-wage amenity value of contract j_t , $v_t(z_t, j_t; n, \eta_t)$ is the disutility of supplying hours, and $\varepsilon_{j_t} \sim \text{T1EV}(0, \sigma_\varepsilon)$ is an iid contract taste shock.

In the terminal promotion period, let $V^{P,n}$ denote the value of being promoted and $V^{NP,n}$ denote the value of not being promoted.²⁸ Therefore, for workers in the firm of type n :

$$V_T^n(h_T, \eta_T) = p_T(h_T)V^{P,n}(\eta_T) + [1 - p_T(h_T)]V^{NP,n}(\eta_T).$$

Consistent with industry evidence that part-time equity partner roles are uncommon (Calvert, Chanow, and Marks 2009), the model treats partnership as a single terminal full-time role rather than differentiating between full-time and part-time partner contracts.

Given this, we can define an n -specific value function $V_t^n(h_t, \eta_t)$ for workers in the firm after observing their latent state η_t , but ex ante with respect to current-period taste shocks. For $t < T$, the worker's problem can be written recursively as:

$$V_t^n(h_t, \eta_t) = \mathbb{E}_{\varepsilon_t} \left[\max_{j_t, z_t} \begin{cases} u_t(j_t, z_t; n, \eta_t) + \beta q_{t+1}(z_t, j_t) \mathbb{E} [V_{t+1}^{\mathcal{L},n}(\eta_{t+1}) | \eta_t] \\ \quad + \beta (1 - q_{t+1}(z_t, j_t)) \mathbb{E} [V_{t+1}^n(h_{t+1}, \eta_{t+1}) | \eta_t] & \text{if } j_t \in \{\mathcal{F}, \mathcal{P}\} \\ u_t(\mathcal{L}, z_o; n, \eta_t) + \beta \mathbb{E} [V_{t+1}^{\mathcal{L},n}(\eta_{t+1}) | \eta_t] & \text{if } j_t = \mathcal{L} \end{cases} \right],$$

where $q_{t+1}(z_t, j_t)$ is defined in (3), $h_{t+1}(h_t, z_t)$ is defined in (4), $z_t = z_o$ when $j_t = \mathcal{L}$, and expectations over η_{t+1} are taken with respect to the transition process for the latent state. The outside option value satisfies: $V_t^{\mathcal{L},n}(\eta_t) = \mathbb{E}_{\varepsilon_{\mathcal{L}t}} [u_t(\mathcal{L}, z_o; n, \eta_t)] + \beta \mathbb{E} [V_{t+1}^{\mathcal{L},n}(\eta_{t+1}) | \eta_t]$, with terminal condition $V_T^{\mathcal{L},n}(\eta_T) = V^{NP,n}(\eta_T)$.

4.2 Parameterization

Contracts. We assume that full-time and outside option work have $\phi_{\mathcal{F}} = \phi_{\mathcal{L}} = 1$, while part-time work has $\phi_{\mathcal{P}} \in (0, 1)$. To align with compensation practices in the industry, we

²⁸These values reflect both pecuniary and non-pecuniary components – i.e., the disutility associated with post-promotion hours and the amenity value of leaving.

also specify:

$$\begin{aligned} w_{\mathcal{F}t}(z) &= s_t + b_t \times \mathbb{I}(z > \underline{z}), \\ w_{\mathcal{P}t}(z) &= \phi_{\mathcal{P}} s_t + \phi_{\mathcal{P}} b_t \times \mathbb{I}(z > \phi_{\mathcal{P}} \underline{z}), \text{ and} \\ w_{\mathcal{L}t} &= s_{ot}, \end{aligned}$$

where s_t is the within-firm salary for associates in year t , b_t is the annual bonus paid when hours exceed the full-time threshold $\underline{z}$, and s_{ot} is the outside option salary.²⁹ Part-time workers receive proportional salary and bonus compensation and face a proportional bonus threshold. We set partner compensation, π_T , equal to the profits-per-partner measure reported in the Am Law 100 data.³⁰ Table 2 reports the values of all fixed model inputs.

Worker Heterogeneity. Worker heterogeneity shows up along two dimensions. First, each worker’s permanent type n is drawn at entry from a lognormal distribution:

$$n \sim \text{Lognormal}(\mu_n, \sigma_n^2).$$

Second, a latent state $\eta_t \in \{\eta_L, \eta_H\}$ captures exogenous life events such as parenthood, illness, or caregiving responsibilities and evolves according to a two-state Markov process with transition matrix:

$$M_{\eta}(p_{LH}, p_{HL}) = \begin{bmatrix} (1 - p_{LH}) & p_{LH} \\ p_{HL} & (1 - p_{HL}) \end{bmatrix},$$

independent of n .

Amenities. We normalize the amenity value of full-time work to zero, so the part-time and outside option amenities are measured relative to the default arrangement of full-time work within the firm. Specifically, $\psi_{\mathcal{F}}(\eta) = 0$, while $\psi_{\mathcal{P}}(\eta)$ and $\psi_{\mathcal{L}}(\eta)$ vary with the latent state. These terms capture residual non-wage values of different contracts, including factors like office culture, schedule flexibility, and the value of working fewer days.³¹

²⁹We calibrate the outside option salary using the average salary change at exit from large law firms in the Revelio Labs data. See Appendix Figure 5.

³⁰We then set an implicit partner hourly wage w^{Pn} such that optimal partner hours satisfy $w^{Pn} z^{Pn*} = \pi_T$.

³¹This is one way of capturing the fixed cost of working additional days. Alternatively, schedules could enter the disutility of work discussed below. However, because we do not observe the outside option schedule $\phi_{\mathcal{L}}$, we assume $\phi_{\mathcal{L}} = 1$ and allow the value of different schedules to be flexibly captured by the amenity term.

Disutility of Work. We model the disutility of work as iso-elastic in FTE hours:³²

$$v_t(z, j; n, \eta_t) = \frac{n/\kappa(\eta_t)}{1 + 1/\epsilon} \left(\frac{z/\phi_j}{n/\kappa(\eta_t)} \right)^{1+1/\epsilon}, \quad (6)$$

where ϵ is the labor supply elasticity, $\kappa(\eta_t)$ shifts the cost of work across latent states, and n captures permanent heterogeneity in the cost of supplying hours. We assume $\kappa(\eta_L) = 1$.

Gender-Specific Primitives. We allow the labor supply elasticity (ϵ), the distribution of permanent types (μ_n, σ_n), the latent state transition matrix $M_\eta(p_{LH}, p_{HL})$, the scale of the contract-specific taste shocks (σ_ϵ), and the mappings from latent states into amenities (ψ_P, ψ_L) and effort costs (κ) to differ by gender. These gender-specific primitives allow the model to match the observed differences in retention, contract choice, and performance distributions documented above.

Other Model Inputs. Table 2 shows model objects that are fixed by institutional features of large law firms and therefore assumed rather than estimated.

Table 2: Fixed Model Objects

Model Object	Source	Value
Salary (\$k)	2023 Cravath Scale	$s_t = \{215, 225, 250, 295, 345, 370, 400, 415\}$
Bonus (\$k)	2023 Cravath Scale	$b_t = \{20, 30, 57.5, 75, 90, 105, 115, 115\}$
Bonus Threshold (norm. hrs)	Data Provider	$\underline{z} = 0.65$
Profits per Partner (\$k)	2024 Am Law 100	$\pi_T = \text{redacted} \in [2000, 4000] \times 10 \text{ yrs}$
Outside Option Salary	Revelio Labs	$s_t^o = (1 - 0.26) \cdot s_L$
Promotion Year	Vault (2009)	$T = 9$
Outside Option Hours	Bloomberg (2025)	$z_o = 40 \text{ hrs} \cdot 50 \text{ wks} \cdot \frac{2}{3} \text{ billable hr share}$
Discount Rate		$\beta = 0.95$
Part-Time Share	Law Firm Dataset	$\phi = 69\%$

Notes: Table reports model inputs that are fixed outside the estimation. Salary and bonus schedules are in thousands of dollars and correspond to associate years 1 through 8. The bonus threshold is reported in normalized annual hours units. Profits per partner are in thousands and enter the terminal value of partnership. Outside option salary is set equal to a fraction of the worker’s within-firm salary in the year of exit, where s_L denotes that within-firm salary. The part-time share is the mean contracted share among associates with a contracted share less than or equal to 80% of full-time work. We calculate the share of billable hours to correct for non-billable attorneys in the Bloomberg sample as: $36 \div \left[\frac{48-40 \cdot 0.42}{0.58} \right] = 0.669$.

³²As shown in Appendix Section B.3, this functional form can be micro-founded as *hours per day* when the value of working fewer *days per week* enters through the amenity term. This setup captures the fact that supplying the same total hours as a part-time worker requires working more hours per workday.

4.3 Estimation

We estimate the model using SMM (McFadden 1989). The full parameter vector is:

$$\Theta = (\Theta^f, \Theta^m, \delta, c_h^p, \gamma_p, c_z^q, \gamma_q),$$

where the gender-specific block is:

$$\Theta^g = (\epsilon^g, \mu_n^g, \sigma_n^g, p_{LH}^g, p_{HL}^g, \psi_P^g(\eta_L), \psi_P^g(\eta_H), \psi_L^g(\eta_L), \psi_L^g(\eta_H), \kappa^g(\eta_H), \sigma_\epsilon^g), \quad g \in \{f, m\}.$$

For each candidate parameter vector Θ , we solve the worker's dynamic problem separately by gender using backward induction at the corresponding $(\Theta^g, \delta, c_h^p, \gamma_p, c_z^q, \gamma_q)$. We then use the resulting policy functions to propagate the distribution of workers over contracts, types, and firm capital forward through time.³³ Finally, we compute model-implied moments. The estimator solves:

$$\hat{\Theta} = \arg \min_{\Theta} [m^{data} - m^{sim}(\Theta)]' W [m^{data} - m^{sim}(\Theta)],$$

where W is a diagonal weight matrix.³⁴ Additional details on the estimation can be found in Appendix Section F.1. The next subsection describes the main sources of variation in the moment vector that discipline different components of Θ .

4.4 Identification

This subsection provides identification intuition for the joint SMM estimator rather than a formal identification proof. The moment vector contains information on hours, bonus bunching, contract shares, contract transitions, next-year departures, and aggregate promotion outcomes, as described in detail in Appendix Section F. These moments jointly discipline the parameter vector Θ , but different features of the data are most informative about different components of the model.³⁵ In discussing these features below, we suppress gender superscripts to simplify notation.

The first set of parameters governs baseline labor supply through the distribution of the latent type n , parameterized by (μ_n, σ_n) , and the curvature of hours costs, ϵ . These

³³Both the T1EV taste shock and the η Markov transitions can be integrated out – the T1EV shock via a closed-form logsumexp inside the value function, and the η transitions by applying the Markov matrix M_η to the η dimension of the distribution each period.

³⁴The 11 gender-specific parameters per gender (Θ^f and Θ^m) and the 5 gender-neutral parameters $(\delta, c_h^p, \gamma_p, c_z^q, \gamma_q)$ are estimated jointly in a single optimization. Minimization is performed using global derivative-free optimization, specifically a differential evolution algorithm.

³⁵Appendix Figure 9 illustrates these relationships through comparative statics.

parameters are disciplined jointly by the empirical full-time hours distribution. Within that distribution, ϵ primarily changes the local response to the bonus threshold, $\underline{z}$, while (μ_n, σ_n) shift the broader level and dispersion of hours, including away from the threshold. The model must fit this local bunching together with the rest of the full-time hours distribution. Thus, variation around the threshold helps separate responsiveness to the bonus incentive from baseline heterogeneity in hours.

The second set of parameters governs state-dependent departures from baseline labor supply. The latent state, $\eta \in \{\eta_L, \eta_H\}$, is not observed, nor do we observe transitions between latent states. Instead, the transition matrix M_η , together with the state-dependent amenities $\psi(\eta)$ and hours cost shifter $\kappa(\eta)$, determines how latent states translate into part-time work, leaving, and low-hours spells, given model-implied baseline labor supply and empirical compensation offers. Within this block, the amenity terms $\psi_{\mathcal{P}}(\eta)$ and $\psi_{\mathcal{L}}(\eta)$ primarily affect the relative value of part-time work and leaving, while the hours cost shifter $\kappa(\eta)$ primarily affects FTE hours conditional on contract. The transition probabilities p_{LH} and p_{HL} then affect the frequency and duration of spells in each state. Thus, we include the empirical levels of part-time work and leaving, differences between part-time and full-time FTE hours, and conditional contract transition probabilities to jointly inform this set of parameters. Any residual contract movements are informative about the scale of idiosyncratic taste shocks, σ_ϵ .

The third set of parameters governs workers' perceptions of the firm policies that map hours into firing and promotion risk. These perceived policies are not directly observed. Instead, they are disciplined jointly by the empirical time path of hours, year-to-year departures from the firm, and aggregate promotion outcomes. The depreciation rate δ determines how observed hours are converted into accumulated firm capital: when δ is higher, past hours decay more quickly, so late career hours matter more for promotion than earlier career hours. The empirical evolution of hours over time is therefore informative about δ . Together with this accumulation process, the promotion parameters (c_h^p, γ_p) map accumulated firm capital into promotion probabilities. The empirical promotion rate helps locate the promotion threshold, c_h^p , while observed bunching around that threshold is informative about the steepness of the perceived promotion policy, γ_p .³⁶ Similarly, the firing parameters (c_z^q, γ_q) map FTE hours into next-year firing risk. The empirical FTE hours of associates who depart in the following year help locate the firing threshold, c_z^q , while observed bunching around the implied firing margin is informative about the steepness parameter, γ_q .

³⁶The empirical bunching moments are calculated at the candidate thresholds for each optimization step.

4.5 Parameter Estimates and Model Fit

Table 3 presents the parameter estimates. To start, the estimates suggest that men and women do not differ meaningfully in baseline labor supply. The estimated labor supply elasticity is nearly identical across genders, with $\hat{\epsilon}^f = 0.013$ for women and $\hat{\epsilon}^m = 0.014$ for men. These estimates are consistent with the prior bunching literature; Kleven (2016) discusses how even sizable bunching can yield small corresponding elasticities – e.g., Chetty, Friedman, Olsen, and Pistaferri (2011) estimate bunching elasticities of 0.01 to 0.02. In the context of the model, these estimates imply that the median associate would be willing to work up to 125 additional billable hours per year, or 2.5 additional billable hours per week, in order to receive the bonus. The estimated distributions of the permanent type n , which capture persistent heterogeneity in baseline labor supply, are also similar across genders, with $\hat{\mu}_n^f = 0.53$ and $\hat{\mu}_n^m = 0.54$. The male distribution is slightly more dispersed, with $\hat{\sigma}_n^m = 0.18$ compared to $\hat{\sigma}_n^f = 0.13$. Overall, though, the model does not attribute the gender promotion gap primarily to large gender differences in baseline labor supply or permanent type.

Instead, the largest gender differences appear in latent state transitions. Women and men enter η_H at similar rates, with $\hat{p}_{LH}^f = 0.21$ and $\hat{p}_{LH}^m = 0.22$, implying that both groups first enter η_H after about 4.5 years on average. However, women are substantially less likely to return to η_L each year, with $\hat{p}_{HL}^f = 0.12$ compared to $\hat{p}_{HL}^m = 0.40$, implying an expected return time of about 2.5 years for men and about 8.3 years for women. That is, once women enter η_H , the state can be persistent enough to extend through the rest of their associate years. This mirrors the part-time dynamics documented in Section 3.3: men and women both transition into part-time work, but women’s part-time spells are substantially more persistent.

The latent state η_H can be understood as capturing child-related demands for time flexibility. As discussed above, the After the JD survey data show that most lawyers working part-time in large law firms cite children as the reason for doing so.³⁷ Given this, the implied motherhood penalty (estimated in Appendix E) is approximately 14% and broadly consistent with prior estimates (e.g., Kleven, Landais, and Egholt-Søgaard 2019, 2021).

The remaining state-dependent parameters describe how η_H affects associate choices. The estimated hours cost shifter is close to one for both genders, with $\hat{\kappa}^f(\eta_H) = 0.97$ and $\hat{\kappa}^m(\eta_H) = 1.00$, implying that the latent state does not meaningfully lower optimal FTE hours.³⁸ Instead, the amenity estimates imply that part-time work and leaving become

³⁷This interpretation is also consistent with prior work documenting gender differences in time spent on childcare and home production (e.g., Ciasullo and Ucciolli 2025). And the transitory nature of these part-time spells is consistent with prior work showing that mothers’ labor supply tends to rebound as children age (Goldin, Kerr, and Olivetti 2022).

³⁸For example, this suggests that an associate who would bill X hours on a full-time contract would

substantially more attractive in η_H . For women, the estimated part-time amenity rises from 19 in η_L to 2269 in η_H , while the leaving amenity rises from 183 to 1499. For men, the estimated part-time amenity rises from 11 to 4931, while the leaving amenity rises from 16 to 1002. In the model, this implies that moving a median associate from η_L to η_H increases the probability of choosing part-time work by 38 percentage points for women and 42 percentage points for men, and increases the probability of leaving by 2 percentage points for women and close to zero percentage points for men.³⁹ Thus, the estimates suggest that associates value the option to reduce days at the firm or to move to non-firm work *specifically* while in η_H . Finally, the scale of idiosyncratic taste shocks is similar across genders, with $\hat{\sigma}_\varepsilon^f = \hat{\sigma}_\varepsilon^m = 19.5$, suggesting that gender differences are not driven by idiosyncratic shocks.

The final set of estimates describes the incentives associates face around promotion and firing. The estimated firm capital depreciation rate, $\hat{\delta} = 0.33$, implies that an hour retains about two-thirds of its weight in firm capital each year. As a result, early-career hours remain relevant for promotion, but later-year hours receive more weight. For example, fourth-year hours retain only about 20% of their firm capital value for promotion compared to eighth-year hours. Given this accumulation process, the estimated promotion threshold is $\hat{c}_h^p = 0.58$ in normalized annual hour equivalents, and the estimated steepness is $\hat{\gamma}_p = 1.68$. These estimates imply that associates perceive promotion probabilities to rise smoothly with accumulated firm capital, with promotion becoming substantially more likely only at higher levels of accumulated hours. This smoothness is consistent with uncertainty about others' performance, which can make a threshold-like promotion rule appear smooth from the worker's perspective, as discussed in Appendix Section B.2. The estimated firing threshold, $\hat{c}_z^q = 0.17$, is located at low FTE hours, with estimated steepness $\hat{\gamma}_q = 2.94$, implying a smooth perceived separation risk.

Taken together, these parameters imply realistic labor supply patterns, as illustrated in the model-implied policy functions in Appendix Figures 10 and 11. For both genders, associates with higher n are generally less likely to exit and more likely to supply long hours. However, the model generates different selection patterns into full-time and part-time work across genders. For men, the probability of full-time work rises much more sharply with n than the probability of part-time work, so men with higher n who remain at the firm mostly select into full-time contracts. In contrast, high-type women who remain at the firm retain a substantial probability of choosing part-time work.⁴⁰

Overall, the estimated model also fits the data well. To assess model fit, we simulate

generally bill $0.8 \times X$ hours on an 80% part-time contract.

³⁹These percentage changes are averaged across associate years 2 through 8.

⁴⁰Appendix Section G discusses associates on the margin of partnership.


associate careers using the estimated parameters and compare the simulated moments to their empirical counterparts. We focus here on the three key objects from Section 3 – the promotion gap decomposition, contracts including exit, and hours – while Appendix Table 5 compares the full set of empirical and simulated moments. Figure 6 shows that the model matches the empirical promotion gap and its retention and performance components almost exactly. In both the data and the model, the performance component accounts for most of the promotion gap, while the retention component accounts for the remainder. This fit reflects the model’s ability to reproduce the contract and hours patterns underlying the gap. Figure 7 compares empirical and simulated contract shares over the associate career. The model captures the main patterns in the data: full-time work declines with tenure, leaving rises with tenure, and part-time work is more prevalent among women than men. The model overpredicts early-career female exit, though the fit improves in later years. Figure 8 compares empirical and simulated full-time hours distributions in selected years, showing that the model reproduces the general location and dispersion of empirical hours.

Table 3: Model Results: Parameter Estimates

Parameter	Female	Male
Labor supply elasticity, ϵ^g	0.013 (0.000)	0.014 (0.000)
Mean type, μ_n^g	0.53 (0.01)	0.54 (0.00)
Dispersion of type, σ_n^g	0.13 (0.00)	0.18 (0.01)
$\eta_L \rightarrow \eta_H$ transition probability, p_{LH}^g	0.21 (0.02)	0.22 (0.04)
$\eta_H \rightarrow \eta_L$ transition probability, p_{HL}^g	0.12 (0.01)	0.40 (0.04)
Hours cost shifter in η_H , $\kappa^g(\eta_H)$	0.97 (0.01)	1.00 (0.01)
Part-time amenity in η_L , $\psi_{\mathcal{P}}^g(\eta_L)$	19 (12)	11 (2)
Part-time amenity in η_H , $\psi_{\mathcal{P}}^g(\eta_H)$	2269 (174)	4931 (17)
Leaving amenity in η_L , $\psi_{\mathcal{L}}^g(\eta_L)$	183 (35)	16 (3)
Leaving amenity in η_H , $\psi_{\mathcal{L}}^g(\eta_H)$	1499 (92)	1002 (118)
Scale of idiosyncratic taste shocks, σ_ε^g	19.5 (7.8)	19.5 (3.7)
Firm capital depreciation, δ	0.33 (0.01)	
Promotion policy threshold, c_h^p	0.58 (0.00)	
Promotion policy steepness, γ_p	1.68 (0.09)	
Firing rule threshold, c_z^q	0.17 (0.00)	
Firing rule steepness, γ_q	2.94 (0.15)	


Notes: Table reports parameter estimates from the SMM procedure. Female and male columns report gender-specific parameters. Parameters without gender-specific columns are common across genders. Standard errors are shown in parentheses and are computed using $B = 80$ bootstrap replications that resample associates with replacement. The policy location terms, c_h^p and c_z^q , are reported in normalized annual hour equivalents. **Data:** Law Firm Dataset and supplemental model objects from Table 2.

Figure 6: Model Fit – Promotion Gap Decomposition




Notes: Figure compares the empirical promotion gap decomposition to the decomposition implied by the estimated model. The hashed bars show the empirical decomposition from Table 1. Solid bars show the corresponding decomposition in simulated data using the parameter estimates in Table 3. The decomposition separates the overall promotion gap into the retention gap and the performance gap, as defined in Section 3.3. Positive gap values indicate female under-representation relative to parity. **Data:** Empirical moments from the Law Firm Dataset and promotion rates constructed in Appendix Table 1; simulated moments from the estimated model summarized in Table 3.

Figure 7: Model Fit – Contract Shares



Notes: Figure compares empirical contract shares to the contract shares implied by the estimated model. Panel A reports results for men and Panel B reports results for women. Solid lines show the empirical shares. Dashed lines show the corresponding simulated shares using the parameter estimates in Table 3. Shares are computed within each gender and associate year cell. **Data:** Empirical moments from the Law Firm Dataset; simulated moments from the estimated model summarized in Table 3.

Figure 8: Model Fit – Full-Time Hours



Notes: Figure compares empirical full-time hours distributions to the distributions implied by the estimated model. Solid lines show the empirical distribution of hours among full-time associates in the specified year. Dashed lines show the corresponding simulated distribution using the parameter estimates in Table 3. Panels report the gender-specific distributions for sixth and eighth-year associates. Hours are normalized to preserve confidentiality. **Data:** Empirical moments from the Law Firm Dataset; simulated moments from the estimated model summarized in Table 3.

5 The Impact of Alternative Promotion Policies

We now use the estimated model to evaluate how counterfactual promotion policies would affect both representation and firm productivity. This joint focus is important because policies that reduce the promotion gap may differ substantially in their costs to firms.

To assess these costs, we use the model to trace two productivity-relevant metrics. The first is total associate billable hours, which approximates short-run revenue generated by associates and captures the output consequences of changing associates' career incentives.

The second is the average permanent type n among promoted associates, which summarizes the productivity of hypothetical partner cohorts in terms of their persistent willingness or ability to supply long hours and captures the selection role of promotion policies.

These metrics do not define the firm’s full objective. Solving the firm’s profit maximization problem would require a much larger model to map promotion policies into not only associate hours and partner type, but also client demand, competition across firms, recruiting, post-promotion labor supply shocks, and the long-run value of the partnership.⁴¹ Accordingly, these counterfactuals should be interpreted as policy frontiers rather than solutions to the firm’s optimization problem.

We consider three main counterfactuals. The first accounts for FTE hours in promotion-relevant firm capital. The second imposes gender-based representation targets among promoted partners. The third caps the hours that an associate may work annually. These three instruments isolate distinct margins through which firms could try to reduce the promotion gap: how performance is measured, how promotions are allocated conditional on performance, and how associates compete. They also correspond to popular policy proposals related to gender promotion gaps in high-hours professions. For each counterfactual promotion policy, we recalibrate the promotion threshold to keep the expected size of the promoted class fixed and trace how the female share of promotions, short-run associate revenue, and partner productivity change relative to the status quo.⁴²

5.1 FTE Hours

The first counterfactual varies the weight placed on FTE hours relative to total billable hours in promotion-relevant firm capital. As discussed in Section 3.2, total billable hours are directly tied to short-run associate revenue and provide a signal of worker type, but total hours mechanically lower measured performance for part-time associates even when they supply high hours relative to their contracted schedule. We therefore consider an FTE hours metric that preserves the hours-based promotion system while accounting for the associate’s contracted share of full-time work.

Let $\lambda \in [0, 1]$ denote the weight placed on FTE hours in firm capital and $(1 - \lambda)$ denote

⁴¹The firm could also change firm size or the structure of partnership, both of which are held fixed here. Appendix H considers counterfactuals that change the structure of partnership.

⁴²Holding the promotion threshold fixed would change the expected size of the promoted class and therefore workers’ promotion incentives. Recalibrating the threshold keeps the expected number of promotion slots fixed, allowing us to focus on responses to the promotion policy change under consideration.

the weight placed on total hours. Under this policy, firm capital from (4) now evolves as:

$$h_{t+1}^{CF} = (1 - \delta)h_t^{CF} + (1 - \lambda)z_t + \lambda \frac{z_t}{\phi_{j_t}},$$

where ϕ represents the share of days worked, as above. When $\lambda = 0$, the policy coincides with the baseline total hours-based approach. When $\lambda = 1$, the policy evaluates associates using only FTE billable hours.


Figure 9 shows how outcomes change as we place greater weight on FTE hours. At low values of λ , the female share of promotions is essentially unchanged, but it rises once FTE hours receive substantial weight – e.g., to 41% at $\lambda = 1$. Higher values of λ shift promotions toward associates with high permanent types but lower total hours because of part-time spells, many of whom are women, thereby reducing the part-time penalty emphasized in prior work (e.g., Goldin and Katz 2016). Under the FTE hours-only policy, where $\lambda = 1$, the female promotion rate increases by 19%, while the male promotion rate falls by 9%.

On the productivity side, total associate billable hours change by less than 0.4% across all values of λ . This occurs because high- λ policies slightly increase female retention but also increase part-time take-up, leaving short-run associate output essentially unchanged on net.⁴³ Partner productivity also improves, as the average n among promoted associates rises by 2% at $\lambda = 0.25$ and 3% at $\lambda = 0.5$ and above. This reflects the same reallocation mechanism discussed above: the FTE hours metric shifts promotion toward associates whose total hours are lower because of part-time spells, but whose permanent types remain high. Thus, promotion policies that account for FTE hours can narrow the promotion gap while preserving short-run associate output and improving selection on n .


Overall, the FTE hours policy performs well on the productivity metrics tracked here. However, as discussed above, other margins that the firm cares about may co-move with these metrics. For example, total hours may screen for dimensions of future partner capacity outside the current model, such as flexibility constraints that arise during the partner years. However, Appendix Figure 15 shows that simulations of partner hours closely track the average n metric.

⁴³See Appendix Figure 13 for details.


Figure 9: Effects of Placing Greater Weight on FTE Hours



(A) Female Share of Promotions



(B) Change in Short-Run Revenue



(C) Change in Partner Productivity

Notes: Figure reports counterfactual effects of placing greater weight on FTE hours in promotion-relevant firm capital. The horizontal axis is the FTE-hours weight, λ . When $\lambda = 0$, the promotion metric coincides with the baseline total hours metric. When $\lambda = 1$, the promotion metric uses only FTE billable hours. For each value of λ , all other model parameters are held fixed and the promotion threshold is recalibrated to preserve the expected size of the promoted class. Each panel reports the percentage change from the baseline simulated status quo. Panel A reports the female share of promotions. Panel B reports the change in short-run associate revenue, measured by total associate billable hours. Panel C reports the change in partner productivity, measured by average permanent type n among promoted associates. **Data:** Simulations from the estimated model summarized in Table 3.

5.2 Representation Targets

The second counterfactual varies the target female share among promoted partners. Gender-based representation targets have often been proposed as a policy response to gender disparities in promotion and leadership (e.g., Smale and Miller 2015). Existing evidence

also suggests that such policies may be able to increase women’s representation while maintaining, or even improving, the quality of selected leaders (Bertrand, Black, Jensen, and Lleras-Muney 2019). We therefore consider a counterfactual policy that sets a target female share, $\bar{q}$, among newly promoted partners.

A parity target, $\bar{q} = 0.5$, is one natural benchmark. However, less stringent targets are common in practice and may have different effects.⁴⁴ We therefore vary $\bar{q}$ over a grid and implement each target by allowing promotion thresholds to differ by gender, while requiring that the expected size of the promoted class remains at its status quo level.

Figure 10 shows how outcomes change as the target female share $\bar{q}$ increases. Representation targets substantially reduce the promotion gap by construction – e.g., the female share of promotions is 35% at $\bar{q} = 0.35$ and 50% at $\bar{q} = 0.5$. Total associate billable hours increase slightly across the range of targets, reflecting a small increase in pre-promotion retention, with female retention gains more than offsetting male retention losses. However, the effects on partner productivity depend on the target. Moderate targets preserve, or slightly improve, average n among promoted associates, while more aggressive targets weaken selection on n – e.g., under the parity target, average n among promoted associates falls by 0.3%.


Representation targets operate through a different margin than the FTE policy. Rather than changing how associates’ hours accumulate into the promotion metric, they change how the firm maps a given level of firm capital into promotion. As $\bar{q}$ rises, the policy lowers the promotion threshold for women relative to men in order to offset retention and performance gaps. At moderate targets, the marginal women promoted under the policy tend to have relatively high permanent types despite lower total hours, so average n is preserved or slightly improved. At stricter targets, the female threshold must be pushed farther down and the male threshold farther up. The policy therefore increasingly promotes women with lower n than the men displaced at the margin.⁴⁵

Thus, representation targets involve a tradeoff. Moderate targets can substantially increase women’s representation while preserving short-run associate revenue and partner productivity. More aggressive targets generate larger representation gains, but do so by changing the allocation of promotion slots rather than improving selection on partner productivity.


⁴⁴E.g., the Norwegian board quota studied by Bertrand, Black, Jensen, and Lleras-Muney (2019) required at least 40% representation of each gender, as did the French law on academic hiring committees considered in Deschamps (2024). Schaefer and Mankiw (2026) also study a 40% quota for *male* primary school teachers.

⁴⁵This may help explain the backlash to quotas discussed in prior work – e.g., Leibbrandt, Wang, and Foo (2018), Ip, Leibbrandt, and Vecchi (2020), Deschamps (2024), and González and Martínez-Marquina (2026).


Figure 10: Effects of Representation Targets



(A) Female Share of Promotions



(B) Change in Short-Run Revenue



(C) Change in Partner Productivity

Notes: Figure reports counterfactual effects of representation targets among promoted partners. The horizontal axis is the target female share of the promoted class, $\bar{q}$. For each value of $\bar{q}$, the policy is implemented by allowing promotion thresholds to differ by gender while recalibrating thresholds to preserve the expected size of the promoted class. All other model parameters are held fixed. Each panel reports the percentage change from the baseline simulated status quo. Panel A reports the female share of promotions. Panel B reports the change in short-run associate revenue, measured by total associate billable hours. Panel C reports the change in partner productivity, measured by average permanent type n among promoted associates. **Data:** Simulations from the estimated model summarized in Table 3.

5.3 Hours Caps

The third counterfactual varies the maximum number of hours that associates can work. Hours caps have been proposed as a policy response to the “rat race” that can arise in careers like private sector law and investment banking (e.g., Saeedy 2024). They may also narrow


gender gaps if they weaken nonlinear returns to long, inflexible hours (Goldin 2014).⁴⁶ Unlike the first two policies, hours caps do not directly change the promotion metric or allocation rule in period T . Instead, they limit the extent to which associates can compete through high billable hours throughout their pre-promotion years.⁴⁷ We consider a range of hours caps, z_{cap} , under which associates may not choose $z_t > z_{\text{cap}}$.

Figure 11 reports the hours caps in weekly hour equivalents and shows how outcomes change as the cap varies. Tight caps reduce the promotion gap substantially – e.g., a 40-hour cap increases the female share of promotions to 47%, while a 60-hour cap only increases it to 35%. However, these reductions come at the cost of lower associate output – partly by construction and partly because of worker responses. For example, the 40-hour cap reduces total associate billable hours by 32%, with retention falling for both male and female associates, and part-time take-up among the associates who remain at the firm rising sharply. By preventing high- n associates from supplying additional hours, the cap also compresses the right tail of the performance distribution and limits associates’ ability to signal their type. As a result, average n among promoted associates falls by 8% under the 40-hour cap and remains below the status quo across the range of cap values. Thus, hours caps can narrow the promotion gap, but they do so by weakening both associate output and selection into partnership.


⁴⁶Wasserman (2023) also considers how hours caps affect gender differences in *entry* into a given profession.

⁴⁷Here, we hold fixed the focus on hours. However, in practice, compressing the hours distribution could induce firms to place greater weight on other performance measures.


Figure 11: Effects of Hours Caps



(A) Female Share of Promotions



(B) Change in Short-Run Revenue



(C) Change in Partner Productivity

Notes: Figure reports counterfactual effects of caps on associate hours. The horizontal axis reports the annual hours cap, z_{cap} , converted to weekly-hour equivalents for readability. Under each cap, associates may not choose hours above z_{cap} . For each value of z_{cap} , all other model parameters are held fixed and the promotion threshold is recalibrated to preserve the expected size of the promoted class. When more associates clear the promotion threshold than there are promotion slots, promotion is allocated randomly among associates above the threshold. Each panel reports the percentage change from the baseline simulated status quo. Panel A reports the female share of promotions. Panel B reports the change in short-run associate revenue, measured by total associate billable hours. Panel C reports the change in partner productivity, measured by average permanent type n among promoted associates. **Data:** Simulations from the estimated model summarized in Table 3.


5.4 Comparison Across Policies

Taken together, these results show that all three counterfactual policies can reduce the promotion gap – but through different margins and with different corresponding effects.

Placing greater weight on FTE hours narrows the promotion gap while preserving short-run associate revenue and improving partner productivity. Representation targets also reduce the gap with little effect on short-run associate revenue, but their effect on partner productivity depends on the target level. Hours caps reduce the promotion gap by directly limiting high-hour work, but this same mechanism lowers associate output and weakens selection into partnership.⁴⁸

Figure 12 plots how the productivity metrics co-move with the female share of promotions across the three policies and highlights the relative strength of the FTE hours policy: although it cannot fully close the promotion gap, it preserves short-run associate revenue and dominates the other policies in terms of partner productivity.

Figure 12: Representation and Productivity Across Counterfactual Policies



Notes: Figure reports how representation and productivity co-move across the three main counterfactual promotion policies. The horizontal axis reports the female share of promotions. Each connected marker corresponds to a different value of the relevant policy parameter: the weight placed on FTE hours (λ), the target female share among promoted partners ($\bar{q}$), or the annual hours cap (z_{cap}). The black star marks the simulated status quo. The vertical dashed line marks gender parity in promotions, and the horizontal dashed line marks the baseline value of each productivity metric. Panel A reports short-run associate revenue, measured by total associate billable hours. Panel B reports partner productivity, measured by average permanent type n among promoted associates. Both productivity metrics are normalized so that the baseline equals one. **Data:** Simulations from the estimated model summarized in Table 3.

Thus, promotion policies that account for part-time schedules may be a particularly promising response to gender promotion gaps in high-hours professions like private sector

⁴⁸Appendix Figure 14 reports the corresponding welfare effects. Weighting FTE hours primarily benefits high- n workers while reducing welfare for mid- n workers. Representation targets shift welfare toward women, especially higher n women, and away from higher n men. Hours caps impose the largest losses on high- n workers, especially high- n men, while their effects on lower n workers depend on how binding the cap is.

law. In the model, such policies improve selection by placing less weight on transitory reductions in associate hours and more weight on permanent worker characteristics relevant to long-term partner productivity.

An important implementation concern beyond the model, though, is whether firms can ensure that such policies are actually applied in promotion reviews and not undermined by discretionary evaluations or promoting managers' continued reliance on older promotion criteria (e.g., Thornton 2008; Leslie, Manchester, Park, and Mehng 2012). In light of this concern, moderate representation targets may provide a useful alternative since compliance with explicit targets may be easier to monitor.⁴⁹

6 Conclusion

We study how firm promotion policies shape gender promotion gaps in large law firms. We find that, in the status quo, differences in total billable hours account for most (69%) of the observed gap. These hours gaps are primarily driven by differences in the persistence of part-time work: men and women transition into part-time work at similar rates, but women remain part-time substantially longer. We then estimate a dynamic labor supply model, which similarly points to persistence as the key gender difference, and use it to evaluate counterfactual promotion policies. We find that several policies can substantially reduce the gap, but with different implications for firm productivity. Accounting for part-time schedules when measuring performance can cut the promotion gap in half with little effect on the productivity margins captured by the model. Moderate representation targets can also substantially reduce the gap with limited productivity effects. By contrast, hours caps and more aggressive representation targets generate meaningful productivity costs.

Large law firms exemplify high-hours professional labor markets in which the existing literature has documented sizable gender promotion gaps. Prior work has identified several mechanisms that may contribute to these gaps, including retention, performance, and differential promotion standards. Our within-firm panel data allow us to quantify the contribution of these components and to evaluate how counterfactual promotion policies would affect outcomes once workers adjust their labor supply. The findings highlight the importance of accounting for the interaction between worker labor supply and firm policies when evaluating personnel policy reforms, and shed light on specific margins through which firms may be able to increase the representation of women in senior positions.

⁴⁹However, these targets should be understood as voluntary rather than strict quotas to avoid violating anti-discrimination laws. E.g., *Johnson v. Transportation Agency, Santa Clara Cty.*, 480 US 616 (1987).

References

- Above the Law. 2015. “Biglaw Mommy: Share And Share Alike.” *Above the Law*.
- Altonji, Joseph. 1986. “Intertemporal Substitution in Labor Supply: Evidence from Micro Data.” *Journal of Political Economy* 94 (3).
- Antecol, Heather, Kelly Bedard, and Jenna Stearns. 2018. “Equal but Inequitable: Who Benefits from Gender-Neutral Tenure Clock Stopping Policies?” *American Economic Review* 108:2420–2421.
- Azmat, Ghazala, Vicente Cuñat, and Emeric Henry. 2025. “Gender Promotion Gaps and Career Aspirations.” *Management Science* 72 (3): 1865–1888.
- Azmat, Ghazala, and Rosa Ferrer. 2017. “Gender Gaps in Performance: Evidence from Young Lawyers.” *Journal of Political Economy* 125 (5): 1306–1355.
- Babcock, Linda, Maria Recalde, Lise Vesterlund, and Laurie Weingart. 2017. “Gender Differences in Accepting and Receiving Requests for Tasks with Low Promotability.” *American Economic Review* 107 (3): 714–747.
- Baker, George, Michael Gibbs, and Bengt Holmstrom. 1994. “The Internal Economics of the Firm: Evidence from Personnel Data.” *The Quarterly Journal of Economics* 109 (4): 881–919.
- Barlevy, Gadi, and Derek Neal. 2019. “Allocating Effort and Talent in Professional Labor Markets.” *The Journal of Labor Economics* 37 (1): 187–246.
- Bastani, Spencer, and Håkan Selin. 2014. “Bunching and non-bunching at kink points of the Swedish tax schedule.” *Journal of Public Economics* 109:36–49.
- Bell, Linda, and Richard Freeman. 2001. “The incentive for working hard: explaining hours worked differences in the US and Germany.” *Labour Economics* 8 (2): 181–202.
- Benson, Alan, Danielle Li, and Kelly Shue. 2019. “Promotions and the Peter Principle.” *Quarterly Journal of Economics* 134 (4): 2085–2134.
- . 2026. “Potential and the Gender Promotions Gap.” *American Economic Review* 116 (2): 375–417.
- Bertrand, Marianne. 2018. “Coase Lecture - The Glass Ceiling.” *Economica* 85 (338): 205–231.
- Bertrand, Marianne, Sandra E. Black, Sissel Jensen, and Adriana Lleras-Muney. 2019. “Breaking the Glass Ceiling? The Effect of Board Quotas on Female Labor Market Outcomes in Norway.” *The Review of Economic Studies* 86 (1): 191–239.
- Bertrand, Marianne, Claudia Goldin, and Lawrence Katz. 2010. “Dynamics of the Gender Gap for Young Professionals in the Financial and Corporate Sectors.” *American Economic Journal: Applied Economics* 2 (3): 228–255.

- Bertrand, Marianne, and Kevin Hallock. 2001. "The Gender Gap in Top Corporate Jobs." *ILR Review* 55 (1): 3–21.
- Blau, Francine D., and Jed Devaro. 2007. "New Evidence on Gender Differences in Promotion Rates: An Empirical Analysis of a Sample of New Hires." *Industrial Relations* 46 (3): 511–550.
- Blomquist, Sören, Whitney Newey, Anil Kumar, and Che-Yuan Liang. 2021. "On Bunching and Identification of the Taxable Income Elasticity." *Journal of Political Economy* 129:2320–2343.
- Bloomberg. 2025. "Attorney Workload and Hours (2024 Wrapup)." *Bloomberg Law Reports*.
- Boone, Ryan. 2021. "Tacit Collusion in Labor Markets: The Case of BigLaw." *UCLA Electronic Theses and Dissertations*.
- Cairo, Sofie, Ria Ivandic, Anne Sophie Lassen, and Valentina Tartari. 2026. "Parenthood and the career ladder: evidence from academia." *Centre for Economic Performance Working Paper*.
- Calvert, Cynthia, Linda Chanow, and Linda Marks. 2009. "Reduced Hours, Full Success: Part-Time Partners in U.S. Law Firms." *The Project for Attorney Retention*.
- Cappelli, Peter, and Wayne F. Cascio. 1991. "Why some jobs command wage premiums: A test of career tournament and internal labor market hypotheses." *Academy of Management Journal* 34 (4): 848–868.
- Chen, Jihui, Qihong Liu, and Myongjin Kim. 2022. "Gender gap in tenure and promotion: Evidence from the economics Ph.D. class of 2008." *Southern Economic Association* 88 (4): 1277–1312.
- Chetty, Raj, John Friedman, Tore Olsen, and Luigi Pistaferri. 2011. "Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records." *The Quarterly Journal of Economics* 126 (2): 749–804.
- Choi, Sheridan, and Delia Violante. 2020. "Overview of Gender Diversity in Business Law, A Research Compendium." *Berkeley Law: Women in Business Initiative*.
- Ciasullo, Ludovica, and Martina Uccioli. 2025. "What Works for Working Couples? Work Arrangements, Maternal Labor Supply, and the Division of Home Production." *IZA Discussion Paper*.
- Comunale, Mariarosaria, Petra de Bruxelles, Kalpana Kochhar, Juliette Raskauskas, and D. Filiz Unsal. 2023. "Who are Central Banks? Gender, Human Resources, and Central Banking." *IMF Working Paper*.
- Cook, Cody, Rebecca Diamond, Jonathan Hall, John List, and Paul Oyer. 2020. "The Gender Earnings Gap in the Gig Economy." *The Review of Economic Studies* 88:2210–238.
- Deschamps, Pierre. 2024. "Gender Quotas in Hiring Committees: A Boon or a Bane for Women?" *Management Science* 70 (11): 7486–7505.

- Doeringer, Peter B., and Michael J. Piore. 1971. *Internal Labor Markets and Manpower Analysis*.
- Farber, Henry. 2005. "Is Tomorrow Another Day? The Labor Supply of New York City Cabdrivers." *Journal of Political Economy* 113 (1): 46–82.
- Fatima, Ma. 2025. "Partner Track Transparency Report Reveals Law Firm Odds." *JD Journal*.
- Ferrer, Rosa, and Danisz Okulicz. 2025. "Adversarial Career Concerns: Theory and Evidence on Lawyers' Career Trajectories." *Working Paper*.
- Galanter, Marc, and Thomas Palay. 1991. "Tournament of Lawyers: The Transformation of the Big Law Firm." *University of Chicago Press*.
- Garicano, Luis, and Thomas Hubbard. 2007. "Managerial Leverage Is Limited by the Extent of the Market: Hierarchies, Specialization, and the Utilization of Lawyers' Human Capital." *The Journal of Law and Economics* 50 (1): 1–43.
- Gibbons, Robert, and Michael Waldman. 1999. "Careers in organizations: Theory and evidence." In *Handbook of Labor Economics*, vol. 3. In: Orley C. Ashenfelter and David Card (eds.) Elsevier.
- Goldin, Claudia. 2004. "Interview with Claudia Goldin." *Federal Reserve Bank of Minneapolis*.
- . 2006. "The Quiet Revolution That Transformed Women's Employment, Education, and Family." *AEA Papers and Proceedings: Richard T. Ely Lecture*.
- . 2014. "A Grand Gender Convergence: Its Last Chapter." *American Economic Review* 104 (4): 1091–1119.
- . 2021. "Career and Family: Women's Century-Long Journey toward Equity." *Princeton University Press*.
- Goldin, Claudia, and Larry Katz. 2016. "A Most Egalitarian Profession: Pharmacy and the Evolution of a Family-Friendly Occupation." *Journal of Labor Economics* 34 (3): 705–746.
- Goldin, Claudia, and Lawrence Katz. 2008. "Transitions: Career and Family Life Cycles of the Educational Elite." *American Economic Review: Papers & Proceedings* 98 (2): 363–369.
- Goldin, Claudia, Sari Pekkala Kerr, and Claudia Olivetti. 2022. "When the Kids Grow Up: Women's Employment and Earnings Across the Family Cycle." *NBER Working Paper*.
- González, Juan B., and Alejandro Martínez-Marquina. 2026. "The Backlash of Gender Quotas."
- Greenwood, Royston, and Laura Empson. 2003. "The Professional Partnership: Relic or Exemplary Form of Governance?" *Organization Studies* 24 (6): 909–933.

- Gupta, Nabanita Datta, Yana Gallen, Kristian Stamp Hedeager, and Kerstin Holzheu. 2026. “Workplace Amenities and the Gender Pay Gap.” *The Rockwool Foundation Study Paper No. 261*.
- Harrington, Mona, and Helen Hsi. 2007. “Women Lawyers and Obstacles to Leadership.” *MIT Workplace Center*.
- Holmström, Bengt. 1999. “Managerial Incentive Problems: A Dynamic Perspective.” *The Review of Economic Studies* 66 (1): 169–182.
- Hospido, Laura, Luc Laeven, and Ana Lamo. 2022. “The Gender Promotion Gap: Evidence from Central Banking.” *The Review of Economics and Statistics* 104 (5): 981–996.
- Hotz, V. Joseph, Finn Kydland, and Guilherme Sedlacek. 1988. “Intertemporal Preferences and Labor Supply.” *Econometrica* 56 (2): 335–360.
- Huang, Ruidi, Erik Mayer, and Darius Miller. 2024. “Gender Bias in Promotions: Evidence from Financial Institutions.” *The Review of Financial Studies* 37 (5): 1685–1728.
- Ip, Edwin, Andreas Leibbrandt, and Joseph Vecchi. 2020. “How Do Gender Quotas Affect Workplace Relationships? Complementary Evidence from a Representative Survey and Labor Market Experiments.” *Management Science* 66 (2): 805–822.
- Keller, Wolfgang, Teresa Molina, and William Olney. 2023. “The gender gap among top business executives.” *Journal of Economic Behavior and Organization* 211:270–286.
- Kleven, Henrik. 2016. “Bunching.” *Annual Review of Economics* 8:435–464.
- Kleven, Henrik, Camille Landais, and Jakob Egholt-Søgaard. 2019. “Children and Gender Inequality: Evidence from Denmark.” *American Economic Journal: Applied Economics* 11 (4): 181–209.
- . 2021. “Does Biology Drive Child Penalties? Evidence from Biological and Adoptive Families.” *American Economic Review: Insights* 3 (2): 183–198.
- Kleven, Henrik, and Mazhar Waseem. 2013. “Using Notches to Uncover Optimization Frictions and Structural Elasticities: Theory and Evidence from Pakistan.” *The Quarterly Journal of Economics* 128 (2): 669–723.
- Landers, Renée, James Rebitzer, and Lowell Taylor. 1996. “Rat Race Redux: Adverse Selection in the Determination of Work Hours in Law Firms.” *American Economic Review* 88 (3): 329–348.
- Lassen, Anne-Sophie, and Ria Ivandić. 2024. “Parenthood and Academic Career Trajectories.” *AEA Papers and Proceedings* 114:238–242.
- Lateral Link. 2026. “Female Representation in Biglaw Partnerships — A Long Way to Go.” *Lateral Link*.
- Lazear, Edward, and Sherwin Rosen. 1990. “Male-Female Wage Differentials in Job Ladders.” *Journal of Labor Economics* 8 (1): S106–S123.

- Lazear, Edward, and Sherwin Rosen. 1997. "Rank-Order Tournaments as Optimum Labor Contracts." *Journal of Political Economy* 89 (5): 473–522.
- Leibbrandt, Andreas, Liang Choon Wang, and Cordelia Foo. 2018. "Gender Quotas, Competitions, and Peer Review: Experimental Evidence on the Backlash Against Women." *Management Science* 64 (8): 3501–3516.
- Leibowitz, Arleen, and Robert Tollison. 1978. "Earning and Learning in Law Firms." *Journal of Legal Studies* 7 (1): 65–81.
- Leslie, Lisa M., Colleen Flaherty Manchester, Tae-Youn Park, and SI Ahn Mehng. 2012. "Flexible Work Practices: A Source of Career Premiums or Penalties?" *Academy of Management Journal* 55 (6): 1407–1428.
- Levin, Jonathan, and Steven Tadelis. 2002. "A Theory of Partnerships." *SSRN Electronic Journal*.
- MaCurdy, Thomas E. 1981. "An Empirical Model of Labor Supply in a Life-Cycle Setting." *Journal of Political Economy* 89.
- Manchester, Colleen Flaherty, Lisa M. Leslie, and Amit Kramer. 2010. "Stop the Clock Policies and Career Success in Academia." *American Economic Review* 100 (2): 219–223.
- Marhoffer, Elizabeth A., Samer Ein-Alshaeba, Alyssa A. Grimshaw, Jürgen L. Holleck, Benjamin Rudikoff, Lori A. Bastian, and Craig G. Gunderson. 2024. "Gender Disparity in Full Professor Rank Among Academic Physicians: A Systematic Review and Meta-Analysis." *Academic Medicine* 99 (7): 801–809.
- Martínez, Isabel, Emmanuel Saez, and Michael Siegenthaler. 2021. "Intertemporal Labor Supply Substitution? Evidence from the Swiss Income Tax Holidays." *American Economic Review* 111 (2): 506–546.
- McFadden, Daniel. 1989. "A Method of Simulated Moments for Estimation of Discrete Response Models Without Numerical Integration." *Econometrica* 57 (5): 995–1026.
- McLymore, Arriana. 2021. "Law Firms Rethink Training, Push Incentives to Close Gaps on Gender and Race." *Reuters*.
- Minni, Virginia, Kieu-Trang Nguyen, Heather Sarsons, and Carla Srebot. 2026. "Managers and the Cultural Transmission of Gender Norms." *NBER Working Paper*.
- NALP. 2024. "2023 Report on Diversity in U.S. Law Firms." *The National Association for Law Placement*.
- Niederle, Muriel, and Lisa Vesterlund. 2007. "Do Women Shy Away from Competition? Do Men Compete Too Much?" *Quarterly Journal of Economics* 122 (3): 1067–1101.
- Nordenflycht, Andrew von. 2010. "What Is a Professional Service Firm? toward a Theory and Taxonomy of Knowledge-Intensive Firms." *The Academy of Management Review* 35 (1): 155–174.

- Okuyama, Yoko, Takeshi Murooka, and Shintaro Yamaguchi. 2025. "Unpacking the Child Penalty Using Personnel Data: How Promotion Practices Widen the Gender Pay Gap." *IZA Discussion Paper*.
- Oyer, Paul, and Scott Schaefer. 2016. "Firm/Employee Matching: An Industry Study of U.S. Lawyers." *ILR Review* 69 (2): 378–404.
- . 2019. "The Returns to Elite Degrees: The Case of American Lawyers." *ILR Review* 72 (2): 446–479.
- Peery, Destiny. 2017. "Annual Survey Report on the Promotion and Retention of Women in Law Firms." *NAWL*.
- Rebitzer, James, and Lowell Taylor. 2007. "When Knowledge Is an Asset: Explaining the Organizational Structure of Large Law Firms." *Journal of Labor Economics* 25 (2): 201–229.
- Rosen, Sherwin. 1992. "The Market for Lawyers." *Journal of Law and Economics* 35 (2): 215–246.
- Rummel, Michael. 2024. "Law Firms Work to Close the Gender Gap in Leadership and Management." *Law Week Colorado*.
- Russian, Anna. 2024. "LSAC's Knowledge Report The 2024 1L Profile." *Law School Admission Council*.
- Saeedy, Alexander. 2024. "Wall Street Curbs Young Bankers' Hours After Overwork Outcry." *Wall Street Journal*.
- Saez, Emmanuel. 2010. "Do Taxpayers Bunch at Kink Points?" *American Economic Journal: Economic Policy* 2 (3): 180–212.
- Schaede, Ursina, and Ville Mankki. 2026. "Quota versus Quality? Long-Term Gains from an Unusual Gender Quota." *American Economic Review* 116 (6): 1996–2037.
- Simons, Hugh. 2018. "What's a Reasonable Ratio of Partner Comp to Associate Starting Salary?" *ALM Intelligence: Legal Compass*.
- Smale, Alison, and Claire Cain Miller. 2015. "Germany Sets Gender Quota in Boardrooms." *New York Times*.
- Taddia, Marialuisa. 2025. "BigLaw's Two-Tier Model Risks Widening Gender Disparities." *Law360*.
- Tannenbaum, Daniel, and Brenden Timpe. 2025. "Job Tasks, Leadership, and the Gender Gap after Parenthood." *AEA Papers and Proceedings* 115:271–275.
- Tesar, Linda. 2025. "CSWEP Annual Report." *American Economic Association*.
- Thornton, Saranna R. 2008. "The Implementation and Utilization of Stop the Tenure Clock Policies in Canadian and U.S. Economics Departments."
- Vault. 2009. "Getting on the Partnership Track."

- Vesterlund, Lise. 2025. “The No Club, an Interview with Lise Vesterlund.” *Robert G. and Sue Douthit O’Donnell Center for Behavioral Economics*.
- Wasserman, Melanie. 2023. “Hours Constraints, Occupational Choice, and Gender: Evidence from Medical Residents.” *The Review of Economic Studies* 90 (3): 1535–1568.
- Williams, Joan, and Cynthia Thomas Calvert. 2002. “Balanced Hours: Effective Part-Time Policies for Washington Law Firms: The Project for Attorney Retention, Final Report, Third Edition.” *William and Mary Journal of Women and the Law* 8 (3): 357–441.
- Wiswall, Matthew, and Basit Zafar. 2018. “Preference for the Workplace, Investment in Human Capital, and Gender.” *The Quarterly Journal of Economics* 133 (1): 457–507.
- Wood, Robert, Mary Corcoran, and Paul Courant. 1993. “Pay Differences among the Highly Paid: The Male-Female Earnings Gap in Lawyers’ Salaries.” *Journal of Labor Economics* 11 (3): 417–441.
- Zhao, Xinyi, Anna I. Thoma, Ralph Hertwig, and Dirk U. Wulff. 2025. “Mapping the gender attrition gap in academic psychology.” *arXiv Working Paper*.

A Supplementary Data

Revelio Labs. We use Revelio Labs employment histories to observe career outcomes for lawyers who leave large law firms. The unit of observation is an employment spell. For each spell, we observe the employer, job title, start and end dates, location, and estimated salary. We restrict attention to lawyers previously employed at Am Law 100 firms and use subsequent job spells to characterize outside options after leaving BigLaw, with a focus on post-exit salaries (as shown in Appendix Figure 5).

Am Law 100. We use law firm-level financial data from Am Law 100 to measure revenue, headcount, and profits for our sample firm and comparable firms in the industry. These data provide annual firm-level statistics for the 100 largest U.S. law firms and allow us to capture key aspects of the firm’s financials.

After the JD. We use longitudinal survey data from After the JD to provide additional information about career dynamics and promotion outcomes. The survey follows a cohort of lawyers first admitted to the bar in 2000, surveying them at multiple points over their first decade in the profession. It includes information on occupation, career progression, compensation, workload, family, and individual expectations and experiences. While these data span multiple firms and rely on self-reported measures of objects like billable hours, they complement our administrative data by providing insight into individual-level promotions and the role of other factors like expectations and aspirations in shaping outcomes.

BigLaw Investor. We use data from BigLaw Investor on historical Cravath scales to quantify compensation schedules in our setting. These data provide the mapping from seniority to salaries and bonuses that underlies associates’ incentives.

Aggregate Promotion Moments. We use public promotion announcements (e.g., “Firm X promotes Y new equity partners”) to observe the typical size of new equity partner classes at the firm level. Combining these counts with the size of the promotion-eligible cohort in our administrative data allows us to construct aggregate promotion rates. Because historical public announcements do not consistently report the gender composition of new partner classes, we supplement these firm-level moments with industry-level evidence from Choi and Violante (2020), which reports the typical gender composition of new partner cohorts by subgroup.⁵⁰ We combine these data with the gender composition of the promotion-eligible

⁵⁰Our firm is similar to industry trends along related dimensions; for example, women make up about 25% of the firm’s equity partner stock, similar to the industry benchmark in NALP (2024).

cohort in our administrative data and the firm-specific aggregate promotion rates to construct gender-specific promotion rates.

Law Student Survey. We are also in the process of conducting a survey of Stanford Law School students who have worked as summer associates to elicit beliefs about how promotion decisions are made in private sector law firms. The survey instrument is shown in Appendix [I](#). Responses will supplement anecdotal evidence of beliefs summarized in Appendix [J](#) and provide evidence on how prospective associates interpret promotion policies.

B Supplemental Discussion

B.1 Estimated Law Firm Promotion Policy

Given the empirical performance distributions, we choose the parameters that minimize the distance between model-implied and empirical gender-specific promotion rates:

$$(\hat{\gamma}, \hat{c}_z^p) = \arg \min_{\gamma, c_z^p} \left[\left(\int p(z; \gamma, c_z^p) dF_f(z) - \psi_f \right)^2 + \left(\int p(z; \gamma, c_z^p) dF_m(z) - \psi_m \right)^2 \right],$$

where the subscripts f and m denote female and male associates, respectively.⁵¹ The empirical partner promotion rates, constructed in Appendix Table 1, are 45% for female eighth-year associates and 72% for male eighth-year associates.

Table 1: Calculating Gender-Specific Promotion Rates

Moment	Data Source	Value
Promotion Rate of Eighth-Year Associates	New Partner Class Announcements + Empirical Eighth-Year Cohorts	60%
Female Share of Eighth-Year Associates	Empirical Eighth-Year Cohorts	44%
Female Share of Promotions	Choi and Violante (2020)	33%
Implied Promotion Rate of Female Eighth-Year Associates = $(0.60/0.44) \times 0.33 =$		45%
Implied Promotion Rate of Male Eighth-Year Associates = $(0.60/0.56) \times 0.67 =$		72%

Notes: Table constructs gender-specific promotion rates among eighth-year associates. The aggregate promotion rate is calculated using new partner class announcements and empirical eighth-year associate cohorts. The female share of newly promoted equity partners is taken from Choi and Violante (2020). The implied female promotion rate equals the aggregate promotion rate divided by the female share of eighth-year associates and multiplied by the female share of promotions. The implied male promotion rate is constructed analogously using the male shares. Although we cannot observe gender-specific promotion rates in our firm, the firm aligns with industry patterns on related observable dimensions – e.g., women make up about 25% of the firm’s equity partner stock, consistent with the industry benchmark in NALP (2024). **Data:** New partner class announcements, Law Firm Dataset, and Choi and Violante (2020).

⁵¹This is a just-identified exercise: the two policy parameters are disciplined by the two promotion rates, conditional on the empirical performance distributions.


Table 2: GMM Estimates for Promotion Policy $p(z; \gamma, c_z^p)$

<i>Panel A: Parameter Estimates</i>	Estimate
Steepness $\hat{\gamma}$	8640
Threshold $\hat{c}_z^p$	0.5737

<i>Panel B: Model Fit</i>	Data	Model
Male Promotion Rate $\bar{p}_m$	0.72	0.70
Female Promotion Rate $\bar{p}_f$	0.45	0.45

Notes: Table reports estimates of the gender-neutral promotion policy used in the static exercise. The policy maps normalized eighth-year billable hours, z , into promotion probabilities according to $p(z; \gamma, c_z^p) = [1 + \exp -\gamma(z - c_z^p)]^{-1}$. Panel A reports the estimated steepness parameter, $\hat{\gamma}$, and threshold, $\hat{c}_z^p$. Panel B compares empirical promotion rates to model-implied promotion rates under the estimated policy. The distance column reports the difference between the empirical and model-implied rates. **Data:** Eighth-year hours from the Law Firm Dataset and promotion rates constructed in Appendix Table 1.

Figure 1: Average Hours Worked in Year 6, by Gender and Partner Outcome



Notes: Figure reports average self-reported billable hours with 95% confidence intervals in the sixth year after law school, separately by gender and subsequent partnership outcome, from the After the JD survey data. Partnership is measured using respondents' reported employment in Wave 3. The sample is limited to private sector associates observed in Wave 2 at firms with at least 200 lawyers and equity partners, and who also responded to the Wave 3 survey. Men are shown in orange and women in gray. **Data:** After the JD Survey – Wave 2 (2007) and Wave 3 (2012).

B.2 Perceived Promotion Policy

Suppose there are N associates who remain in the firm through period T , and the firm wants to promote a fixed share $q \in (0, 1)$. That is, they use a rank-based promotion policy:

$$\text{promote associate } i \text{ if and only if } \text{rank}_{iT} \leq qN,$$

where rank_{iT} is i 's rank based on firm capital (h_{iT}). This is equivalent to saying that there exists a threshold (c_h^p) at the $(1 - q)$ sample quantile of $\{h_{iT}\}_{i=1}^N$ such that the policy is:

$$\text{promote associate } i \text{ if and only if } h_{iT} \geq c_h^p.$$

If we assume that associate i observes her own h_{iT} but does not observe others' performance, then the precise cutoff c_h^p is not known to her.⁵² As a result, the associate faces uncertainty about the effective cutoff and perceives it as

$$Z^* = c_h^p + \eta,$$

where η is a mean-zero random variable independent of h_{iT} that captures uncertainty about her realized rank.⁵³ Under this belief, the associate is promoted when $h_{iT} \geq Z^*$, so her perceived promotion probability is:

$$p^{\text{perc}}(h_{iT}) = \Pr(h_{iT} \geq c_h^p + \eta \mid h_{iT}) = \Pr(\eta \leq h_{iT} - c_h^p) = F_\eta(h_{iT} - c_h^p),$$

where F_η is the CDF of η .

To connect this object to the promotion policy in the model, assume η is logistic with scale parameter $1/\gamma$. Then the perceived promotion probability takes the same form as in (5):

$$p^{\text{perc}}(h_{iT}) = [1 + \exp\{-\gamma(h_{iT} - c_h^p)\}]^{-1}.$$

Thus, even if the underlying firm policy is a deterministic fixed-quota threshold rule, uncertainty about peers' performance (and hence about the realized cutoff) makes promotion probabilistic in the associate's decision problem.

⁵²The firm's policy is deterministic conditional on the full cross-section of firm capital, but the realized cutoff depends on the distribution of peers' h_{jT} , which the associate does not observe.

⁵³That is, the associate's beliefs are correctly centered on the firm's standard, c_h^p , but η captures uncertainty regarding peers' performance.

B.3 FTE-Based Disutility

Let z_t denote total hours supplied in period t , ϕ_j denote the fraction of full-time workdays associated with contract j , and D denote the number of workdays in a full-time schedule over the period. Then a worker on contract j works $\phi_j D$ days, so that $z_t/(\phi_j D)$ is average hours per workday. Since D is common across contracts, z_t/ϕ_j is proportional to daily work intensity. Thus, holding total hours z_t fixed, a lower value of ϕ_j implies that the same total hours must be compressed into fewer workdays.


We model the disutility of work as convex in this daily intensity measure:

$$v_{it}(z_t, j) = \frac{n_i/\kappa(\eta_{it})}{1 + 1/\epsilon} \left(\frac{z_t/\phi_j}{n_i/\kappa(\eta_{it})} \right)^{1+1/\epsilon}.$$

This setup captures the idea that supplying a given number of total hours is more costly when those hours are concentrated into fewer workdays. Other non-hour factors – including the value of working fewer days, workplace culture, etc. – are captured by the amenity term.


C Supplementary Figures

Figure 2: After the JD Data on Promotion and 6th-Year Hours




Notes: Figure shows the relationship between self-reported billable hours in the sixth year after law school and subsequent partnership from the After the JD survey data. The horizontal axis reports Wave 2 self-reported billable hours. The vertical axis reports the share of respondents who report being partners in Wave 3. The sample is limited to private sector associates observed in Wave 2 at firms with at least 200 lawyers and equity partners, and who also responded to the Wave 3 survey. **Data:** After the JD Survey – Wave 2 (2007) and Wave 3 (2012).

Figure 3: After the JD Data on 6th-Year Hours and Business Development




Notes: Figure plots self-reported billable hours in the sixth year after law school against new client revenue from the After the JD survey data. Both variables are measured in Wave 2, with new client revenue reported in log dollars. The sample is limited to private sector associates observed in Wave 2 at firms with at least 200 lawyers and equity partners. **Data:** After the JD Survey – Wave 2 (2007).

Figure 4: After the JD Data on 6th-Year Hours and Relationships




Notes: Figure plots self-reported billable hours in the sixth year after law school against respondents' reported number of mentors from the After the JD survey data. Both variables are measured in Wave 2. The sample is limited to private sector associates observed in Wave 2 at firms with at least 200 lawyers and equity partners. **Data:** After the JD Survey – Wave 2 (2007).

Figure 5: Revelio Labs Comparison of Firm and Outside Option Salaries




Notes: Figure compares salaries before and after lawyers leave Am Law 100 firms for jobs outside the Am Law 100. The dashed horizontal line at zero marks no salary change. The red horizontal line reports the average salary change across observed exits, and the blue markers report the observed salary changes by year at the firm before exit. The sample includes lawyers who left Am Law 100 firms for non-Am Law 100 jobs in 2010 or later. **Data:** Revelio Labs Individual Professional Profiles.

Figure 6: Bin Scatter of Age on Associate Year



Notes: Figure plots average associate age by associate year. Each point reports the average age among associates observed in that associate year in the Law Firm Dataset. The figure is used to characterize the ages at which associates approach the standard partnership decision. **Data:** Law Firm Dataset.

Figure 7: Contract Transition Probabilities by Gender



Notes: Figure reports one-year contract transition probabilities by gender, pooled across associate years. Transitions are defined from a worker's contract in year t to their contract or exit status in year $t + 1$. The origin contracts are full-time (FT) and part-time (PT); destination states are FT, PT, and leaving the firm.

Data: Law Firm Dataset.

D Appendix Tables

Table 3: Job Shares Over Time

Survey Year:	Graduation Year: 2000		
	2002-3	2007-8	2012-13
Law Firm	76.1%	54.0%	41.2%
Federal Gov't (Incl. Judiciary)	5.2%	9.6%	10.3%
State or Local Gov't (Incl. Judiciary)	4.4%	6.5%	6.1%
All Other	14.3%	30.0%	42.4%

Notes: Table reports employment-sector shares over time for the After the JD sample. Columns correspond to survey waves: Wave 1 in 2002–2003, Wave 2 in 2007–2008, and Wave 3 in 2012–2013. Rows group respondents into law firm jobs, federal government jobs including clerkships, state or local government jobs including clerkships, and all other jobs. The sample is limited to respondents who graduated from a top 20 law school. **Data:** After the JD Survey – Wave 1, Wave 2, and Wave 3.

Table 4: Reasons for Job Change by Gender

Reason:	Gender		
	Overall	Female	Male
For a Better Job	15.4%	10.7%	19.6%
Dissatisfaction with Practice Area	10.9%	12.9%	9.2%
For Better Work Environment	9.2%	10.0%	8.5%
Moved From Area Because of Spouse Job	6.8%	7.1%	6.5%
Desire to Change Sector	6.5%	7.9%	5.2%
Was Fired	1.0%	0.0%	2.0%
All Other	50.2%	51.4%	49.0%

Notes: Table reports stated reasons for job changes, overall and separately by gender, from the After the JD survey data. Entries are shares of respondents reporting each reason for leaving. The sample is limited to respondents who graduated from a top 20 law school and had a job change. **Data:** After the JD Survey – Wave 1, Wave 2, and Wave 3.


E Motherhood Penalty Regression

Using simulated model data, we also conduct a child penalty exercise in the spirit of Kleven, Landais, and Egholt-Søgaard (2019). We interpret the arrival of a worker’s first child as the first period in which the worker enters η_H . We use workers who never enter η_H as controls. Given this, we can calculate the difference between each parent i in period t ’s earnings and mean earnings among never-parents of the same gender in the same period, Δ_{it} . We then estimate the following regression separately by gender:

$$\Delta_{it} = \sum_{j \neq -1} \beta_j \cdot \mathbb{I}[t - b_i = j] + \varepsilon_{it},$$

where b_i denotes the period of first birth. Each β_j therefore captures the change in parents’ earnings relative to never-parents at event time j , normalized to the period before birth. We cluster standard errors by worker. In Appendix Figure 8, we scale each β_j by mean earnings among never-parents in the corresponding periods to express the estimates in proportional terms. The child penalty at event time j is then the difference between the proportional earnings effects for mothers and fathers. The estimated 14% penalty is consistent with estimates in the existing literature – e.g., 19% in Kleven, Landais, and Egholt-Søgaard (2019) and 17-18% in Kleven, Landais, and Egholt-Søgaard (2021).

Figure 8: Child Penalty Regression



Notes: Figure reports estimates of the child penalty in a panel of simulated workers. Entry into parenthood is defined as the worker’s first entry into η_H . Earnings include the outside-option salary after exit and partner compensation after promotion, so impacts reflect the percent change in total earnings. Standard errors are clustered by worker. **Data:** Simulated panel from the estimated model summarized in Table 3.

F Identification and Estimation Moments

F.1 Moment Definitions

Hours Bunching for the Bonus. We summarize the local hours response to the bonus threshold using a bunching statistic around $\underline{z}$, pooling full-time worker-year observations across associate years. Let $\mathcal{I}^{\mathcal{F},g}$ denote the set of full-time worker-year observations for gender g , and let $N^{\mathcal{F},g} = |\mathcal{I}^{\mathcal{F},g}|$. For bandwidth $\Delta = 0.017$ normalized hours, define:

$$\widehat{B}^g(\Delta) = \frac{1}{N^{\mathcal{F},g}} \sum_{(i,t) \in \mathcal{I}^{\mathcal{F},g}} \mathbb{1}\{z_{it} \in [\underline{z}, \underline{z} + \Delta)\} - \frac{1}{N^{\mathcal{F},g}} \sum_{(i,t) \in \mathcal{I}^{\mathcal{F},g}} \mathbb{1}\{z_{it} \in [\underline{z} - \Delta, \underline{z})\}.$$

This moment informs the labor supply elasticity, ϵ^g . A larger ϵ^g makes hours more responsive to the bonus, shifting more marginal workers from below to above the threshold and increasing $\widehat{B}^g(\Delta)$.

Hours Distribution. We summarize the level and spread of full-time hours using the median of hours among full-time workers:

$$\widehat{z}_t^{\mathcal{F},g} = \widehat{Q}_{0.5} \left(\{z_{it} : i \in \mathcal{I}_t^{\mathcal{F},g}\} \right),$$

and the dispersion of hours:

$$\widehat{D}_t^{\mathcal{F},g} = \left\{ \frac{1}{N_t^{\mathcal{F},g}} \sum_{i \in \mathcal{I}_t^{\mathcal{F},g}} \left(z_{it} - \widehat{z}_t^{\mathcal{F},g} \right)^2 \right\}^{1/2}.$$

These moments inform the latent permanent type distribution, (μ_n^g, σ_n^g) . A higher μ_n^g shifts hours upward, while a higher σ_n^g increases dispersion in hours.

Contract Shares Over Time. We summarize the level and time path of contract choices using contract shares. For $j \in \{\mathcal{F}, \mathcal{P}, \mathcal{L}\}$, define:

$$\widehat{S}_{j,t}^g = \frac{1}{N_t^g} \sum_i \mathbb{1}\{j_{it} = j, g_i = g\}.$$

These moments inform the contract amenity parameters, $\psi_{\mathcal{P}}^g(\eta)$ and $\psi_{\mathcal{L}}^g(\eta)$, which shift the relative value of each contract. Higher values of $\psi_{\mathcal{P}}^g(\eta)$ raise part-time take-up, while higher values of $\psi_{\mathcal{L}}^g(\eta)$ raise leaving to the outside option.

Conditional Contract Transition Rates. We calculate empirical transition rates as follows. For $j \in \{\mathcal{F}, \mathcal{P}\}$ and $k \in \{\mathcal{F}, \mathcal{P}, \mathcal{L}\}$, define:

$$\hat{T}_{jk,t}^g = \frac{1}{N_t^{j,g}} \sum_i \mathbb{1}\{j_{it} = j, g_i = g\} \mathbb{1}\{j_{i,t+1} = k\}.$$

These moments inform the latent state transition probabilities, p_{LH}^g and p_{HL}^g . More persistent latent states generate more persistent contract choices, while more frequent latent state transitions generate more movement between contracts.⁵⁴

Part-Time FTE Hours. We summarize part-time workers' choice of FTE hours as:

$$\hat{z}_t^{\mathcal{P},g} = \hat{Q}_{0.5} \left(\left\{ \frac{z_{it}}{\phi_{\mathcal{P}}} : i \in \mathcal{I}_t^{\mathcal{P},g} \right\} \right),$$

This moment informs the high-state effort cost shifter, $\kappa^g(\eta_H)$. A larger $\kappa^g(\eta_H)$ lowers FTE hours among part-time workers relative to full-time workers.

Terminal-Year Hours Spread. We summarize the gap between full-time hours in the terminal associate year and full-time hours earlier in the career as:

$$\hat{G}^g = \hat{z}_{T-1}^{\mathcal{F},g} - \frac{1}{T-3} \sum_{t=2}^{T-2} \hat{z}_t^{\mathcal{F},g}.$$

This moment informs the career-capital depreciation rate, δ . When δ is low, early hours remain valuable for promotion and effort is spread more evenly over the career. When δ is high, recent hours receive more weight, increasing the return to supplying hours near the promotion decision and raising $\hat{G}^g$.

Aggregate Promotion Rate. We summarize promotion outcomes using the aggregate conditional promotion rate among workers who remain at the firm through the promotion period. Let $\mathcal{W}$ denote the promoted set of workers. Define:

$$\hat{R} = \frac{\sum_i \mathbb{1}\{i \in \mathcal{W}, j_{iT} \neq \mathcal{L}\}}{\sum_i \mathbb{1}\{j_{iT} \neq \mathcal{L}\}}.$$

This moment informs the promotion threshold, c_h^p , because the threshold shifts the overall level of promotion probabilities. A higher c_h^p lowers the share of retained associates whose

⁵⁴Contract transitions that are not accounted for by persistent latent states or permanent heterogeneity discipline the taste shock variance, σ_ε^g . Greater residual movement across contracts requires a larger σ_ε^g .

firm capital is high enough for promotion.

Bunching Above the Promotion Threshold. For sufficiently large δ , cumulative firm capital at year T is dominated by year- T hours, so c_h^p corresponds closely to the relevant year- T hours threshold. Using the same bandwidth $\Delta = 0.017$ normalized hours as above, define:

$$\begin{aligned}\widehat{M}_T^g(c_h^p, \Delta) &= \frac{1}{N_T^{\mathcal{F},g}} \sum_i \mathbb{1}\{j_{iT} = \mathcal{F}, g_i = g\} \mathbb{1}\{z_{iT} \in [c_h^p, c_h^p + \Delta)\} \\ &\quad - \frac{1}{N_T^{\mathcal{F},g}} \sum_i \mathbb{1}\{j_{iT} = \mathcal{F}, g_i = g\} \mathbb{1}\{z_{iT} \in [c_h^p - \Delta, c_h^p)\}.\end{aligned}$$

This moment informs the promotion steepness parameter, γ_p , because γ_p determines how strongly promotion probabilities respond to firm capital around the threshold. A larger γ_p raises the return to moving just above the perceived promotion threshold, shifting mass from below to above c_h^p and increasing $\widehat{M}_T^g(c_h^p, \Delta)$.

Hours Among Low-Hour Departing Workers. We summarize the level of FTE hours among departing workers in the lower part of the hours distribution using the median. Specifically, for a fixed cutoff z^{low} , define:

$$\widehat{\mathcal{Z}}_{\text{low}}^{\mathcal{L},g} = \widehat{Q}_{0.5} \left(\left\{ \frac{z_{it}}{\phi_{jit}} : j_{it} \in \mathcal{J}^I, e_{i,t+1} = \mathcal{L}, g_i = g, \frac{z_{it}}{\phi_{jit}} \leq z^{\text{low}} \right\} \right).$$

This moment informs the firing threshold, c_z^q . A higher c_z^q shifts the firing-risk schedule to the right, increasing FTE hours of workers and raising $\widehat{\mathcal{Z}}_{\text{low}}^{\mathcal{L},g}$.⁵⁵

Bunching Above the Firing Threshold. We summarize the local hours response to the firing threshold using a bunching statistic around c_z^q , pooling inside-firm worker-year observations. Using the same bandwidth $\Delta = 0.017$ normalized hours as above, define:

$$\begin{aligned}\widehat{M}_q^g(c_z^q, \Delta) &= \frac{1}{N^{\mathcal{F},g} + N^{\mathcal{P},g}} \sum_{i,t} \mathbb{1}\{j_{it} \in \mathcal{J}^I, g_i = g\} \mathbb{1}\left\{ \frac{z_{it}}{\phi_{jit}} \in [c_z^q, c_z^q + \Delta) \right\} \\ &\quad - \frac{1}{N^{\mathcal{F},g} + N^{\mathcal{P},g}} \sum_{i,t} \mathbb{1}\{j_{it} \in \mathcal{J}^I, g_i = g\} \mathbb{1}\left\{ \frac{z_{it}}{\phi_{jit}} \in [c_z^q - \Delta, c_z^q) \right\}.\end{aligned}$$

This moment informs the firing steepness parameter, γ_q , because γ_q determines how sharply firing risk responds to FTE hours around the threshold. A larger γ_q raises the return to

⁵⁵We assume $z^{\text{low}} = 0.5$

moving just above the firing threshold, shifting mass from below to above c_z^q and increasing $\widehat{M}_q^g(c_z^q, \Delta)$.

Table 5: Targeted Moments

Moment group	Moment	Cell	Data	Model
<i>Bonus bunching</i>	$\widehat{B}^g$	$g = \text{male}$	0.106	0.080
		$g = \text{female}$	0.126	0.077
<i>Median full-time hours</i>	$\widehat{z}_t^{\mathcal{F},g}$	$t = 2, g = \text{male}$	0.668	0.665
		$t = 3, g = \text{male}$	0.663	0.666
		$t = 4, g = \text{male}$	0.671	0.666
		$t = 5, g = \text{male}$	0.672	0.667
		$t = 6, g = \text{male}$	0.681	0.667
		$t = 7, g = \text{male}$	0.703	0.668
		$t = 8, g = \text{male}$	0.682	0.668
		$t = 2, g = \text{female}$	0.665	0.618
		$t = 3, g = \text{female}$	0.664	0.665
		$t = 4, g = \text{female}$	0.672	0.666
		$t = 5, g = \text{female}$	0.673	0.666
		$t = 6, g = \text{female}$	0.664	0.667
		$t = 7, g = \text{female}$	0.682	0.667
		$t = 8, g = \text{female}$	0.680	0.667
<i>Full-time hours dispersion</i>	$\widehat{D}_t^{\mathcal{F},g}$	$t = 2, g = \text{male}$	0.079	0.109
		$t = 3, g = \text{male}$	0.076	0.112
		$t = 4, g = \text{male}$	0.068	0.110
		$t = 5, g = \text{male}$	0.101	0.119
		$t = 6, g = \text{male}$	0.087	0.120
		$t = 7, g = \text{male}$	0.088	0.119
		$t = 8, g = \text{male}$	0.092	0.127
		$t = 2, g = \text{female}$	0.073	0.085
		$t = 3, g = \text{female}$	0.095	0.086
		$t = 4, g = \text{female}$	0.078	0.079
		$t = 5, g = \text{female}$	0.085	0.082

Continued on next page

Moment group	Moment	Cell	Data	Model
<i>Contract shares</i>	$\widehat{S}_{j,t}^g$	$t = 6, g = \text{female}$	0.068	0.077
		$t = 7, g = \text{female}$	0.102	0.085
		$t = 8, g = \text{female}$	0.084	0.075
		$t = 2, g = \text{male}, j = \mathcal{P}$	0.293	0.321
		$t = 2, g = \text{male}, j = \mathcal{L}$	0.243	0.215
		$t = 3, g = \text{male}, j = \mathcal{P}$	0.256	0.228
		$t = 3, g = \text{male}, j = \mathcal{L}$	0.351	0.379
		$t = 4, g = \text{male}, j = \mathcal{P}$	0.214	0.174
		$t = 4, g = \text{male}, j = \mathcal{L}$	0.465	0.505
		$t = 5, g = \text{male}, j = \mathcal{P}$	0.197	0.137
		$t = 5, g = \text{male}, j = \mathcal{L}$	0.557	0.604
		$t = 6, g = \text{male}, j = \mathcal{P}$	0.146	0.110
		$t = 6, g = \text{male}, j = \mathcal{L}$	0.652	0.682
		$t = 7, g = \text{male}, j = \mathcal{P}$	0.128	0.088
		$t = 7, g = \text{male}, j = \mathcal{L}$	0.729	0.743
		$t = 8, g = \text{male}, j = \mathcal{P}$	0.093	0.072
		$t = 8, g = \text{male}, j = \mathcal{L}$	0.783	0.791
		$t = 2, g = \text{female}, j = \mathcal{P}$	0.318	0.283
		$t = 2, g = \text{female}, j = \mathcal{L}$	0.148	0.346
		$t = 3, g = \text{female}, j = \mathcal{P}$	0.234	0.252
		$t = 3, g = \text{female}, j = \mathcal{L}$	0.343	0.488
		$t = 4, g = \text{female}, j = \mathcal{P}$	0.222	0.220
		$t = 4, g = \text{female}, j = \mathcal{L}$	0.520	0.593
		$t = 5, g = \text{female}, j = \mathcal{P}$	0.218	0.186
		$t = 5, g = \text{female}, j = \mathcal{L}$	0.642	0.674
		$t = 6, g = \text{female}, j = \mathcal{P}$	0.162	0.156
		$t = 6, g = \text{female}, j = \mathcal{L}$	0.714	0.738
		$t = 7, g = \text{female}, j = \mathcal{P}$	0.130	0.127
		$t = 7, g = \text{female}, j = \mathcal{L}$	0.770	0.791
		$t = 8, g = \text{female}, j = \mathcal{P}$	0.113	0.104
		$t = 8, g = \text{female}, j = \mathcal{L}$	0.835	0.830

Continued on next page


Moment group	Moment	Cell	Data	Model
<i>Contract transition rates</i>	$\hat{T}_{jk}^g$	pooled, $g = \text{male}, j = \mathcal{F}, k = \mathcal{F}$	0.620	0.623
		pooled, $g = \text{male}, j = \mathcal{F}, k = \mathcal{P}$	0.124	0.173
		pooled, $g = \text{male}, j = \mathcal{F}, k = \mathcal{L}$	0.256	0.204
		pooled, $g = \text{male}, j = \mathcal{P}, k = \mathcal{F}$	0.487	0.328
		pooled, $g = \text{male}, j = \mathcal{P}, k = \mathcal{P}$	0.284	0.468
		pooled, $g = \text{male}, j = \mathcal{P}, k = \mathcal{L}$	0.230	0.204
		pooled, $g = \text{female}, j = \mathcal{F}, k = \mathcal{F}$	0.592	0.622
		pooled, $g = \text{female}, j = \mathcal{F}, k = \mathcal{P}$	0.148	0.145
		pooled, $g = \text{female}, j = \mathcal{F}, k = \mathcal{L}$	0.260	0.232
		pooled, $g = \text{female}, j = \mathcal{P}, k = \mathcal{F}$	0.381	0.099
		pooled, $g = \text{female}, j = \mathcal{P}, k = \mathcal{P}$	0.351	0.656
		pooled, $g = \text{female}, j = \mathcal{P}, k = \mathcal{L}$	0.268	0.244
<i>Part-time FTE hours</i>	$\widetilde{z}_t^{\mathcal{P},g}$	$t = 2, g = \text{male}$	0.598	0.603
		$t = 3, g = \text{male}$	0.650	0.656
		$t = 4, g = \text{male}$	0.660	0.660
		$t = 5, g = \text{male}$	0.675	0.664
		$t = 6, g = \text{male}$	0.663	0.668
		$t = 7, g = \text{male}$	0.678	0.672
		$t = 8, g = \text{male}$	0.667	0.676
		$t = 2, g = \text{female}$	0.635	0.643
		$t = 3, g = \text{female}$	0.651	0.644
		$t = 4, g = \text{female}$	0.657	0.668
		$t = 5, g = \text{female}$	0.663	0.669
		$t = 6, g = \text{female}$	0.671	0.669
		$t = 7, g = \text{female}$	0.685	0.670
		$t = 8, g = \text{female}$	0.690	0.671
<i>Terminal-year hours spread</i>	$\hat{G}^g$	$g = \text{male}$	0.006	0.002
		$g = \text{female}$	0.010	0.009
<i>Aggregate promotion rate</i>	$\hat{R}$	aggregate	0.600	0.581
		$g = \text{male}$	0.72	0.674

Continued on next page

Moment group	Moment	Cell	Data	Model
		$g = \text{female}$	0.45	0.446
<i>Promotion-threshold bunching</i>	$\widehat{M}_T^g$	$g = \text{male}$	0.028	0.143
		$g = \text{female}$	0.000	0.223
<i>Departing worker hours</i>	$\widehat{z}_t^{\mathcal{L},g}$	$g = \text{male}$	0.441	0.467
		$g = \text{female}$	0.447	0.467
<i>Firing-threshold bunching</i>	$\widehat{M}_q^g$	$g = \text{male}$	0.003	0.000
		$g = \text{female}$	-0.000	0.000

Notes: Table compares the empirical moments targeted in estimation to the corresponding model-implied moments. The data column reports moments computed from the empirical data; the model column reports simulated moments generated using the parameter estimates in Table 3. Hours are reported in normalized units. Full-time contract shares are omitted because full-time, part-time, and leaving shares sum to one within each gender and associate year. Contract transition rates and bunching moments are pooled across associate years. **Data:** Empirical moments from the Law Firm Dataset and promotion rates constructed in Appendix Table 1; simulated moments from the estimated model summarized in Table 3.


Figure 9: Identification Intuition



Notes: Figure illustrates identification intuition by varying one or two model parameters at a time while holding the remaining parameters fixed at their estimated values. **Data:** Simulations from the model.

F.2 Policy Functions

Figure 10: Contract Choice Probabilities by Type and Associate Year




(A) Female


(B) Male

Notes: Figure reports model-implied contract choice probabilities as a function of permanent type n . Panel A reports results for women and Panel B reports results for men. Within each panel, sub-panels correspond to associate years 2, 4, 6, and 8. The blue line shows the probability of full-time work, the red line shows the probability of part-time work, and the yellow line shows the probability of leaving the firm, including firing risk. Gray shading shows the estimated gender-specific density of permanent type n . **Data:** Simulations from the estimated model summarized in Table 3.

Figure 11: Hours by Type and Associate Year



(A) Female




(B) Male

Notes: Figure reports model-implied optimal hours choices as a function of permanent type n . Hours are normalized for confidentiality. Panel A reports results for women and Panel B reports results for men. Within each panel, sub-panels correspond to associate years 2, 4, 6, and 8. The blue line shows hours chosen under full-time work. The red line shows FTE hours chosen under part-time work, equal to total part-time hours divided by the contracted share of full-time work. The dashed horizontal line marks the bonus threshold. Gray shading shows the estimated gender-specific density of permanent type n . **Data:** Simulations from the estimated model summarized in Table 3.

G Associates on the Margin

Figure 12 illustrates the firm capital paths for three simulated associates with the same permanent type. The marginal partner works full-time in all eight years, while the marginal non-partner enters η_H in years four through six and works part-time during those years. The inframarginal non-partner enters η_H in year two and remains in that state thereafter. The estimated amenity terms imply that the marginal non-partner would need about \$2 million per year to be willing to work full-time while in η_H . Ex ante, her expected payoff to working full-time is only about \$1 million in present value at the start of her part-time spell (reflecting a relatively small increase in her expected promotion probability). However, ex post, working full-time would have pushed her over the promotion threshold, generating a payoff of approximately \$16 million in present value. Thus, uncertainty about the promotion threshold is important for generating near-misses like this.⁵⁶

Figure 12: Simulated Firm Capital Paths




Notes: Figure traces cumulative firm capital over the associate years for three simulated female associates who follow the model's estimated policy functions. All three associates have the same permanent type ($n = 1.67$) and differ only in their latent state paths: the marginal non-partner enters η_H in years 4–6 and works part-time during the spell; the inframarginal non-partner enters η_H in year 2 and remains in it. The dashed line is the promotion cutoff. The promotion cutoff is the cumulative equivalent of the annual threshold reported in Table 3. **Data:** Simulations from the estimated model summarized in Table 3.

⁵⁶Given this, we consider a counterfactual in Section H.5 that reduces promotion policy uncertainty.


H Counterfactuals

Figure 13: Contract Share Effects of Main Counterfactuals




Notes: Figure reports contract share effects of the three main counterfactual policies from Section 5. Panel A varies the FTE-hours weight, λ ; Panel B varies the target female share among promoted partners, $\bar{q}$; and Panel C varies the annual hours cap, z_{cap} , reported in weekly hour equivalents. Outcomes are reported as percentage point changes from the baseline simulated status quo, separately by gender. For each counterfactual value, all other model parameters are held fixed and the promotion threshold is recalibrated to preserve the expected size of the promoted class. **Data:** Simulations from the estimated model summarized in Table 3.

Figure 14: Welfare Effects of Main Counterfactuals



Notes: Figure reports contract share effects of the three main counterfactual policies from Section 5. Panel A varies the FTE-hours weight, λ ; Panel B varies the target female share among promoted partners, $\bar{q}$; and Panel C varies the annual hours cap, z_{cap} , reported in weekly hour equivalents. Outcomes are reported as percentage changes from baseline simulated welfare. Welfare changes are shown separately by gender and tercile of permanent type n . For each counterfactual value, all other model parameters are held fixed and the promotion threshold is recalibrated to preserve the expected size of the promoted class. **Data:** Simulations from the estimated model summarized in Table 3.

Figure 15: Partner Hours Effects of Main Counterfactuals



Notes: Figure reports projected partner-stage billable hours for the promoted classes generated by the three main counterfactual promotion policies from Section 5. Panel A varies the weight placed on FTE hours, λ ; Panel B varies the target female share among promoted partners, $\bar{q}$; and Panel C varies the annual hours cap, z_{cap} , reported in weekly-hour equivalents. For each policy value, we take the resulting promoted class, including its composition over permanent type n and latent state η_T at the promotion date, and project billable hours over the next 30 years, allowing partners to choose part-time contracts as in the associate years. Partner hours are summed over those 30 years and discounted back to period T using $\beta = 0.95$. Outcomes are reported as percentage changes relative to the status quo simulation. The two lines give bounds on projected partner hours under alternative assumptions about the latent state process after promotion. The blue line resets the latent state to η_L at T , eliminating post-promotion differences in η , while the red line allows the latent state process to persist throughout the partner career. The shaded region shows the range of projected outcomes between these two bounds. **Data:** Simulations from the estimated model summarized in Table 3.

H.1 Additional Counterfactual: Persistent Firm Capital


We also consider a counterfactual that varies the persistence of promotion-relevant firm capital. Under the estimated promotion policy, past hours depreciate at the estimated rate $\hat{\delta} = 0.33$, so promotion probabilities place less weight on early-career performance than on more recent hours.⁵⁷ Figure 5A shows that the promotion gap arises relatively late in the associate career, suggesting that placing greater weight on early-career performance could narrow the overall promotion gap. To study this margin, we vary the firm capital depreciation parameter δ . Lower values of δ make firm capital more persistent and therefore place greater weight on hours supplied earlier in the associate career.

Appendix Figure 16 shows how outcomes change as δ varies. Making firm capital more persistent does not close the promotion gap. For lower values of δ , the female share is generally weakly *lower* than the status quo. Increasing the persistence of firm capital also reduces short-run revenue, with total associate billable hours falling below the status quo when δ is below its estimated value. Nonetheless, persistent firm capital can improve partner productivity, with average n among promoted associates above the status quo for lower values of δ .


The limited effect of this counterfactual on the promotion gap reflects offsetting worker responses. More persistent firm capital reduces the weight placed on late-career hours, which weakens the penalty for late part-time spells. However, it also makes early-career shortfalls matter for longer. For associates with low early-career hours, this lowers future promotion prospects and reduces the value of staying at the firm. This retention response is concentrated among women: in the $\delta = 0.1$ case, female retention through the eighth associate year falls by 14%, while male retention is essentially unchanged. Average part-time take-up among women who remain at the firm also falls by 8%, consistent with the policy making low-hours periods more costly for career progression. Thus, making firm capital more persistent modestly improves selection into partnership, but does not reduce the promotion gap and lowers associate output.

⁵⁷Benson, Li, and Shue (2019) similarly find that firms often prioritize current job performance in promotion decisions even when other observables are better predictors of post-promotion performance.


Figure 16: Effects of Firm Capital Persistence




(A) Female Share




(B) Change in Short-Run Revenue




(C) Change in Partner Productivity



(D) Change in Contract Shares



(E) Change in Welfare



(F) Change in Partner Hours

Notes: Figure reports counterfactual effects of varying the depreciation rate of promotion-relevant firm capital, δ . Lower values of δ correspond to more persistent firm capital and therefore place more weight on earlier career hours. For each value of δ , all other model parameters are held fixed and the promotion threshold is recalibrated to preserve the expected size of the promoted class. Panel A reports the female share of promotions. Panels B–C report percentage changes from the baseline simulated status quo in the short-run associate revenue (total associate billable hours) and and partner productivity (average permanent type n). Panel D reports percentage point changes in contract shares by gender. Panel E reports percentage changes in welfare by gender and tercile of permanent type n . Panel F reports simulated percentage changes in partner hours. **Data:** Simulations from the estimated model summarized in Table 3.

H.2 Additional Counterfactual: Non-Equity Partnership


In the baseline model, associates who are not promoted transition to the outside option. We now consider a counterfactual policy in which associates who are not promoted to equity partner at T may instead remain at the firm as salaried non-equity partners. This policy is motivated by the increasing use of non-equity partnership and related senior roles as alternatives to traditional equity partnership (Peery 2017; Lateral Link 2026). In the simulation, non-equity partners earn a fixed premium over their eighth-year associate salary and remain eligible for the bonus offered to eighth-year associates. As above, we hold the number of equity partners fixed. To isolate the effect of introducing a new within-firm job level, we do not allow non-equity partners to be promoted later to equity partner; delayed equity-promotion opportunities are considered separately in Appendix Section H.3.

Appendix Figure 17 shows how outcomes change as the non-equity partner premium rises. The policy has essentially no effect on the female share of equity partner promotions or on partner productivity. However, total associate hours increase modestly as the premium rises.

These patterns reflect the margin through which the policy operates. Introducing non-equity partnership increases the value of remaining at the firm through the promotion stage, but it does not change how equity partner slots are allocated conditional on performance. Higher non-equity partner premiums reduce exit, especially among women, and increase both full-time and part-time work.

Thus, non-equity partnership is not a strong policy for reducing the equity partner promotion gap in this model. Its main effect is to improve retention and generate small gains in associate output, while leaving the gender composition and average type of equity partners unchanged. That said, by reducing exit – especially among women – the policy may still improve the gender balance of the firm overall and may help close the *pay* gap within the firm. This is consistent with real-world concerns that the increased use of non-equity partnership may raise female representation while also creating a “parking lot” for women short of equity partnership (Taddia 2025).

Figure 17: Effects of Introducing Non-Equity Partnership



Notes: Figure reports counterfactual effects of varying the fixed premium that non-equity partners earn over their eighth-year associate salary. For each premium, all other model parameters are held fixed and the promotion threshold is recalibrated to preserve the expected size of the equity partner class. Panel A reports the female share of equity partner promotions. Panels B–C report percentage changes from the baseline simulated status quo in the short-run associate revenue (total associate billable hours) and and equity partner productivity (average permanent type n). Panel D reports percentage point changes in contract shares by gender. Panel E reports percentage changes in welfare by gender and tercile of permanent type n . Panel F reports simulated percentage changes in partner hours. **Data:** Simulations from the estimated model summarized in Table 3.

H.3 Additional Counterfactual: Delayed Equity Partner Review

Under the typical promotion timing in large law firms, associates face the partnership decision after their eighth associate year, at an average age of 37.⁵⁸ The years leading up to promotion overlap with an important period for fertility and family formation – a well-known convergence of the career and family clocks (Goldin 2004). Delaying the partnership decision may allow associates who temporarily reduce hours, including around childbirth, to return to full-time work and accumulate additional firm capital before being evaluated for equity partnership. We therefore consider a counterfactual policy that delays the partnership decision from T to $T + d$ for all associates, where $d \in \{1, 2, 3\}$. During the intervening years, associates who remain at the firm temporarily become non-equity partners under the 20%-premium contract described in Section H.2. We assume that workers do not enter η_H after period $T - 1$, so the delay allows existing high-cost spells to resolve but does not introduce new shocks during the delay period. This policy captures the broader logic of arrangements that give workers a longer period to build a record before a final promotion decision, including tenure-clock policies in academia (Manchester, Leslie, and Kramer 2010; Antecol, Bedard, and Stearns 2018).⁵⁹ As above, we hold fixed the total number of equity partner promotions.


Appendix Figure 18 reports the effects of delaying the partnership decision. The policy has essentially no effect on the female share of equity partner promotions. However, it substantially reduces total associate billable hours (measured over the standard pre-promotion associate years) and weakens selection into equity partnership. Delaying the review gives associates more time to accumulate firm capital, but it also pushes the equity partner prize further into the future and requires additional years of pre-promotion effort. This second force dominates in the simulations. Women’s exit share rises as the delay grows, while men’s contract choices are largely unchanged. Thus, rather than retaining women through a temporary period of lower hours, the delayed review discourages continued firm attachment among the same associates the policy is intended to help.

This result is consistent with evidence from academia that delay-based policies can have unintended gendered effects. For example, Antecol, Bedard, and Stearns (2018) find that gender-neutral tenure-clock stopping policies increased men’s tenure rates and reduced women’s tenure rates. Although the policy considered here differs because it delays review for all associates rather than extending the clock after an individual event, the logic is similar: extending the evaluation horizon need not reduce gender gaps if the additional time is more costly for the more constrained workers.


⁵⁸See Appendix Figure 6.

⁵⁹Unlike a tenure-clock pause, however, the delayed review applies to all associates rather than being triggered by an individual event or contract choice.


Figure 18: Effects of Delaying Equity Partner Review




(A) Female Share




(B) Change in Short-Run Revenue




(C) Change in Partner Productivity



(D) Change in Contract Shares



(E) Change in Welfare



(F) Change in Partner Hours

Notes: Figure reports counterfactual effects of delaying equity partner review by d years. All other model parameters are held fixed and the promotion threshold is recalibrated to preserve the expected size of the equity partner class. Panel A reports the female share of equity partner promotions. Panels B–C report percentage changes from the baseline simulated status quo in the short-run associate revenue (total associate billable hours) and and equity partner productivity (average permanent type n). Panel D reports percentage point changes in contract shares by gender. Panel E reports percentage changes in welfare by gender and tercile of permanent type n . Panel F reports simulated percentage changes in partner hours. **Data:** Simulations from the estimated model summarized in Table 3.

H.4 Equalized Non-Promotable Task Assignment

In addition to differing in the number of hours they bill, associates may differ in the share of total work time that is billable. In particular, if women perform a disproportionate share of non-promotable tasks (Babcock, Recalde, Vesterlund, and Weingart 2017), then observed billable hours may understate women’s total work output. Although we do not observe total work hours separately from billable hours, we can use the model to consider a counterfactual policy that increases the share of women’s work time that is billable, implemented as a proportional rightward shift in women’s promotion-relevant billable hours.

Specifically, firm capital accumulates according to:

$$h_{t+1}^{CF} = (1 - \delta)h_t^{CF} + \tilde{z}_t,$$


where $\tilde{z}_t = \lambda_f z_t$ for female associates and $\tilde{z}_t = z_t$ for male associates, with $\lambda_f > 1$ capturing the increase in women’s billable share of work time.⁶⁰ When $\lambda_f = 1$, the policy coincides with the status quo. We vary λ_f over a grid and, as above, recalibrate the promotion threshold to keep the expected size of the promoted class fixed.⁶¹

Appendix Figure 19 reports the effects of these adjustments. The female share of promotions rises monotonically with λ_f , reaching 43% when $\lambda_f = 1.1$ and 52% when $\lambda_f = 1.2$. Short-run associate revenue also increases slightly, driven by higher female retention. The effects on partner productivity depend on the size of the adjustment. For moderate values of λ_f , average n among promoted associates is essentially unchanged or slightly higher. For larger adjustments, partner productivity falls. This pattern suggests that moderate adjustments could improve representation with limited effects on productivity but that larger across-the-board adjustments may weaken selection into equity partnership. That said, since we do not observe total work hours or task assignment, this exercise should not be interpreted as an estimate of the effect of a specific policy change.


⁶⁰This implementation abstracts away from effects that a change in the billable share would have on bonus awards and firing risk, though.

⁶¹Vesterlund (2025) suggest that non-promotable tasks take up about one additional month of women’s work time each year. If that month were reallocated to billable work, this would correspond to an adjustment of approximately $\lambda_f = 12/11 = 1.09$.


Figure 19: Effects of Increasing Female Share of Promotable Tasks




(A) Female Share




(B) Change in Short-Run Revenue




(C) Change in Partner Productivity



(D) Change in Contract Shares



(E) Change in Welfare



(F) Change in Partner Hours

Notes: Figure reports counterfactual effects of adjusting female billable hours up by a constant λ_f . All other model parameters are held fixed and the promotion threshold is recalibrated to preserve the expected size of the equity partner class. Panel A reports the female share of equity partner promotions. Panels B–C report percentage changes from the baseline simulated status quo in the short-run associate revenue (total associate billable hours) and equity partner productivity (average permanent type n). Panel D reports percentage point changes in contract shares by gender. Panel E reports percentage changes in welfare by gender and tercile of permanent type n . Panel F reports simulated percentage changes in partner hours. **Data:** Simulations from the estimated model summarized in Table 3.

H.5 Reducing Promotion Policy Uncertainty

We also consider a counterfactual that reduces associates’ uncertainty about the promotion threshold.⁶² Greater transparency about associates’ relative rank, for example, could reduce uncertainty, even if the realized threshold is only determined after the firm observes the full distribution of period- T firm capital. To study this margin, we increase the steepness of workers’ perceived promotion probabilities such that:


$$\gamma_p^{CF} = \frac{\hat{\gamma}_p}{1 - u},$$

where u is a proportional reduction in uncertainty. We hold all other model parameters fixed and recalibrate the promotion threshold to preserve the expected size of the promoted class.

Appendix Figure 20 reports the effects of reducing uncertainty. Even under a substantial 50% reduction in uncertainty, the female share of promotions rises only modestly – by about 3 percentage points – while short-run associate revenue falls by 4% and partner productivity improves by 3%. These patterns reflect how uncertainty shapes associate incentives. Under high uncertainty, even associates far from the eventual promotion threshold perceive some return to staying at the firm and supplying additional hours. However, reducing uncertainty concentrates these incentives among associates closer to the promotion margin.

⁶²As discussed in Appendix Section B.2, uncertainty about the threshold can arise when associates do not observe the full distribution of peer performance. This uncertainty weakens the perceived return to additional hours for associates near the promotion margin, as illustrated in Appendix Section G.

Figure 20: Effects of Reducing Promotion Policy Uncertainty



Notes: Figure reports counterfactual effects of reducing promotion policy uncertainty through γ_p . All other model parameters are held fixed and the promotion threshold is recalibrated to preserve the expected size of the equity partner class. Panel A reports the female share of equity partner promotions. Panels B–C report percentage changes from the baseline simulated status quo in the short-run associate revenue (total associate billable hours) and equity partner productivity (average permanent type n). Panel D reports percentage point changes in contract shares by gender. Panel E reports percentage changes in welfare by gender and tercile of permanent type n . Panel F reports simulated percentage changes in partner hours. **Data:** Simulations from the estimated model summarized in Table 3.

I Law Student Survey

Q1. Have you been a summer associate at a large law firm (e.g., Am Law 100)?

(1) Yes

(2) No

Q2a [displayed if “Yes” to Q1]. During your most recent summer, how often did you:

	Never	Sometimes (10–40%)	About half of the time	Often (60–90%)	Always
Work evenings or weekends	☐	☐	☐	☐	☐
Feel pressure to be “always available”	☐	☐	☐	☐	☐
Have work reassigned to others close to deadlines	☐	☐	☐	☐	☐
Receive feedback on pace/availability	☐	☐	☐	☐	☐
Have difficulty finding work	☐	☐	☐	☐	☐

Q2b [displayed if “Yes” to Q1]. Compared with other summer associates, your workload felt:

(1) Much lighter (20+% less work)

(2) Lighter (5–20% less work)

(3) Similar (5% less to 5% more work)

(4) Heavier (5–20% more work)

(5) Much heavier (20+% more work)

Q2c [displayed if “No” to Q1]. Based on your general understanding, how often do summer associates in large law firms:

	Never	Sometimes (10–40%)	About half of the time	Often (60–90%)	Always
Work evenings or weekends	☐	☐	☐	☐	☐
Feel pressure to be “always available”	☐	☐	☐	☐	☐
Have work reassigned to others close to deadlines	☐	☐	☐	☐	☐
Receive feedback on pace/availability	☐	☐	☐	☐	☐
Have difficulty finding work	☐	☐	☐	☐	☐

Q3a [displayed if “Yes” to Q1]. During your most recent summer, how often did full-time associates:

	Never	Sometimes (10–40%)	About half of the time	Often (60–90%)	Always
Work evenings or weekends	☐	☐	☐	☐	☐
Feel pressure to be “always available”	☐	☐	☐	☐	☐
Have work reassigned to others close to deadlines	☐	☐	☐	☐	☐
Receive feedback on pace/availability	☐	☐	☐	☐	☐
Have difficulty finding work	☐	☐	☐	☐	☐

Q3b [displayed if “No” to Q1]. Based on your general understanding, how often do full-time associates in large law firms:

	Never	Sometimes (10–40%)	About half of the time	Often (60–90%)	Always
Work evenings or weekends	☐	☐	☐	☐	☐
Feel pressure to be “always available”	☐	☐	☐	☐	☐
Have work reassigned to others close to deadlines	☐	☐	☐	☐	☐
Receive feedback on pace/availability	☐	☐	☐	☐	☐
Have difficulty finding work	☐	☐	☐	☐	☐

Q4a [displayed if “Yes” to Q1]. Does the firm from your most recent summer allow reduced-hours/part-time schedules for associates?

- (1) Yes, formally
- (2) Yes, informally
- (3) No

Q4b [displayed if “Yes” to Q1]. At the firm from your most recent summer, how common are reduced-hours/part-time schedules for associates?

- (1) Very rare (less than 1% of associates)
- (2) Somewhat rare (1–5% of associates)
- (3) Somewhat common (5–10% of associates)
- (4) Very common (more than 10% of associates)

Q4c [displayed if “No” to Q1]. Based on your general understanding, do large law firms allow reduced-hours/part-time schedules for associates?

- (1) Yes, formally
- (2) Yes, informally
- (3) No

Q4d [displayed if “No” to Q1]. Based on your general understanding, how common are reduced-hours/part-time schedules for associates in large law firms?

- (1) Very rare (less than 1% of associates)
- (2) Somewhat rare (1–5% of associates)
- (3) Somewhat common (5–10% of associates)
- (4) Very common (more than 10% of associates)

Q5. How likely are you to aim for each of the following career outcomes?

	Not likely (<10%)	Somewhat likely (10–40%)	Likely (40–60%)	Very likely (60–90%)	Almost certainly (>90%)
Equity partner role at a large firm	☐	☐	☐	☐	☐
Non-equity partner role at a large firm	☐	☐	☐	☐	☐
Equity partner role at a smaller firm	☐	☐	☐	☐	☐
Non-equity partner role at a smaller firm	☐	☐	☐	☐	☐
In-house role	☐	☐	☐	☐	☐
Government/public interest role	☐	☐	☐	☐	☐
Academic role	☐	☐	☐	☐	☐
Other role	☐	☐	☐	☐	☐

Q6. How strongly do you agree with each of the following statements?

	Strongly disagree	Somewhat disagree	Neither agree nor disagree	Somewhat agree	Strongly agree
I value work-life balance highly in deciding between career paths.	☐	☐	☐	☐	☐
I'm willing to trade personal time for faster advancement.	☐	☐	☐	☐	☐
I expect my goals to change after the first few years of practice.	☐	☐	☐	☐	☐

Q7. In your opinion, how much do the following matter for advancement toward equity partner in a large law firm?

	Not at all important	Slightly important	Moder- ately important	Very important	Extremely important
Visibility and availability	☐	☐	☐	☐	☐
Work quality	☐	☐	☐	☐	☐
Total hours billed	☐	☐	☐	☐	☐
Efficiency per hour billed	☐	☐	☐	☐	☐
Business development	☐	☐	☐	☐	☐
Face time / being seen working late	☐	☐	☐	☐	☐
Willingness to take last-minute work	☐	☐	☐	☐	☐

Q8. The following questions will ask you to think about hypothetical scenarios. For each question always assume the following.

1. Adam and Bob are senior associates in a large law firm.
2. They are identical in every way (same practice, office, work quality, professional relationships, etc.).
3. Except: the only difference between them is that Adam is on a full-time schedule (earning standard compensation) while Bob is on a reduced-hours schedule (earning 80% compensation).

Q8a. If Adam bills 2,000 hours and Bob bills 1,400 hours, who is more likely to be promoted to equity partner?

- (1) Adam
- (2) Bob
- (3) Equal chances

Q8b. If Adam bills 2,000 hours and Bob bills 1,600 hours, who is more likely to be promoted to equity partner?

- (1) Adam
- (2) Bob
- (3) Equal chances

Q8c. If Adam bills 2,000 hours and Bob bills 1,800 hours, who is more likely to be promoted to equity partner?

- (1) Adam
- (2) Bob
- (3) Equal chances

Q8d. If Adam bills 2,000 hours and Bob bills 2,000 hours, who is more likely to be promoted to equity partner?

- (1) Adam
- (2) Bob
- (3) Equal chances

Q8e. If Adam bills 2,000 hours and Bob bills 2,200 hours, who is more likely to be promoted to equity partner?

- (1) Adam
- (2) Bob
- (3) Equal chances

Q9. How confident are you in your answers to this survey?

	Not very confident (<20%)	Slightly confident (20–50%)	Somewhat confident (50–80%)	Very confident (>80%)
How confident?	○	○	○	○

Q10. What sources of information have shaped your answers? (Select all that apply)

- ☐ Direct experience
- ☐ Observation of associates or partners
- ☐ Conversations with associates or partners
- ☐ Career services / recruiting materials
- ☐ Conversations with peers
- ☐ General reputation of firms

Q11. Is there anything else about law firm workload and career planning that you'd like to share? [free response]

J Qualitative Evidence on How to Make Partner

From posts in online forums about working in large private sector law firms:

- “Having no soul, white and grey haired man. Also having a family of four but never talking about their children or wife. Billable hours are the way to live: one 0.1 at a time.”
- “At my firm, 2,000 billable were a functional mandatory minimum. My 8th year, when I was up for partner, I billed 2,400 or so.”
- “At this point all you can do is be a damn good associate and bust your ass. If you don’t do that first the chances are literally zero. Even if an associate brings in business independently, it’s all about scale. . . \$100k a year in fees is not going to move the needle in a way that makes a partner case.”
- “Totally numbers driven. It’s very hard to make equity billing less than 1900 or so.”
- “A former classmate of mine just got a ‘time to focus on partnership’ pitch from group heads to seniors. It was conveyed you are expected 2400-500 of true client billables to be on the path to consideration. You are a profit center. Nothing more.”
- “Sell your soul. Say goodbye to work life balance.”