

achieving pay equity

How analytics has evolved to
support true progress

Authors:

Brian Levine, PhD

Linda Chen

Alex Grecu, PhD

welcome to brighter



Over the last few years, the pressure on organizations to report on pay equity has grown substantially. Activist investors have been challenging companies with shareholder proposals — focusing submissions first on the technology sector, which has since seen proactive disclosure become commonplace, and more recently on financial services firms, for whom disclosure is on the rise. On top of that, new UK legislation also requires companies to disclose average pay differences between women and men regularly. Even as these and other global regulations continue to exert pressure, organizations have moved to prioritize pay equity as part of due diligence and to ensure access to top talent. Our *When Women Thrive* research, initiated in 2014, shows that organizations engaging in rigorous pay-equity review based on statistical analysis are more successful in achieving diverse representation.

Under such heightened scrutiny and attention, organizations are focused on firm-wide, global analysis to assess their circumstances proactively, support action where necessary and inform their communications to various constituencies, including shareholders, customers and employees, even when not required by law. To meet a broadening set of demands, we argue that such analyses need to be more focused on achieving organization-level change than simply ensuring employee-level alignment with norms. Analysis should afford decision makers the opportunity to test the impact of different pay-adjustment strategies on pay gaps, to achieve as much progress as possible within budget constraints. In this paper, we show how pay-equity assessments have evolved — first, reviewing the enduring attributes of effective analysis and second, presenting critical process refinements to drive greater impact.



Defining all-else-equal “adjusted” pay gaps

A standard pay equity analysis compares pay for similar work — that is, it accounts for legitimate factors intended to drive differences in pay, such as experience, location and role, before calculating an all-else-equal “adjusted” pay gap between women and men or between people of color and white employees.

By contrast, the UK disclosure requirement focuses on a “raw” pay gap, an all-in number that does not account for these factors and therefore compares dissimilar workers. Analysis to identify potential pay discrimination would rely on the adjusted pay gap and not the raw figures, though analyses should also be considered to assess equity across related employment dimensions, including opportunity for promotion to higher roles and fairness in performance management.

Broader analyses could further reveal, for example, differences in opportunities to develop experiences and skills that are more favorably rewarded.

Figure 1 below represents both raw and adjusted pay rates, by gender, for one organization. As the top chart shows, nearly 60% of women are paid less than \$50,000, whereas only 40% of men feature in this lowest pay band. This skew in representation contributes heavily to the overall raw pay gap and largely reflects differences in roles occupied by women and men. By developing statistical models to help us account for legitimate factors and properly estimate expected levels of pay for each individual, we can truly assess whether a pay disparity exists on an apples-to-apples basis, as the bottom chart illustrates. In this particular example, the raw pay gap of more than 20% falls to less than 2% after applying models; the shapes of the “standardized” pay distributions are clearly more similar than those of the raw pay distributions.

Figure 1: Actual and “standardized” pay distributions for one organization

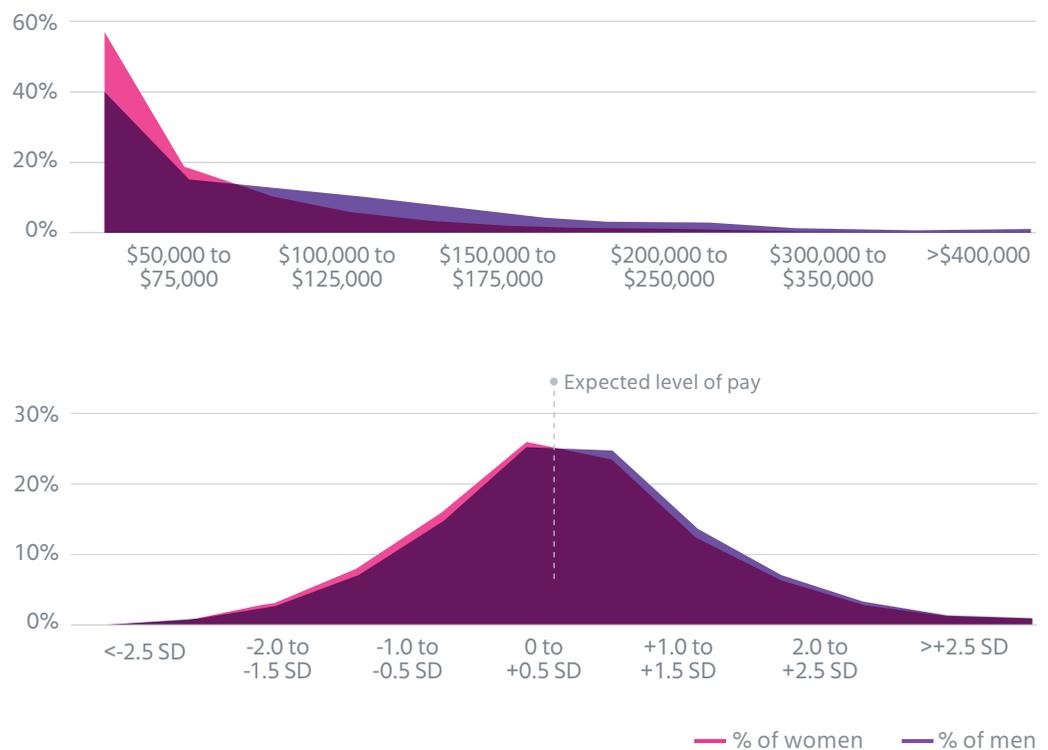
Disguised client example

Actual pay

The chart on the right represents the distribution of employees, by gender, across FTE pay bands. Nearly 60% of women have base salaries of less than \$50,000 versus 40% of men. But these actual pay levels do not yet account for any legitimate factors that may explain differences in pay.

Standardized pay

To account for legitimate factors, our statistical models enable us to “standardize” pay and calculate how far, in terms of standard deviations, actual pay levels deviate from model expectation. The chart on the right shows the standard deviation (SD) distribution of employees’ actual pay.



Enduring elements of effective pay-equity analysis

The process that Mercer relies on to assess pay equity consists of three primary steps (summarized in Figure 2):

Step 1: Data collection

Standard elements include gender and, in the US, race/ethnicity, relevant pay constructs that generally include base salary and total compensation plus a broad set of legitimate employee factors that potentially drive differences in pay. Legitimate factors captured should reflect the organization’s compensation philosophy and the reality of what organizational information systems capture.

Step 2: Running statistical models

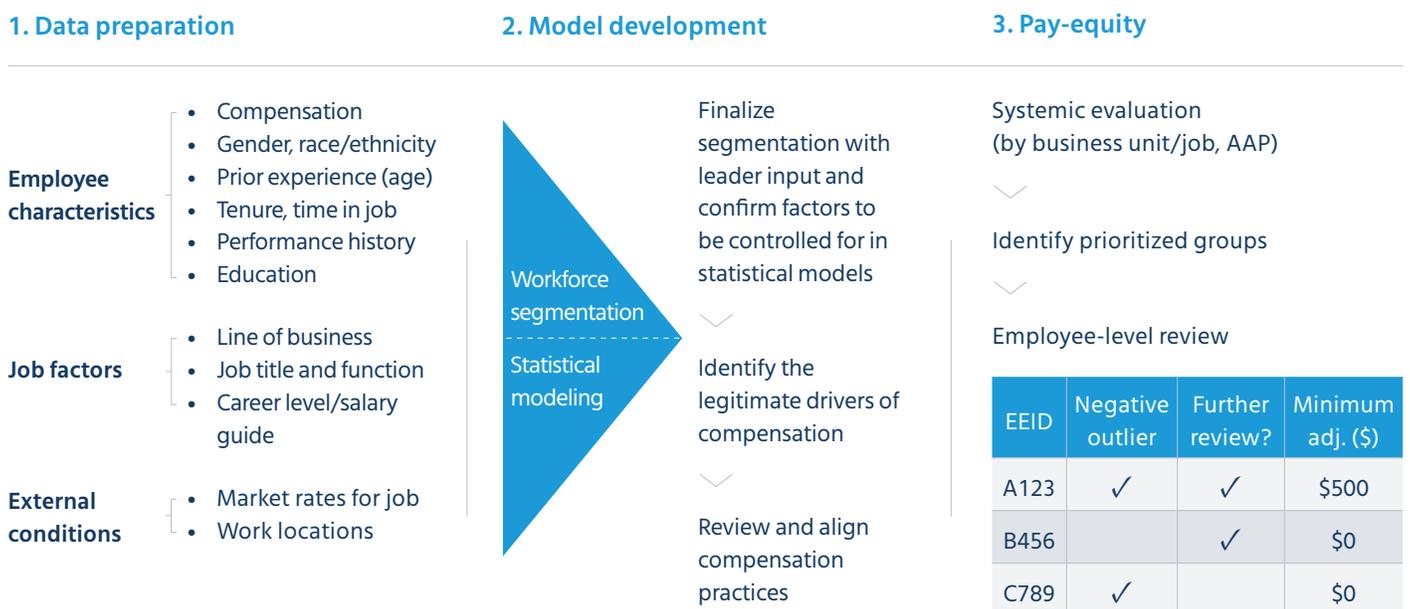
The workforce is broken into segments of employees subject to similar pay practices. In each of these segments, regression models, including legitimate factors, are then estimated — at

this point, models exclude gender and race/ethnicity. These models are validated with critical stakeholders (for example, compensation leaders and legal counsel) to ensure they represent compensation norms that the organization seeks to build upon (for example, pay for organization-specific over general work experience, pay driven by geographic differentials, pay for performance).

Step 3: Identifying areas of risk and remediation steps

The regression models are used to calculate systemic pay differences across employee groups (for example, women and men) for the overall enterprise, separate businesses/regions and distinct jobs. In areas where there are statistically significant differences between employee groups (that is, where factors in the model cannot explain meaningful differences in pay), the analysis points to specific “outliers” — employees who may need pay adjustments. It also calculates the actual adjustments required to bring pay levels into the range of expectation.

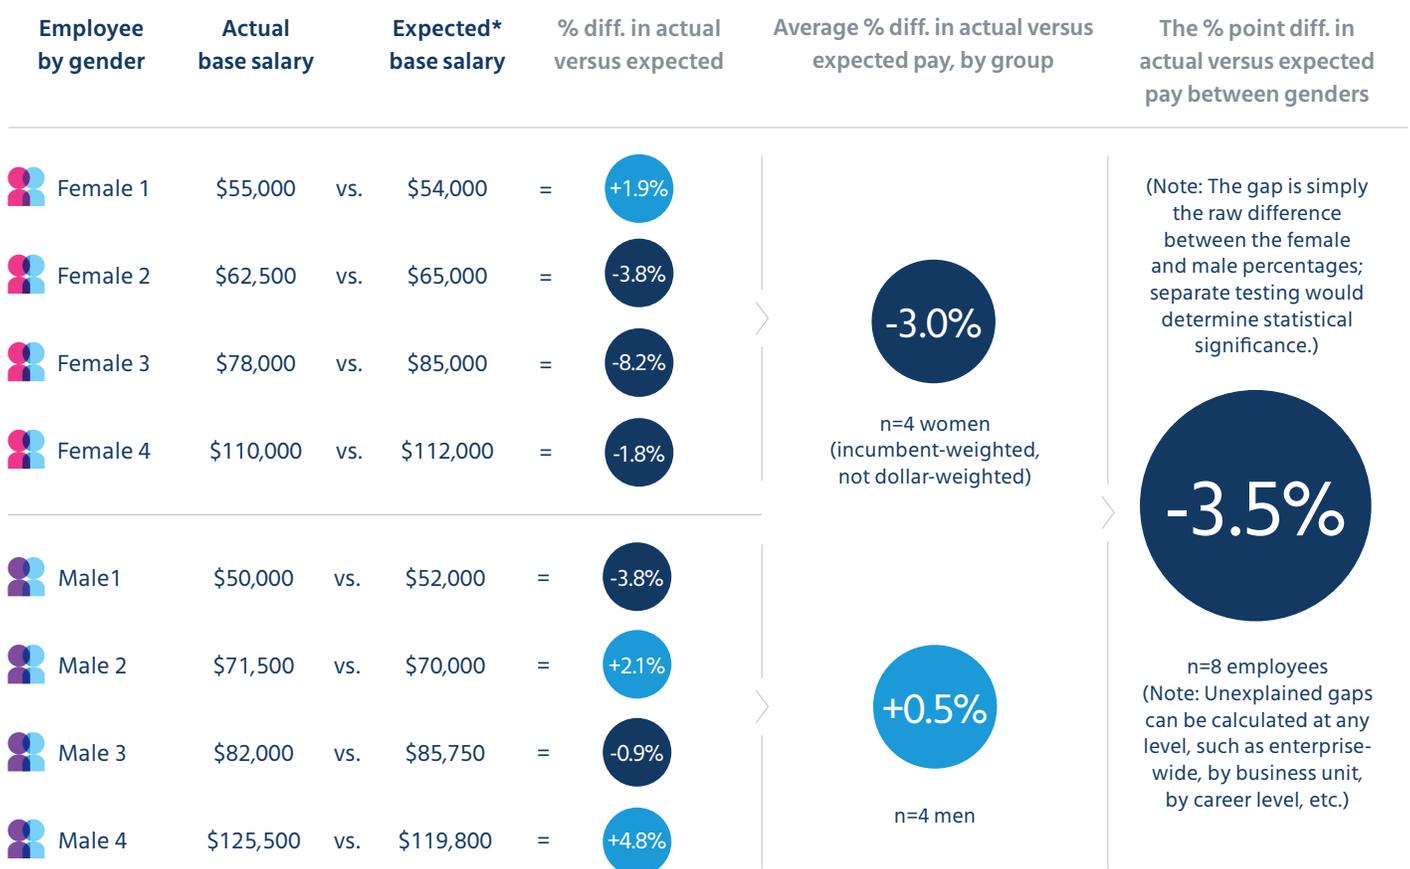
Figure 2: Mercer’s pay equity process



To calculate the adjusted pay gap for any workforce unit, one first calculates the pay gap for each employee, defined as the percentage difference between actual pay and expected pay from the relevant statistical model. These individual pay gaps

are then averaged for each employee group (for example, women and men) separately. Finally, the adjusted pay gap is calculated as the difference in the average gaps between the groups. The calculation logic is shown in Figure 3.

Figure 3: Calculating the gender pay gap



* Expected pay levels are estimated from statistical models that account for legitimate factors that differentiate pay (for example, career levels, experience, performance, location, etc.).

From process fairness to closing gaps

While much of Mercer's recommended process remains the same, it has nonetheless evolved to meet the changing needs of the organizations we support. Up until a few years ago, the priority for most was to assess whether there were any areas of systemic difference or risk and to consider pay increases in those areas for employees who appeared to be paid below the model-based norm. For these "outlier" employees, we recommended increasing pay so that it would no longer be statistically significantly different from expectation¹ (see Figure 4). We associate this priority with **process fairness**, as every employee in a risk area would be considered for a pay increase, regardless of gender or race/ethnicity.

The tradeoff for maintaining process fairness, however, is a muted impact on addressing systemic pay gaps. For example, if a particular area shows a statistically significant pay difference where women are paid less than men, a disproportionate amount of women may be concentrated in the lower tail of the normal distribution, but there may be male outliers found in these negative tails as well. If pay levels for these outliers are similarly adjusted upwards, their increases would counter the adjustments made for women and reduce the net impact on the pay gap.

Increasingly, our clients are focused on making more progress in **closing pay gaps**, in identified areas of risk and for the organization overall. To address the tradeoffs described

previously, closing gaps generally requires focusing on a different or larger set of outliers and, possibly, a greater expenditure on pay adjustments as well. It also requires the ability to rapidly test the impact of different adjustment strategies.

To close gaps, we generally recommend a "bottom-up" approach — building up from individual outliers to close gaps in areas of risk. But we refine the conventional "process fair, ± 1.96 SDs" methodology in two ways:

- First, we consider diverting adjustments to the group that is significantly underpaid relative to the other. Directing dollars to that group can dramatically increase the impact of limited budgets and is "fair" in the sense that any underpaid group would be considered for an adjustment. For example, in areas where men are paid statistically significantly below women, pay for men would be adjusted upwards.
- Second, we consider narrowing the confidence interval; that is, reducing the number of SDs used for outlier identification. As the number of SDs declines, the number of outliers, and the adjustment for each outlier, increases. Figure 5 depicts the exercise of reducing the number of SDs.



¹ Mercer's standard has been to consider increasing pay to be within 1.65 standard deviations (SDs) of the expected pay level; this approach is more aggressive than the traditional focus on ± 1.96 SDs, identifying more outliers and allocating more budget to counter identified risks.

Alternatively, one can consider a “top-down” approach, calculating the total budget that would eliminate the statistical significance of the “adjusted” pay gap if allocated to the “disadvantaged” (that is, lower-paid) group. Although the ultimate impact of this remediation budget would depend on the specific allocation of pay adjustments between employees, a “best estimate” is to calculate the budget assuming a flat

percentage increase for each employee in the disadvantaged group. During the review process, such a budget would ultimately be allocated to employees who are high performers or low in their pay ranges. The complexity and potential subjectivity of such employee-specific considerations is why we tend to favor a “bottom-up” approach.

Figure 4: The distribution of expected pay for a specific employee

Conventional remediation has focused on “minimum” adjustments that ensure employees are paid no less than 1.96 SDs below expectation.

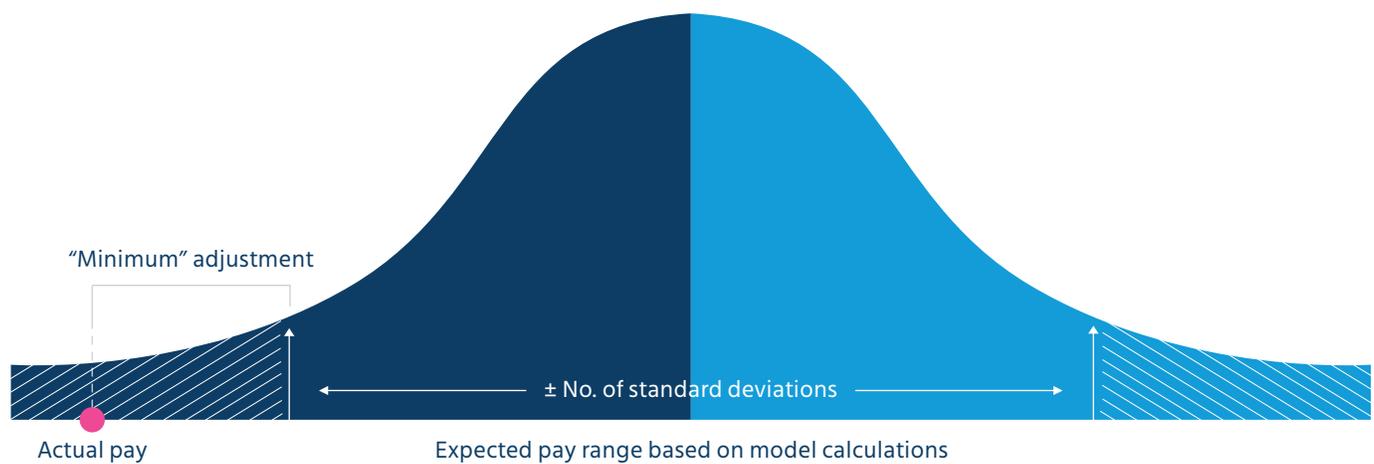
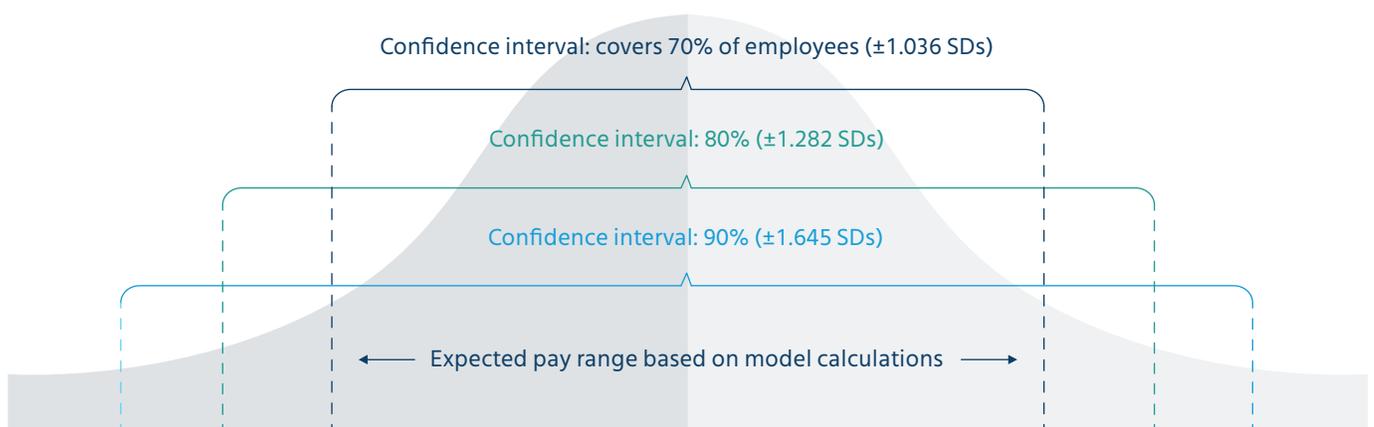


Figure 5: Various confidence intervals around the expected pay level

Reduce the number of SDs — and the size of the confidence interval that bounds the range of expected pay — to identify more outliers and increase potential adjustments.



Note that while the approaches described tend to distribute budgets in favor of women and people of color — as systemic pay differences still tend to negatively impact these groups —

organizations frequently also consider adjustments for men and white employees in areas where they are disadvantaged.

Figure 6 below shows a standard output that summarizes different adjustment strategies, the remediation budgets associated with each and the remaining post-adjustment gaps, assuming budgets are spent. Organizations choosing one strategy over another must consider the relevant tradeoffs.

A further opportunity to accelerate progress would be to address the positive outliers. In some units, it may be the case that pay differences are driven primarily by those in the advantaged group who are paid significantly more than expectation. Containment of positive outliers (for example, treating them similarly to cases where employees are paid above the maximum of their pay range) should be considered as part of a holistic strategy to close gaps.

The “residual” distributions in Figure 7, which depict the differences (in terms of SDs) between actual and expected pay levels for women and men, side by side, show that the optimal strategy to reduce gaps might well vary across units. In the first unit, women are indeed concentrated among negative

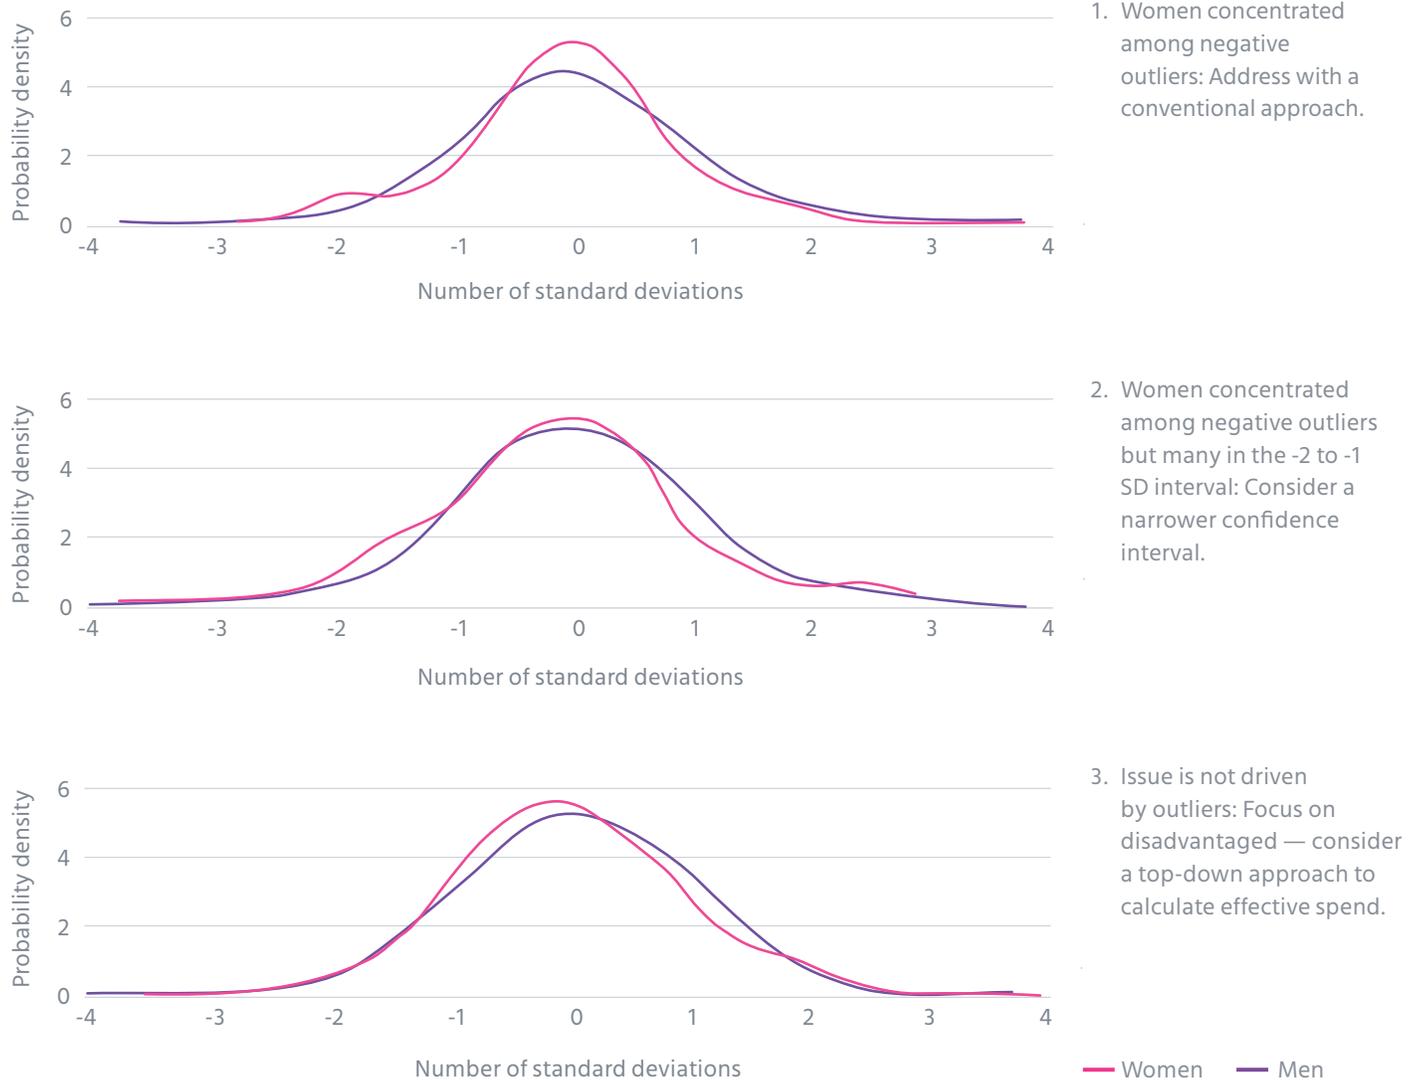
outliers; here, a focus on all residuals below two SDs might be very effective toward diminishing the adjusted gender pay gap. In the second unit, it would be more impactful to explore adjustments for those who are below one SD; that is, narrowing the confidence interval to pick up more outliers in the lower tail. In the third unit, a focus on women only and consideration of broader adjustments than just a set targeted to the negative outliers might be required.

In every case, we advise that outliers be separately reviewed before pay adjustments are finalized, as there may be reasons for differentiated pay rates that are not reflected in analysis data. The better the quality of the data and the lower the number of exclusions, the greater the potential for the actual impact to match the impact anticipated. Furthermore, the benefit of the analysis cannot be realized if recommended adjustments are effectively countered by reductions in merit and related increases. Recommended adjustments need to be processed on top of other changes that are regularly considered.

Figure 6: The impact of various remediation strategies on the “adjusted” gender pay gap

Disguised client example

Primary focus	Description of adjustment option				Adjustment actions		Post-adjustment results (starting gap = 2.6%)
	#	Who gets adjusted	Where to apply adjustment	Conf. interval	# of EEs adj.	Total budget (\$)	Overall gap %
Conventional — focused on process fairness ; all EEs eligible for an adjustment	1	Adjust all negative outliers (female and male)	In any work group with a significant pay difference between women and men	90%	900	800,000	-2.6%
	2			80%	1,800	1,900,000	-2.5%
	3			70%	2,700	2,700,000	-2.5%
Focused on most efficient use of budget ; smaller budgets make more progress	4	Adjust all negative outliers (female and male)	Women in work groups where women are disadvantaged, men in work groups where men are disadvantaged	90%	500	500,000	-2.3%
	5			80%	900	1,000,000	-2.1%
	6			70%	1,400	1,400,000	-1.8%
Focused on reducing gaps by addressing outliers across the enterprise	7	Adjust female negative outliers only	Across the entire company	90%	2,000	1,800,000	-1.9%
	8			80%	3,800	3,600,000	-1.5%
	9			70%	5,600	7,900,000	-0.7%

Figure 7: Residual distributions for three business units and the implications for the remediation strategy**Disguised client example****Broader analysis to support closing gaps once and for all**

Of course, as already stated, the “raw” pay gap will not be addressed through pay equity analysis alone. It is primarily driven by differences in the roles occupied by women and men and by employees of color and white employees, as well as other individual and situational factors, such as experience, performance, business affiliation and location.

Statistical modeling of pay drivers can identify root causes of pay gaps; for example, revealing the impact of role, employee experience and performance management processes.

Root-cause analytics should be conducted as part of a standard pay-equity analysis (see the “decomposition” analysis in Figure 8) and should be followed by further examination of equity in hiring, promotion and employee retention. See the Internal Labor Market (ILM[®]) map, shown in Figure 9.

Figure 8: A “decomposition” that shows reductions in the “raw” gap achieved by different sets of controls

When examining the gender pay gaps, differences in roles account for 70% of the raw gap (that is, the 20% gender pay gap reduces to 6% after accounting for career levels), suggesting that women’s lower pay is largely due to a higher concentration of women in lower-paying jobs. For the race/ethnicity pay gaps, differences in location account for 64% of the raw difference, suggesting that employees of color tend to be concentrated in lower-paid work locations.

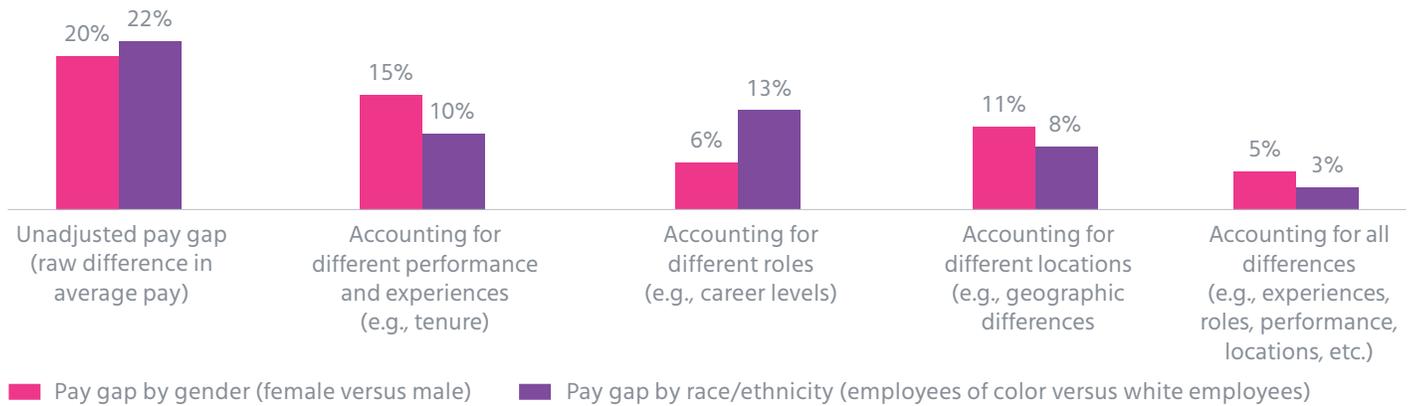
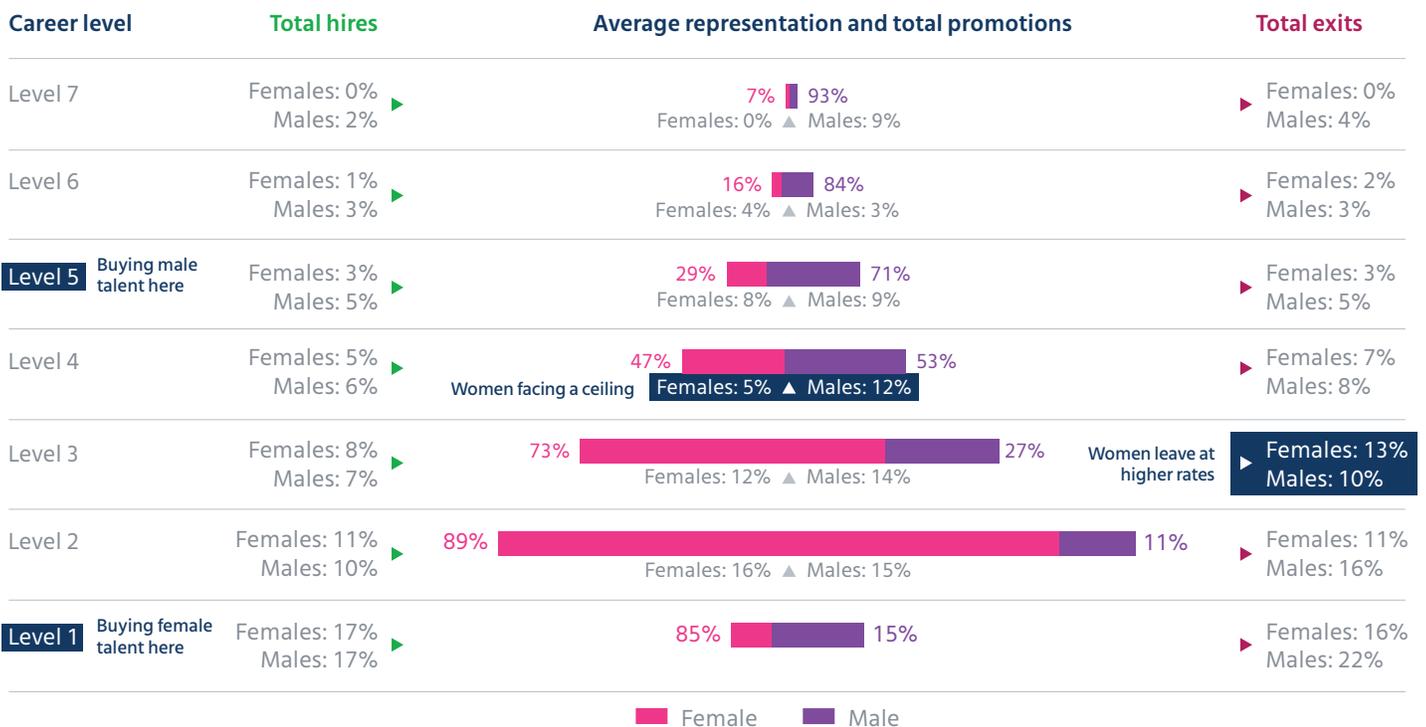


Figure 9: A sample internal labor market (ILM®) map depicting differences in gender representation across career levels, as well as differences in hiring, promotion and turnover rates

The skew in representation substantially impacts the raw pay gap, and the observed inequities in the rates of entry, exit and upward progression hinder this organization’s ability to build their female talent beyond Level 3. ILM maps can be created for any two-group comparison (for example, employees of color versus white employees) and can aid in broader review of root causes.

Disguised client example



Conclusion

Pay equity has long been a standard consideration of compensation programs. Commitments, however, have often been met through ad hoc, sometimes cursory reviews of those sitting in the same job and/or confidential analyses via legal counsel linking to a set of pay adjustments processed with little understanding of the impact. Greater scrutiny — and demands from business leaders for progress — requires that organizations raise the bar on their efforts. In our work with large, global organizations, across industries, we have found that effective pay equity analyses are characterized by three elements:

- 1. Thorough review of statistical models by compensation leaders** — Models should reflect how compensation is actually determined in the business and reinforce legitimate factors that should drive differentiation in pay.
- 2. Remediation strategies informed by estimates of impact** — Organizations should focus on a set of adjustments that reflect budget realities but also demonstrably drive progress.
- 3. Learning from the statistical models and related examination of data to expand equity efforts** — Pay equity can be best achieved through changes in practices that prevent gaps from arising in the first place.

For pay equity, the ultimate goal may be “100 cents on the dollar,” but the exercise should not be turned into a numbers game. The learnings discovered along the way — from what uniquely drives pay inside your organization to how outliers are distributed above and below model expectation — will not only inform how best to allocate adjustments but also identify what else to consider to more holistically address root causes. To get to equity, companies need to pursue more aggressive strategies than conventional approaches and, once there, maintain a vigilant focus.

Please reach out to the authors at the addresses below with any questions, comments or curiosities. We would love to hear from you!

brian.levine@mercer.com

alex.grecu@mercer.com

www.mercer.com



