Wage Signaling, Salary History Bans, and Equality

Jeffrey Meli & James C. Spindler[1]

We develop a model of employee mobility and wage formation suitable for analyzing labor market interventions that target wages and discrimination. By augmenting the standard asymmetric learning model with costly effort, such that wages are necessary for production, wages become fully informative of employee ability rather than merely products of public signals and employer competition for workers, as in prior models. When salary disclosure is allowed, strategies designed to mask employee quality, commonly seen in standard models, are untenable; instead, the equilibrium has symmetric learning, optimal employee assignment, and efficient production. When salary disclosure is prohibited – actual policy in many jurisdictions – the equilibrium reverts to asymmetric learning, with strategic under-assignment due to manipulation of noisy signals. Equal pay mandates that restrict intra-title wage variation, as may result from anti-discrimination law, prevent firms from fully utilizing employee skill. In the presence of discrimination, both salary history bans and equal pay mandates may result in reduced wage gaps, though at the expense of reduced production efficiency and returns to skill.

[1] Jeffrey Meli is Clinical Professor at the Stern School of Business (jm11241@stern.nyu.edu). James C. Spindler is a Professor at University of Texas School of Law and the McCombs School of Business (jspindler@law.utexas.edu.). The views of the authors are their own, and do not necessarily reflect the views of their respective institutions.

I. Introduction

In this article we develop a theoretical model of wage formation in which wage history is fully revealing of employee skill, rather than merely reflective of observable employee characteristics, such as title. We use our framework, which is an augmented version of the standard asymmetric learning model, to provide the first theoretically grounded comparison of three different legal regimes governing wage formation: wage disclosure, in which outside firms can inquire about prior wages, a "salary history ban" prohibiting such inquiry, and equal pay laws that restrict variation in wages across employees. We consider both an unbiased environment and one in which certain employees (women) are subject to discrimination in the form of a glass ceiling. In an unbiased market, wage disclosure results in an efficient equilibrium equivalent to symmetric learning, in which all employees are paid their full (outside) marginal product. In the discriminatory setting, wage disclosure only partially mitigates the effect of bias; employees subject to early-career discrimination benefit from wage disclosure but are still paid less on average over the course of their careers. Policy makers have adopted new labor market interventions to combat such wage inequities (Cullen 2023), including salary history bans and equal pay reforms. However, we demonstrate that while salary history bans and equal pay laws may have some success in reducing wage gaps, they come with significant costs. First, they each introduce inefficiencies which reduce overall wages. Second, these interventions have distributional consequences that fall most heavily on high skill employees, including those most impacted by discrimination (i.e., high skill women). In other words, these interventions reduce both overall welfare and the welfare of the very employees they are intended to help.

Our model is an improvement over existing models of wage formation in labor markets, in which wages reflect publicly available information (such as title), and thus are an output of competition, rather than an input into production. As a result, wage disclosure is either irrelevant or assumed to be banned, and thus these models are poorly suited to analyze legal interventions restricting wage formation or disclosure.[2] In contrast, we establish a link between wages, productivity, and skill, by making two adjustments to the standard asymmetric learning model. First, we add disutility of effort, with disutility increasing along the intensive margin of

[2] In some models, wage disclosure affects the relative bargaining power of employers and employees, but not overall production. See the literature review for a comprehensive comparison of our framework to existing models.

employee production. This simple but fundamental import from contract theory implies that higher wages are required for increased production, because higher skilled employees must be paid sufficiently to motivate harder work. Second, we assume that task assignments are continuous, rather than discrete as is typical in the literature. Outside firms cannot observe the precise assignments, but they can observe a noisy signal of assignment, akin to the "promotion" threshold that is standard in the literature.

With these adjustments, we obtain a tractable model in which wages are informative of employee skill. In fact, when wages can be disclosed, the disutility of effort ensures that they fully reveal employee ability. Unlike in the standard literature, strategic underpromotion or other masking strategies are untenable. The difference between our conclusion and the standard result is driven by the fact that employees maximize utility net of effort, rather than just their wage. An incumbent employer cannot allow a negative utility gradient to exist with respect to skill, as the highest-skill employees would be easiest to poach. An incumbent firm that paid a pooling wage to mask the skill of certain employees loses its highest skilled (but lowest utility) employees to competing firms that match the pooled wage but pair it with a job assignment that is utility-increasing only for those employees, who will earn the same wage but exert less effort. Lower skilled employees reject the offer because it pays the same wage but requires more effort. Outside firms would earn profits from the highly skilled employees they poached, and the incumbent firms would be left with highly paid but low skill employees.

As a result, a viable equilibrium requires that incumbent firms pay their more skilled employees higher wages. Competing firms induce employee quality from the wage, and they bid up wages to marginal product.[3] The result is equivalent to a competitive equilibrium with symmetric information: employees earn their full (outside) marginal product every period, firms make no profits (excepting those linked to firm-specific human capital), wages start low and rise for skilled employees over time, and the full productive capacity of the economy is realized each period. This occurs even in absence of any extrinsic public signal; informative wages imply that noisy signals of employee assignment are irrelevant. Our analysis provides a theoretical

[3] Disutility of effort presents another challenge to pooling equilibria: if pooling were viable, lower skilled employees would earn higher utility than higher skilled employees. This would create disincentives for high-ability employees to reveal their ability in the first place, which would require them to accept utility-decreasing promotions. Although we do not rely on such a revelation mechanic in our model, we do think it illustrates another reason why costly effort renders masking and wage pooling strategies untenable.

grounding for the empirical observation that, in the status quo ante, many employers asked about prior wages when making lateral hires. Hall & Krueger (2012) report that "about half" of workers reported that their employers learned their salary histories before extending job offers, and Barach & Horton (2017) find that for workers whose employers learned their histories, such learning occurred prior to extending the job offer in more than 80% of cases.[4]

We then analyze the effect of the new class of labor market interventions. One prominent intervention is the salary history ban, a recent regulatory reform in important labor markets in the United States.[5] Bans seek to suppress the information content in wages by prohibiting firms from asking prospective employees about their prior salaries. They are intended to benefit workers, particularly women and minorities, by reducing "anchoring," whereby unfairly low wages follow employees throughout their careers.[6]

In our framework, a salary history ban jams the wage signal and leads to an equilibrium with inefficient utilization of employees. Since outside firms cannot utilize the wage signal, they rely instead on the noisy signal of employee assignment. This reintroduces the associated inefficiencies linked to manipulation of these signals. Incumbent firms strategically underpromote some skilled employees, in keeping with the standard asymmetric learning models. Underpromotion reduces aggregate productivity and thus aggregate wages. Bans also result in significant distributional consequences. Jamming the wage signal reduces the value of the outside option for more skilled employees, and thus their second period wage. Lower second period wages increase employer profits, which are reflected in higher wages for new entrants to the labor market. Employees with skill at the low end of each title band benefit from this tradeoff; they receive a second period wage close to their outside marginal product and a higher entry wage. In contrast, bans reduce the welfare of employees at the high end of the skill distribution within each title because of the significant decline in their second period wage.

---

[4] Bessen, Meng & Denk (2020) posit an alternative explanation: that employers request salary history because it reveals heterogeneous reservation wages among otherwise identical workers.

[5] Enacting jurisdictions include New York City, California, and Massachusetts. One source counts 22 states as having enacted some form of salary history ban. See https://www.hrdive.com/news/salary-history-ban-states-list/516662/.

[6] While one might expect informational unraveling leading to voluntary disclosure, it is the stated purpose of these laws to prevent the use of such information. Accordingly, some of these laws, such as those of California and Massachusetts, render voluntarily disclosure problematic. Voluntary disclosure is also potentially subject to later claims of illegal "prompting." As a result, certain large employers, such as Google, Facebook, and Amazon, have stated that they will no longer utilize salary histories.

We then consider what we broadly term equal pay mandates, which require equal pay among persons performing the same job. These include anti-discrimination laws (such as the federal Equal Pay Act) and recent court interpretations that require employers to demonstrate that pay disparities are based on acceptable factors (i.e., not sex or race) and which disallow past pay as an acceptable factor.[7] We note that modern reforms promoted by pay equity advocacy groups include lockstep compensation, limiting managerial discretion in setting pay, and standardizing compensation.[8] Further, many jobs, such as those in the public sector or in unionized settings are already subject to such practices. These policies cabin potential wages and provide transparency to existing employees about the prevailing market wage for their role.

We model equal pay mandates as limiting the within-title variation of wages. This leads to underutilization of employees, but through a subtler (and new) channel. Because tasks are continuous, but the "promotion" threshold is discrete, employees with a range of ability can have the same title. Equal pay mandates force firms to pool employees, such that workers with different ability levels earn the same wage. However, costly effort prevents firms from fully utilizing the higher-skilled employees within a given job title because requiring them to exert more effort would create a negative utility gradient in skill, allowing outside firms to poach the more-skilled employees. Conceptually, by constraining wages within titles, an equal pay mandate constrains the maximum feasible assignment of each title. The result is within-title underutilization of employees, a novel source of inefficiency, and a corresponding reduction in average wages and average productivity.

The distributional consequences are similar to a salary history ban: employees at the upper end of the skill distribution within each title are made worse off, because they cannot be assigned (and paid) up to their true ability. This conclusion accords with recent empirical analysis of wage-equality initiatives (Cullen 2023). However, employees are paid wages (approximately) equal to their realized, as opposed to potential, outside marginal product. This

[7] A recent 9th Circuit opinion interpreted the Equal Pay Act to required employers to prove that pay differentials are based upon valid factors, such as productivity, as opposed to sex, and disallowed the use of salary histories as such a valid factor; the minority dissent noted that this rule would "demand a lockstep pay system." Rizo v. Yovino, Case No. 16-15372, 9th Circuit, February 27, 2020.

[8] These recommendations come from the California Governor's Commission on the Status of Women and Girls. See, What Can I do to Promote a Culture of Pay Equity?, https://women.ca.gov/californiapayequity/employers-resources/what-can-i-do-to-promote-a-culture-of-pay-equity/.

limits the lifetime profits and the associated increase in initial wages paid to new entrants to the labor market. Although promotion decisions are generally inefficient under an equal pay mandate, both under-assignment and over-assignment are possible. The latter is driven by the "wedge" between feasible assignments at each title. When the maximum feasible assignment for employees who are not promoted is low, it may be worthwhile to over-assign employees with skill close to, but below, the public promotion threshold, if the losses from over-assignment are below the gains from greater utilization of their skill.

Of course, these interventions were not created in a vacuum. They are intended to combat discrimination, and advocates of these policies would likely accept some aggregate efficiency loss in exchange for greater fairness in labor markets. To better understand this tradeoff, we introduce a realistic and tractable form of discrimination: some women employees are subject to a "glass ceiling". Their perceived skill (as observed, say, by an immediate supervisor) is incorrectly capped just below the promotion threshold.

When wage disclosure is allowed, some firms can successfully poach women who are subject to discrimination. However, successful poaching does not require an offer equal to their full marginal product; the outside offer need only provide incrementally greater utility to be accepted. As a result, mobility reduces, but does not eliminate, the gender pay gap. Some women are paid less than men with the same assignment, and the effects of early-career bias have long-term effects on women's wages (and utility). Put differently, market forces only partially mitigate the effects of discrimination, and the labor market is characterized by a pay gap, anchoring, and lesser pay for equal work, all features that the interventions discussed above are intended to address.

However, both interventions have negative effects. A salary history ban reduces the welfare of employees just below the promotion threshold, which includes all women subjected to the glass ceiling. Bans increase the reliance on noisy promotion signals; so long as discrimination reduces the number of women who are promoted, the distributional consequences fall most heavily on the victims of discrimination. Equal pay mandates can have a similar effect: under some parameters, they reduce the welfare of employees just below the promotion threshold, which includes all employees subject to discrimination. Interestingly, however, equal pay mandates can have a salutary effect in which they cause firms to promote high-skilled women despite the glass ceiling; this occurs when there is a large gap between the productivity

of promoted and non-promoted employees, effectively causing the firm to lower the threshold for promotion. That said, from a distributional perspective, wage disclosure is generally the best scenario for high skill women.

We make a number of empirical predictions regarding the implications of bans and pay transparency laws based on this analysis. First, we expect that the evidence for asymmetric learning will be stronger in jurisdictions that enact salary history bans. Our assumptions regarding costly effort and continuous tasks are plausible for any jobs where the intensive margin of employees varies significantly (e.g., deal volume of an investment banker), as opposed to occupations where the work is largely commoditized (e.g., instructors teaching standardized course loads). We predict less evidence of asymmetric learning in jobs with significant within-title wage variation when wages can be disclosed, and greater evidence after the imposition of a ban. That said, we expect bans to reduce ex-post income inequality, due to the reduction of the outside option and associated reduction in the return to skill earned by more skilled employees. Again, this effect will be more pronounced in roles where there is significant intra-title variation in tasks and in pay. Bans will also lead to increased wages for new entrants into the labor market. This is not true under equal pay mandates, where the second period wages reflect realized (as opposed to potential) productivity; we expect the reduced wage inequality will come from reduced overall wages. Finally, bans and equal pay mandates create opposing incentives for firms' organizational structure. Bans allow firms to mask their employees' skill, and thus they lead firms to prefer fewer titles (which are observable to outside firms). In contrast, the limits on within-title utilization associated with equal pay mandates incentivize firms to create more titles.

The rest of the article proceeds as follows. In Section II, we relate our model to the prior literature on employer learning and emerging research on the effect of salary history bans and other equality initiatives. In Section III, we develop a formal model of wages and mobility with asymmetric learning and costly effort. In Section IV, we demonstrate that wage disclosure leads to a unique equilibrium equivalent to symmetric learning, where wages equal marginal product. In Section V, we demonstrate the inefficiencies that arise from suppressing the wage signal via salary history bans and equal pay mandates. In Section VI we assess the effect of these interventions on fairness in labor markets, using a tractable form of discrimination. Section VII discusses our results and paths for future work.

II. Relation to the prior literature

A. Theoretical models of employer learning

This article builds on the significant theoretical labor economics literature on asymmetric learning. [9] This literature seeks to explain the role of adverse selection in employee job mobility and the potential negative effects of competitive hiring pressure, which create incentives to deter poaching of employees (Greenwald 1986, Waldman 1984). In the standard approach, public signals are noisy and endogenous, and inefficiencies result from manipulation of the endogenous signal.[10] A typical assumption is that outside firms can observe job assignments (Waldman, 1984) or promotions (DeVaro and Waldman, 2012). In equilibrium, employers underpromote their high skill or high productivity workers to deter poaching by competitor firms (Waldman & Zax (2016) demonstrate the robustness of this result to a range of modelling assumptions). Wages do not generally reflect an employee's marginal product, because competitor firms fail to learn about employee ability. This results in high wages for new entrants to the labor market, as firms compete away future profits they expect to earn from new hires.

The existing models provide limited insight into the effects of labor market interventions designed to regulate or impinge on wage formation, or that seek to suppress the information contained in wages. In the class of models that closely follows the standard approach, wage disclosure (or lack thereof) is irrelevant: wages reflect public signals and are not themselves incrementally informative (Waldman 1984). Equal pay mandates are also irrelevant because these models feature no intra-title wage variation. There are more recent studies in which wages are assumed to have signaling value, and banning disclosure can affect the distribution of bargaining power between firms and employees. For example, in Golan (2009) wages serve as a signal of productivity; differential wages arise from an extrinsic bargaining game over surplus between employer and worker. Pinkston (2009) presents an asymmetric learning model with wage disclosure, where wages are a function of exogenous signals and labor market competition. However, suppressing the information content in wages does not lead to efficiency losses,

[9] If learning was symmetric, then none of these laws would have any effect on wages.
[10] In some studies, the signals are assumed to be exogenous but noisy (see Schonberg (2007), where greater noise generates more adverse selection in worker mobility).

because the productivity of any individual employee is unaffected by her wage.[11] Without competition, employers in these models would need not pay differential wages. Further, in both the original and the more recent models, policy goals designed to achieve wage equality could be achieved by government fiat. For example, any biases or unfairness in labor markets could be resolved by mandating a set wage for all workers in the economy, without any loss of efficiency. Similarly, outlawing competition for workers would equalize wages without a negative consequence. In fact, some such policies create efficiency gains in some of these models. For example, in Waldman (1984), the government could mandate either a wage of zero for all employees or forbid competition for employees. This improves efficiency: Waldman's firms otherwise mismatch employees to jobs in order to limit poaching.[12] If such drastic, and obviously destructive, policies have no aggregate consequences in these models, we believe they are not suited to analyze the more sophisticated recent interventions, such salary history bans and various wage equality initiatives (nor were they designed to do so).

To that end, we put forth a parsimonious model in which wages are an input into production, by which we mean higher wages are necessary to increase output. We demonstrate that wages then serve as signals of quality. As in Waldman (1984) and much of the literature that follows, in our model firms divide their employees' time between two tasks: an unskilled task and a skilled task that is more productive for high ability employees. As is also standard in these models, incumbent firms learn about their own employees' capabilities at the skilled task over time, whereas outside firms do not. Our addition is that the skilled task requires costly effort, increasing with the amount of the task assigned to the employee, and we model the wage as publicly observable.

To be clear, ours is not the first study to consider costly effort in the context of asymmetric learning. Gibbs (1995) analyzes a model in which production is determined by the product of effort and ability. More recent articles assume that production is a function of the sum of effort and ability, including Ghosh and Waldman (2010) and Ekinci, Kauhanen, and Waldman (2018), both of which analyze how optimal contracting relates to learning, effort, and public

[11] More specifically, once the reservation wage is met, employees will produce at any feasible level regardless of the wage. In contrast, the contract theory literature considers settings in which compensation schemes are necessary to promote optimal employee production. Bolton and Dewatripont (2004) provide a survey of such models.

[12] Eliminating all returns to skill could result in other inefficiencies, such as reducing the incentive for workers to invest in their human capital. Banning wage disclosure has a similar effect in our model.

signals. However, these studies either implicitly (e.g., Gibbs, 1995) or explicitly (e.g., Ekinci, Kauhanen, and Waldman, 2018) assume that wages cannot be disclosed. For example, in the equilibria computed in Ekinci, Kauhanen, and Waldman (2018), the first period compensation (wage plus bonus) would be fully informative of employee ability, could it be disclosed to outside firms. Ours is the first article to explicitly study how costly effort affects the information content of wages, how this interacts with wage formation and learning, and how various interventions into wage formation and disclosure affect job assignments, wages, and efficiency.

B. Early empirical work on salary history bans equal pay mandates

Our model also helps explain recent empirical findings regarding salary history bans, requirements for wage ranges in job listings, and pay transparency. We predict that salary history bans will lead to greater evidence of asymmetric learning, particularly for jobs in which higher wages are required to motivate greater productive effort. We expect that pooling and masking strategies will become relatively more prevalent in those jurisdictions, and that mobility will increase among lower wage workers. The pay of low skill workers will increase, and that of high skill workers will decline; in addition, bans reduce overall productivity. Equal pay mandates have similar effects, with low skill, low pay workers gaining at the expense of higher skill workers and overall productivity.

These predictions are broadly consistent with empirical work to date. Initial studies of salary history bans have largely explored demographic cleavages, such as the impact on the gender and minority wage gaps. Agan, Cowgill & Lee (2021) find that salary history bans lead to less inequality in salary offers, with lower salary offers overall. Bessen, Meng & Denk (2020) find that salary history bans generally increased wages, with disproportionate benefits accruing to white females, non-white males, and white male job-switchers. Davis, Ouimet & Wang (2021) find indications that salary history bans made employers more cautious in their offers to new hires, offering lower wages as a result. Hansen & McNichols (2020) detect positive effects concentrated on women over 35 or with children. Mask (2020) examines “scarred” workers (those who employment history is hampered by a recession) and observes that they enjoy increased mobility due to salary history bans. Sinha (2019) provides evidence that salary history bans benefited females and hurt males, while an update (Sinha (2022)) finds effects only on white females (positive). Khanna (2020) provides experimental evidence that the ability to

volunteer salary history data frustrates their intended purpose.[13] Data from Sran, Vetter & Walsh (2020) show new hiring declining after the implementation of a salary history ban. Cullen & Pakzad-Hurson (2021) find that that salary transparency rules (specifically, allowing workers to share salary information with one another) lead to lower overall salaries, due to employer incentives to reduce salary inflation. Cullen (2023), in a survey of recent literature, reports similar results, noting that "'horizontal' pay transparency policies that reveal pay gaps between co-workers at the same firm… lower worker bargaining power and wages," while noting that providing information to workers about job opportunities across the landscape of firms appears to raise wages.

Overall, these results suggest increased wage equality but at the cost of efficiency losses, reflected by lower overall wages (though we note that there is not uniformity in these findings). Further, the measured effects appear to be largely distributional, with transfers from higher wage to lower wage workers (such as from white to non-white and male to female).

III. The model: wage formation with costly effort

Our approach is based on standard asymmetric learning models, with two major modifications. Firstly, and most importantly, we assume that that effort is costly, and that increased production requires greater effort. This differs from other models with costly effort, in which effort is additive to production but not required for it.[14] In order to abstract from the monitoring or contracting necessary to incentivize effort, we assume that firms can observe the effort exerted by their employees.[15] The second major departure from the standard approach is that we assume that task assignments are continuous: employees split their time between the different tasks, rather than perform only one or the other task. We assume that the precise assignments cannot be credibly disclosed to outside firms. This is motivated by the continuous

[13] As stated above, there are reasons to believe that such volunteering is limited. Davis, Ouimet & Wang (2021) provide evidence that salary history bans may change employer norms and practices, such that salary histories may not be utilized even when technically allowed (as it often is in the public sphere).
[14] Our assumption is similar to Gibbs (1995), in which production is the product of effort and ability.
[15] We have developed an alternative version of the model in which effort is only imperfectly observable. If firms identify employees who are "shirking" with some probability then they must gross up their wages to induce effort accordingly (such that employees prefer to exert effort rather than risk being caught and getting fired, which is a version of an efficiency wage). This results in a proportional increase in the equilibrium wage, but otherwise generates identical conclusions so long as the productivity of the skilled task is high enough to sustain the efficiency wage. This enhancement adds parameters and algebra but no economic intuition, and thus we present the version with observable effort.

nature of the assignments, as outside firms would struggle to verify how exactly an employee's time was allocated across tasks. However, we do assume that there is a "promotion" threshold that is observable by outside firms, akin to the promotions considered in the standard models. If an employee's assignment to the higher value task exceeds this threshold, outside firms recognize her as having been promoted.

All our other assumptions are based on those in the standard approach, albeit with minor adjustments to their precise structure as required by the two major departures. We discuss the two major departures and how they interact with the other assumptions to drive our main result in more detail in Section IV, after first solving the model under wage disclosure.

A. The players

We assume there exists a continuum of profit maximizing firms, which have the same constant returns to scale production function that spans two tasks, both of which require only labor as an input. We assume a mass of employees have a unit of time to spend at work each period, which they divide between the two tasks.

One task is "unskilled." It requires no effort and generates the same expected production (per unit of time) for every employee. Without loss of generality, we normalize the expected productivity of the unskilled task to 0.[16] The second task is "skilled," which requires (observable) effort *e* per unit of time to be successfully completed (i.e., we assume a linear cost of effort). Each employee *i* is endowed with a capacity for the skilled task of $\alpha_i \in [0,1]$. The probability distribution function of skill in the population is given by $f(\alpha)$, with expected value $E[\alpha]$. This distribution is known to the players. The per-unit of time expected production at the skilled task is $y > e$, so long as the employee spends no more than α_i of her total time at work that period on the skilled task. Any allocation beyond α_i generates negative expected production. We assume that employees develop firm-specific human capital that increases production if the employee stays with the same firm for more than one period. Production per unit of time in the unskilled task rises by $s_0 > 0$, and in the skilled task by a factor $s_1 > 0$. We assume that $y * s_1 > s_0$; firm

[16] The necessary assumption is that both the effort level and the productivity (when performed up to the capability of an employee) of the skilled task are above those of the unskilled task. Then we normalize both effort and productivity accordingly.

specific human capital does not change the ordinal relationship between the productivity of the two tasks.

Formally, an employee *i* endowed with skill α_i, employed by firm *j*, and assigned to spend a fraction $\beta_{i,t}$ of her work hours in period *t* at the skilled task (and the rest at the unskilled task), generates expected production $E[Y_{i,j,t}]$ of:

$$E[Y_{i,j,\mathrm{t}}] = \left(1 - \beta_{i,t}\right) * s_0 * I_{i,j,t-1} + \min\left(\beta_{i,t}, \alpha_i\right) * y * \left(1 + s_1 * I_{i,j,t-1}\right) + \max\left(\beta_{i,t} - \alpha_i, 0\right) * L \quad [1]$$

where $L < 0$ is the loss associated with assigning an employee the skilled task beyond her capacity, and $I_{i,j,t-1}$ is a dummy variable equal to 1 when employee *i* worked at firm *j* in period *t-1*.

In keeping with the standard asymmetric learning literature (e.g., Waldman 1984), we assume that the distribution of skill is such that the unconditional expected value of assigning any time to the skilled task is negative. Therefore, absent any information about an employee, she would only be assigned the unskilled task; this implies (for example) that $\beta_{i,1} = 0$.

Employees are risk neutral and maximize expected utility, which is equal to lifetime wages minus the cost of any effort they expend (we assume a discount rate of zero): $U_i = \sum_t w_{i,t} - \sum_t \beta_{i,t} e$. Maximizing utility rather than income implies that employees may strictly prefer a job that pays the same or even a lesser wage than an alternative if the difference in required effort is large enough (this will be important when we consider pooling equilibria). In other words, when comparing offers, employees consider the wage and assignment pair, not just the wage, and choose the package with the highest utility. We assume that firms are restricted from making negative utility offers.[17]

B. Timeline, wage formation, and learning

We model two periods of worker production. At the start of each period, a worker is made an offer by her incumbent firm. This offer consists of a wage and an assignment, where the

[17] Employees have an endogenous outside option due to the presence of a lateral hiring market for experienced workers; it binds when it provides utility that is greater than 0. By assuming employees are risk neutral, all wage contracts are formed in the spot market (i.e., there is no role for insurance via long-term contracts). The restriction on negative utility offers constrains employers from forcing employees out of the firm, which is only relevant only under equal pay mandates.

latter is an allocation to the skilled task, with the balance spent on the unskilled task. Each worker then receives an offer from one outside firm, also consisting of a wage and an assignment to the skilled task. The worker switches jobs if doing so strictly increases expected lifetime utility (i.e., there is an infinitesimal switching cost). Workers are randomly assigned to an "incumbent" firm in the first period. Since employees are ex-ante identical, the first period offer from the outside firm ensures Bertrand competition obtains for new entrants to the labor market; the outside firm and incumbent make equivalent offers, equal to the lifetime profits that they expect from a new hire conditional on expected lifetime production, wages, and mobility.[18]

Initially, neither firms nor employees know the employee's type. Following the standard assumption in the asymmetric learning literature, α_i is revealed to both the employee and to her incumbent firm after one period of work, but not to other firms in the market. Firms will use their knowledge of α_i to inform future wage offers and job assignments.

We assume that precise job assignments cannot be credibly disclosed to outside firms. However, we assume that there is a noisy signal of job assignment, defined as a "promotion". Outside firms can observe if an employee's assignment to the skilled task is above a threshold β'; employees above this threshold are observed to have been promoted, equivalent to having a different title. As is standard in the literature, firms strategically determine a promotion criterion that maximizes their profits knowing how the information will be used by outside firms.

We also include an additional signal: when wage disclosure is allowed, employees can make credible disclosure of the incumbent wage offer to the outside firm, from which it will make inferences about the employee's ability when making its offer. Incumbent firms know that their offers will be disclosed and consider the implications of salary disclosure when making wage and job assignments. When a salary history ban is enacted, such disclosure is prohibited, and firms respond accordingly.

Once the worker has selected a firm, she works, wages are paid, and production occurs.

C. Equilibrium wages, mobility, and salary disclosure

We define an equilibrium as a set of wages, job assignments, and mobility decisions that satisfy a Perfect Bayesian Nash equilibrium. Incumbent firms offer their employees a wage-

[18] Alternatively, firms could compete for new hires, who randomize between identical offers.

assignment pair of (w_t, β_t). We denote outside offers with a * superscript: outside firms offer prospective lateral hires a wage-assignment pair of (w_t^*, β_t^*).

In the first period, the same wage is offered to all new hires (since firms have no information about employee skill), and we have assumed that they can assign new hires none of the skilled task. Therefore, $(w_1, \beta_1) = (w_1^*, \beta_1^*) = (w_1, 0)$ for some $w_1 \geq 0$.

In the second period, incumbent firms know the skill of their employees, and will vary their wages and assignments accordingly: $(w_2, \beta_2) = (w_2(\alpha), \beta_2(\alpha))$. Under certain circumstances (such as under a salary history ban) these wage and assignment functions will have discontinuities at a specific skill threshold associated with a change in promotion status. In order to align the notation with that used for outside firms, when wages and assignments depend on promotion status we separately denote $(w_p(\alpha), \beta_p(\alpha))$ for promoted employees and $(w_{np}(\alpha), \beta_{np}(\alpha))$ for employees that are not promoted (the domains of the two sets of functions will be non-overlapping). Given that promotion status is only relevant in the second period, we suppress the time subscript.

Outside firms cannot observe α and the wage-assignment pairs they offer to prospective new hires cannot be conditioned on skill. Instead, they are conditioned on salary (if it can be disclosed) and on promotion status. When salary history is relevant, we write $(w_2^*, \beta_2^*) = (w_2^*(w_2), \beta_2^*(w_2))$. When promotion status is relevant, outside firms offer one wage-assignment pair to each promoted employee and another to each employee who is not promoted. We use the shorthand of (w_p^*, β_p^*) for promoted employees and (w_{np}^*, β_{np}^*) for employees who have not been promoted (under a salary history ban, each of these wages and assignments is a constant, whereas under an equal pay mandate they can be conditioned on the incumbent wage).

Finally, the mobility decisions are characterized by accepting the wage offered by the incumbent firm or accepting the offer from the outside firm.

In all of the rational (i.e., nondiscriminatory) equilibria that we compute, the outside offers are as attractive as possible, a principle we term *maximal outside offer*. The maximal outside offer is the wage-assignment pair that provides the greatest utility affordable for the outside firm conditional on the information available to it about the employee. [19] As such it

[19] We constrain the outside firm from making offers that would result in negative profits if accepted. This eliminates trivial equilibria, such as those with an extremely high outside offers for employees with a positive wage, forcing the incumbent to pool all employees at a wage of 0.

would generate zero profits if accepted. To see why this is necessary in equilibrium, consider what happens if the outside firm instead made an offer that would result in positive profits. The profit-maximizing strategy for the incumbent would be to match that utility, such that it just retains its employees. But then the optimal strategy for the outside firm is a slightly better offer (i.e., a higher wage and/or a lower assignment), which would be accepted by the employee and result in positive profits for the outside firm. Of course, this in turn changes the optimal strategy for the incumbent; the only steady state is at the maximal outside offer. Incumbents have an incentive to at least match this utility, to retain the profits associated with firm-specific human capital, such that there is no mobility in equilibrium. The outside firm does not deviate from this strategy, despite the lack of mobility, because lower (and potentially profitable) offers will not be accepted, and more attractive offers would be accepted but lead to negative profits. This principle is only violated in the presence of discrimination. An employee who receives a biased offer based on a misperception of her skill may be poached by an offer that is "discrimination aware", meaning it is grossed up to reflect the biased assessment; this outside offer may be below the maximal outside offer.

D. Baseline equilibrium: e = 0

To demonstrate the equilibrium concept and to establish that our framework encompasses the standard approach as a special case, we compute the equilibrium that obtains when effort is costless.

Proposition 1: With costless effort, the unique equilibrium is characterized by an inefficient promotion threshold $\alpha' > \beta'$ *and a promotion wage* $w_p^* > 0$ *such that:*

i) *Employees with* $\alpha < \alpha'$ *are not promoted; incumbents offer a second period wage-assignment pair of* $\left(w_{np}(\alpha), \beta_{np}(\alpha)\right) = (0, \min(\alpha, \beta'))$*, and outside firms offer* $(w_{np}^*, \beta_{np}^*) = (0,0)$*;*

ii) *Employees with* $\alpha \geq \alpha'$ *are promoted; incumbents offer a second period wage-assignment pair of* $\left(w_p(\alpha), \beta_p(\alpha)\right) = \left(w_p^*, \alpha\right)$ *and outside firms offer* (w_p^*, β_p^*)*, where* $\beta_p^* = \underset{\beta}{\operatorname{argmax}}\, E[Y_2^* \mid \alpha \geq \alpha']$*, the assignment that maximizes the production of employees with skill above the promotion threshold when employed by the outside firm.*

No workers change firms, and new entrants to the labor market receive positive wages that reflect the expected lifetime profits from new hires but are not assigned any of the skilled task.
Proof: See Appendix.

The equilibrium in *Proposition 1* has all the characteristics of the standard asymmetric learning models. First, the combination of asymmetric learning and the noisy signal of job assignment results in the classic Waldman result: incumbent firms strategically underpromote high skill employees to frustrate outside offers. Incumbents utilize an inefficiently high promotion threshold because the offers made by outside firms to employees who are not promoted are severely constrained. Recall that we assume that the unconditional expectation of assigning an employee to any amount of the skilled task is negative, and the employees who are not promoted have (weakly) less skill than the full population. Therefore, these employees cannot be assigned any of the skilled task by the outside firm, and the maximal outside offer is a wage of 0. Incumbent firms only need match the wage of 0 to retain these employees (when effort is costless, wage and utility are equal), incentivizing them to cap the assignment at β' for some employees with skill above that threshold. At a wage of 0, employees that are not promoted are generally paid below their marginal product in the second period and earn no return to skill.

The second feature that matches the standard results is that incumbents may promote the most highly skilled employees if the production gains from the increased assignment outweigh the extra wage that must accompany a promotion. In equilibrium, the threshold α' is the skill level at which the profits generated from promoting the employee (and paying the wage w_p^*) equal the profits from not promoting, capping the assignment at β', and paying a wage of 0. In the appendix, we derive the equations that determine the promotion threshold, and the conditions under which some employees are promoted. The structure of the outside offer is important in this calculation, as is the level of firm-specific human capital. The maximal outside offer is the full expected production that the outside firm can generate from employees knowing only that they have skill above α', which of course excludes any contribution from firm-specific human capital. The optimal assignment may be above α', depending on the distribution of skill, if the gains from better utilizing more skilled employees outweighs the losses from over-assigning employees with skill close to α'.

Incumbent firms must match this wage to retain their employees. Therefore, promoted employees do earn positive wages (and utility). However, they are also generally paid below

their marginal product, and wage (and utility) is equalized across all promoted employees regardless of their individual productivity.

Third, by virtue of paying wages below marginal product, incumbents earn positive second period profits. These are paid to employees in the first period, when workers earn wages in excess of their production, as is typical in models with asymmetric learning.

Finally, the equilibrium is inefficient because $\alpha' > \beta'$: there are some under-utilized employees that do not produce at their full potential. The inefficiency results from the noise in the signal. Either a fully informative signal or no signal at all is preferable from an efficiency standpoint, although these would have very different implications for the distribution of wages across employees. This inefficiency reduces aggregate productivity and thus wages.

Several policy implications of this baseline equilibrium are worth noting. First, salary disclosure is irrelevant. Wages simply reflect publicly available information (in this case, promotion status), and banning disclosure has no effect on the equilibrium. Second, equal pay mandates are also irrelevant. Within each title, workers are paid identical wages regardless of their skill; this is a standard result in the asymmetric learning literature. Without costly effort there is no intra-title wage variation. The extra production that more skilled employees generate flows directly to firm profits, and thus to the wages paid to new entrants in the labor market.

Finally, although those two policy interventions are irrelevant, there are policy options that would reduce or eliminate the inefficiency caused by strategic underpromotion, including banning disclosure of job history, reducing or eliminating competition for experienced employees, and banning mobility. Mandating a second period wage of 0 also eliminates the inefficiency. While some of these are infeasible on their face (e.g., banning disclosure of job history would invalidate every resume ever written), the point is that such interventions come at no cost, and in fact imply an aggregate gain.[20]

IV. Costly effort, wage disclosure

We now parameterize the model with costly effort, and allow employees to disclose the incumbent wage offer to the outside firm. Specifically, the outside firm observes w_2 before

[20] Many models in the literature consider channels other than in-period production through which such interventions could have deleterious effects, such as investments in human capital through education or job training. Our framework highlights that the additional potential effect on production.

making its offer of (w_2^*, β_2^*)). Incumbent firms know this and make their offers accordingly. We will see that the wage necessary to induce effort in the skilled task reveals each employee's capacity α to the outside firm, which drives up the wage incumbents must pay. Outside firms are willing to match incumbent offers and assign poached employees to the amount of the skilled task that corresponds to that wage.

Proposition 2: With costly effort and wage disclosure, the unique equilibrium is:

$$(w_1, \beta_1) = (w_1^*, \beta_1^*) = (s, 0)$$
$$(w_2(\alpha), \beta_2(\alpha)) = (\alpha * y, \alpha)$$
$$(w_2^*(w_2), \beta_2^*(w_2)) = (w_2, \frac{w_2}{y})$$

No workers change firms and $s = E[(1-\alpha) * s_0 + \alpha * y * s_1]$*, the expected second period production (and profits) linked to firm-specific human capital. The noisy signal of promotion is irrelevant. The equilibrium is efficient.*

Proof: a) *Existence*. The incumbent's second period wage offer reveals the employee's ability since it is a linear function of α. Outside firms cannot make strictly profitable deviations from choosing to match this wage ($w_2^*(w_2) = w_2$) and offering the corresponding assignment of $\beta_2^*(w_2) = \frac{w_2}{y}$. This is the maximal outside offer because it fully utilizes the employee and pays a wage equal to her production. Since $y > e$, utility increases with assignment so long as the wage equals the full production; therefore, the specified outside offer provides the highest utility the outside firm can deliver with nonnegative profits. Decreasing the utility of the offer (with a lower wage and/or a higher assignment) would not attract any lateral hires. Increasing the utility of the offer (via a higher wage and/or lower assignment) would lead to hiring, but negative profits, since the incumbent wage is equal to the maximum (outside) productivity of the employee.

Incumbents cannot make profitable deviations because of the structure of this maximal outside offer. For example, consider an incumbent that offered wages $w_2(\alpha) = \alpha * e$, the minimum necessary to induce effort, resulting in a utility of 0. Any employee with $\alpha > 0$ would accept the outside offer because it pays the same wage but requires less effort, and thus generates positive utility:

$$U^*{}_2(w_2) = w^*{}_2(w_2) - e * \beta^*{}_2(w_2) = w_2\left(1 - \frac{e}{y}\right) > 0 \qquad [3]$$

In fact, this outside offer increases the utility for any employee who is paid less than her full (outside) marginal product, because it pays the same wage but requires less effort:

$$U^*_2 - U_2 > 0 \rightarrow w_2 - \frac{w_2}{y} * e - (w_2 - \beta_2 * e) > 0 \rightarrow w_2 < \beta_2 * y \quad [4]$$

In any strategy that is profitable pre-attrition, including pooling and mixed strategies, incremental profits (above those linked to firm-specific human capital) must be generated by employees assigned an amount of the skilled task above $\frac{w_2}{y}$. But those employees accept the outside offer, and employees who are assigned below this level reject it. Therefore, after attrition, any such strategy reduces profits and incumbent firms have no incentive to deviate from paying wages equal to $\beta * y$. Further, because wages equal $\beta * y$ (the marginal the product for outside firms), the only profits that incumbents earn in the second period are due to firm-specific human capital. Firms maximize these profits by assigning employees their full capacity in the skilled task (because $y * s_1 > s_0$). These profits are reflected in the initial wage.

b) *Uniqueness*. See Appendix.

Uniqueness is proved by eliminating alternative equilibria. We eliminate equilibria with alternative one-to-one maps of ability to wage by the incumbent by demonstrating that the outside firm can infer ability and profitably poach employees from the incumbent. The elimination of potential pooling strategies is more complex. The key intuition is that constructing a pool that deters poaching requires the incumbent firm to inefficiently utilize some employees, vis-à-vis alternative assignments that are available to them. However, pool construction and assignments are not observable and the incumbent cannot commit to a pooling strategy that involves a less profitable use of some employees than an available alternative. Therefore, the incumbent deviates from pooled assignments that would deter poaching, invalidating equilibria with deterrent pools.

In equilibrium, then, wages fully reveal employee skill, and employees capture their full (outside) marginal product. Wage inequality is high, owing to the high returns to skill. Further, wages are not constant within each title; instead, there is significant intra-title wage variation that reflects the full range of abilities. Because salary history is fully informative, any noisy signal about employee ability is irrelevant. Finally, the equilibrium is efficient because each employee is fully utilized. Put differently, all the features in the baseline considered in *Proposition 1* that are linked to asymmetric learning and the noisy signal disappear, and instead the equilibrium is equivalent to one with symmetric learning.

A. Interpretation and key assumptions

The equilibrium in *Proposition 2* differs from the standard equilibria in two important ways. First, wages are informative, rather than reflections of publicly observable information. In fact, they are so informative that when they can be disclosed they are the ultimate signal. Outside firms must inquire about the incumbent wage offer, or else they will suffer from adverse selection and earn negative profits. By itself, this insight is new and our framework can explain why firms do in fact inquire about salary history. Second, wages vary within each title, reflecting the continuous nature of skill and assignments rather than the discrete promotion threshold.

Therefore, from a policy perspective, it is clear that neither wage disclosure nor equal pay mandates are irrelevant in our framework. Wage disclosure is central to wage formation, and suppressing the wage signal can have aggregate and distributional consequences. Similarly, equal pay mandates, which compress intra-title variation in wages, restrict the ability of firms to take advantage of within-title differences in employee ability, which could also have significant consequences. These conclusions distinguish our framework from the standard approach, and feature heavily in the policy analysis that comprises the remainder of this article. Therefore, we analyze how our various assumptions interact to drive these results.

We start by noting that most of our assumptions are purposefully chosen to closely track those in the standard model, and they serve the same function as they do in that framework. The basic framing with two tasks, one of which generates greater productivity, but only when performed by certain employees, is straight from the standard model, as is the learning process, through which the incumbent firm observes the skill of its employees. Implicit in both our framework and the standard model is the assumption that firms can commit to a wage and an assignment. This may be more important in our framework, given the cost of effort, but nonetheless it is not a stretch to assume that firms cannot sustainably subject their employees to a "rug pull" where they dump extra work on them unexpectedly. Like in the standard model, we assume that output is not observable and therefore not contractable. Many factors could contribute to unobservable output, including team-based work and tenuous connections between specific tasks and eventual profits. If output were observable, high skill employees in either framework would seek out contracts with output-specific wages. Of course, there are jobs where output is observable, and those often do have such contracts (e.g., commission-based sales jobs); neither our framework nor the standard one applies in those situations.

As in the standard model, firm-specific human capital gives firms an incentive to fully utilize their employees. In both our framework and the standard one, the absence of firm-specific human capital leads to trivial equilibria in which firms never assign any employees to the higher productivity task. Competition for employees ensures that wages equal outside marginal product; if that is also the incumbent productivity then the firm has no reason to promote employees and/or increase their assignments.[21] Of course, there are other modelling choices that can drive promotions, such as requiring a minimum number of high-level employees for production. We use firm-specific human capital for two related reasons. First, it is a well-documented, heavily studied, and realistic, aspect of employment. Second, it is the more common modelling choice (likely for these reasons), and we purposefully limit the deviations from the standard model, to better emphasize the importance of costly effort.

Therefore, we turn to the two ways we diverge from the standard assumptions: the continuous nature of assignments and costly effort. The continuous nature of task assignments is a realistic enhancement, reflecting the true nature of many jobs. For example, we can think of the higher value task as management of other employees. An employee assigned a small amount of management might have the same title as some employees on her team (e.g., "Associate" at a law firm, or "Vice President" at an investment bank). This employee could not credibly disclose her precise management responsibilities to outside firms. However, management responsibilities above a certain threshold come with a different title (e.g., "Partner" or "Managing Director") that can be disclosed. Importantly, this is meant to capture incremental responsibility within the same profession. We are not comparing manual labor with investment banking, but rather the career arc withing a given field.

Continuous assignments are necessary but not sufficient to generate intra-title wage variation. They are necessary because they allow firms to fully utilize their employees' skill. They are insufficient because, absent costly effort, firms can accomplish this without paying differential wages. This is clear from *Proposition 1,* where firms successfully assign each

[21] Specifically, the standard Waldman result and *Proposition 1* would both revert to a straight cap on employee assignments. In our model, assignments would be capped at the public promotion threshold; in Waldman, no employees would be assigned the higher task (i.e., no employee is promoted). With costly effort, the equilibrium in *Proposition 2* is no longer unique. While any assignment to the skilled task must still be accompanied by an increased wage, the assignments we detail are the maximum feasible for each employee in equilibrium.

employee up to her ability (with the exception of those employees who are underpromoted) but have one wage for each title.

It is also clear from *Proposition 1* that the continuous nature of assignments does not itself alter the character of the equilibrium vis-à-vis the standard Waldman underpromotion result. With costless effort the identical reasoning from Waldman applies to our framework: when outside firms can only observe title, they suffer from potential adverse selection. Outside offers that would attract employees are not sustainable in equilibrium because the incumbent, who knows the skill of its employees, would retain those that are profitable at the higher wage. The outside firm only makes unprofitable hires. Therefore, the equilibrium outside offer is tied to the (outside) marginal product of the lowest skill of promoted employees. This incentivizes the incumbent to underpromote; the firm-specific human capital of the least skilled employee who is promoted must be large enough to justify her higher wage. In fact, the original Waldman model could be reinterpreted as having continuous assignments; higher skill employees in that model generate more production, even when assigned the same task. We simply link that increased production to a greater assignment to the skilled task.

The continuous nature of assignments requires us to refine the production function. We must define production for any assignment to the skilled task, whilst retaining the notion that over-assignment vis-à-vis an employee's skill is inefficient. Equation [1] is a natural extension of the standard approach; production increases with assignment up to the employee's inherent ability, and then declines with further assignment to the skilled task. Using the example of management, a bad manager (i.e., an employee over-assigned to the management task) reduces productivity by reducing moral, or inefficiently utilizing the resources or abilities of her team. This justifies the assumed loss from overpromotion.

Like with the standard model, we assume that the unconditional distribution of skill is such that assignment to the skilled task is inefficient in the first period. In our framework, this assumption is somewhat stronger, since it implies that even an incremental assignment would be damaging, although we could alternatively assume that there is some baseline assignment to the skilled task that can be given to new entrants, and normalize assignments and wages accordingly. This would have no effect on our conclusions. That said, our assumption is not unrealistic. Referring again to the management example, firms do not hire new college graduates into management jobs; instead, they learn about their skills over time and incrementally assign more

responsibility as appropriate. Analytically, this assumption requires a left-skewed distribution of skill. In particular, there must be a sufficient mass of employees with no ability to execute the skilled task. The most important implication of this distribution is that outside firms cannot earn nonnegative profits when offering positive utility to any set of employees that includes those with no skill. This constraints outside firms when they cannot distinguish between employees, as in *Proposition 1*.

Costly effort is necessary for both of the policy-relevant characteristics of the equilibrium in *Proposition 2*. It forces incumbents to pay their productive employees a non-zero wage, with higher productivity garnering increasingly higher wages. Any effort to exploit the continuous nature of skill and assignments will require intra-title variation in wages. By extension, wages are informative, and outside firms can infer skill if they can observe the wage. Masking strategies based on pooling employees, which are typical in the standard approach, are untenable, because they require that firms inefficiently utilize their employees, which leads to lower profits.

We believe this is also a realistic assumption. Some variation in wages must be compensation for additional effort, rather than purely the result of competition driven by observable signals. As mentioned in the Literature Review, if this were not the case, then the policy prescription is simple: mandating a flat wage across all employees in the economy will eliminate any bias or unfairness in the labor market without creating any distortions. We are not suggesting that competition does not factor into wage formation. Instead, we assume that wages are necessary for production, and then allow the market dynamics to play out; notably, competition interacts with the wage signal to generate the equilibrium in *Proposition 2*. In fact, due to competitive pressure any non-zero cost of effort is sufficient to generate our results. Although costly effort is realistic, our modelling is necessarily simplistic. For example, we assume that effort is (at least partially) observable. There are likely examples where that is not true, which would require more sophisticated contracting. Similarly, we assume that the cost of effort is linear and uniform across employees. If instead we assumed that the cost of effort varied across employees, it is possible that firms could exploit that variation (if they could observe it), reducing the value of the wage signal. This could lead to both wages and titles being informative; notably, the latter appears on every resume, and thus must have some signaling value. We discuss these potential extensions in Section VII.

Costly effort also requires that we alter the timing of the job offers. Although there are counterexamples, most asymmetric learning models assume that the outside firm makes an offer based on the observable employee characteristics, and then the incumbent responds. This is not feasible with costly effort because the incumbent cannot simply give the employee an assignment; any assignment must be paired with a wage that is at least high enough to induce the required effort. So long as that minimum threshold is met, the specifics of the ensuing sequence are irrelevant in our framework. The outside firm can make an offer, and then the incumbent can counter, or the outside firm can make an offer and the employee can choose between the two. These are equivalent because the incumbent can infer the optimal outside offer and would respond strategically. Any choice is a pure modelling artifice; ours is at least as realistic as the standard one. Reality evolves in continuous time, and at all times employees have both a title and a wage, which they can disclose (when allowed) to outside firms when seeking alternative offers (i.e., we know of no instance where employees are offered a title, and then later a wage based on whatever outside offer they managed to secure).

V. Salary history bans and equal pay mandates

We now analyze how various interventions in the labor market targeting wage disclosure and wage formation affect task assignments, wages, and efficiency.

A. Salary history bans

With a salary history ban, the outside firm cannot use the second period wage offer to infer worker type. It must instead rely solely on the noisy promotion signal. Therefore, incumbent firms can mask the assignments of their employees, up to that threshold. The result is similar to *Proposition 1*: incumbent firms under-promote employees in order to reduce their outside option and allow them to pay the minimum wage possible.

Proposition 3: With a salary history ban, the unique equilibrium has an inefficient promotion threshold $\alpha' > \beta'$ *and a promotion wage* $w_p^* > 0$ *such that:*

i) Employees with $\alpha < \alpha'$ *are not promoted; incumbents offer a second period wage-assignment pair of* $\left(w_{np}(\alpha), \beta_{np}(\alpha)\right) = (e * \min(\alpha, \beta'), \min(\alpha, \beta'))$, *and outside firms offer* $(w_{np}^*, \beta_{np}^*) = (0,0)$;

ii) *Employees with $\alpha \geq \alpha'$ are promoted; incumbents offer a second period wage-assignment pair of $\left(w_p(\alpha), \beta_p(\alpha)\right) = \left((w_p^* - e * \beta_p^* + e * \alpha\ ,\ \alpha\right)$ and outside firms offer (w_p^*, β_p^*), where $\beta_p^* = \underset{\beta}{\operatorname{argmax}}\, E[Y_{i,2}^* - \beta * e \mid \alpha \geq \alpha']$, the assignment that maximizes the utility of employees with skill above the promotion threshold when employed by the outside firm.*

No workers change firms, and new entrants to the labor market receive positive wages that reflect the expected lifetime profits from new hires but are not assigned any of the skilled task.

Proof: See appendix.

The proof follows that of *Proposition 1* closely, as the equilibria are very similar. The main difference is the need to compensate for the disutility of effort. For example, the flat incumbent wage of 0 from *Proposition 1* paid to employees that are not promoted is replaced with the minimum wage necessary to ensure utility of 0: employees who are not promoted are paid wages that just compensate for their effort. Therefore, wage (but not utility) of these employees increases with skill.

Similarly, because of the cost of effort, outside firms no longer compete for workers purely on the basis of the wage; they offer utility, equal to wage less the cost of effort. The maximal outside offer is the wage and assignment pair with the greatest utility outside firms can offer to promoted employees and still earn nonnegative profits, based on the promotion threshold (rather than simply the maximum wage, as in *Proposition 1*). Incumbent firms must match that utility to retain their employees. Therefore, although the wages of promoted employees vary with their precise assignment, they all earn equivalent utility, equal to that offered by the outside firm. This is above the utility of employees who are not promoted, as in *Proposition 1*.

Firms earn high profits in the second period, as they generally pay wages below marginal product, although the profits are lower than was the case in *Proposition 1* because second period wages must compensate for effort. The profits are reflected in first period wages.

We conclude that jamming the wage signal introduces noise into the learning process. Because they are used to motivate effort, wages are the ultimate signal; without them, outside firms are forced to rely on promotion signals. Incumbent firms strategically adjust the promotion threshold to maximize their profits; this limits the ability of outside firms to make lateral hires.

Banning wage disclosure has both aggregate and distributional effects on welfare. It reduces aggregate welfare because it introduces an inefficient promotion threshold. Employees

with skill between β' and α' are underutilized, and the associated loss of production is reflected in lower aggregate wages.

Corollary 1: A salary history ban reduces aggregate wages, utility, and efficiency.

From a distributional standpoint, banning disclosure improves the welfare of employees with the lowest skill within each title band. For example, employees with skill of 0 receive the same second period wage under a ban as they do when wage disclosure is allowed, but a ban increases their initial wage. Similarly, a ban increases the welfare for employees with skill just above α'. Under a ban, their second period wage is at least as high as it is under disclosure, and they too receive the elevated initial wage.

Given *Corollary 1*, a ban must therefore reduce the welfare of one or both sets of employees at the upper end of the skill distribution for their title (depending on the parameter values and distribution of skill). Banning wage disclosure increases the bargaining power that firms have over these employees because they cannot distinguish themselves from lower skilled employees with the same title. This allows the incumbent to pay them wages below their potential (outside) marginal product, with the benefits dispersed across employees in the form of higher entry wages. For example, consider employees with skill between β' and α'. They receive a second period wage close to $y * \alpha'$ under disclosure but a much lower wage of $e * \beta'$ under a ban. A parameter set such that the average employee is not promoted (i.e., $\beta' < E(\alpha)$) is sufficient to ensure that this decline in the second period wage is larger than the increase in the entry wage, implying that a ban reduces the welfare of these employees.

Corollary 2: A salary history ban increases the welfare of employees with skill at the low end of each title, but reduces the welfare of employees with skill at the upper end of one or both titles.

B. Equal pay mandates

We model an equal pay mandate as a requirement that firms pay all employees with the same title the same wage. In other words, firms must pay one wage (w_{np}) to all employees who are assigned below the observable promotion threshold β', and another to those that are "promoted" past that threshold (w_p). This modelling is intended to capture the intent of rules designed to compress the distribution of wages across employees, such as the Equal Pay Act, and

there is reason to believe that this is their ultimate effect.[22] Further, many jobs already have lockstep compensation based on title, such as roles covered by public sector unions, and this section applies to those roles as well.

The constraint imposed by an equal pay mandate is very different from the constraint imposed by a salary history ban. A ban reduces the information available to an outside firm, and thus reduces the outside option of the more skilled employees, with the benefits accruing to the incumbent firm (in the second period) and to less skilled employees (in the first period). In equilibrium, costly effort still leads to within-title salary differentiation, as the incumbent differentiates assignments within titles and so pays each employee a wage that just covers her effort. In contrast, an equal pay mandate binds the incumbent firm; by limiting the range of possible wages it can pay, it limits the range of assignments it can utilize.

If there is an employee with skill sufficiently high to promote and pay the promotion wage w_p, then any employee with skill above that threshold must also be promoted, since they can be given at least the same assignment. Therefore, the equilibrium under an equal pay mandate takes the familiar form: a promotion threshold α'. We will demonstrate that any equilibrium under an equal pay mandate must have several other characteristics of economic interest (in the appendix we discuss the conditions under which such equilibria exist).

The incumbent will choose wages w_p and w_{np}, the promotion threshold α', and assignment functions $\beta_p(\alpha)$ and $\beta_{np}(\alpha)$, that maximize its profits given the outside offers. Outside firms make offers that are conditioned on both promotion status and wages: $\left(w_p^*, \beta_p^*\right) = \left(w_p^*(w_p), \beta_p^*(w_p)\right)$ and $\left(w_{np}^*, \beta_{np}^*\right) = \left(w_{np}^*(w_{np}), \beta_{np}^*(w_{np})\right)$. There are two types of poaching

[22] This analysis applies to equality-based initiatives that make it costly to pay differential wages based on unverifiable characteristics such as skill or productivity. For instance, a recent 9th Circuit opinion interpreted the Equal Pay Act to required employers to prove that pay differentials are based upon valid factors such as productivity as opposed to sex, and disallowed the use of salary histories as such a valid factor; the minority dissent noted that this rule would "demand a lockstep pay system." Rizo v. Yovino, Case No. 16-15372, 9th Circuit, February 27, 2020. Pay equity advocacy groups also recommend practices such as lockstep compensation, limiting managerial discretion in setting pay, standardizing compensation, and creating a culture of equality. See California Commission on the Status of Women and Girls, What Can I do to Promote a Culture of Pay Equity?, https://women.ca.gov/californiapayequity/employers-resources/what-can-i-do-to-promote-a-culture-of-pay-equity/. Other policies can help enforce the intent of these decisions. For example, requiring publication of wages in job postings pre-commits the employer to wages within the band and limits discretion based on productivity. Pay transparency laws make it easier to discover and allege the impropriety of pay disparities, and both new and existing laws make such disparities costly.

strategies that an outside firm may attempt. The first is the familiar one: an employee assigned more than $\frac{w_2}{y}$ by the incumbent can be poached by the outside firm. The second is a consequence of the limitation on assignments: if the incumbent "bunches" high-skill employees at an assignment of $\frac{w_2}{y}$, it is possible that the outside firm can profitably assign this group of employees an assignment of $\frac{w_2}{y} + \varepsilon$, if the losses from any employees that would be over-assigned (such as those with $\alpha = \frac{w_2}{y}$) are outweighed by the gains from better utilizing those employees with skill above $\frac{w_2}{y}$. Equilibrium requires that neither type of firm can increase its profits by varying these offers, conditional on the strategy being deployed by the other type of firm. In particular, the incumbent strategy must be robust to both poaching strategies.

We derive a series of constraints on the equilibria that together characterize the effect of an equal pay mandate. The first set of constraints involves mobility; unsurprisingly, there cannot be mobility in equilibrium.

Lemma 1: There can be no mobility amongst the set of employees who is not promoted.

Proof: Assume we have an equilibrium with mobility amongst this set of employees. The outside firm cannot profitably offer positive utility to any set of employees that includes those with skill of 0. Therefore, only employees with skill above some threshold α^* accept the offer. If $\alpha^* > 0$, then it must be the case that employees with skill equal to α^* have productivity above the equilibrium wage w_{np}, because these employees have the maximum productivity amongst the remaining employees, and the incumbent must earn nonnegative profits. This gives the incumbent the incentive to deviate; it would assign employees who would accept the outside offer to α^*, which retains them and thus increases its profits. Alternatively, if $\alpha^* = 0$ (such that all employees with positive skill are poached), the incumbent can deviate to a strategy where it pays a wage of 0, assigns none of the skilled task, and earns profits of s_0 per employee. Therefore, this could not be an equilibrium. QED

Knowing that there cannot be mobility for employees who are not promoted implies a lower bound of the profitability of those employees.

Lemma 2: The minimum average profitability of unpromoted employees is s_0.

Proof: The incumbent can at a minimum assign none of the skilled task and pay a wage of 0 to employees who are not promoted. Outside firms cannot poach any employees under this

strategy, and it generates profits of s_0 per employee. Therefore, this is the minimum average profitability of those employees in equilibrium. QED

The fact that the employees who are not promoted generate a profit in equilibrium constrains the outcomes for the employees who are promoted.

Lemma 3: There can be no mobility for employees who are promoted.

Proof: Any employee who is poached by an outside firm generates no second period profits for the incumbent, by definition. Therefore, the incumbent would reassign any employee subject to poaching such that she was not promoted, which generates positive profits. QED

The lack of mobility in equilibrium implies that the utility of employees with the greatest assignment to the skilled task within each title is equal to the utility of the outside offer (with mobility it is possible that they had utility below the outside offer and accepted it). If the employee with the greatest assignment had utility above the outside offer the incumbent could raise its profits by increasing her assignment and/or reducing her wage.

With that insight in hand, we can further constrain the form any equilibrium must take.

Proposition 4: Under an equal pay mandate, the equilibrium wages and assignments of employees who are not promoted are bound above:

i) $w_{np} \in [0, s_0]$

ii) $\beta_{np}(\alpha) \in [0, s_0/y]$

Proof: See Appendix.

Consider first the upper bound on the wage w_{np}. If it is above s_0 then employees with $\alpha = 0$ are unprofitable for the incumbent. If the outside firm made an offer with positive utility, the incumbent could assign those employees enough of the skilled task that they accept the outside offer, which increases its profits. However, the outside firm cannot profitably hire those employees, and would deviate to a lower utility offer. Alternatively, if the outside firm offered a utility of 0, then the incumbent could utilize a maximum assignment that generated a utility of 0 (equal to w_{np}/e). But then the outside firm could deviate to a positive utility offer that would be accepted by those more skilled employees. Therefore, the wage cannot exceed s_0. The upper bound on assignments detailed in Part (ii) follows directly from the maximal outside offer. As in the other equilibria we consider, any employee assigned more than w_{np}/y of the skilled task can be poached by the outside firm, which can make an offer of $(w_{np}, \frac{w_{np}}{y} + \varepsilon)$ that is utility-

improving for those employees at some $\varepsilon > 0$ and generates profits of $\varepsilon * y$ for the outside firm. Given the result in Part (i), this caps the assignment at s_0/y.

This is an important constraint imposed by the equal pay mandate. By virtue of needing to pay the lowest skilled employees the same wages as their higher skilled (but not promoted) colleagues, wages are capped, which translates into an upper bound on possible assignments. This introduces a new form of efficiency loss: within-title underutilization. Some employees who are not promoted cannot be fully utilized because an assignment equal to their skill would reduce their utility and expose them to poaching. Firms cannot in general afford to pay wages sufficiently high to protect all employees when fully utilized because that generates losses from overpaying the lower skilled employees.

Equal pay mandates impose a similar constraint on employees who are promoted.

Proposition 5: Under an equal pay mandate, the equilibrium assignments made by incumbents to employees who are promoted are bound above and below:

i) $\beta_p(\alpha) \leq max\ (\frac{w_p}{y}, \beta')$

ii) $\beta_p(\alpha) \geq max[\beta', (w_p + w_{np} * K)/C]$

where $K = s_1 + s_0/y$ *and* $C = y * (1 + s_1) - s_0$

Proof: See Appendix.

The upper bound on the assignment is derived in the same way as the upper bound in *Proposition 4*. In general, at any assignment above w_p/y the employee can be profitably poached by the outside firm. The only exception is if the max assignment is equal to β'. As we will see below, it is possible that some employees are over-assigned; if that occurs and the maximum assignment is equal to β', then the outside firm cannot offer a wage as high as $y * \beta'$ to the promoted employees, and thus the incumbent need only match the lower wage.

There are two components to the lower bound. First, by definition, promotion requires that the assignment must be greater than β'. Second, the profits generated by the least skilled promoted employee must at least equal the profit that she would generate if not promoted (and thus paid w_{np}) and given the maximum assignment outlined in *Proposition 4*. The lower bound on $\beta_p(\alpha)$ is an assignment level that generates that profit, given the equilibrium wages.

An important implication of *Proposition 5* is that in general (i.e., except when $w_p = y$), the most skilled promoted employees are underutilized. This is analogous to the underutilization

of the most skilled employees who are not promoted. The wage necessary to retain the most skilled employees when assigned up to their capacity is not affordable because it reduces the profits generated from the lower skilled but still promoted employees.

Together, these *Propositions* lead to several general conclusions regarding the effect of equal pay mandates. The first follows directly from the underutilization within each title; overall productivity and thus average wages must fall under an equal pay mandate.

Corollary 3: Equal pay mandates reduce aggregate wages, utility, and efficiency due to within-title underutilization.

Second, just as with a salary history ban, there are distributional consequences; the efficiency losses associated with an equal pay mandate are not uniformly distributed across employees. Unlike a ban, utility is reduced because more skilled employees are not fully utilized, and thus their wage (and utility) cannot reflect their full potential production. Any underutilized employee has a wage that reflects her realized production, excepting that linked to firm-specific human capital; this is ensured by the maximal outside offer. Their utility loss is driven by the fact that utility is highest when both fully assigned and fully paid. In other words, an equal pay mandate does not reduce the outside option available to an employee conditional on her assignment (i.e., unlike under a salary history ban, she is still "fully paid"). It restricts the assignment, which in turn restricts her ability to fully realize her potential. As a result, the losses to more skilled employees are not reflected in higher salaries to new entrants to the labor market. In fact, starting salaries fall, because employees who are not underutilized are overpaid. They receive wages above their (outside) marginal product because they are paid just like more skilled employees with the same title.

Therefore, an equal pay mandate reduces the welfare of more skilled employees within each title. It does not necessarily increase the welfare of less skilled employees; they may receive a higher second period wage but that is balanced by a lower entry wage.

Corollary 4: An equal pay mandate reduces the welfare of the most skilled employees within each title.

A further implication of underutilization of employees is that equal pay mandates potentially create incentives for employers to force out their least skilled workers. Recall that we assume that firms cannot offer negative utility wage-assignment pairs to their employees. This assumption is not binding in the other settings we consider; firm-specific human capital implies

that it is worthwhile employing even those workers with $\alpha = 0$. That potentially changes under an equal pay mandate. While the least skilled workers still generate production from firm specific human capital, their presence in the firm limits the equilibrium assignments (and thus production) of the more skilled workers. This drag could outweigh their production, such that the net effect of the lower skilled employees on profits is negative. One possible effect of equal pay mandates is that firms will seek to "weed out" productive but low skilled workers, possibly through the strategic use of unappealing assignments.

Another possible implication is that firms will have an incentive to create more titles. With more titles, the incumbent creates more ranges of feasible assignments, and thus more fully realizes the profits from firm-specific human capital. This distinguishes an equal pay mandate from a salary history ban, where a relative paucity of titles better obscures the true ability of the incumbent employees and thus better frustrates outside offers.

A final implication of *Proposition 5* is that equal pay mandates raise the prospect of over-assignment. The possibility arises because of the significant wedge between the assignments that can be given to employees at each title. The range of possible assignments for employees who are not promoted is quite narrow, and close to 0 (particularly when s_0 is small). This implies that employees with skill just below β' are severely underutilized if not promoted. Promoting these employees allows for much greater utilization of their skill, and it is possible doing so increases profits, if the gains from more fully utilizing them, net of the losses linked to over-assignment, outweigh the extra wage that accompanies promotion.

An example helps illustrate the point. Assume that the promotion threshold β' is high at 0.95 but that $s_0 \approx 0$. With such a small s_0, the maximum assignment for employees who are not promoted ($\beta_{np}(\alpha) = s_0/y$) is also close to 0. Therefore, all employees who are not promoted generate (nearly) no profits. In this sense, an employee with a skill level of, say, $\alpha = 0.9$, is severely underutilized when not promoted. However, this employee can generate positive profits if promoted, so long as the wage w_p and the loss from over-assignment L are small relative to the production associated with such a high skill level:

$$0.9 * y * (1 + s_1) + 0.05 * L - w_p > 0 \rightarrow$$
$$w_p - 0.05 * L < 0.9 * y * (1 + s_1) \qquad [5]$$

It is only feasible to over-assign employees with skill below the public promotion threshold β'; these employees are the only ones for which over-assignment is ever necessary. An

employee with skill above β' can be fully utilized without being over-assigned. In fact, increasing the assignment of an employee with skill above β' has the counterproductive effect of reducing her utility (and thus making her easier to poach) and reducing firm profits. As a result, over-assignment only occurs when the minimum assignment for promoted employees is pegged at the public threshold β'.

Corollary 5: In general, the threshold α' is inefficient. When $\beta' < (w_p + w_{np} * K)/C$, with C and K defined as in Proposition 5, firms under-promote in the classic Waldman sense. When $\beta' > (w_p + w_{np} * K)/C$, some employees with skill below β' are over-assigned.

Although this analysis describes the key features of an equilibrium under an equal pay mandate, for completeness we outline an example equilibrium in the appendix, along with the conditions under which it obtains. This requires detailing the off-equilibrium outside offers that keep the incumbent from deviating to alternative wages.

VI. Discrimination

Of course, the labor market interventions we analyze are intended to mitigate the effects of discrimination. If they do so, then the resulting inefficiencies may be acceptable or even desirable tradeoffs in exchange for more equitable labor markets. To better understand how discrimination interacts with the wage and assignment dynamics above and to illustrate the flexibility and utility of our general framework, we introduce a simple and tractable model of discrimination. We assume that some employees are subjected to a "glass ceiling", defined as an upper bound on their perceived skill.

Formally, we assume one class of employees, labelled "men", are always correctly perceived by their employers; all the equilibria computed above apply. Another class of employees, labelled "women", are subject to misperception: a proportion τ are observed to have skill $\tilde{\alpha} = \min(\alpha, \beta')$. This glass ceiling of β' implies that these women are not promoted. Setting the glass ceiling at the promotion threshold has the intuitive feature that women are capped in title, rather than assignment (although we discuss the implications of other choices for the glass ceiling below). The remaining (1-τ) women are correctly perceived by their employers. We assume this discrimination is uniformly distributed across firms rather than concentrated in specific employers, which would make it more difficult to identify and root out.

This glass ceiling is a form of unconscious bias, whereby some managers systematically but unknowingly underutilize their female employees. This could arise through awarding more stretch assignments to men or underestimating the management potential of women. We assume that these managers do not recognize what they are doing, nor do they recognize the existence of discrimination in the broader market. From a modelling standpoint, this implies that the same proportion τ of outside offers are made without an adjustment for the possible presence of discrimination. These offers can involve high wages and high assignments, as justified by the incumbent wage or title, but they will not "gross up" wages or assignments to reflect the possible presence of discrimination.

We assume that the other (1- τ) proportion of offers are made by managers who do acknowledge the existence of discrimination and strategically adjust any outside offers they make accordingly. Because employees receive only one outside offer, this implies that women subjected to the glass ceiling have a (1- τ) probability of receiving such a "rational" offer. It is important to note that discrimination is irrational. While it is obviously harmful to women, it also reduces firm profits, because they fail to fully utilize a subset of employees (and fail to learn about their error). We do not find this troubling as a modelling device, as discrimination itself is inherently irrational, and the proponents of bans and mandates clearly believe that it persists despite its pernicious effects. However, the irrationality effects the nature of the outside offers, which will no longer necessarily satisfy the maximal outside offer principal. Specifically, a woman subject to the glass ceiling can be profitably poached; if the offer is "discrimination aware", it can profitably induce mobility by grossing up her assignment and wage. In other words, discrimination can result in mobility, which otherwise does not occur in our model because unbiased firms seek to retain the production associated with firm-specific human capital. We label these outside offers as "rational".

Below, we assess the welfare of women under the three legal regimes we consider. We consider both the welfare of women subjected to the glass ceiling (i.e., those with high skill who are discriminated against) and the overall gender pay gap.

A. Wage Disclosure

Women who work for discriminatory managers are assigned $\tilde{\alpha}$ and paid $y * \tilde{\alpha}$. This implies that all women with skill above β' receive $w_2 = y * \beta'$. Given the continuous

distribution of skill, a rational outside offer will strategically adjust for the effect of discrimination.[23] These women can potentially be profitably poached; they will accept an offer that delivers even incrementally more utility than the incumbent.[24] Therefore, the rational outside offer has a discontinuity at $w_2 = y * \beta'$. It will involve the assignment that maximizes profits, given the range of skill for women paid that wage and the incremental wage necessary to ensure effort at the new assignment (whilst maintaining the utility of the incumbent offer):

$$\beta_2^*(y * \beta') = \underset{\beta}{\operatorname{argmax}}\, E[Y_{i,2}^* - \beta' * y + (\beta - \beta') * e \mid \alpha \geq \beta'] \qquad [6]$$

where $Y_{i,2}^*$ is the production function defined in [1]. Implicit in [6] is a wage offer designed to match the utility of the incumbent offer (plus an infinitesimal increment which we suppress), given the additional effort required:

$$w_2^*(y * \beta') = \beta' * y + (\beta^* - \beta') * e \qquad [7]$$

Proposition 6: With wage disclosure and discrimination, the equilibrium follows that in Proposition 2, except that (1-τ) of the women with wages of $y * \beta'$ are offered a wage and assignment pair defined by [7] and [6]. If $\beta_2^*(y * \beta') > \beta'$ women accept this offer.

If $\beta_2^*(y * \beta') > \beta'$ then market forces reduce the impact of discrimination. Some women subjected to the glass ceiling are profitably hired by outside firms, which increases their wage. The profits are driven by the departure from the maximal outside offer. Mobility also punishes the incumbent firm, which loses an employee and thus the expected profits that motivated the initial wage.

However, market forces only partially mitigate discrimination. First, not all women subject to the glass ceiling will receive a rational offer. Second, "anchoring" occurs, whereby the effect of prior discrimination affects women's future outcomes. In wage terms, the artificially low incumbent wage is reflected in the outside offer; women are paid their incumbent wage plus an increment that is just high enough to induce acceptance of the new assignment and wage pair, rather than their full outside marginal product, which is what other employees are paid.

[23] With a discrete distribution of skill, the probability of being subject to the glass ceiling conditional on a wage of $y * \beta'$ would be elevated but less than unity. This has no qualitative effect on our analysis.
[24] We could alternatively assume rational offers are magnanimous, meaning they are the maximal outside offer, with no qualitative change to the conclusions. The potential for over-assignment of women with skill just above the glass ceiling is the source of anchoring in wage and utility terms.

Formally, high skill women subject to the glass ceiling have a post-mobility average wage of $y * \beta' + (1 - \tau) * e * (\beta^* - \beta')$. The last term is the gross-up for the additional effort associated with the higher assignment for those who receive a rational offer. This remains below the male wage of $y * \alpha$ for all women except those with skill between $\beta'$and β^*. Mobility also does not result in greater utility; women who change jobs are just barely better off.

In assignment terms, high skill women who change firms are still subject to a glass ceiling (equal to β^*), albeit one that is above the promotion threshold. Mobility simply replaces one ceiling with another. Finally, on average women are paid less for the same work than men. Specifically, those who accept outside offers are paid below similarly assigned men.

Putting these pieces together, we demonstrate that in aggregate a gender pay gap still exists, even after mobility.

Corollary 6: With wage disclosure and discrimination, there is an aggregate gender pay gap, anchoring, and some women are paid than men less for the same work.

Proof: The aggregate wage for women who accept outside offers is below their aggregate productivity (or the outside firm would not be increasing its profits). But the aggregate productivity of those women is below the outside marginal product of similarly skilled men, because some women are underutilized and others are over-assigned. Male wages fully reflect their outside marginal product; thus, aggregate women's wages are below aggregate men's wages. High skilled women subject to the glass ceiling have lower lifetime wages than similarly skilled men (and women not subject to it). Women who accept an outside offer are paid less than similarly assigned men at the same firm. QED

We conclude that under disclosure, mobility reduces but does not eliminate the effect of discrimination on the welfare of female employees. It only partially reduces the gender pay gap, and even after mobility, anchoring ensures that early-career discrimination has lifetime effects on women's wages and assignments, and on average women are paid less for the same work. Taken together, these features of the labor market closely match those cited by proponents of bans and other interventions designed to combat discrimination.

Finally, the same qualitative conclusions apply with any level of the glass ceiling.

B. Salary History Ban

Under a ban, women subject to the glass ceiling are not promoted. They receive a second period assignment of β' paired with a wage of $e * \beta'$. Outside firms cannot distinguish between these women and low-skilled women (at either type of firm). Therefore, $(w_2^*, \beta_2^*) = (0,0)$, as the population of unpromoted women has less skill on average than the overall population, and there is no mobility.

Proposition 7: With a salary history ban and discrimination, high skill women subject to the glass ceiling are trapped; they are pooled with low skill women and outside firms cannot offer them positive utility. The equilibrium follows that in Proposition 3, except that these high skill women receive a wage assignment pair of $(e * \beta', \beta')$.

A ban has two deleterious effects on women subject to the glass ceiling. First, it traps them; they can no longer change employers and increase their wage. Second, the glass ceiling implies that their incumbent firm treats them as if they are just below the promotion threshold. Per *Corollary 2*, this is the worst possible category from a distributional standpoint, because the imposition of a ban decreases the second period wage of these employees by more than it increases their initial wage.

However, this does not mean that the imposition of a ban increases the gender pay gap. While it reduces the welfare of the specific women who are subject to discrimination, a ban also generally reduces the return to skill, and with it the wages of high skilled men. These two forces work in opposite directions, such that the gender pay gap can increase or decrease. Of course, any aggregate reduction in the pay gap comes at the cost of inefficiency (and thus lower aggregate wages) and reduced wages and utility for the victims of discrimination. That others might suffer more from a ban seems like poor succor for those employees.

Corollary 7: A salary history ban reduces the welfare of women subject to discrimination, but may reduce the gender pay gap.

Intuitively, a ban ties wages and utility to titles. So long as discrimination results in certain classes of employees being underpromoted, a ban reduces the welfare of the victims of discrimination. They are particularly affected because their assignments are likely clustered just below the promotion threshold, which is the range of skill that is most negatively affected by a ban. It is important to note that we would draw a different conclusion if discrimination was not closely linked to title. For example, if the glass ceiling was close to 0, then women subjected to it are perceived to have very low skill. Low skill employees benefit from a ban due because it

raises the entry wage. In contrast, disclosure results in these employees earning low wages and utility, even after the potential gross-up by an outside firm, due to anchoring. A ban could be beneficial in this case. However, we prefer our baseline approach, as it is more intuitive that discrimination effect promotions as well as assignments.

C. Equal pay mandate

First, we note that, unlike in the other settings we consider, the equilibrium under an equal pay mandate depends on the distribution of skill across workers in the firm, rather than on the skill of the individual employee. In particular, the equilibrium wage for promoted employees will be weakly lower with discrimination than without; firms will observe fewer high skill employees, and thus may determine that a wage necessary to facilitate high assignments is not worth paying. In other words, discrimination shifts downwards (observed) aggregate skill, which may result in lower wages and assignments for highly skilled men and women not subject to the glass ceiling, leading to greater underutilization of high skill employees. The assumption that discrimination is distributed across firms implies that this effect is symmetric. Nonetheless, one effect of an equal pay mandate in the presence of discrimination is weakly lower welfare for highly skilled men and highly skilled women not subject to the glass ceiling.

The particular effect of the equal pay mandate on women who are subject to the glass ceiling will depend upon whether the promotion threshold α' is above or below β' (recall from *Corollary 5* that the promotion threshold α' may be inefficiently high or low due to the constraints of the equal pay mandate). Where $\alpha' > \beta'$ (inefficient underpromotion), no women subject to the glass ceiling firms will be promoted, their perceived ability is truncated at β'. If $\alpha' \leq \beta'$ (inefficient overpromotion), then all women with actual skill in excess of the promotion threshold will be promoted.

In the $\alpha' > \beta'$ case, women subject to the glass ceiling are pooled with the low-skilled men and women at the non-promoted wage of $w_{np} \in [0, s_0]$ as described in *Proposition 4*. While some firms recognize that some of these unpromoted women were discriminated against, wage pooling prevents learning about individual skill levels and, as assumed, the unconditional distribution of skill is such that positive skill assignments are unprofitable. There is no mobility for these workers (as in *Lemma 1*), and these women are trapped. In this case, the equal pay mandate reduces the welfare of women subject to discrimination; it relegates them to the low

wage and assignment pairs associated with the lower title, and thus they are severely underutilized. However, as in the case with the salary history ban, the overall effect on the gender pay gap is uncertain; it depends on what range of assignments (and the associated wage) can be given to promoted employees. For example, if the maximum assignment for promoted employees is close to β' then the overall returns to skill are low, which could reduce the gender pay gap. This is more likely the case because of the distributional issue noted above.

Corollary 8: Under discrimination, when $\alpha' > \beta'$ an equal pay mandate reduces the welfare of women subject to discrimination, but may reduce the gender pay gap.

If, on the other hand, $\alpha' \leq \beta'$, then women subject to the glass ceiling are promoted; they have perceived skill above the promotion threshold α'. In this case, the promotion results are identical across women, regardless of the status of their employers. Assignments, however, will be different. Promoted women who are subject to the glass ceiling have perceived skill of β', and thus will generally have lower assignments than similarly skilled (and promoted) men. However, assignments are not relevant for the gender pay gap; all employees with the same title receive the same wage. In fact, those with lower assignments earn higher utility.

Therefore, an equal pay mandate eliminates the gender pay gap in this circumstance. Ironically, the net effect of discrimination is slightly higher female utility, driven by the biased cap on their assignment. The equal pay mandate still reduces the welfare of employees just below the promotion threshold. However, that no longer includes the women who are subject to discrimination, who are now above the promotion threshold.

In this case, the precise effect of a mandate on the welfare of high skill women is indeterminant. It depends on the equilibrium wage for promoted employees, which could be above or below the wage these women earn without the mandate.

Corollary 9: Under discrimination, when $\alpha' \leq \beta'$, an equal pay mandate has an indeterminant effect on the welfare of women subject to discrimination but eliminates the gender pay gap and creates a negative utility gap.

Corollary 9 is more likely to apply when the promotion threshold is very high, such that the prospect of severe underutilization induces over-assignment. That said, the benefits it generates vis-à-vis the gender pay gap are not necessarily reflected in the welfare of the employees subject to discrimination, and come at the cost of an aggregate efficiency loss. Further, as noted above, an equal pay mandate gives firms an incentive to create more titles, and

thus reduce the inefficiency associated with over-promotion. This incentive is particularly strong when the extant promotion threshold is very high, such that skilled but unpromoted employees are underutilized. In other words, firms have an incentive to create more titles (and undo the effects of over-promotion) in exactly the circumstance where an equal pay mandate is likely to eliminate the gender pay gap.

VII. Discussion

In this article, we augment the standard asymmetric learning framework with costly effort, such that employers must pay high ability workers a greater wage to induce the optimal level of effort. This simple but powerful adjustment, which is in keeping with the basic fundamentals of contract theory, completely changes the potential signaling value of wages. Wages are an output in the standard approach, reflective of publicly observable extrinsic signals of employee quality, such as title, and thus are irrelevant to potential outside firms. In our approach, wages are informative about employee quality and play an important role in determining equilibrium pay and promotions. In fact, they are perfectly informative. The signaling value of wages is so powerful that wage disclosure leads to an equilibrium equivalent to one of symmetric learning: firms pay their employees their full marginal product each period and the full productive capacity of the economy is utilized each period. Wage disclosure prevents strategic manipulation of noisy public signals by underpromotion, which is a standard feature and source of inefficiency in asymmetric learning models, because wages are fully informative. This explains why many employers ask about wages in the hiring process when they are allowed to do so.

Our framework is relevant because policy makers are increasingly enacting reforms that intervene in wage disclosure and wage formation in efforts to combat historical inequities in the distribution of wages across employees. The standard models are poorly suited to analyze these interventions. By modelling wages as an output, they imply that the extant interventions are irrelevant. In addition, they lead to a paradoxical result, in that a large class of seemingly drastic interventions in the labor market are either costless or even welfare improving, such as banning competition for workers or mandating equal wages for all workers. In contrast, by modeling wages as necessary to generate increased production, our framework is equipped to evaluate these interventions, and ours is the first theoretically grounded analysis of their likely effects.

We demonstrate that both salary history bans and equal pay mandates introduce aggregate inefficiencies that reduce overall wages and have distributional consequences across employees. A ban introduces inefficient reliance upon noisy signals, which reduces aggregate productivity and thus welfare. A ban also reduces the outside option of the highest skilled employees within each title, because they can no longer use their higher wage to signal their ability to outside firms. Low skill workers within each title benefit as the profits from this shift in bargaining power are distributed to new hires. This is consistent with the emerging empirical literature on salary history bans, which finds conflicting results across population subsamples. This literature (Agan, Cowgill & Gee (2020, 2021), Bessen, Meng & Denk (2020), Hansen & McNichols (2020), Mask (2020), Sinha (2019), (Sinha (2022)) is roughly consistent with our prediction that salary history bans will benefit those with less-developed career skills, given the significant data limitations.

Wage equality initiatives include measures such as legal rules that place the onus on employers to justify pay disparities and that mandate disclosure of internal pay disparities. We posit that such measures generally serve to make pay disparities costly, and that the effect is to standardize wages around observable, verifiable, non-suspect traits such as job titles. While this is not the only possible way to conceptualize such measures, it does accord with the intent of the policymakers and statements of equality activists. In our model, constraining wages by title also leads to production inefficiencies and lower aggregate wages: firms that must pay employees of different skill levels the same wage must also assign them (approximately) the same level of the skilled task in order to retain them, due to the cost of effort involved in higher assignments. This results in under-assignment of the higher skilled employees within each title band. Productivity and average wages are reduced, with the welfare losses concentrated in the higher skilled employees within each title. The welfare loss does not come from reduced returns to realized production; employees are paid fairly for what they produce. Instead, higher skilled employees cannot realize (and get compensated for) their full potential. These results are consistent with some emerging empirical literature, which find that highlighting internal disparities leads to more equal pay but overall lower wages (Cullen 2023).

Of course, these efficiency losses may be acceptable if the interventions reduced the pernicious effects of discrimination, as they are intended. To assess these tradeoffs, we introduce a simple and tractable model of discrimination, in which women employees are subject to a glass

ceiling. With wage disclosure, market forces do not fully address discrimination. Although wage disclosure enables mobility that reduces somewhat the gender pay gap, a form of anchoring implies that women subject to early-career discrimination have lower lifetime wages and utility. However, neither salary history bans nor equal pay mandates are more effective regimes. Although they may reduce (or even eliminate) the gender pay gap, they do so by generally reducing the welfare of the women subject to the glass ceiling. The pay gap declines because the welfare of skilled men declines by even more. Reducing the returns to skill is a poor method of addressing the effect of discrimination on highly skilled women.

Our analysis generates a significant number of testable empirical predictions. First, we predict that the evidence for asymmetric learning will be greater in jurisdictions that impose a salary history ban. Wage disclosure mitigates asymmetric learning and bans reimpose a reliance on other noisy signals that can be manipulated by incumbent firms. This will be particularly true in professions with significant variation in within-title wages. This variation is a sign that effort and skill vary within a given title, and suggest a role for wage disclosure in learning. Reduced within-title wage variation, fewer promotions, lower aggregate wages but higher entry wages would all be suggestive of greater asymmetric learning.

We predict that bans and equal pay mandates will both reduce aggregate wages, the returns to skill, and the incentives for discretionary investments in human capital. Both interventions will reduce mobility, particularly for the highest paid employees within a given title, who can no longer distinguish themselves from their lesser skilled peers. However, only bans will increase entry wages. These interventions also give firms opposing incentives regarding titles. Bans encourage firms to use fewer titles, whereas equal pay mandates encourage firms to use more. Once again, these effects will be stronger in professions which have greater intra-title wage variation under wage disclosure. Equal pay mandates also potentially give firms an incentive to weed out their least productive employees; lower skilled employees would exhibit greater voluntary and involuntary mobility.

Our framework can also be used to indirectly test for the presence of discrimination. Under wage disclosure, we predict that women (and other disadvantaged groups) will have greater mobility than men but benefit less from job switching because mobility motivated by escaping discrimination does not require a wage that fully reflects outside marginal product. We

predict that both bans and equal pay mandates will reduce or eliminate this additional mobility, as they "trap" high skilled employees at discriminatory firms.

Our work raises a number of avenues for future research. First, as suggested above, the differential incentives related to titles can be fully explored; one possibility is to endogenize these choices and examine the optimal organizational structure under different legal regimes. This can be combined with the potential for involuntary separations, which could increase under equal pay mandates. Second, we posit one form of discrimination based on the justification for the interventions we study, and which leads to disparities in the status quo ante. However, discrimination can take many forms, and it is possible that these, or other, interventions are effective at mitigating some different forms of bias. That said, we stress that our results will extend to any discrimination that results in underpromotion of certain classes of employees, as the interventions we analyze tie wages closely to noisy promotion signals.

Further, our analysis of wage disclosure suggests that there is room for improvement in labor markets. Market forces are an imperfect salve; possibly, different legal regimes could improve on the outcomes. These may involve greater information, as opposed to constraints on the information available to prospective employers. For example, in the UK, firms must disclose the ratio of male to female pay. While this disclosure is crude (it does not adjust for title or responsibility), it is possible that outside firms could use this to better identify prospective lateral hires who are underpaid and/or underpromoted. Determining the optimal disclosure regime is an interesting avenue for future work.

In addition, the asymmetric learning literature is over 40 years old, and includes a large variety of models and enhancements. These include slot constraints (whereby firms have a limited number of promotions available), various bargaining schema, variation across employees in the cost of effort, multiple dimensions of skill (e.g., employees could be endowed with two types of skill, relevant for different tasks), variations in the observability (and thus contractability) of different inputs and outputs, etc. These have implications for wages and productivity, as in our model, but also address a wider range of concerns, such as investments in education, training, and firm-specific human capital, job mobility amidst costly search, monopsony power, etc. Our intention in this article is not to replicate a half century of the literature. Instead, we hewed as closely as possible to the most standard assumptions, to demonstrate how meaningfully the conclusions change when costly effort is introduced to the

framework. The various legal regimes we study likely affect these other outcomes as well, and there is much room for further study along these lines.

References

1. Amanda Agan,, Bo Cowgill, & Laura Gee, Do Workers Comply with Salary History Bans? A Survey on Voluntary Disclosure, Adverse Selection, and Unraveling, in AEA Papers and Proceedings Vol. 110, pp 215-19, (2020)
2. Amanda Agan, Bo Cowgill, & Laura K. Gee, Salary History and Employer Demand: Evidence from a Two-Sided Audit, NBER Working Paper 29460 (2021) (available at http://www.nber.org/papers/w29460)
3. Moshe A. Barach & John J. Horton, How Do Employers Use Compensation History?: Evidence from a Field Experiment, CESifo Working Paper Series No. 6559 (2017) (available at SSRN: https://ssrn.com/abstract=3014719)
4. James Bessen, Meng Chen & Erich Denk, Perpetuating Inequality: What Salary History Bans Reveal About Wages, working paper (2020) (available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3628729)
5. Zoe B. Cullen and Bobak Pakzad-Hurson, Equilibrium Effects of Pay Transparency, Technical Report, Working Paper 2021.
6. Zoe B. Cullen, Is Pay Transparency Good?, NBER Working Paper No. 31060 (http://www.nber.org/papers/w31060), March 2023
7. Jesse Davis, Paige Ouimet, and Xinxin Wang, "Do Salary History Bans Limit Discrimination?," Working Paper, 2021
8. J. DeVaro & M. Waldman, The Signaling Role of Promotions: Further Theory and Evidence, Journal of Labor Economics, Vol. 30 no. 1 (2012)
9. L. Golan, Wage Signaling: A Dynamic Model of Intrafirm Bargaining and Asymmetric Learning, 50 International Economic Review 831 – 854 (2009)
10. B. C. Greenwald, Adverse Selection in the Labor Market, Review of Economic Studies, 53, 325–347 (1986)
11. Robert E. Hall & Alan B Krueger, Evidence on the Incidence of Wage Posting, Wage Bargaining, and On-the-job Search, 4 American Economic Journal: Macroeconomics 56-67 (2012)
12. J. Mask , Salary History Bans and the Wage Scarring Gap, working paper, (2020)

13. Meli, Jeffrey, and James C. Spindler. 2019. "Salary History Bans and Gender Discrimination." SSRN, March. https://papers.ssrn.com/abstract=3361431
14. J. Pinkston, A Model of Asymmetric Employer Learning with Testable Implications, The Review of Economic Studies, Vol. 76, pp 367-94, (2009)
15. Uta Schonberg, Testing for Asymmetric Employer Learning, 25 Journal of Labor Economics 651 - 691 (2007)
16. G. Sran,, F. Vetter, & M. Walsh, Employer Responses to Pay History Inquiry Ban, working paper (2020) (available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3587736)
17. Sourav Sinha, Salary History Ban: Gender Pay Gap and Spillover Effects, SSRN working paper, 2019
18. Sourav Sinha, US Salary History Bans: Strategic Disclosure by Job Applicants and the Gender Pay Gap, 2022 SSRN (ssrn.com/abstract=4025580)
19. Michael Waldman, Job Assignments, Signaling, and Efficiency, 15 RAND Journal of Economics 255 – 267 (1984)
20. Michael Waldman and Ori Zax, An Exploration of the Promotion Signaling Distortion, 32 Journal of Law Economics & Organization 119 (2016)

Appendix

Proposition 1: With costless effort, in the unique equilibrium there exists an inefficient promotion threshold $\alpha' > \beta'$ *and a promotion wage* $w_p^* > 0$ *such that:*

i) *Employees with* $\alpha < \alpha'$ *are not promoted; incumbents offer a second period wage-assignment pair of* $(w_{np}, \beta(\alpha)) = (0, \min(\alpha, \beta'))$*, and outside firms offer* $(w_{np}^*, \beta_{np}^*) = (0,0)$

ii) *Employees with* $\alpha \geq \alpha'$ *are promoted; incumbents offer a second period wage-assignment pair of* $(w_p, \beta(\alpha)) = (w_p^*, \alpha)$ *and outside firms offer* (w_p^*, β_p^*)*, where* $\beta_p^* = \underset{\beta}{\operatorname{argmax}}\, E[Y_2^* \mid \alpha \geq \alpha']$*, the assignment that maximizes the production of employees with skill above the promotion threshold when employed by the outside firm*

No workers change firms, and new entrants to the labor market receive positive wages that reflect the expected lifetime profits from new hires, but are not assigned any of the skilled task.

Proof: We assume that the equilibrium is characterized by a promotion threshold α', and derive the conditions under which is it in fact an equilibrium. We start with the pay and assignments of employees that are not promoted. Because we assume that the unconditional expected value of assigning an employee to any of the skilled task is negative, employees who are not promoted can only receive offers $(w_{np}^*, \beta_{np}^*) = (0,0)$ from the outside firm (this population is of lower skill that the full population). With the outside offer capped at 0, the incumbent need only pay these employees a wage of 0. At that wage, the incumbent maximizes its profits by fully utilizing the capacity of these employees, up to the public promotion threshold. Therefore, it assigns these employees $\beta(\alpha) = \min(\alpha, \beta')$. Neither incumbents nor outside firms have an incentive to deviate. We conclude that, with a given threshold α', the wages and assignments detailed in Part (i) meet the criteria for an equilibrium.

Next we turn to the threshold α' and the wages and assignments for promoted employees. We first consider the optimal wage and assignment strategy for the incumbent conditional on the outside offer w_p^*. In order to retain promoted employees, incumbent firms must match that wage. Although we derive the exact form of w_p^* below, for now it is enough to realize that $w_p^* > 0$. Outside firms know that they can assign new hires at least α' of the skilled task (this is a minimum because it is possible that there is a greater level of assignment that increases the average production), and that these hires will generate average production of at least $\alpha' * y$. If the prevailing wage for these employees was 0, outside firms would have an incentive to deviate to some small positive wage, hire those employees, and earn positive profits.

This positive wage creates a wedge between $\alpha'$and the public promotion threshold β'. From the above analysis, the maximum profit generated by an employee who is not promoted is $\beta' * y * (1 + s_1) + (1 - \beta') * s_0$. The incumbent will only promote employees if doing so generates profits at least that high. But promoting entails a positive wage. At the threshold α', the incumbent must be indifferent between promoting (and paying a wage of w_p^*) and not promoting (and paying a wage of 0). We compare the profits from promotion and no promotion to compute the relationship between α' and β':

$$\alpha' * y * (1 + s_1) + (1 - \alpha') * s_0 - w_p^* = \beta' * y * (1 + s_1) + (1 - \beta') * s_0 \rightarrow$$

$$\alpha' = \beta' + w_p^*/K \qquad \text{[A1]}$$

where $K = y * (1 + s_1) - s_0 > 0$. First, it is clear from [A1] that $\alpha' > \beta'$. Second, with a fixed wage but production that increases with assignment (up to ability), incumbent profits are maximized by fully utilizing all employees with skill greater than α'. Therefore, in equilibrium incumbents offer promoted employees $(w_p, \beta(\alpha)) = (w_p^*, \alpha)$.

We now turn to the outside firm, which simultaneously solves its own maximization problem, choosing the wage w_p^* conditional on the promotion threshold α'. Knowing that only employees with $\alpha > \alpha'$ are promoted, the outside firm can determine β_p^*, the optimal assignment for promoted employees:

$$\beta_p^* = \underset{\beta}{\operatorname{argmax}}\, E[Y_2^* \mid \alpha \geq \alpha'] \qquad \text{[A2]}$$

It is possible that $\beta_p^* = \alpha'$. It is also possible that $\beta_p^* > \alpha'$; some over-assignment of the lowest skilled promoted employees is optimal if the gains from better utilizing the more skilled employees outweigh the losses from over-assignment of the less skilled employees.

Outside firms cannot offer wages above the production of promoted employees at that optimal assignment (doing so would guarantee negative profits). Therefore, $w_p^* \leq E[Y_2^* | \beta = \beta_p^*, \alpha \geq \alpha']$. However, in equilibrium the incumbent will match the outside offer, which implies that it also cannot be the case that $w_p^* < E[Y_2^* | \beta = \beta_p^*, \alpha \geq \alpha']$. This would give the outside firm an incentive to deviate to a wage of $w_p^* + \varepsilon$ for some $\varepsilon > 0$, which would be accepted by all promoted employees and generates profits of $E[Y_2^* | \beta = \beta_p^*, \alpha \geq \alpha'] - (w_p^* + \varepsilon)$. The increment ε can always be chosen small enough that this sum is positive. We conclude that:

$$w_p^* = E[Y_2^* | \beta = \beta_p^*, \alpha \geq \alpha'] \qquad \text{[A3]}$$

Together, [A1] and [A3] give two relationships between α' and w_p^*. Each can be used to express w_p^* as an increasing function of α'. In the case of [A1], this is a linear function. For [A3], it is not necessarily linear, but we do know that $w_p^* \geq \alpha' * y$: the outside firm can at least assign lateral hires an amount of the skilled task equal to the minimum skill of the set of promoted employees. Further, when $\alpha' = 1$, $w_p^* = y$: if only the most skilled employees are promoted, outside firms make the highest possible offer.

The unique equilibrium is defined by the crossing of the functions [A1] and [A3] in the feasible domain for α' ($\beta' < \alpha' \leq 1$). Evaluated at β', the function in [A1] yields $w_p^* = 0$,

which is below the level of [A3], which requires that $w_p^* \geq \beta' * y$. Therefore, these functions cross in this domain iff [A1], evaluated at $\alpha' = 1$, results in $w_p^* \geq y$. Then the linear [A1] has at least one point in the domain of α' that is weakly greater than the curve in [A3], and thus the two (increasing) functions must cross. We rewrite [A1] to express w_p^* as a function of α' and solve for the parameter set that ensures a crossing:

$$w_p^*(\alpha')|\alpha' = 1 \geq y \rightarrow$$
$$(1 - \beta') * K \geq y \rightarrow$$
$$\beta' \leq (y * s_1 - s_0)/[y * (1 + s_1) - s_0] \qquad \text{[A4]}$$

Note that [A4] is expressed as an upper bound on the public promotion threshold β'. If the public threshold is low, then high skilled employees are severely underutilized when not promoted, and thus there is more scope to promote them despite the higher salary that promotion requires. If the public promotion threshold is very high, then there is little additional production to be gained from fully utilizing highly skilled employees, and it is possible that there is no crossing that defines an equilibrium. If there is no crossing, then in equilibrium no employees are promoted (i.e., all employees with skill above β' have assignments capped at β'), which is equivalent to $\alpha' > 1$.

Regardless of the exact level of α', at least some employees are inefficiently utilized. This results in positive second period profits for the incumbent firm, which are reflected in the first period wages. QED.

Proposition 2 (uniqueness): With costly effort and wage disclosure, the following equilibrium is unique:

$$(w_1, \beta_1) = (w_1^*, \beta_1^*) = (s, 0)$$
$$(w_2(\alpha), \beta_2(\alpha)) = (\alpha * y, \alpha)$$
$$(w_2^*(w_2), \beta_2^*(w_2)) = (w_2, \frac{w_2}{y})$$

Where no workers change firms and $s = E[(1 - \alpha) * s_0 + \alpha * y * s_1]$.

Proof: We demonstrate uniqueness by eliminating potential alternative equilibria. First, any 1-to-1 mapping of ability to wage by the incumbent is revealing of α. If the incumbent wage was below $\alpha * y$, the outside firm could profitably poach the employee by making an offer between the incumbent wage and $\alpha * y$. This would reduce incumbent profits vis-a-vis the

equilibrium, because of lost production linked to firm specific human capital. Therefore, no other 1-1 mapping is a viable equilibrium.

Second, we consider pooling strategies, in which the incumbent pays employees of different skill levels the same wage. To describe a potential pooling equilibrium, we must characterize the pools as constructed, paid, and assigned by the incumbent, and the offers made by the outside firm. The incumbent pooling strategy is defined by a set of triples $[w_2, \lambda(\alpha), \beta(\alpha)]$, where w_2 denotes the wage paid to employees in the particular pool, $\lambda(\alpha)$ denotes the measure of employees of ability α included in the pool, and $\beta(\alpha)$ denotes the assignment to the skilled task to each member of the pool (the incumbent can tie assignments to ability).[25] Every employee must be assigned to a pool, there is no requirement that all pools have multiple skill levels, and for some feasible wages $\lambda(\alpha)$ may be the null set.

Outside firms cannot observe $\lambda(\alpha)$ or $\beta(\alpha)$, and the incumbent cannot commit to them. Therefore, the incumbent will deviate from $\lambda(\alpha)$ and/or $\beta(\alpha)$ in any pool if it is profitable to do so. These deviations in pool construction and assignment will prove critical to invalidating prospective pooling equilibrium.

The strategy of the outside firm is characterized by a set of wages and assignments that are functions of the incumbent wage: $(w_2^*(w_2), \beta_2^*(w_2))$. This needs to be specified for all feasible incumbent wages, to ensure that the incumbent cannot deviate from a proposed equilibrium.

First, we show that there can be no mobility in equilibrium. The incumbent always has an incentive to pay its employees with ability of 0 a wage of 0. The outside firm cannot poach these employees and they generate incumbent profits of s_0. Therefore, in equilibrium, it must be that $(w_2^*(0), \beta_2^*(0)) = (0,0)$. Any employee that accepts an outside offer generates no profits for the incumbent. The incumbent would increase its profits from that employee to s_0 by reassigning the employee to $(0,0)$. Therefore, a viable pooling equilibrium must have no mobility.

The lack of mobility has several implications for any prospective pooling equilibria. First, without mobility, incumbents will adjust $\lambda(\alpha)$ and $\beta(\alpha)$ within each pool such that no employee is over-assigned to the skilled task (i.e., $\beta(\alpha) \leq \alpha$ for all α with $\lambda(\alpha) > 0$). This must be true because, without mobility, an employee that is over-assigned reduces profitability from the

[25] Note that $\beta(\alpha)$ need not be 1-1. The triples must be defined for all feasible levels of w_2, i.e., those between 0 and y.

losses associated with over-assignment (if there were mobility, then it is possible that an incumbent will not realize the losses from over-assignment). The incumbent increases profits by reducing the employee's assignment. This increases the utility of the employee, by requiring less effort, and so reassignment does not affect the assumption of no mobility.

Second, if there is a pool in which the profits associated with a given skill α are above the profits employees with that skill would generate in all other pools, then it must be the case that $\lambda(\alpha) = f(\alpha)$ for that pool, and $\lambda(\alpha) = 0$ for all other pools. For example, all employees with skill of 0 will be assigned a wage and assignment of 0. In any other pool they generate lower profits.

A third implication is that the maximal outside offer places an upper bound on the assignments that can be given to pooled employees. The outside offer is a wage and assignment pair based on the incumbent wage: $(w_2^*(w_2), \beta_2^*(w_2))$. Since the outside firm cannot distinguish between pooled employees, this offer must be made to every member of the pool, although the outside firm knows that only employees for whom this offer is utility-improving will accept it. Knowing that no employee is over-assigned, the minimum equilibrium outside offer is $(w_2^*, \beta_2^*) = (w_2, \frac{w_2}{y})$. As above, any offer with a utility below this cannot be a stable point. At any less attractive offer, the incumbent would include in the pool some employees $\alpha > \frac{w_2}{y}$ and fully utilize them. But then the outside firm could poach those employees with an offer of the form $(w_2^*, \beta_2^*) = (w_2, \frac{w_2}{y} + \varepsilon)$, and generate a positive profit of $\varepsilon * y$ (they were not over-assigned at the incumbent, and so generate production of $w_2 + \varepsilon * y$).[26] Note that this holds for pools with a null set of employees: the outside offer must be $(w_2, \frac{w_2}{y})$. If this were not the case, then the incumbent would deviate and profitably construct a pool at that wage.

If the outside offer takes the form $(w_2^*, \beta_2^*) = (w_2, \frac{w_2}{y})$ for all feasible levels of w_2, then the incumbent maximizes profits by assigning each employee up to her ability and paying a wage equal to $\alpha * y$. In other words, it would set $\lambda(\alpha) = 0$ in every pool except that with $w_2 = \alpha * y$ and would set $\beta(\alpha) = \frac{w_2}{y}$, effectively dissolving any pool. Doing so maximizes profits because the only profits are driven by firm specific human capital: employees are paid wages equal to

[26] If over-assignment was possible, then incumbents could deter these offers using the over-assigned employees.

their realized outside marginal product. We assume that $y * s_1 > s_0$, and thus the profits linked to firm-specific human capital are highest when employees are fully utilized. A lower wage for any employee would require a lesser assignment, and thus lower profits.

However, it is possible that the maximal outside offer is even more attractive. Consider a pool with $w_2 > 0$ that includes employees with skill above w_2/y, who are under-assigned. It is possible that the outside firm can increase the utility of its offer by assigning more than w_2/y, if the losses from over-assigning those employees with skill at or below w_2/y who will accept the offer are outweighed by the gains from better utilizing employees with skill above w_2/y. In equilibrium, the incumbent would have to at least match the utility of the outside offer (or its employees would accept the outside offer). To do so, it must reduce the assignment for employees in that pool to below w_2/y; at the same wage but a sufficiently low assignment, their utility would be too high for the outside firm to poach them.

Note that it is not feasible to construct pools that meet this criterion for all $w_2 > 0$. In particular, the outside offer for $w_2 = y$ is $(w_2^*, \beta_2^*) = (y, 1)$; there are no more skilled employees that can be in included in the pool to facilitate a more attractive offer. As a consequence, all employees with $\alpha = 1$ are assigned to the pool with $w_2 = y$ and are assigned $\beta(1) = 1$. At all other wages, they will be assigned at or below w_2/y and thus generate lower profits.

To demonstrate that pools that meet this criterion are not viable in equilibrium, consider the alternatives available to the incumbent for the under-assigned employees included in the pool. Assume we have an equilibrium with a pool with a wage w_2 where the maximal assignment is below $\frac{w_2}{y}$. Define the maximum profits employees in this pool generate as π_{max}. These profits are below those that would be generated if the most skilled employees in this pool were assigned $\frac{w_2}{y}$:

$$\pi_{max} < w_2(1 + s_1) + \left(1 - \frac{w_2}{y}\right) * s_0 - w_2 = w_2 * s_1 + \left(1 - \frac{w_2}{y}\right) * s_0 \qquad [A5]$$

Per the above, this pool must contain more skilled employees who are under-assigned. These employees will all generate π_{max}. Let the most skilled employee in the pool have skill $\alpha' > \frac{w_2}{y}$. This employee can feasibly be assigned to any pool where the maximum assignment is equal to or below her ability. In order for this to be an equilibrium, reassignment to those alternative pools cannot increase incumbent profits. Therefore, the profits π_{max} are weakly

greater than the analogous maximum profit in those alternative pools. At the same time, the most skilled members of those other pools can be feasibly reassigned to this pool (which has a lower maximum assignment); therefore, it must be the case that all feasible alternative pools generate the same maximum profits per employee (or they would be reassigned to this pool). This implies that every feasible alternative pool must contain employees that are under-assigned: pools with no under-assigned employees, including those with a null set of employees, allow greater utilization of their employees and thus greater profits.

In fact, this implies that all pools with greater assignments must contain under-assigned employees. If there was a maximum level of skill at which employees were under-assigned, that employee would have feasible alternatives at which she could be more fully utilized: those pools with maximum assignments up to her ability, which allow for full utilization because they do not have under-assigned employees by assumption. But we know that employees with $\alpha = 1$ are paid a wage of y, assigned 1 of the skilled task, and generate profits of $y * s_1 > \pi_{max}$. In other words, this pool has no under-assigned employees. Therefore, it is not the case that all pools with greater assignments contain under-assigned employees, and there must be feasible alternatives for the under-assigned employees that increase profits. This invalidates the only remaining form of prospective pooling equilibria. QED

Proposition 3: With a salary ban, the unique equilibrium has an inefficient promotion threshold $\alpha' > \beta'$ and a promotion wage $w_p^* > 0$ such that:

i) *Employees with $\alpha < \alpha'$ are not promoted; incumbents offer a second period wage-assignment pair of $\left(w_{np}, \beta(\alpha)\right) = (e * \min(\alpha, \beta'), \min(\alpha, \beta'))$, and outside firms offer $(w_{np}^*, \beta_{np}^*) = (0,0)$*

ii) *Employees with $\alpha \geq \alpha'$ are promoted; incumbents offer a second period wage-assignment pair of $\left(w_p(\alpha), \beta(\alpha)\right) = \left((w_p^* - e * \beta_p^* + e * \alpha\ ,\ \alpha\right)$ and outside firms offer (w_p^*, β_p^*), where $\beta_p^* = \underset{\beta}{\operatorname{argmax}}\, E[Y_{i,2}^* - \beta * e \mid \alpha \geq \alpha']$, the assignment that maximizes the utility of employees with skill above the promotion threshold when employed by the outside firm*

No workers change firms, and new entrants to the labor market receive positive wages that reflect the expected lifetime profits from new hires but are not assigned any of the skilled task.

Proof: The reasoning follows that of *Proposition 1* very closely, with small adjustments to account for the cost of effort. Like in *Proposition 1*, outside firms cannot make positive offers to employees who are not promoted. That allows the incumbent to pay them the minimum possible wage. That minimum wage is no longer 0, however, it is $e * \beta$. Similarly, the outside offer to promoted employees is now a wage-assignment pair that delivers utility equal to $w_p^* - e * \beta_p^*$. Incumbent firms must match this utility in order to retain their promoted employees. Therefore, in order to assign these employees their full capacity α, they must pay wages of $w_p^* - e * \beta_p^* + e * \alpha$, which equates utility with the outside offer.

We derive the analog to [A1]:

$$\alpha' = \beta' + (w_p^* - e * \beta_p^*)/K \qquad \text{[A6]}$$

where $K = y * (1 + s_1) - s_0 - e > 0$. Similarly, the analog to [A3] is:

$$w_p^* = E[Y_2^* | \beta = \beta_p^*, \alpha \geq \alpha'] \qquad \text{[A7]}$$

Like in *Proposition 1*, the equilibrium is defined by the crossing of the functions in [A6] and [A7] in the feasible domain of α' ($\beta' < \alpha' \leq 1$). Evaluating [A5] at $\alpha' = \beta'$ yields $w_p^* = e * \beta_p^*$. This would result in a utility of 0. Yet we know that the minimum utility that the outside firm can offer when $\alpha' = \beta'$ is positive (it must be at least $\beta' * (y - e)$). Therefore, the wage in [A7] must be above $e * \beta_p^*$. As in *Proposition 1*, [A7] yields $w_p^* = y$ when $\alpha' = 1$. Therefore, the functions cross iff [A6] yields a wage of at least y when evaluated at $\alpha' = 1$:

$$w_p^*(\alpha') | \alpha' = 1 \geq y \rightarrow$$

$$(1 - \beta') * K + e \geq y \rightarrow$$

$$\beta' \leq (y * s_1 - s_0)/[y * (1 + s_1) - s_0 - e] \qquad \text{[A8]}$$

The denominator in [A8] is greater than that in [A4], implying that the constraint on β' is looser; it is more likely to have some promotions when effort is costly. This is because wages rise with ability even amongst the employees who are not promoted (in order to just compensate for effort) and thus the penalty of needing to pay a wage upon promotion has less of an effect on profits. As in *Proposition 1*, it is possible that there is no feasible solution, in which case no employees are promoted. Regardless, incumbent firms generate positive second period profits, and these are reflected in the first period wages. QED

Proposition 4: Under an equal pay mandate, the equilibrium wages and assignments of employees who are not promoted are bound above:

i) $w_{np} \in [0, s_0]$

ii) $\beta_{np}(\alpha) \in [0, s_0/y]$

Proof: For Part (i), assume that we have an equilibrium where $w_{np} > s_0$. This implies that at least some employees are unprofitable for the incumbent, absent mobility, including those with $\alpha = 0$. If the outside offer generates positive utility, then the incumbent would assign those low skilled employees enough of the skilled task that they accept the outside offer. However, those employees cannot be profitably hired by the outside firm. Therefore, the outside firm would deviate from its offer.

If instead the outside firm offered a utility of 0, then the maximum assignment given by incumbent would generate a utility of zero:

$$\beta_{np}^{MAX} = w_{np}/e \qquad \text{[A9]}$$

However, the outside firm could then make profitable hires by deviating to a more attractive offer. For example, an offer of $(w_{np}^*, \beta_{np}^*) = (w_{np}, \beta_{np}^{MAX} - \varepsilon)$ is accepted by all employees with $\beta_{np} \in (\beta_{np}^{MAX} - \varepsilon, \beta_{np}^{MAX}]$, and these hires generate profits of $\beta_{np}^{MAX}(y - e) - \varepsilon * y$. For a small enough choice of ε, these profits are positive. Therefore, the outside firm would deviate and offer positive utility. We conclude that we cannot have an equilibrium where $w_{np} > s_0$.

For Part (ii), the upper bound on the assignment follows similar logic to *Proposition 2*: if an employee is paid below her outside marginal product, she can be poached. For any employee assigned more than w_{np}/y, there exists an $\varepsilon > 0$ small enough that an outside offer of $(w_{np}^*, \beta_{np}^*) = (w_{np}, \frac{w_{np}}{y} + \varepsilon)$ is accepted by the employee, and generates profits of $\varepsilon * y$. Because $w_{np} \leq s_0$, it is worth retaining all employees, and the incumbent would deviate to a lower assignment or a higher wage and retain the employee (and her production linked to firm-specific human capital). QED

Proposition 5: Under an equal pay mandate, the equilibrium assignments made by incumbents to employees who are promoted are bound above and below:

i) $\beta_p(\alpha) \leq \max(\beta', \frac{w_p}{y})$

ii) $\beta_p(\alpha) \geq \max(\beta', (w_p + w_{np} * K)/C]$

where $K = s_1 + s_0/y$ *and* $C = y * (1 + s_1) - s_0$

Proof: The reasoning on the upper bound in Part (i) exactly follows that of the analogous upper bound in *Proposition 4*. Given a wage of w_p, any employee assigned above w_p/y can be profitably poached by the outside firm with an offer of the form $(w_p, \frac{w_p}{y} + \varepsilon)$. The only caveat is that, if the maximum assignment is β', and some employees with skill below β' are promoted, then the outside firm cannot necessarily make an offer as attractive as $y * \beta'$, if the losses from over-assigning some employees are too large.

For Part (ii), we first recognize that a promoted employee must have an assignment greater than β' by definition. However, it is possible that an even tighter constraint applies. An employee who is not promoted but has skill above β', would be assigned $\beta = \frac{w_{np}}{y}$ if not promoted. That results in profits of:

$$\pi_{np} = \beta * y * (1 + s_1) + (1 - \beta) * s_0 - w_{np} = w_{np} * s_1 + \left(1 - \frac{w_{np}}{y}\right) * s_0 \qquad [A10]$$

Only employees who generate at least that level of profits will be promoted. The profits upon promotion are:

$$\pi_p = \beta_p * y * (1 + s_1) + \left(1 - \beta_p\right) * s_0 - w_p \qquad [A11]$$

Promotion requires that $\pi_p \geq \pi_{np}$. We solve for the constraint on β_p:

$$\beta_p \geq [w_p + w_{np}\left(s_1 + \frac{s_0}{y}\right)]/[y * (1 + s_1) - s_0] \qquad [A12]$$

QED

To demonstrate that equilibria of the proposed for exists under an equal pay mandate, we give one particular example.

Under an equal pay mandate, an equilibrium exists with a promotion threshold $\alpha' > 0$ *such that:*

i) *Employees with* $\alpha < \alpha'$ *are not promoted; incumbents offer a second period wage-assignment pair of* $\left(w_{np}(\alpha), \beta_{np}(\alpha)\right) = (0,0)$*, and outside firms offer* $(w^*_{np}(0), \beta^*_{np}(0)) = (0,0)$*;*

ii) *Employees with* $\alpha \geq \alpha'$ *are promoted; incumbents offer a second period wage-assignment pair of* $\left(w_p, \beta_p(\alpha)\right) = \left((y-e) + e * \beta_{max},\ \min\left(\max\left(\alpha, \beta'\right), \beta_{max}\right)\right)$ *for some* $\beta_{max} \leq 1$, *and outside firms offer* $(w_p^*, \beta_p^*) = (y, 1)$;

iii) *Outside firms make (off-equilibrium) offers to any employee who is not promoted but paid a positive wage equivalent to the offers they make to promoted employees:* $(w_{np}^*(w), \beta_{np}^*(w)) = (y, 1)$ *for all w > 0.*

No workers change firms, and α' *is the skill level at which incumbents are indifferent between not promoting (and assigning none of the skilled task) and promoting.*

Proof: First we note that the outside offer for any employee paid a positive wage generates a utility of *(y-e)*. The incumbent must match this utility to retain any employee paid a positive wage. This implies that the promoted employee assigned the highest amount of the skilled task (who has the lowest utility) will have exactly this level of utility. Assume that employee is assigned β_{max}. This level of assignment implies the wage necessary to pay promoted employees:

$$w_p = (y-e) + e * \beta_{max} \qquad \text{[A13]}$$

For example, if the incumbent chooses to assign the most skilled employee to $\beta_{max} = 1$, then it must pay a wage of y. If it settles for a lesser assignment (we will determine the optimal β_{max} below) it can reduce the wage. This wage then determines the optimal promotion threshold α'. It is the level of skill at which the incumbent firm is indifferent about the promotion status of the employee. If not promoted, the employee generates profits of s_0. Since promotion entails paying a wage of w_p, the employees who are worth promoting generate enough production to make it worth paying them that wage. We start by assuming that the promotion threshold is above the public threshold of β'.

$$\pi_p = \pi_{np} \rightarrow$$
$$\alpha' y(1+s_1) + (1-\alpha')\, s_0 - w_p = s_0 \rightarrow$$
$$\alpha'[y * (1+s_1) - s_0] = w_p \rightarrow$$
$$\alpha' = \frac{w_p}{[y*(1+s_1) - s_0]} \qquad \text{[A14]}$$

If [A14] is in fact above β', then this is the optimal promotion threshold. However, if [A14] is below β', then we need to adjust the threshold for over-assignment, since promotion requires an assignment of at least β':

$$\pi_p = \pi_{np} \rightarrow$$

$$\alpha' y(1+s_1) + (\beta' - \alpha') * L + (1-\beta')\, s_0 - w_p = s_0 \rightarrow$$

$$\alpha'[y * (1+s_1) - L] = w_p - \beta'(L - s_0) \rightarrow$$

$$\alpha' = \frac{w_p - \beta'(L-s_0)}{[y*(1+s_1) - L]} \qquad \text{[A15]}$$

Therefore, either [A14] or [A15] define the threshold. If it is [A15], then some employees are over-assigned. Note that both the numerator and denominator of [A15] are positive, and thus $\alpha' > 0$ for any choice of β_{max}.

The incumbent maximizes its profits by using its employees to the fullest extent possible, conditional on two constraints. At most, it can assign β_{max}. At least, it must assign β'. Therefore, for all $\alpha > \alpha'$, we have $\beta_p(\alpha) = (\min(\max(\alpha, \beta'), \beta_{max}))$.

The incumbent will choose the profit-maximizing level of β_{max}, based on these assignments to the skilled task. As it reduces β_{max}, it can reduce the wage it must pay its promoted employees, and thus is able to promote a wider range of employees. The optimal choice will satisfy the following:

$$\beta^*_{max} = \underset{\beta_{max}}{\operatorname{argmax}}\, \pi(\beta_{max}) \rightarrow$$

$$\beta^*_{max} = \underset{\beta_{max}}{\operatorname{argmax}} \int_{\alpha'(\beta_{\max})}^{1} \Big\{ Min(\alpha, \beta_p(\alpha)) * y * s_1 + \max\big(0, \beta_p(\alpha) - \alpha\big) L + (1 - \beta_p(\alpha)) * s_0 - w_p(\beta_{max}) \} f(\alpha) d\alpha \qquad \text{[A16]}$$

The outside firm cannot make greater offers to the employees who are not promoted. Any such offer would be accepted by all those employees, who have lower skill than the general population and thus cannot be assigned any of the skilled task. Similarly, the outside firm cannot raise its offer to promoted employees; it is already offering the highest possible utility, equal to *(y-e)*. Outside firms have no incentive to change the off-equilibrium offer made to employees who are not promoted but are paid a positive wage, as these offers are not accepted by any employees. A justification for this offer is that the outside firm is itself subject to the equal pay mandate, and thus should in principle be restricted to only two wages. The off-equilibrium offer to employees who are not promoted preserves its compliance with the equal pay mandate, since it is the same wage offered to promoted employees. We will see below that a high offer is

necessary: it provides the equal pay mandate version of a steady state. It must provide sufficient utility such that the incumbent has no incentive to raise its wage and better utilize its employees.

Incumbent firms choose the threshold α' optimally, and so are already maximizing profits conditional on assigning none of the skilled task to employees who are not promoted and providing utility of at least *(y-e)* to promoted employees. The equal pay mandate prevents them from reducing the wage of employees with skill below β^*_{max}, who are overpaid vis-à-vis the outside offer (i.e., they earn utility above that implied by the outside offer, because they are paid the maximum wage but assigned less than β^*_{max} of the skilled task).

The only remaining deviation to consider is that the incumbent may prefer to raise the wage of the unpromoted employees and more fully utilize those with some skill. The off-equilibrium outside offer ensures that doing so requires awarding employees a utility of at least *(y-e)* to retain them. Specifically, if the incumbent wished to assign an employee β of the skilled task, w_{np} must be at least $(y-e)+\beta*e$. The profits assigning a given β are:

$$\pi(\beta) = \beta * y(1+s_1) + (1-\beta)\, s_0 - w_{np} \rightarrow$$
$$\pi(\beta) = \beta * [(y-e) + (y*s_1 - s_0)] + s_0 - (y-e) \qquad \text{[A17]}$$

where we substitute the required salary for that assignment. Both $(y-e)$ and $(y*s_1-s_0)$ are positive by assumption, implying that [A17] is increasing in β. Therefore, a sufficient condition for deviations of this kind to be unprofitable is that even the maximum assignment to unpromoted employees of β' is unprofitable. Note that this is more strict than necessary because raising the wage sufficiently high to allow for an assignment of β' entails overpaying many employees (and possibly over-assigning some with low skill such that they accept the outside offer), and the average profits from this type of deviation will be much lower than the profits from the most skilled employees. That said, we can substitute β' into [A17] and compare to the equilibrium profits of s_0:

$$\beta' * y(1+s_1) + (1-\beta')\, s_0 - w_{np} \leq s_0 \rightarrow$$
$$y*s_1 - s_0 \leq (y-e)*(1-\beta')/\beta' \qquad \text{[A18]}$$

We conclude that if the firm specific human capital associated with the skilled task s_1 is not too high, then [A18] is satisfied and the proposed equilibrium is robust to this type of deviation. Although the necessary condition is a looser constraint on s_1, it is the case that a pure strategy equilibrium is not possible if s_1 is too high. The maximum utility the outside firm can ever offer is *(y-e)*. If s_1 is very high relative to s_0, then deviations to a wage above s_0 for unpromoted

employees can be impossible to deter, and per *Proposition 4* a wage above s_0 cannot be an equilibrium.

Average wages are lower than in the unrestricted case considered in *Proposition 2* because employees with $\alpha \in (0, \alpha')$ are underutilized. This reduces both the second period wages earned by those employees, and the second period profits from firm-specific human capital, and thus reduces the first period wage as well. If $\beta^*_{max} < 1$ then some promoted employees are also underutilized. QED