

REORGANIZING WORK INSIDE THE FIRM: TASK CHANGE AFTER TECHNOLOGY ADOPTION*

Eleonora Brandimarti[†] Virginia Minni[‡] Raffaella Sadun[§]
Jorge Tamayo[¶]

October 2026

Abstract

We study task reorganization after technological change using worker–task-level data from a large Latin American bank. Two sequential interventions, automation of routine transactions and a later redesign of performance evaluation, separately alter production technology and organizational design without reducing employment. Automation raises output per worker but generates only limited task reallocation: workers expand routine production rather than reallocating effort toward higher-value tasks. Task reallocation remains limited until the firm redesigns performance evaluation. After the redesign, activity shifts toward the newly prioritized task across the workforce, while gains in non-targeted complex work, and the associated wage gains, concentrate among workers with low adjustment costs, proxied by tenure. The results highlight how automation-induced task complementarities require deliberate organizational redesign, with gains shaped by the distribution of adjustment costs within the firm.

Keywords: technology adoption; task reallocation; organizational design; managerial quality; firm productivity

JEL: J24; M54; O33

*We thank our partners at the Bank for sharing the data and essential context knowledge. This manuscript has not been subject to prior review by any party, except to verify that it does not contain any confidential data or information, as per the research collaboration contract signed at the outset of this project.

[†]Adam Smith Business School, University of Glasgow.

[‡]The University of Chicago Booth School of Business, CEPR & NBER, J-PAL.
virginia.minni@chicagobooth.edu.

[§]Harvard Business School, Digital Reskilling Lab, CEPR & NBER. rsadun@hbs.edu.

[¶]Harvard Business School, Digital Reskilling Lab, NBER. jtamayo@hbs.edu.

1. Introduction

Technology adoption drives productivity growth, but its returns depend on how firms reorganize work around it (Bresnahan et al., 2002; Bloom et al., 2012; Brynjolfsson et al., 2021). Task-based models formalize this process as reallocation across tasks: technology changes the relative returns of tasks, and relative prices draw workers toward those tasks in which they retain a comparative advantage (Autor et al., 2003; Acemoglu and Autor, 2011; Acemoglu and Restrepo, 2018, 2022).

Inside firms, however, the relative return to a task is not a market price. It is set by what the organization asks, measures, and rewards (Coase, 1937; Simon, 1951; Holmström and Milgrom, 1991). These organizational incentives are embedded in routines and costly to revise (Arrow, 1974; Nelson and Winter, 1982). When technology changes the relative value of tasks, incentives may not adjust immediately, creating a wedge between what the firm values and what workers are rewarded for.

We study this wedge in a Latin American bank that digitizes routine transactions in 2018 to free branch capacity for complex products, but does not revise task-specific performance priorities until 2021. The sequence and setting let us observe directly a stable workforce across three regimes: old technology with old performance priorities, new technology with old priorities, and new technology with redesigned priorities. The data cover the universe of financial trades from 2017 to 2023, each linked to the worker who completes it, together with personnel records, wages, training histories, and manager assessments for roughly 10,000 employees across 450 branches.

Two centrally coordinated interventions unfold in sequence within the bank and are implemented in a staggered fashion across branches. In 2018, the bank introduces a customer-facing mobile app that digitizes account openings, consumer loans, credit-card applications, and other routine retail transactions. The “Tech Roll-out” aims to broaden the bank’s customer base and to free branch capacity for the more complex products that the bank regards as strategically valuable. In January 2021, the bank redesigns how branch work is directed and evaluated. Under this “Task Intervention,” SME lending (the most analytically and in-

terpersonally intensive activity in branches) becomes the central performance priority. The new priorities are communicated and implemented through branch-level visits, KPI explanations, and training, while the production technology and the base wage schedule remain unchanged.

Throughout both interventions, the bank maintains an explicit no-layoff policy. Hiring and separations remain stable, making labor quasi-fixed in the sense of Oi (1962). Adjustment occurs primarily through changes in incumbent workers' task portfolios rather than turnover.

The first finding is that the Tech Roll-out raises output substantially but does not produce the broad shift toward complex products that the bank anticipates. Following app adoption, worker-level transaction volume increases by approximately 19 percent. Total clients, first-time clients, active accounts, and account balances also rise. These realized outcomes are consistent with an expansion in client demand alongside productive capacity.

Nevertheless, under the old performance system, the additional production is concentrated in the routine products enabled by the app. Simple trades rise sharply, other complex trades increase much less, and aggregate SME lending does not change significantly. The complex-to-simple output ratio falls. The app therefore expands the bank's production frontier, but production does not reallocate toward complex products. This pattern reflects the priorities embedded in the existing performance system. It is also consistent with what the technology alone would produce if the app makes routine tasks more productive and expands the client demand pool for them.

Average results, however, mask a pronounced tenure gradient. Low-tenure workers exhibit the largest increase in routine output and are the only group to show modest increases in SME lending and other complex products after the technology rollout. Responses decline with tenure, and the most experienced workers change little on either complex margin. Thus, under the old organizational incentives, low-tenure workers move toward complex work even without a formal change in what the organization rewards, whereas high-tenure workers change little. The tenure gradient remains present within age and baseline-ability groups and is not explained by initial product specialization, the tenure composition of coworkers or managers, differential changes in days worked, or worker turnover. We therefore interpret

the tenure gradient as consistent with adaptation costs associated with established routines, task-specific experience, and customer relationships.

The subsequent performance redesign, which changes priorities while leaving the technology in place, generates a sharp but highly targeted reallocation toward the activity it prioritizes. Following the intervention, SME lending increases sharply and persistently — by approximately 10 percent over the subsequent year, excluding the implementation week. Other complex activities, which are not targeted by an explicit KPI, respond markedly less. The ratio of SME lending to other complex products rises, while the overall complex-to-simple ratio remains broadly flat. The SME response is positive and broadly similar across tenure groups, closing the tenure gap on the directly prioritized product. Yet, the gap persists in other complex activities: output on these untargeted margins continues to rise primarily among low-tenure workers.

The reprioritization also redistributes earnings. Wage gains are concentrated among low-tenure workers, the same workers whose task portfolios change the most. Workers with eight or more years of tenure receive little or no average wage gain. This result shows that preserving employment does not necessarily equalize the gains from organizational change. Even within a stable workforce, the workers who adjust most capture a disproportionate share of the associated earnings gains.

Finally, we examine branch managers' role in coordinating resource allocation within the bank. While headquarters sets the administered returns to tasks, managers allocate coaching and training across workers, thus shaping how these incentives translate into task reallocation. To study their role, we estimate branch manager branch productivity fixed effects using pre-intervention variation from branch-manager rotations across branches, with branch and week fixed effects. These manager effects predict subsequent internal performance evaluations and subordinate assessments of teamwork, trust, adaptability, and inspiration, and are largely unrelated to manager demographics.

Before the performance redesign, the pattern of adjustment across tasks and workers is similar regardless of manager quality. After the redesign, however, manager quality becomes important, especially for activities that are not directly rewarded. Higher manager productiv-

ity is associated with larger increases in untargeted complex activities, almost entirely among low-tenure workers. By contrast, the increase in SME lending, the activity directly targeted by the new performance system, does not vary with manager quality.

These results point to two complementarities in the implementation of organizational change. First, the change in organizational priorities increases the value of managerial coaching. Second, managerial coaching is more productive for workers who are already better positioned to adapt, leading branch managers to concentrate their efforts on these workers. Managerial implementation therefore contributes to widening, rather than closing, the tenure gap in adaptation. Training patterns reinforce this interpretation. Total training hours change little across manager and worker groups, but branch managers allocate training differently: low-tenure workers receive relatively more analytical training, while high-tenure workers receive relatively more routine-oriented training.

The economic magnitude of these adjustment frictions is meaningful. Had all workers adjusted to the new technology as low-tenure workers did, the bank would have generated 18 percent more simple loans and 29 percent more complex loans through the end of 2019. Since complex loans are substantially larger, sluggish adjustment translates into approximately \$27 million in foregone complex lending and \$1.6 million in associated profits. Thus, heterogeneity in workers' ability to adjust can materially reduce the returns to technological change.¹

We organize these findings with a model of task allocation inside a three-tier firm. Technology changes production opportunities by increasing routine task productivity and expanding the client demand pool. The model takes as given the performance weights before and after the Task Intervention; these weights determine the relative private returns workers receive from different outputs. Branch managers observe workers more closely than headquarters and allocate scarce coaching between analytical and routine content. Workers allocate effort across routine, other complex, and SME products and differ in the cost of adapting to new technologies and tasks, which we proxy with worker tenure.

The framework separates three margins of adjustment. Technology expands *capacity* by

¹A regime more willing to separate workers could attenuate the tenure gradient we document through selection, by inducing exit among workers with the highest adjustment costs. Our setting instead isolates within-worker adaptation in an internal labor market, identifying tenure-based frictions in task reallocation without confounding selection through hiring and firing.

increasing what can be produced with existing resources. Performance priorities provide *direction* by changing which outputs are measured and rewarded. Worker heterogeneity and managerial implementation shape the *incidence* of the adjustment: who moves toward newly prioritized tasks and who captures the gains.

First, under the new technology with old priorities in place, routine output expands for every worker, with the largest increase among adaptable, lower-tenure workers. Whether capacity is released toward complex work depends on whether demand for routine products has been exhausted and any released capacity flows disproportionately to the same adaptable workers.

Second, the redesign disproportionately increases the prioritized activity and raises every worker's output in that activity. The tenure gap in responses can therefore close on the prioritized task while persisting on untargeted complex tasks, which continue to depend on workers' adaptability. Third, managerial coaching becomes more valuable once formal priorities are aligned and managers tilt coaching toward analytical content for workers for whom it has the highest expected return, that is the workers with lower adaptation costs.

The paper contributes to three literatures. First, it provides worker–task-level evidence on the organizational adjustment underlying technology–organization complementarity. Existing research documents that technology and organizational practices are correlated with productivity and that organizational barriers can limit the returns to adoption (Bresnahan et al., 2002; Bloom et al., 2012, 2014; Brynjolfsson et al., 2021; Atkin et al., 2017; Juhász et al., 2024; Feigenbaum and Gross, 2024; Houeix, 2026; Verhoogen, 2023; McElheran et al., 2026; Kastis and Vipond, 2025). We observe technology adoption, aggregate output, the subsequent change in performance priorities, and the associated within-worker, product-specific responses. The results connect this literature to the organizational-economics view of the firm as a system of administratively directed allocation and to theories in which performance measures act as internal prices (Coase, 1937; Simon, 1951; Arrow, 1974; Williamson, 1975; Holmström and Milgrom, 1991; Baker, 1992).

Second, the paper contributes to research on managers and productivity by showing that the relationship between manager quality and organizational adjustment is state-contingent

(Ichniowski et al., 1997; Lazear et al., 2015; Adhvaryu et al., 2023; Friebe et al., 2022; Metcalfe et al., 2023; Minni, 2026; Diaz et al., 2025; Chen et al., 2026). Branch managers facilitate adaptation only after formal performance priorities change, and primarily among workers who also display greater responsiveness to technology. The findings suggest that management is a complement to organizational direction rather than a substitute for it.

Third, the paper contributes to task-based research on technological change by documenting economically meaningful adjustment within existing jobs. Most empirical measures of technological exposure emphasize employment, occupations, or broad job categories. In our setting, hiring, separations, and occupational assignments remain largely unchanged even as product-level output, specialization, training, and earnings respond (Brynjolfsson et al., 2025; Eloundou et al., 2024; Acemoglu and Autor, 2011; Autor and Handel, 2013; Ocampo, 2022; Adenbaum, 2023; Autor and Thompson, 2025; Demirel et al., 2026; Freund and Mann, 2026; Humlum and Vestergaard, 2026). The findings therefore highlight a measurement challenge: adoption and occupational exposure alone may reveal neither the direction of within-job adjustment nor its distribution across incumbent workers.

The paper is organized as follows: Section 2 describes the setting and Section 3 presents the data and empirical strategy; Section 4 examines the technology roll-out; Section 5 develops the conceptual framework; Section 6 studies the task intervention and the role of managers; and Section 7 assesses the economic importance of worker adjustment frictions.

2. Context

This section describes the bank, the task structure of branch work, and the two sequential changes that underpin our empirical strategy.

2.1. The bank

We partner with one of the largest and longest-established banks in a Latin American country. After more than 50 years of operation, the bank serves roughly 3 million individual, SME, and rural clients and employs over 10,000 workers.

In 2017, the bank launches a mobile-first digital transformation strategy. Existing online services are limited to basic payments and transfers. The new strategy seeks to create a fully digital customer journey, expand access to banking services, raise productivity in routine transactions, and redirect branch workers toward higher-value products. The customer-facing app launches in June 2018. Between 2017 and 2023, the bank facilitates more than 15 million new trades and becomes one of the country’s largest financial institutions by assets. The strategy targets two frictions at once: an external one — customers’ costly access to the branch network — and an internal one — the time cost of producing routine transactions in branches. Account managers remain essential to both margins: they onboard clients onto the app, retain responsibility for advisory-intensive products, and receive performance-based commissions on the products they sell.

The bank also commits to retaining incumbent workers and pairs this commitment with internal training and redeployment. The policy is explicitly communicated to employees:

“The bank operates on the premise that “we do not fire anyone.” We achieve efficiency [...] not by dismissing people, but because they move into another role or through natural turnover.”

This constraint makes within-worker adaptation the central adjustment margin in our setting. Consistent with the policy, we find no aggregate or tenure-specific changes in hiring or separations around either intervention (Figure 5 and Table A.4).

2.2. Workers, managers, and tasks

A typical branch is led by a branch manager and staffed primarily by account managers, who are the focus of our worker-level analysis. Account managers perform most client-facing trades and receive performance-based commissions. We group their work into three categories: simple trades, such as basic account openings and consumer loans; other complex trades, such as investment products and mortgage loans; and SME lending, which requires relationship management and credit evaluation. Throughout, we treat these three categories as bundles of tasks: each product combines routine, analytical, and interpersonal task content in different intensities, and these bundles map directly into the three tasks of the framework in Section 5. Branch managers, by contrast, do not primarily sell products: they supervise

account managers, allocate coaching and training, and are evaluated on branch performance.

Figure 1, Panel (a), maps banking products into O*NET-based task measures following [Autor et al. \(2003\)](#). Relative to simple products, other complex products rely more on analytical and interpersonal skills and less on routine cognitive work; SME lending is the most analytically and interpersonally intensive category.² SME lending also involves a distinct production process. Account managers must source local firms, build relationships with owners and managers, assess business models, cash-flow cycles, and supply-chain exposure, and draw on local information that standardized product features and transaction histories do not capture.

2.3. Technology roll-out and performance redesign

The transformation proceeds through two centrally coordinated interventions. The Tech Roll-out changes how routine transactions are produced and expands operational capacity. The later Task Intervention changes roles and performance evaluation, shaping how that capacity is allocated across tasks. Their sequence lets us distinguish a change in production technology from a change in organizational direction. Figure 1, Panel (b), summarizes this sequence: simple retail products dominate before either intervention, the app expands routine capacity, and the later performance redesign prioritizes SME lending. We describe both interventions and their rationale below.

1. Tech roll-out: introduction of the mobile app. The first intervention introduces a customer-facing mobile app that digitizes account opening, small consumer loans, and credit card applications. For eligible app-based loan applications, it also removes several documentation requirements, including reference letters, employment letters, and asset documentation. These changes make routine transactions substantially faster and less time-intensive for account managers.³ The rollout combines the app with customer onboarding, branch-level support, and active adoption efforts. Branch employees help clients enter the platform, complete initial

²The combined analytical–interpersonal score rises from about 0.23 standard deviations for low-complexity products to roughly 0.4–0.5 for medium-complexity products and to around 0.6 or more for SME lending.

³A branch manager reports: “Opening an account used to take about an hour in the analog process; today we open it in five minutes. [...] Starting the month after we launched the mobile account, we doubled the number of accounts opened each month, and it has remained at that level.”

transactions, and integrate the app into existing service workflows. The intervention changes how routine tasks are performed while leaving performance priorities unchanged. Simple products are available through the app, whereas SME lending and other advisory-intensive products remain branch-based. The app itself diffuses rapidly: nearly all branches adopt it within four months, allowing us to study a rapid technology shock under the old evaluation system.

2. *Performance redesign: reallocation toward SME lending.* The redesign of roles and performance evaluation is more demanding and proceeds gradually over more than two years. The contrast between the speed of technical deployment and the difficulty of organizational redesign is explicit in the bank’s own account: *“Digital adoption itself was not difficult: people are practical, they learn to use the tool and see that it is fast and efficient. [...] What has been harder is transforming people’s roles. [...] That has been a challenge because people become accustomed to what they have done for years, to what works for them, and to where they can earn money most easily.”*

The later redesign includes changes to performance incentives.⁴ As one senior manager explains: *“People gravitate toward the businesses they can close quickly [...] retail banking is ‘sexy,’ but SME banking is paperwork. [...] We have had to work extensively on the incentive system to prevent people from meeting their targets only through the easiest, digital products.”*

Beginning in January 2021, the bank implements what we define as the “Task Intervention”: headquarters asks branch managers to shift attention away from routine volume and toward complex products, with SME lending as the centerpiece. For account managers, SME lending becomes the primary performance priority, while other complex activities are not reprioritized and the explicit monetary compensation schedule remains unchanged. The reform therefore changes how the remaining tasks are bundled and evaluated after automation (Demirer et al., 2026; Gans and Goldfarb, 2026).

Human Resources (HR) staff communicate the new objectives, explain the SME-related KPI change, and deliver branch-level training through in-person visits. We define a branch’s Task Intervention date as the week of this training and explanation. The roll-out is staggered and

⁴Interview evidence also reports centralized changes in tasks rebundling, target setting, training, certification, and branch-level implementation support.

reaches 99% of branches by March 2023.

For intuition, write the account manager’s performance score as:

$$Sc_i = p_R y_{Ri} + p_O y_{Oi} + p_S y_{Si} \quad (1)$$

The Tech Roll-out makes routine tasks easier to scale but leaves the performance weights unchanged. The Task Intervention raises p_S , reduces p_R , and leaves p_O unchanged, increasing the relative return to SME lending while reducing the weight on routine volume.

In contrast, the Task Intervention raises the marginal return to SME effort and lowers the reward to routine volume, while leaving the reward to other complex work unchanged; the explicit base-wage schedule itself remains unchanged. The operational contrast between routine and SME tasks is apparent to account managers themselves: automation transforms routine work but leaves SME lending largely untouched.

The two interventions are separately timed and identified. The app lowers customer and processing costs, expands the client base, and creates operational capacity. Under the old KPI system, workers use most of that capacity to expand routine output. The later performance redesign redirects attention toward SME lending.

The redesign therefore operates on capacity that the new technology creates. Our estimates capture the redesign after technology adoption and do not identify its counterfactual effect without the technology.

3. Data and empirical strategy

Our empirical strategy relies on detailed administrative records that allow us to trace, at high frequency, how branches and workers respond to the introduction of new technology and task assignments. Because the structure and granularity of the data shape both the definition of outcomes and the timing of treatment exposure, we begin by outlining the construction of the datasets before turning to the empirical strategy.

3.1. Data and samples

We combine four data sources: trade-level production records, monthly branch balance sheets, annual personnel files, and survey-based assessments of managers. We link these sources into branch-week and worker-week panels for the two interventions.

Trades and balance-sheet data. Our primary data contain the universe of financial trades completed between 2015 and 2023. Table 1 panels A and B report descriptive statistics at the branch and worker levels across both interventions, while Panel C summarizes product coverage. Each trade is linked to a product, date, branch, and worker. Because the bank evaluates production primarily through trade volume, trade counts are our main output measure; sales values provide a complementary measure of economic value.⁵

Monthly branch balance-sheet data from 2018 to 2022 provide total balances and active accounts. Because these data are aggregated across products, they cannot distinguish simple, other complex, and SME activity.

Personnel records. Annual HR files from 2017 to 2022 provide job titles, tenure, wages, gender, work locations, and entry and exit information. We use them to identify worker roles, track mobility, and determine the timing of worker exits by combining roster dates with the last observed trade.

Manager surveys. The manager-survey data begin in 2020 and combine two sources. First, the bank’s internal manager performance instrument provides assessments of branch managers’ performance by their supervisors, as well as managers’ self-reported alignment with the bank’s values, for 2020–2022. Second, an external subordinate survey, administered from 2020 through 2023, provides workers’ assessments of their managers’ leadership. Table 2 lists the underlying items and index definitions. These measures are collected after the January 2017–June 2018 window used to estimate manager fixed effects and after the Tech Roll-out;

⁵Sales values are unavailable for some product types (e.g., checking accounts and credit cards), and financial gains accrue over time in ways that the data do not directly capture. We therefore cannot isolate product-specific balance-sheet effects at the same granularity or frequency as the other outcomes.

the 2021–2022 waves also overlap the Task Intervention. Because intervention exposure and subsequent branch performance may themselves affect both internal ratings and subordinate assessments, we use these measures only for ex post validation of the manager fixed effects, rather than as independent measures of manager quality.

Branch-week and worker-week panels. The Tech Roll-out sample covers January 2017 through December 2019. The Task Intervention sample covers January 2019 through March 2023, after completion of the app roll-out; the first treated branch is observed in January 2021.

Treatment is assigned at the branch level. We construct branch-week and worker-week panels over these windows. Worker panels include customer-facing employees who complete at least one trade in a branch-week, excluding headquarters and back-office staff. For the Task Intervention, we generally restrict the worker sample to account managers, the employees directly subject to the redesign. Appendix Figure A.10 is an exception: both panels use the full sample of branch workers to summarize aggregate heterogeneity by manager quality.

Outcome variables. Trade counts are our primary output measure, and branch productivity is trades per worker. Worker-level outcomes include total trades, sales values, and task composition across simple, other complex, and SME activities. Trade and sales outcomes are transformed using the inverse hyperbolic sine.

Branch-level outcomes include client acquisition, total balances, active accounts, hiring, exits, and worker transfers across branches or into client-facing work. We date a move into client-facing work at the first observed trade after at least three months of tenure, reducing contamination from onboarding.⁶

3.2. Empirical strategy

We exploit the staggered implementation of each intervention using cohort-specific event studies and difference-in-differences specifications that compare earlier and later-treated branches under the parallel trends assumption. The event studies trace the adjustment around each

⁶Worker exits are identified using the last observed trade week combined with HR separation records.

roll-out, while the static specifications summarize average effects and heterogeneity by worker, manager, and branch characteristics.

Treatment timing and roll-out. The app reaches 99 percent of all branches within approximately four months. The Task Intervention spreads more gradually and reaches 99 percent of eligible branches within two years. Appendix Figure A.1 shows the distribution of treatment timing across branches.⁷

A branch is treated by the Tech Roll-out in the first week it records a product sold through the mobile app; treatment begins in June 2018. For the Task Intervention, a branch is treated in the week when HR staff deliver branch-level training and explain the SME-related KPI change; treatment begins in January 2021. Because training is concentrated in the treatment week, the contemporaneous SME coefficient includes a training-week response, which we discuss with Figure 7.

Both interventions are coordinated by headquarters, with HR determining roll-out timing to manage implementation. Appendix Table A.1 collects characteristics for branches and workers across the implementation window. Early treated branches are smaller and less likely to be urban. While branches vary in their characteristics and product composition with the timing of the interventions, workers are observationally identical across groups. Our event-study specifications include branch fixed effects in branch-level analyses and worker fixed effects in worker-level analyses, which account for these differences in levels to the extent that they are time-invariant, and week fixed effects, which absorb shocks common to all units. Our identification strategy therefore does not require earlier and later-treated branches to be identical before treatment. Rather, it relies on the assumption that, absent the interventions, their outcomes would have followed parallel trends, and that branches did not adjust in anticipation of treatment. We assess the plausibility of these assumptions by examining the pre-treatment coefficients in the event-study specifications.

⁷ All 449 branches receive access to the banking app, whereas 355 receive the Task Intervention, which targets SME lending. SME lending is not offered by small, retail-oriented branches without an entrepreneurial client base, although all branches conduct both simple and complex retail trades.

Event-study design. For each intervention, we estimate dynamic effects at the branch-week or worker-week level. For expositional simplicity, write:

$$y_{it} = \sum_{s \in \mathcal{S}_j, s \neq -1} \beta_{j,s} D_{it}^{j,s} + \alpha_i + \zeta_t + \epsilon_{it}, \quad (2)$$

where i indexes a branch or worker, and $D_{it}^{j,s}$ indicates that unit i is s weeks from its branch’s treatment date for intervention j . The omitted period is $s = -1$; α_i and ζ_t denote unit and week fixed effects. We bin the endpoints of each event window. The Tech Roll-out window runs from eight weeks before to twenty-six weeks after treatment, and the Task Intervention window runs from eight weeks before to fifty-six weeks after treatment.⁸

We complement the dynamic estimates with static two-way fixed-effects models that replace the event-time indicators with a post-treatment indicator. To study heterogeneity, we interact the post-treatment indicator with baseline worker, manager, and branch characteristics. These specifications summarize average post-treatment effects and characterize the tenure and manager-quality analyses. We cluster standard errors at the branch level.

4. Technology Roll-out

This section examines the impact of the Tech Roll-out. We present branch-level evidence first — client activity, output, and workforce stability — and then turn to task composition and how the response varies across workers. Because employment and worker transfers across branches remain stable, the patterns we document reflect adjustment within the incumbent workforce.

⁸Because adoption is staggered and there are no never-treated units, TWFE coefficients need not recover a convex average of treatment effects when effects vary across cohorts or over time. We assess this empirically with the diagnostics proposed by [De Chaisemartin and d’Haultfoeuille \(2020\)](#). We decompose the static TWFE coefficient into a weighted sum of cohort-week treatment effects. If effects are allowed to vary over time, negative weights sum to -0.36 (branch) and -0.33 (worker) for the Tech Roll-out, and to -0.13 (branch) and -0.24 (worker) for the Task Intervention. If instead each cohort’s effect is constant over time, all weights are non-negative except for the Tech Roll-out at the worker level, where they sum to -0.07 . Even though contamination is small, we additionally estimate the main event studies with the interaction-weighted estimator of [Sun and Abraham \(2021\)](#).

4.1. Client activity, output, and task composition

Client activity and operational capacity. Figure 2 shows that realized client activity rises alongside branch activity. Total and first-time clients increase, and trades rise among both recent and established clients. Client-base growth differs across tasks, in particular, the number of new SME-loan clients decreases (panel 2e). This pattern is consistent with workers shifting effort away from acquiring new SME clients toward meeting app-generated demand for simple products, and relying on existing client networks to originate any new SME loans. Figure 3 shows that total balances and active accounts also rise after the app roll-out.⁹ Together, these outcomes indicate that the expansion involves both more clients and greater activity within existing client relationships.

Output and productivity. Figure 4, Panels (a) and (b), shows that worker output rises steadily after app adoption in both the full sample and the sample of account managers, with no evidence of differential pre-trends. The effect builds over time, consistent with progressive adoption and integration of the app into branch workflows. The corresponding static estimate is a 19.73 percent increase in worker-level total trades (Table A.2, Panel A).¹⁰

Adjustment within the incumbent workforce. Figure 5, Panels (a)–(b), shows no changes in hiring and separations. Appendix Figure A.4, Panels (a)–(b), further shows no variation in transfers across branches, or transitions into customer-facing roles. The app therefore changes output through incumbent workers rather than workforce replacement.¹¹

The app pushes the bank’s production frontier outward by expanding routine production, but it redirects little work toward the higher-value complex tasks. Figure 6 shows that simple trades rise sharply after adoption in both the full worker sample and the account-manager

⁹Recent clients are those observed within six months of their first trade; established clients are all others. Balances and accounts are only available at the branch level and they aggregate client portfolios managed by branch teams.

¹⁰The pattern is similar at the branch level, where total trades rise by 14.47 percent, and in the account-manager sample, where the event-study response is larger (Appendix Figure A.2 and Figure 4b). Percentage changes are computed from the IHS coefficients using the Bellemare and Wichman (2020) formula, evaluated at the sample mean of the dependent variable.

¹¹Pre-intervention turnover averages 0.011 exits and 0.011 hirings per branch-week, equivalent to roughly one separation or hire per branch every two years.

sample (Panels (a) and (c)), while complex trades increase much less (Panels (b) and (d)). The similarity across the two samples indicates that the pattern is not driven by occupational composition. Appendix Figure A.5 confirms that the complex-to-simple ratio falls, the share of SME loans in complex trades remains broadly unchanged, and specialization increases.

4.2. Heterogeneity across workers

Figure 8 (top row) illustrates the worker-level heterogeneity analyzed in more detail Section 6: during the Tech Roll-out the increase in routine output declines steeply with tenure, while complex work rises only among low-tenure workers.¹²

The uneven task upgrading of low-tenure workers is reflected in compensation (Figure 10a). The figure reports non-parametric worker-level TWFE estimates of annual log wages by tenure group and shows a clear, monotonic decline in wage responses with tenure. Despite this gradient, all workers experience an increase in wages driven by the productivity expansion. Appendix Figure A.7 further documents that the negative tenure gradient persists within age and baseline-ability groups and across workers with different exposure to high-tenure peers.

5. Conceptual framework

The evidence in the previous section raises two questions: why does the new capacity flow into the routine products the app automates rather than into the complex products the bank values, and why is adjustment, where it occurs, concentrated among low-tenure workers? One interpretation is that the app changes production opportunities and expands the client base without changing the task-specific private returns workers face. We formalize this intuition in a simple three-tier model of the firm. The model first interprets the technology phase and then derives predictions on how a change in task incentives—the “Task Intervention”

¹²Figure 8 compares long-run responses to the Tech Roll-out across the four tenure groups. Here, “long-run” denotes the tenure-group-specific coefficient on an indicator equal to one from week 13 onward after the intervention; weeks 0–12 are captured by a separate short-run indicator, and the full pre-treatment period is the reference category.

in our terminology—affects the direction of reallocation across products, its incidence across workers, and the value of managerial coaching.

5.1. Model set-up

Setting. Consider a bank with a fixed workforce of account managers. Workers choose effort to maximize performance pay net of effort costs. Branch managers allocate coaching to maximize branch surplus. We take the technology state m and the performance-weight vectors before and after the Task Intervention, p^0 and p^1 , as given, abstracting from headquarters' choice of technology and performance weights.

Worker i allocates effort $e_i = (e_{Ri}, e_{Oi}, e_{Si})$ across simple products, measured as simple trades (R), other complex products (O), and SME lending (S). Effort is the time and intensity devoted to a task; output $y_i = (y_{Ri}, y_{Oi}, y_{Si})$ is the number of trades or loans produced. We normalize $y_{Oi} = e_{Oi}$ and $y_{Si} = e_{Si}$ and write $y_{Ri} = D_i(m)g(A_R(m, \theta_i, \tau_i^R)e_{Ri}/D_i(m))$, so effort and output need not move one-for-one. Here, A_R is routine task productivity, $D_i(m)$ is worker i 's client demand pool, and g is increasing and concave: as the pool gets exhausted, additional routine effort produces fewer additional trades. Worker i has adaptation cost θ_i , which we proxy with tenure in the data.

Within each regime, branch managers and workers take the performance-weight vector $p = (p_R, p_O, p_S)$ as given. Branch managers observe worker types and allocate analytical and routine coaching, τ_i^A and τ_i^R ; branch manager quality determines effective coaching capacity T_M . Pay is

$$W_i = \bar{W}_i + \bar{\lambda}(p_R y_{Ri} + p_O y_{Oi} + p_S y_{Si})$$

Here, W_i is worker i 's total pay, $\bar{W}_i$ is fixed pay, and $\bar{\lambda} > 0$ is the common sensitivity of pay to the weighted performance score. The weights p_R , p_O , and p_S apply to simple, other complex, and SME output, respectively. Effort costs are defined as:

$$C_i = \frac{e_{Ri}^2}{2} + \frac{c_O(\theta_i, \tau_i^A)e_{Oi}^2}{2} + \frac{\tilde{c}(\theta_i, \tau_i^A)e_{Si}^2}{2} + \frac{\chi}{2}(e_{Ri} + e_{Oi} + e_{Si})^2$$

Here, C_i is worker i 's total effort cost. The coefficients $c_O(\theta_i, \tau_i^A)$ and $\tilde{c}(\theta_i, \tau_i^A)$ are the effort-

cost coefficients for other complex work and SME lending, respectively, while $\chi \geq 0$ captures time trade-offs across tasks within the shared workday. The coefficient on routine effort is normalized to one. When $\chi > 0$, effort in one task raises the marginal cost of effort in the others. Analytical coaching lowers c_O and $\tilde{c}$ by helping workers interpret client information, match clients to products, and evaluate SME borrowers; routine coaching raises A_R and makes simple trades easier to complete.

Branch surplus is the sum across workers of $\alpha_R y_{Ri} + \beta_O y_{Oi} + \beta_S(m) y_{Si} - C_i$. Here, α_R , β_O , and $\beta_S(m)$ are the branch values per unit of simple, other complex, and SME output, respectively. We let $\beta_S(m) = \beta_{S0} + \gamma m$, where β_{S0} is the baseline component of the branch value of SME lending and $\gamma > 0$ captures how that value rises with the technology. This formulation reflects the bank's stated strategy of using the app's expansion of the client base and creation of additional capacity to support growth in SME lending; see Section 2.3.

Regimes and solution. We study the continuation game under the three regimes observed in the data: old technology with p^0 , new technology with p^0 , and new technology with p^1 . The timing of the shift from p^0 to p^1 is taken from the data. Section 2.3 documents the institutional basis for this timing: redesigning roles and performance evaluation is substantially slower and more demanding than deploying the app. Within each regime, branch managers allocate coaching before workers choose effort. We solve this continuation game by backward induction. The main text studies the technology, performance-weight, and coaching margins separately; Appendix B presents the full setup and solution.

5.2. Assumptions

Assumption 1 (Technology and simple-product demand). Technology m raises routine task productivity A_R and expands worker i 's client demand pool $D_i(m)$. The demand pool is finite, and the productivity gain is larger for workers with lower adaptation cost θ_i .

Assumption 2 (Multitasking and limited time). The shared-workday parameter χ in the cost of effort function, C_i , is nonnegative. When $\chi > 0$, tasks compete for time; when $\chi = 0$, task-specific effort costs are separable.

Assumption 3 (Adaptation and coaching). Higher adaptation cost θ_i raises the effort-cost coefficients for other complex work, $c_O(\theta_i, \tau_i^A)$, and SME lending, $\tilde{c}(\theta_i, \tau_i^A)$. Analytical coaching τ_i^A lowers both cost coefficients and is more effective for workers with lower θ_i , while routine coaching τ_i^R raises routine task productivity A_R .

Assumption 4 (Old SME reward). Before the redesign, the old SME performance weight satisfies $p_S^0 < p_S^{\text{align}}(m) := \beta_S(m)/\bar{\lambda}$. Here, $\beta_S(m)$ is the branch value per unit of SME output under technology m , and $\bar{\lambda}$ is the common sensitivity of pay to the weighted performance score. Thus, $p_S^{\text{align}}(m)$ is the SME weight that equates the worker's private marginal return, $\bar{\lambda}p_S$, with the branch value of SME output, $\beta_S(m)$. Because the app raises $\beta_S(m)$ while p_S^0 remains fixed, this wedge widens after adoption.

Appendix B gives the full production and cost functions, the optimization problems, the game sequence, the precise conditions, and the proofs.

5.3. Implications

Technology under old performance priorities. Before the Task Intervention, the bank uses $p^0 = (p_R^0, p_O^0, p_S^0)$. The app changes the technology m while leaving these weights unchanged. Assumption 1 implies that the app raises routine task productivity A_R and expands worker i 's client demand pool $D_i(m)$. Simple-product output y_{Ri} therefore rises. Whether routine effort e_{Ri} also rises depends on whether simple-product demand has been exhausted.

The performance weights determine the direct private return to each task and, together with effort costs and the shared-workday parameter χ , the level of SME effort e_{Si} . The app leaves the direct private return to SME effort, $\bar{\lambda}p_S^0$, unchanged but raises the branch value per unit of SME output, $\beta_S(m)$. The surplus-aligned SME weight, $p_S^{\text{align}}(m) = \beta_S(m)/\bar{\lambda}$, therefore rises. This widens both the wedge $p_S^{\text{align}}(m) - p_S^0$ and the gap between the SME effort induced by p_S^0 and the level aligned with branch surplus. This formulation reflects the institutional strategy described in Section 2.3: the app is intended to expand the client base and free branch capacity for complex products; the later performance redesign seeks to direct that capacity toward SME lending.

We denote this the *client demand pool channel*. When $\chi > 0$, two cases arise depending on

whether the simple-product demand pool $D_i(m)$ has been exhausted.¹³ If simple-product demand is not exhausted, the expanded demand pool absorbs the time saved by the app, simple work occupies more of the workday, and complex work falls. If simple-product demand is largely exhausted, the more time on simple products generates few additional trades, so the app frees time for SME lending and other complex work.

Among workers facing comparable simple-product demand pools $D_i(m)$, the model predicts that workers with lower adaptation cost θ_i exhaust simple-product demand sooner because their routine task productivity A_R rises more. They therefore move toward complex work sooner, while workers with higher θ_i are more likely to face simple-product demand that is not yet exhausted. Earlier exhaustion can coincide with a larger increase in simple output y_{Ri} . The app expands $D_i(m)$ for all workers, and the gain in A_R is larger for low- θ_i workers. They can therefore produce more simple output y_{Ri} while reaching the flatter part of g sooner.

Performance redesign. The Task Intervention changes the performance weights from $p^0 = (p_R^0, p_O, p_S^0)$ to $p^1 = (p_R^1, p_O, p_S^1)$, where $p_R^1 < p_R^0$ and $p_S^1 > p_S^0$, while technology remains fixed. The higher SME weight raises workers' direct private return to SME lending, allowing them to shift toward it before exhausting simple-product demand.

Let $\Delta p_S := p_S^1 - p_S^0$ denote the increase in the SME performance weight and $\Delta y_{Si} := y_{Si}^1 - y_{Si}^0$ denote worker i 's change in SME output. When tasks do not compete for time ($\chi = 0$),

$$\frac{\Delta y_{Si}}{y_{Si}^0} = \frac{\Delta p_S}{p_S^0} \quad (3)$$

for every worker. The proportional response is the same for every worker, while the number of additional loans can differ because workers begin at different output levels. The tenure gap in the proportional SME response, $\Delta y_{Si}/y_{Si}^0$, can therefore close. The weight on other complex work, p_O , does not change, and the average response of other complex output is ambiguous: the decline from p_R^0 to p_R^1 frees time for other complex work, while the increase from p_S^0 to p_S^1 draws time away from it.

¹³When $\chi = 0$, SME lending and other complex work do not respond directly to the app.

Wages. Because fixed pay $\bar{W}_i$ is unchanged, the change in worker i 's pay, $\Delta W_i := W_i^1 - W_i^0$, reflects two margins, both scaled by the performance-pay sensitivity $\bar{\lambda}$: the change from p^0 to p^1 applied to the worker's pre-redesign task portfolio y_i^0 , and the output response $y_i^1 - y_i^0$ evaluated under the new weights p^1 . Workers who shift toward SME lending and other complex work receive larger gains. Workers who continue to do mostly simple work may lose variable pay because the simple-output weight falls from p_R^0 to p_R^1 . Remark A.2 in Appendix B gives the exact decomposition.

Branch Managers. Assumptions 3 and 4 imply a complementarity between the performance regime and branch manager quality; branch manager quality enters the model through effective coaching capacity T_M . Over the relevant range, raising the SME weight from p_S^0 to p_S^1 increases the return to analytical coaching τ_i^A . When the marginal unit of coaching is used on the analytical margin, additional effective coaching capacity is therefore more valuable under p^1 than under p^0 . Because analytical coaching lowers the effort-cost coefficients $c_O(\theta_i, \tau_i^A)$ and $\tilde{c}(\theta_i, \tau_i^A)$ more for workers with lower adaptation cost θ_i , higher-quality branch managers direct analytical coaching toward these workers. The same coaching can also raise other complex work by lowering $c_O(\theta_i, \tau_i^A)$. This complementarity underlies Prediction 5 and concerns coaching content; it does not require differences in total coaching hours.

Predictions. Five empirical predictions organize the empirical analysis.

Prediction 1. (*Technology and simple-product output*). Under old performance priorities p^0 , the app raises routine task productivity A_R and expands worker i 's client demand pool $D_i(m)$, so simple-product output y_{Ri} rises for every worker. Among workers facing comparable $D_i(m)$, workers with lower adaptation cost θ_i , proxied in the data by lower tenure, adapt to the app more quickly and therefore exhaust simple-product demand and move toward complex work earlier.

Prediction 2. (*Client demand pool channel*). Before the Task Intervention, other complex and SME output, y_{Oi} and y_{Si} , can respond to the app through the client demand pool $D_i(m)$ only when the shared-workday parameter satisfies $\chi > 0$. Complex

work falls when simple-product demand is not exhausted and can rise once demand is exhausted; workers with lower adaptation cost θ_i make the latter shift more readily.

- Prediction 3. (*Performance regime*). When the SME performance weight rises from p_S^0 to p_S^1 , SME output y_{Si} rises across workers. When there are no time trade-offs ($\chi = 0$), every worker has the same proportional response, $\Delta y_{Si}/y_{Si}^0 = \Delta p_S/p_S^0$, although the number of additional loans can differ. The tenure gap can close in SME lending while persisting in other complex work.
- Prediction 4. (*Wage incidence*). The change in worker i 's pay, $\Delta W_i := W_i^1 - W_i^0$, reflects the change in performance weights from p^0 to p^1 and the worker's output response $y_i^1 - y_i^0$, with both components scaled by the performance-pay sensitivity $\bar{\lambda}$. Workers who adjust their output more in response to the new weights gain more; those who do not adjust can lose variable pay.
- Prediction 5. (*Performance regime–manager quality complementarity*). Over the relevant range, raising the SME performance weight from p_S^0 to p_S^1 increases the return to analytical coaching τ_i^A . Effective coaching capacity T_M , which depends on branch manager quality in the model, is therefore more valuable under p^1 than under p^0 when marginal coaching is used on the analytical margin. Higher-quality branch managers should have larger effects after the redesign, especially for workers with lower adaptation cost θ_i . The same analytical coaching can also increase other complex work.

Section 4 directly examines Prediction 1's implications for simple and complex output: simple-product output rises after the app, with the largest increase among low-tenure workers (Figure 8). Regarding Prediction 2, we do not observe latent demand or whether simple-product demand is exhausted, precluding us from classifying workers by demand status or directly testing its mechanism. However, two patterns are consistent with Prediction 2: total and first-time clients rise after the app (Figure 2), and other complex and SME work rise among low-tenure workers but change little among workers with 12 or more years of tenure

(Figure 8). Section 6 examines Predictions 3–5.¹⁴

6. Task Intervention

The Task Intervention changes what account managers are rewarded for, holding production technology fixed. This section examines how work moves toward the newly incentivized task, which workers move, the margins on which they adjust, and who captures the gains.

6.1. Task reallocation

Figure 4, Panel (c), and Table A.2, Panel B, show that total output rises at both the worker and branch level. Figure 7 shows a large increase in SME trades in the HR training week followed by a persistent rise thereafter (Panel 7c), while other complex trades change much less (Panel 7b). The response is therefore concentrated on the task whose performance weight increases. The unusually large contemporaneous SME coefficient reflects the implementation week: HR staff train branch employees and explain the KPI change during that week. The subsequent coefficients remain positive, showing that SME activity stays elevated beyond the training week.¹⁵

Adjustment within existing jobs. Figure 3, Panels 3c and 3d, shows no detectable change in total balances or active accounts. Because the balance sheet data are aggregated across products, these panels cannot identify SME-specific portfolio effects. Figure 5, Panels (c)–(d), and Appendix Figure A.4, Panels (c)–(d), show no changes in hiring, exits, branch transfers, or moves into customer-facing roles. The reform therefore reallocates activity within existing jobs and targets a product that remains a small share of the bank’s aggregate portfolio.

Overall, the redesign of performance incentives shifts the composition of complex work toward SME lending without generating uniform growth across all complex tasks, as sum-

¹⁴Prediction 1 and Prediction 2 follow from Proposition A.1 and Corollaries A.1–A.2 in Appendix B. Prediction 3 follows from Proposition A.2, Prediction 4 from Remark A.2, and Prediction 5 from Proposition A.3 and Corollary A.3.

¹⁵Appendix table A.2 presents static TWFE estimates with and without observations from the training week, showcasing almost no difference between the two specifications. For consistency and transparency across interventions, we keep all observations in our main specifications.

marized in Figure 7. Panels 7e and 7f show that the targeted shift is also visible in sales. SME sales increase markedly and persistently, while sales of other complex products rise less and the corresponding estimates are less precise.¹⁶ Appendix Figure A.5 shows that the share of SME loans to complex trades and task specialization rises (panels A.5d and A.5f), while the complex-to-simple ratio remains broadly flat (A.5b).

6.2. Tenure and task reallocation

We use worker tenure as a proxy for adaptation cost arising from accumulated routines, established work practices, and task-specific knowledge. In our sample, median tenure is either seven or eight years across the two interventions, varying from 0 to 30 years.¹⁷

Figure 9 compares long-run responses across the four tenure groups. Here, “long-run” denotes the tenure-group-specific coefficient on an indicator equal to one from week 31 onward after the intervention; weeks 0–30 are captured by a separate short-run indicator, and the full pre-treatment period is the reference category.

During the Tech Roll-out, responses to both simple and complex work decline monotonically with tenure, with high-tenure workers showing no change in complex trades (Figure 8). Following the Task Intervention, the performance redesign closes the tenure gap on the directly targeted task but not in broader complex work (Figure 9): SME lending responds similarly across tenure groups, whereas the response in other complex work remains concentrated among low-tenure workers. This contrast is consistent with the model’s Prediction 3.

The corresponding sales results in Figure 9 show a similar pattern. SME sales responses do not vary across tenure groups, whereas sales of other complex products increase only for lower-tenure workers. This pattern weighs against the interpretation that high-tenure workers substitute toward higher-value trades at the expense of trade volume.

Appendix Table A.3 shows that these patterns are robust to restricting the sample to urban

¹⁶Publicly available financial reports from the bank and its rating agency show a stable NPL rate of about 3 percent across commercial, consumer, and mortgage loans around the intervention (2021–2022), providing no indication of a contemporaneous deterioration in loan quality.

¹⁷Tenure is measured at the treatment date and held fixed throughout the panel. The four non-parametric groups are 0–3, 4–7, 8–11, and 12+ years. In the binary age specifications, age is also measured at treatment and split at the sample median.

branches, indicating that the tenure gradient is not driven by differences between urban and rural locations.

Wages. The redesign changes both workers' task allocation and the outputs rewarded by their pay. Prediction 4 implies that wage gains should be concentrated among workers whose task portfolios change the most—in our setting, the low-tenure workers. Figure 10, Panel (b), shows a clear decline in wage responses with tenure. After the Task Intervention, workers with 0–3 years of tenure receive the largest wage gains; the response becomes smaller across the remaining groups and is close to zero or negative among workers with eight or more years of tenure.

Panel (c) shows the corresponding wage distributions. By 2022, the upper tail expands most visibly among low tenure workers, while the distribution for high tenure workers shifts more uniformly. This pattern is consistent with the largest wage gains accruing to the workers who adjust their task allocation the most.

The task and wage results point to a common adjustment margin: low-tenure workers reallocate more toward other complex work and receive larger wage gains, while long-tenured workers show smaller changes on both margins. Heterogeneous adaptation costs thus shape both behavioral responses and the distribution of returns from organizational change.

6.3. What does tenure capture?

We interpret tenure as a proxy for the extent to which workers have established routines, task-specific knowledge, and client relationships. Interview evidence provides institutional context for the relationship channel. Commercial-transformation staff report that *“ninety-seven percent of the customers who self-serve today were cultivated by the branch network”* and that employees continue to manage many clients through one-to-one contact, including through WhatsApp. This account is consistent with digital activity continuing to rely on relationships built by branch employees.

The gradient could, however, reflect other worker attributes that correlate with tenure. We organize the main alternative explanations into three groups: age and remaining career hori-

zon; ability, specialization, and work environment; and effort and employment risk.

Age and remaining career horizon. Age is correlated with tenure and may reduce the horizon over which workers recover the cost of learning new tasks. Figure A.9 shows that adjustment also declines with age. However, Figures A.7 and A.8 show that the tenure gradient remains within both younger and older groups. Age and remaining career horizon thus do not account for the pattern on their own.

Ability, specialization, and work environment. High-tenure workers are more specialized in complex activities at baseline, which could indicate comparative advantage in those tasks. Yet Figures 8b and 9b show that responses in non-incentivized complex work are weaker for high-tenure workers under both interventions. If comparative advantage were the main explanation, these workers would be expected to respond at least as strongly as others.

The sales evidence points in the same direction since the smaller quantity response among high-tenure workers is not offset by a larger response on the value margin. Figures 8e and 9e show that gains in sales associated with other complex work are largest among low-tenure workers.

Selection on observable baseline ability also does not account for the gradient. Figures A.7 and A.8 show that adjustment declines with tenure among workers both below and above the median of baseline ability.

Another related possibility is that lower-tenure workers have steeper underlying learning curves, generating faster growth in productivity and wages independently of the interventions. However, short term coefficients of the heterogeneous responses by tenure mirror the magnitude and ordering of the long term coefficients reported in Figures 8 and 9. Moreover, this explanation is difficult to reconcile with the closing of the tenure gap in SME lending following the Task Intervention. If differences in underlying learning trajectories were driving the tenure gradient, we would expect low-tenure workers to continue responding more strongly.

The pattern is similarly robust to the work environment. The tenure gradient remains visible in branches with both low and high shares of high-tenure peers, and worker responses

are similar under lower- and higher-tenure managers.

Effort and employment risk. Figure 10d shows that days worked increase across all tenure groups after the Task Intervention, with relatively small differences by tenure. The evidence on employment risk points to the same conclusion. Table A.4 shows no tenure-specific changes in hiring or separations around either intervention. The observed gradient is therefore not generated by differential exit, replacement, or restructuring across tenure groups.¹⁸

Alternative proxies for adaptation costs. We complement tenure with an alternative proxy for worker adaptation costs: baseline worker productivity in other complex trades. Figure A.6 shows how workers who are more productive in other complex work adjust after the Tech Roll-out. Tenure gradients for simple trades are flatter for workers with higher baseline complex productivity (A.6b). For complex work, SME loans account for most of the difference: more productive workers increase their output, whereas less productive workers exhibit a negative tenure gradient and reduce their overall production. These results suggest that workers who are more productive in the targeted tasks adapt more uniformly across tenure groups and shift immediately toward high-value trades such as SME loans.

6.4. Managers and task reallocation

We examine whether branch managers help workers adapt to the new priorities. According to Prediction 5, manager quality should matter more after the performance redesign raises the SME weight giving branch managers greater scope to influence task reallocation. Interview evidence illustrates this role. As an HR leader explains: “*The manager accompanies the employee through the process. [...] What are you assigning this person to do, and how is the person progressing? [...] Your leadership carries a major responsibility to develop them.*”

Measuring and validating manager quality. We estimate manager value added (ψ_j) from pre-intervention branch productivity using:

¹⁸Termination costs increase linearly with tenure in this setting: employers must compensate workers for the equivalent of 20 days of work for every additional year with the company.

$$\text{asinh}(\text{branch productivity})_{bjt} = \psi_j + \eta_b + \tau_t + \varepsilon_{bjt} \quad (4)$$

where branch productivity is total trades per employee; b , j , and t index branches, managers, and weeks. Manager fixed effects capture persistent productivity differences as managers rotate across branches, conditional on branch and week fixed effects. Their estimation relies on manager mobility across branches and the largest connected set covers more than 90% of observations.¹⁹ We use ψ_j as an empirical proxy for the manager-quality component of effective coaching capacity T_M in the model.

Table 2 provides suggestive validation of our manager-quality measure against independently collected manager survey and performance ratings. These measures were collected in 2020–2022, after the 2017–2018 period used to estimate manager fixed effects, and may therefore reflect subsequent interventions and changes in branch performance. For this reason, we do not use them in the treatment-effect analysis. Instead, we examine whether manager fixed effects estimated before the interventions predict these later measures of managerial quality.

Manager fixed effects are essentially unrelated to age, tenure, gender, or fixed pay, but are positively associated with internal performance measures and with subordinate ratings of teamwork, trust, embracing change, and inspiration. They also show no association with manager self-reported alignment with company values. Overall, the measure appears to capture persistent leadership and implementation capacity rather than observable demographics or general values alignment.

Manager quality and worker tenure. Appendix Figure A.10, which uses the full sample of branch workers in both panels, shows a modest positive gradient in total output across manager-productivity quartiles under both interventions. This aggregate pattern does not reveal whether branch managers redirect workers toward newly prioritized tasks more effectively or whether this ability varies with worker tenure. First, Appendix Figure A.11 shows that the negative tenure gradient after the Tech Roll-out is nearly identical under low- and high-quality managers. When performance priorities remain unchanged, manager quality does not appear to

¹⁹We exclude 39 managers who never move and demean fixed effects within the connected set.

affect workers' responses to the technological change. Instead, after the performance redesign Figure 11 shows that more productive branch managers are associated with larger increases in other complex activity, primarily among low-tenure workers, while the corresponding difference is small among high-tenure workers.

Managerial support therefore appears to matter more after priorities are aligned. It also accentuates the underlying adaptation gap. Looking at sales, the gains in total sales and other complex sales are largest among low-tenure workers supervised by more productive branch managers, whereas SME sales are more similar across worker-tenure and manager-quality groups. This indicates that manager quality adds the most value on the complex margin that receives no direct KPI increase. These results are consistent with Prediction 5 whereby the return to effective coaching capacity T_M rises with the SME weight, so differences in manager quality matter the most after the redesign raises that weight (Proposition A.3 in Appendix B).

Training composition. One mechanism through which managers may support adjustment is targeted skill investment. We examine whether the Task Intervention changes the amount and content of training differently across worker-tenure and manager-quality groups. Using module titles and descriptions, we manually classify training into analytical, interpersonal, and routine cognitive content, following the same skill framework used for products in Section 2. Training records contain sessions, hours, and participation indicators; Figure 12 uses weekly training hours.²⁰

Figure 12 provides little evidence of systematic differences in total training hours. The composition of training changes more clearly: under more productive branch managers, low-tenure workers receive relatively more analytical training, while high-tenure workers receive relatively more routine cognitive training. Interpersonal training shows smaller and less systematic differences.

²⁰Routine cognitive training includes compliance, occupational safety, and digital-account procedures and averages 0.274 hours per worker-week, with participation in 18.5% of worker-weeks. Analytical training includes data analysis, database management, and SME advisory content and averages 0.155 hours, with participation in 9.3% of worker-weeks. Interpersonal training includes communication, customer service, and public speaking and averages 0.057 hours, with participation in 4.2% of worker-weeks. Creative training averages 0.016 hours and 3.3% participation; language and physical training are rare. Training is one observable implementation margin; we do not directly observe informal coaching, routines, or coordination.

This pattern is consistent with targeted managerial support and with the allocation rule in Corollary A.3 in Appendix B. In particular, more productive branch managers direct analytical investment toward workers who are better able to adjust and steer harder-to-move workers toward routine-oriented content, changing the composition of coaching without changing its total amount. Thus, training composition provides one observable implementation channel consistent with the selective manager effects.

7. Economic value of reorganization

We conclude the paper with an assessment of the impact of the two interventions on the bank's profitability. So far, we have focused on trade volumes and task complexity because they capture within-job reallocation and are available at the worker-task level. We complement the analysis with financial reports compiled by the bank and its credit rating agency to understand the aggregate impact of the changes on the bank's profits. Finally, we combine the empirical estimates with the model's interpretation to estimate the aggregate cost of sluggish reallocation toward complex work.

7.1. Profitability assessment

We draw on the bank's own financial statements and the reports of its credit rating agency (Fitch Ratings) to examine whether the reallocation patterns documented above leave a discernible mark on profitability, and whether any such mark can plausibly be attributed to the two interventions.

Aggregate profitability rises over the sample period. Net profit increases by COP 333 billion between 2017 and 2022, but the path is not linear: profits plateau around COP 1,200 billion between 2017 and 2021 before jumping to over COP 1,600 billion by the end of 2022.

Net interest margins – our proxy for the profitability of individual loan products – remain stable above 6% throughout 2020–2022, indicating that this later acceleration in profits is driven by volume rather than by a repricing of the bank's products.²¹ Throughout the period,

²¹Loan issuance grows from 11.2% to 22.1% year-on-year between 2021 and 2022 alone.

the bank's performance tracks the sector closely, both in profitability and in asset quality.

Defaults are also fairly stable over time. The share of non-performing loans more than 90 days past due (NPL) rises from 2.8% in 2017 to 3.9% in 2018, before falling back to 3.13% by 2022. However, the bank's credit rating agency attributes the 2018 increase to macroeconomic conditions bearing on large commercial borrowers, not to the quality of lending originated at the branch level.²²

Between 2017 and 2022, the bank's retail franchise expands markedly. Its customer base more than doubles between late 2018 and 2022, and consumer lending grows faster than both the commercial and mortgage book in 2022. These expansions are documented both in our analysis and in the bank's financial reports.

7.2. Counterfactuals

We conduct a simple counterfactual exercise that estimates how adjustment would differ if all workers had the same adjustment cost (θ) as low-tenure workers. Our starting point is the baseline specification in Figure 8, which estimates trade volumes for each outcome as a function of the Tech Roll-out, allowing the response to vary across four tenure groups (0–3, 4–7, 8–11, and 12+ years) and across two post-roll-out horizons: a short-run window (within three months of roll-out) and a long-run window (three months and beyond). We then predict trade flows using this empirical specification and compare the resulting predictions with the data to assess model fit.²³

From this baseline specification we construct two sets of predicted trade levels for the long-run post-roll-out period. The first assigns each worker her own estimated tenure-group response, reproducing the model's fit to the observed pattern of reallocation. The second imposes the estimated low-tenure response on every worker, regardless of her actual tenure, asking how much output each worker would have generated had she faced the lower adaptation cost.²⁴ We sum the resulting worker-level differences by tenure group and in aggregate

²²In 2018, one of the bank's largest clients, a country-wide infrastructure company, defaults outright, leaving the bank directly exposed.

²³In Figure 8 and in this analysis, only the long-term response is displayed. All estimated models include worker and week fixed effects.

²⁴For both predictions, we further impose that time effects are constant and set to May 2018, the final common

to obtain the total output foregone because of slower reallocation among higher-tenure workers. To express this magnitude in economic terms, we multiply the foregone trade counts for each loan product by that product’s average sales value per loan.

We make two main changes to the specification used in Figure 8. First, we focus on loan products for which we observe the transaction amount in pesos and exclude other products (e.g., accounts, insurance products, and bonds). Loan products are grouped into simple loans (available on the banking app throughout the Tech Roll-out), and all complex loans (e.g., revolving loans, mortgages, SME loans). Second, we use a Poisson pseudo-maximum-likelihood (PPML) model with the outcome in levels. We prefer this specification to a linear model with the outcome in inverse hyperbolic sine because it lets us aggregate fitted values and compute differences directly in levels, rather than in a rescaled transformation of the outcome. This specification also accommodates the substantial share of worker-week observations with zero trades in these narrower product categories.²⁵

Figure 13 summarizes the counterfactual implications of adaptation costs for the number of loans originated by workers. Under the counterfactual without the tenure-related adaptation friction, the decline in predicted loan counts with tenure disappears. Under the low-adaptation-cost counterfactual, high-tenure workers would originate as many simple loans as low-tenure workers (panel 13a) and even more complex loans (panel 13b). The baseline predicts weekly loan counts falling from 0.35 to 0.08 across tenure groups, whereas the simulation predicts them rising from 0.35 to 0.67. The counterfactual separates experience-related

pre-treatment period for the banking app. This restriction removes changes in trade volume driven by calendar-time factors unrelated to θ (seasonality, aggregate demand conditions, or other shocks common to all workers) that would otherwise be conflated with the tenure-driven adjustment effect.

²⁵One limitation of the PPML model is that it excludes workers who never complete a trade in a given category before or after the Tech Roll-out. This mainly affects the analysis of SME loans, which are relatively rare at the time. We therefore group all complex loans together to reduce the resulting information loss. Specifically, we estimate the following equation with PPML:

$$\mathbb{E}[y_{it} \mid i, t] = \exp \left(\sum_{k=1}^4 \beta_k^S \text{PostShort}_{it} \times \text{Tenure}_i + \sum_{k=1}^4 \beta_k^L \text{PostLong}_{it} \times \text{Tenure}_i + \alpha_i + \gamma_t \right)$$

where y_{it} is the number of loans originated by worker i in week t . PostShort_{it} indicates whether the observation falls within a 12-week window after the Tech Roll-out, and PostLong_{it} indicates periods more than 12 weeks after the Tech Roll-out. Tenure is a set of indicators for workers with 0–3, 4–7, 8–11, or 12+ years of tenure at the time of the shock.

productivity from adaptation costs, as worker fixed effects absorb persistent differences in productivity across workers. It implies a sizable role for sluggish adaptation: the increase in predicted loans for high-tenure workers suggests that they possess relevant skills and client portfolios, but their experience-related productivity gains are offset by higher adaptation costs.

Missed loan opportunities also result in foregone revenue to the firm. Table 3 performs back-of-the-envelope calculations to quantify the overall loss. The model predicts loan counts rather than loan values, so the valuation uses pre-roll-out average loan values and the bank's reported profit margin on loans. We sum loan production over all branches and report predicted differences by loan types in column (1). Lower adjustment costs would have resulted in 1,401 additional simple loans and 821 additional complex loans in one week, equivalent to an 19 percent increase in total simple loans and a 2 percent increase in total complex loans. Simple loans, however, have a lower average value than complex loans: USD 6,150 versus USD 20,633 (column 2). Foregone revenue is therefore concentrated in complex loans, at an estimated USD 27 million, compared with USD 8 million for simple loans (column 3). The corresponding foregone weekly profits are USD 517 thousand for simple loans and USD 1.6 million for complex loans.

8. Conclusion

This paper shows how the returns to technology depend on the firm's capacity to reorganize work around it. As the bank's leadership put it: *"The issue was not simply to transform products digitally. [...] Over time, we realized that we had to transform the organization itself. And that meant changing every process in the organization."*

For firms, the results imply that technology investment, job design, and performance evaluation should be considered jointly. New technology can expand productive capacity while leaving task priorities unchanged. Realizing gains from reallocation requires evaluation systems and managerial practices that direct this capacity toward valued activities.

For policymakers, the no-layoff setting shows that the costs and benefits of technological

change can be uneven even when employment is preserved. Adjustment support may thus need to reach incumbent workers as well as displaced workers. Employees with established routines and task-specific knowledge may face especially high adaptation costs and may require targeted training, coaching, and transition support to benefit from new technologies.

Holding employment approximately fixed limits the scope for selection through hiring and separations, making within-worker adaptation the central margin of adjustment. In firms that rely more heavily on workforce replacement, observed tenure gradients may be smaller because workforce composition can change. Hence, the estimates speak most directly to organizations seeking to redeploy an incumbent workforce, and the within-worker frictions we document may be larger than those observed in settings with more extensive turnover.

The rapid diffusion of artificial intelligence highlights the importance of understanding how new technologies reshape work within existing jobs. Firms may substantially reorganize tasks without changing employment or worker flows, leaving much of this adjustment invisible in occupation-level measures of technological exposure. Our results speak directly to how tasks can be reallocated within existing occupations. In addition, they indicate that understanding the effects of new technologies requires observing not only whether firms adopt them, but also whether and how they redesign performance measures, task assignments, and managerial practices around them. In our setting, the substantial shift toward SME lending appears only after the firm redefines its performance criteria, so technological exposure alone reveals neither the direction of task adjustment nor its distributional effects across workers. More broadly, our results indicate that understanding the productivity and distributional effects of AI and other workplace technologies requires studying not only technological exposure, but also how firms reorganize work around their existing workforce.

References

Acemoglu, Daron and David Autor (2011) “Skills, Tasks and Technologies: Implications for Employment and Earnings,” in Ashenfelter, Orley and David Card eds. *Handbook of Labor Economics*, 4, 1043–1171: Elsevier, [10.1016/S0169-7218\(11\)02410-5](https://doi.org/10.1016/S0169-7218(11)02410-5).

- Acemoglu, Daron and Pascual Restrepo (2018) “The Race Between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment,” *American Economic Review*, 108 (6), 1488–1542, [10.1257/aer.20160696](#).
- (2022) “Tasks, Automation, and the Rise in U.S. Wage Inequality,” *Econometrica*, 90 (5), 1973–2016, [10.3982/ECTA19815](#).
- Adenbaum, Jacob (2023) “Endogenous Firm Structure and Worker Specialization,” working paper.
- Adhvaryu, Achyuta, Anant Nyshadham, and Jorge Tamayo (2023) “Managerial Quality and Productivity Dynamics,” *The Review of Economic Studies*, 90 (4), 1569–1607, [10.1093/restud/rdac038](#).
- Arrow, Kenneth J. (1974) *The Limits of Organization*, New York: W. W. Norton.
- Atkin, David, Azam Chaudhry, Shamyla Chaudry, Amit K Khandelwal, and Eric Verhoogen (2017) “Organizational Barriers to Technology Adoption: Evidence from Soccer-Ball Producers in Pakistan,” *The Quarterly Journal of Economics*, 132 (3), 1101–1164, [10.1093/qje/qjx010](#).
- Autor, David H. and Michael J. Handel (2013) “Putting Tasks to the Test: Human Capital, Job Tasks, and Wages,” *Journal of Labor Economics*, 31 (S1), S59–S96, [10.1086/669332](#).
- Autor, David H, Frank Levy, and Richard J Murnane (2003) “The skill content of recent technological change: An empirical exploration,” *Quarterly Journal of Economics*, 118 (4), 1279–1333.
- Autor, David H. and Neil Thompson (2025) “Expertise,” *Journal of the European Economic Association*, 23 (4), 1203–1271, [10.1093/jeea/jvaf023](#).
- Baker, George P. (1992) “Incentive Contracts and Performance Measurement,” *Journal of Political Economy*, 100 (3), 598–614, [10.1086/261831](#).


- Bellemare, Marc F and Casey J Wichman (2020) “Elasticities and the inverse hyperbolic sine transformation,” *Oxford Bulletin of Economics and Statistics*, 82 (1), 50–61.
- Bloom, Nicholas, Luis Garicano, Raffaella Sadun, and John Van Reenen (2014) “The Distinct Effects of Information Technology and Communication Technology on Firm Organization,” *Management Science*, 60 (12), 2859–2885, [10.1287/mnsc.2014.2013](https://doi.org/10.1287/mnsc.2014.2013).
- Bloom, Nicholas, Raffaella Sadun, and John Van Reenen (2012) “Americans Do IT Better: US Multinationals and the Productivity Miracle,” *American Economic Review*, 102 (1), 167–201, [10.1257/aer.102.1.167](https://doi.org/10.1257/aer.102.1.167).
- Bresnahan, Timothy F., Erik Brynjolfsson, and Lorin M. Hitt (2002) “Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence,” *The Quarterly Journal of Economics*, 117 (1), 339–376, <http://www.jstor.org/stable/2696490>.
- Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond (2025) “Generative AI at Work,” *The Quarterly Journal of Economics*, 140 (2), 889–942, [10.1093/qje/qjae044](https://doi.org/10.1093/qje/qjae044).
- Brynjolfsson, Erik, Daniel Rock, and Chad Syverson (2021) “The Productivity J-Curve: How Intangibles Complement General Purpose Technologies,” *American Economic Journal: Macroeconomics*, 13 (1), 333–372, [10.1257/mac.20180386](https://doi.org/10.1257/mac.20180386).
- Chen, Yuyu, Miguel Espinosa, Shuo Liu, and Jingchen Zhu (2026) “Who Benefits from Digital Monitoring? The Role of Managerial Task-Allocation Skill,” Unpublished manuscript.
- Coase, Ronald H. (1937) “The Nature of the Firm,” *Economica*, 4 (16), 386–405, [10.1111/j.1468-0335.1937.tb00002.x](https://doi.org/10.1111/j.1468-0335.1937.tb00002.x).
- De Chaisemartin, Clément and Xavier d’Haultfoeuille (2020) “Two-way fixed effects estimators with heterogeneous treatment effects,” *American economic review*, 110 (9), 2964–2996.
- Demirer, Mert, John J. Horton, Nicole Immorlica, Brendan Lucier, and Peyman Shahidi (2026) “Chaining Tasks, Redefining Work: A Theory of AI Automation,” NBER Working Paper 34859, National Bureau of Economic Research, [10.3386/w34859](https://doi.org/10.3386/w34859).

- Diaz, Brayan S., Andrea Neyra Nazarrett, Julian Ramirez, Raffaella Sadun, and Jorge A. Tamayo (2025) “Training within Firms,” NBER Working Paper 33670, National Bureau of Economic Research, [10.3386/w33670](#).
- Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock (2024) “GPTs Are GPTs: Labor Market Impact Potential of LLMs,” *Science*, 384 (6702), 1306–1308, [10.1126/science.adj0998](#).
- Feigenbaum, James and Daniel P. Gross (2024) “Organizational and Economic Obstacles to Automation: A Cautionary Tale from AT&T in the Twentieth Century,” *Management Science*, 70 (12), 8520–8540.
- Freund, Lukas B. and Lukas F. Mann (2026) “Job Transformation, Specialization, and the Labor Market Effects of AI,” Working paper, first version August 16, 2025; this version March 18, 2026.
- Friebel, Guido, Matthias Heinz, and Nikolay Zubanov (2022) “Middle Managers, Personnel Turnover, and Performance: A Long-Term Field Experiment in a Retail Chain,” *Management Science*, 68 (1), 211–229, [10.1287/mnsc.2020.3905](#).
- Gans, Joshua S. and Avi Goldfarb (2026) “O-Ring Automation,” NBER Working Paper 34639, National Bureau of Economic Research, [10.3386/w34639](#).
- Goldstein, Harvey and Michael JR Healy (1995) “The graphical presentation of a collection of means,” *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 158 (1), 175–177.
- Holmström, Bengt and Paul Milgrom (1991) “Multitask Principal–Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design,” *Journal of Law, Economics, and Organization*, 7 (Special Issue), 24–52, [10.1093/jleo/7.special.issue.24](#).
- Houeix, Deivy (2026) “Asymmetric Information and Digital Technology Adoption: Evidence from Senegal,” pedl working paper, Private Enterprise Development in Low-Income Countries, Latest version February 22, 2026.


- Humlum, Anders and Emilie Vestergaard (2026) “Still Waters, Rapid Currents: Early Labor Market Transformation under Generative AI,” NBER Working Paper 33777, National Bureau of Economic Research, [10.3386/w33777](https://doi.org/10.3386/w33777), Latest version March 20, 2026.
- Ichniowski, Casey, Kathryn Shaw, and Giovanna Prennushi (1997) “The Effects of Human Resource Management Practices on Productivity: A Study of Steel Finishing Lines,” *American Economic Review*, 87 (3), 291–313.
- Juhász, Réka, Mara P. Squicciarini, and Nico Voigtländer (2024) “Technology Adoption and Productivity Growth: Evidence from Industrialization in France,” *Journal of Political Economy*, 132 (10), 3215–3259.
- Kastis, Yannis and Hillary Vipond (2025) “Organizational Practices and Technology Adoption: Evidence from Jewish Immigration and the Tailoring Industry in England,” October, https://jkastis.github.io/yanniskastis/Kastis_JMP.pdf, Working paper.
- Lazear, Edward P, Kathryn L Shaw, and Christopher T Stanton (2015) “The Value of Bosses,” *Journal of Labor Economics*, 33 (4), 823–861, [10.1086/679517](https://doi.org/10.1086/679517).
- McElheran, Kristina, Mu-Jeung Yang, Zachary Kroff, and Erik Brynjolfsson (2026) “The Adoption of Industrial AI in America,” *AEA Papers and Proceedings*, 116, 20–25.
- Metcalf, Robert D., Alexandre B. Sollaci, and Chad Syverson (2023) “Managers and Productivity in Retail,” NBER Working Paper 31192, National Bureau of Economic Research, [10.3386/w31192](https://doi.org/10.3386/w31192).
- Minni, Virginia (2026) “Making the Invisible Hand Visible: Managers and the Allocation of Workers to Jobs,” *The Quarterly Journal of Economics*, 141 (3), 1871–1920, [10.1093/qje/qjag017](https://doi.org/10.1093/qje/qjag017).
- Nelson, Richard R. and Sidney G. Winter (1982) *An Evolutionary Theory of Economic Change*, Cambridge, MA: Harvard University Press.

- Ocampo, Sergio (2022) “A Task-Based Theory of Occupations with Multidimensional Heterogeneity,” working paper, University of Western Ontario, Human Capital and Economic Opportunity Working Group.
- Oi, Walter Y (1962) “Labor as a Quasi-Fixed Factor,” *Journal of Political Economy*, 70 (6), 538–555, [10.1086/258468](https://doi.org/10.1086/258468).
- Simon, Herbert A. (1951) “A Formal Theory of the Employment Relationship,” *Econometrica*, 19 (3), 293–305, [10.2307/1906815](https://doi.org/10.2307/1906815).
- Sun, Liyang and Sarah Abraham (2021) “Estimating dynamic treatment effects in event studies with heterogeneous treatment effects,” *Journal of Econometrics*, 225 (2), 175–199, <https://doi.org/10.1016/j.jeconom.2020.09.006>, Themed Issue: Treatment Effect 1.
- Verhoogen, Eric (2023) “Firm-Level Upgrading in Developing Countries,” *Journal of Economic Literature*, 61 (4), 1410–1464, [10.1257/jel.20221655](https://doi.org/10.1257/jel.20221655).
- Williamson, Oliver E. (1975) *Markets and Hierarchies: Analysis and Antitrust Implications*, New York: Free Press.

9. Figures




(a) Task complexity by product




(b) Visual summary of the interventions


Figure 1: **Interventions and the complexity of tasks.** Panel (a) shows standardized differences in task content for specific complex products relative to average simple products using composite O*NET measures (Autor et al., 2003). High-complexity products show higher analytical and lower routine-cognitive content than simple products; medium-complexity products (checking and payroll accounts) differ mainly in interpersonal content, with analytical content slightly below and routine content above that of simple products. SME lending and pre-mortgages show the largest analytical shift. Panel (b) summarizes the Tech Roll-out and Task Intervention: the Tech intervention automates simple products (blue), while the Task intervention targets SME lending via KPI changes (red). As the app matures, customers increasingly initiate purchases independently (blue bubble), and we focus on how complex tasks respond (purple).




(a) Total clients




(b) Unique first-time clients




(c) Trades by new clients



(d) Trades by old clients




(e) SME Loans by new clients




(f) SME Loans by old clients


Figure 2: Effects of Tech Roll-out on new client acquisition and activity by client type at the branch level. All outcomes are transformed using the inverse hyperbolic sine (IHS). Panel (a) reports the effect of the Tech Roll-out on the total number of weekly clients. Panel (b) reports the effect on unique new clients (first-time visits). Panels (c) and (d) report the effects on the number of trades performed by new and old clients, where client status is defined based on a six-month window since the first observed trade. Panels (e) and (f) further report SME loans by client type. Standard errors are clustered at the branch level.




(a) Tech: Total Balances (Branch)



(b) Tech: Number of Accounts (Branch)




(c) Task: Total Balances (Branch)




(d) Task: Number of Accounts (Branch)


Figure 3: Effects of Tech Roll-out and Task Intervention on balances and number of accounts at the branch level. All outcomes are transformed using the inverse hyperbolic sine (IHS). Panel (a) reports total balances (in millions of pesos) and Panel (b) reports the number of active accounts for the Tech Roll-out. Panels (c) and (d) report corresponding outcomes for the Task Intervention. Due to data availability and reporting, the outcomes are only available at the branch and monthly level. Standard errors are clustered at the branch level.



(a) Tech Roll-out: Total trades




(b) Tech Roll-out: Total Trades (only Account Managers)




(c) Task Intervention: Total trades


Figure 4: Effects of Tech Roll-out and Task Intervention on total trades at the worker level. All outcomes are transformed using the inverse hyperbolic sine (IHS). Panel (a) shows the Tech Roll-out on the full sample of workers, Panel (b) shows the Tech Roll-out on the sample of account managers, and Panel (c) shows the Task Intervention on the sample of account managers. Standard errors are clustered at the branch level. Corresponding branch-level estimates are reported in Appendix Figure A.2.




(a) Tech: Total Hiring (Branch)



(b) Tech: Total Exits (Branch)




(c) Task: Total Hiring (Branch)




(d) Task: Total Exits (Branch)


Figure 5: Effects of Tech Roll-out and Task Intervention on worker turnover at the branch level. All outcomes are in levels. Row 1 reports the Tech Roll-out and Row 2 reports the Task Intervention. Panels (a) and (c) report total hirings, and Panels (b) and (d) report total exits. Hirings measure the number of workers entering the company and assigned to a branch in a given week. Exits count workers who permanently leave the company in a given week; each worker is assigned to the last branch in which the worker is observed. The analysis includes all workers. Standard errors are clustered at the branch level.




(a) Full sample: Simple Trades




(b) Full sample: Other Complex Trades




(c) Full sample: SME Loans



(d) Account Managers: Simple Trades



(e) Account Managers: Other Complex Trades



(f) Account Managers: SME Loans

Figure 6: Effects of Tech Roll-out on task composition. All outcomes are transformed using the inverse hyperbolic sine (IHS). Panels report worker-level task composition for simple trades, other complex trades, and SME loans. Row 1 uses the full sample of branch workers, whereas Row 2 is restricted to account managers. Standard errors are clustered at the branch level.







Figure 7: Effects of Task Intervention on task composition. All outcomes are in inverse hyperbolic sine (IHS) transformation, estimated on the sample of account managers. Panels report worker-level task composition by trade type: simple trades, other complex trades excluding SMEs, and SME loans. Row 1 reports effects of the intervention on trade volumes, and Row 2 reports effects on sales (in millions of pesos). The unusually large treatment-week coefficients for SME trades and sales in panels 7c and 7f coincide with the week when HR staff deliver branch-level training and explain the KPI change. They therefore include a short-run training-week response. The subsequent coefficients remain positive, showing that SME activity stays elevated after the training week. Standard errors are clustered at the branch level.




(a) Trades: Simple Trades




(b) Trades: Other Complex Trades




(c) Trades: SME Loans



(d) Sales: Simple Trades




(e) Sales: Other Complex Trades




(f) Sales: SME loans


Figure 8: Heterogeneous responses to the Tech Roll-out by worker tenure. All outcomes are transformed in inverse-hyperbolic-sine (IHS). Row 1 reports long-run coefficients for simple trades, other complex trades, and SME loans. Row 2 reports the corresponding long-run coefficients for sales outcomes. The sample includes all workers. The long-run indicator equals one from week 13 onward after the intervention; weeks 0–12 are captured by a separate short-run indicator, and the full pre-treatment period is the reference category. Tenure is measured at the treatment date and grouped into 0–3, 4–7, 8–11, and 12+ years. Standard errors are clustered at the branch level. Panels show 83% confidence intervals, such that non-overlapping intervals indicate a statistically significant difference between means at the 5% level (Goldstein and Healy, 1995).




(a) Simple Trades




(b) Other Complex Trades




(c) SME Loans



(d) Sales: Simple trades




(e) Sales: Other Complex Trades




(f) Sales: SME loans


Figure 9: Heterogeneous responses to the Task Intervention by worker tenure. All outcomes are transformed in inverse-hyperbolic-sine (IHS). Row 1 reports long-run coefficients for simple trades, other complex trades, and SME loans. Row 2 reports the corresponding long-run coefficients for sales outcomes. The sample includes account managers. The long-run indicator equals one from week 31 onward after the intervention; weeks 0–30 are captured by a separate short-run indicator, and the full pre-treatment period is the reference category. Tenure is measured at the treatment date and grouped into 0–3, 4–7, 8–11, and 12+ years. Standard errors are clustered at the branch level. Panels show 83% confidence intervals, such that non-overlapping intervals indicate a statistically significant difference between means at the 5% level (Goldstein and Healy, 1995).




(a) Tech Roll-Out: Log wages (yearly)



(b) Task Intervention: Log wages (yearly)




(c) Wage distribution by tenure group




(d) Task Intervention: Worker Days (levels)

Figure 10: Non-parametric tenure heterogeneity in wages and worker days after the Tech Rollout and Task Intervention. The figure reports worker-level TWFE estimates of annual log wages. Panel (a) reports the effect of the Tech Roll-out on annual log wages, with wages measured in millions of pesos, for the worker sample; tenure is measured at the treatment date. Panel (b) reports the effect on log wages (in millions of pesos) for the Task Intervention, with tenure years computed at the time of the shock over the sample of account managers in the yearly panel. This is due to wage information being available at the yearly level. Panel (c) reports the yearly distribution of wages (in millions of pesos) by tenure group among account managers. Worker median tenure is split within each year. Panel (d) reports the effect on days worked, with tenure years computed at the time of the shock over the sample of account managers in the weekly panel. The plotted coefficients correspond to effects from week 31 onward after the intervention, relative to the full pre-treatment period. Standard errors are clustered at the branch level. Panels (a), (b) and (d) show 83% confidence intervals, such that non-overlapping intervals indicate a statistically significant difference between means at the 5% level (Goldstein and Healy, 1995).




(a) Trades: Total Trades (b) Trades: Simple Trades (c) Trades: Other Complex (d) Trades: SME Loans



(e) Sales: Total Trades (f) Sales: Simple Trades (g) Sales: Other Complex (h) Sales: SME Loans

Figure 11: Heterogeneity in responses to the Task Intervention by manager quality and worker tenure. Row 1 refers to trades outcomes; Row 2 refers to sales outcomes. Columns report total trades, simple trades, other complex trades, and SME loans. All outcomes are transformed in inverse hyperbolic sine (IHS), "MQ" denotes manager quality. Manager quality is assigned based on an AKM specification with branch and manager fixed effects, and branch productivity as the outcome. Manager fixed effects are then split at the median of the distribution and are assigned to workers based on the branch in which they are observed at the intervention date. Worker tenure is measured at baseline and split at the median in the cross-section of workers. The sample includes all workers. Standard errors are clustered at the branch level. The figure shows 83% confidence intervals, such that non-overlapping intervals indicate a statistically significant difference between means at the 5% level (Goldstein and Healy, 1995).




(a) Total training hours

(b) Analytical


(c) Interpersonal

(d) Routine Cognitive

Figure 12: Effects of the Task Intervention on training by worker tenure and manager quality. Panels report treatment effects of the task intervention on training by worker tenure and manager quality at the worker level. Training outcomes are measured in hours and refer to total (a), analytical (b), interpersonal (c), and routine cognitive (d) training activities. "MQ" denotes manager quality. Manager quality is assigned based on an AKM specification with branch and manager fixed effects, and branch productivity as the outcome. Manager fixed effects are then split at the median of the distribution and are assigned to workers based on the branch in which they are observed at the intervention date. Worker tenure is measured at baseline and split at the median in the cross-section of workers. The sample includes all workers. Standard errors are clustered at the branch level. The figure shows 83% confidence intervals, such that non-overlapping intervals indicate a statistically significant difference between means at the 5% level (Goldstein and Healy, 1995).



(a) Simple Loans



(b) Complex Loans (incl. SME)

Figure 13: Simulated loan counts after the Tech Roll-out, assuming low adaptation costs. The figure reports average weekly differences in predicted counts between the baseline model and a simulation in which all workers have the lowest observed adaptation costs, as represented by workers with 0–3 years of tenure. Only loan products are considered. Panel 13a includes simple loans available through the banking app during the Tech Roll-out (fixed-rate and variable-rate loans). Panel 13b includes all complex loans, including SME, mortgages, revolving, and vehicle loans. Predicted counts are estimated from a Poisson pseudo-maximum likelihood model where number of loans are expressed in levels, interacted with tenure bands set at the time of the intervention, and short-run (up to 12 weeks) and long-run (over 12 weeks) adjustments to the Tech Roll-out. Here, only long-run predictions are showcased and compared with the full pre-treatment period. In both predictions, the time period is fixed to the last week prior to the Tech Roll-out to disregard seasonality and trends. In the simulation, the interaction between tenure and the Tech Roll-out is set to the lowest band (0-3 years) to simulate a lower adaptation-cost response. Throughout the analysis, worker fixed effects remain unchanged. The specification includes worker and week fixed effects, and the sample includes all workers.

10. Tables

	Tech Roll-out		Task Intervention	
	Avg.	Std. Dev.	Avg.	Std. Dev.
Panel A: Branch Level Descriptive Statistics				
Average No. Workers	7.528	5.442	8.263	4.773
Total Trades	72.755	73.197	70.336	46.416
Complex Trades (Excl. SME)	28.789	38.565	19.499	22.774
Simple Trades	41.675	44.362	48.724	31.759
SME Loans	2.291	3.990	2.114	4.058
Productivity (Trades/Workers)	9.703	6.692	8.952	6.299
Observations	69,465		77,830	
Panel B: Worker Level Descriptive Statistics				
Total Trades	11.769	11.982	14.629	10.233
Complex Trades (Excl. SME)	4.823	6.219	4.038	4.591
Simple Trades	6.946	7.893	10.517	7.615
SME Loans	0.368	1.504	0.077	0.438
Observations	382,447		158,701	
Panel C: Sample Summary				
Time period	Jan 2017–Dec 2019		Jan 2019–Mar 2023	
Number of treated branches	449		355	
Number of total workers	6,948		8,129	
Number of account managers	2,152		2,018	
Number of simple products	4		8	
Number of complex products (incl. SME)	12		8	

Table 1: Branch- and worker-level descriptive statistics and sample summary. Panels A and B: All statistics refer to average weekly values. Panel C: additional summary information. Panel B: Worker-level descriptive statistics are reported for all workers in the Tech Roll-out sample and for account managers in the Task Intervention sample. Statistics for the sample of account managers during the Tech Roll-out are comparable to those for the Task Intervention (total trades: 15.395; complex trades excl. SME: 6.078; simple trades: 9.270; SME loans: 0.046). The average number of SME loans is higher in the Tech Roll-out sample because that sample includes workers specialized in SME lending who are not covered by the Task Intervention. Panel C: Summary of the data used to assess the Tech Roll-out and Task Intervention. Simple products available during the Tech Roll-out include savings accounts (approx. 50% of trades), fixed and variable rate loans, and credit cards. The Task Intervention additionally includes all products that are added to the mobile app by January 2021, such as deposit certificates, other saving accounts, simple mortgages and car loans (these products are classified as complex during the Tech Roll-out). Complex products include payroll accounts, checking accounts, life insurances, pension funds, investment funds, and more complex loans and mortgages. SME loans are the most complex product as they also require skills from commercial banking.

Prod. Manager FEs	Correlation	p-value
<i>Leadership</i>		
Teamwork	0.255	0.000***
Trust	0.172	0.001***
Embrace Change	0.217	0.000***
Inspire	0.214	0.000***
Avg. Leadership	0.242	0.000***
<i>Performance Measures</i>		
Performance	0.127	0.013**
Objectives	0.138	0.007***
Skills	0.026	0.611
Avg. Performance	0.118	0.020**
<i>Company Values</i>		
Principles	-0.005	0.928
Big Picture	0.032	0.529
Context	0.033	0.520
Client Focus	0.010	0.843
Workforce Development	0.086	0.090*
Co-creation	0.006	0.901
Agility	0.061	0.234
Adventure	-0.002	0.972
Avg. Values	0.036	0.485
<i>Demographics</i>		
Age	-0.0277	0.432
Tenure	-0.0264	0.452
Adj. Wage	-0.0479	0.173
Gender	0.0351	0.319


Table 2: Correlation between manager AKM fixed effects, demographics, and leadership measures. The internal performance and values measures and the leadership observations used in this table cover 2020–2022; the external subordinate leadership survey is available through 2023. These observations postdate the January 2017–June 2018 window used to estimate manager AKM fixed effects, and the 2021–2022 waves overlap the Task Intervention. Because the survey ratings may reflect intervention exposure and subsequent branch performance, they are used only for ex post validation and do not define manager quality in the treatment-effect analysis. Performance measures come from the bank’s internal manager-evaluation instrument, while company values reflect managers’ self-reported alignment with firm values. “Performance” captures overall managerial performance, “Objectives” reflects the ability to meet objectives and deadlines, and “Skills” indicates whether the manager meets the required skill level; “Avg. Performance” is the mean of these three items. “Avg. Values” is the mean across value dimensions. Leadership measures are based on annual evaluations conducted by an external consulting firm and reflect worker ratings of their branch manager; “Avg. Leadership” is the mean of these items. Demographic characteristics are drawn from personnel records; wages are inflation-adjusted in millions of pesos.

Product Type	Predicted Difference (Counts) Simulation – Baseline (1)	Average Prod. Value (USD) Baseline (2)	Aggregate Prod. Value (USD) Counts $\times$ Avg. Value (3)	Profit Margin (USD) Agg. Value $\times$ 6% (4)
Simple Loans	1,401	6,150	8,618,810	517,129
Complex Loans (incl. SME)	821	20,633	27,436,788	1,646,207


Table 3: Counterfactual magnitudes by loan type based on low adjustment cost simulation. Back-of-the-envelope estimate of aggregate weekly losses due to adaptation costs after the Tech Roll-out, expressed in 2017 USD. Worker-level differences in predicted counts between the prediction model and a simulation where all workers have the lowest observed adaptation costs (as determined by workers with 0-3 years of tenure) are aggregated over the available post-period until December 2019. Only loan products are considered. The first row includes simple loans available on the banking app during the Tech Roll-out (fixed rate and variable rate loans), the second row includes all complex loans (SME loans, revolving loans, mortgages, vehicle loans). Column (1): Predicted counts are estimated at the worker level from a Poisson pseudo-maximum likelihood model where number of loans are expressed in levels, interacted with tenure bands set at the time of the intervention, and short-run (up to 12 weeks) and long-run (over 12 weeks) adjustments to the Tech Roll-out. Here, we showcase the average long-run prediction aggregated over one week. In both the prediction model and the simulation with low adjustment costs, the time period is fixed to the last week prior to the Tech Roll-out to disregard seasonality and trends. In the simulation, the interaction between tenure and the Tech Roll-out is set to the lowest band (0-3 years) to simulate a lower adaptation-cost response. Throughout the estimation, worker fixed effects remain unchanged, worker and week fixed effects are included, and the sample includes all workers. Figure 13 reports average predictions by worker type. Column (2): average product value is estimated based on the average loan value of all products in a category at baseline, and converted into 2017 USD (this is a conservative estimate as average loan value increases slightly after the Tech Roll-out). Column (3): Aggregate product value is estimated by multiplying the predicted difference in counts (column 1) by the average product value (column 2). Column (4): the profit margin is set at 6%, reflecting the average profit margin across loan products reported by the bank.

A. Additional figures and tables

A.1. Appendix figures




(a) Tech Roll-out




(b) Task Intervention

Figure A.1: **Roll-out of the Tech and Task Intervention.** The figure shows the timing and distribution of treatment across units for the Tech Roll-out and the Task Intervention. The Tech Roll-out date is the first week in which a branch records an app-based sale. The Task Intervention date is the week in which HR staff deliver branch-level training and explain the SME-related KPI change.




(a) Tech: Total trades




(b) Tech: Productivity


Figure A.2: Effects of Tech Roll-out on total trades and productivity at the branch level. All outcomes are transformed using the inverse hyperbolic sine (IHS). Panel (a) reports total trades at the branch level, and Panel (b) reports trades per employee at the branch level.




(a) Total Trades



(b) Simple Trades




(c) Other Complex Trades




(d) SME Loans


Figure A.3: Effects of Tech Roll-out at the worker level on the sample of branches also treated by the Task Intervention. All outcomes are transformed using the inverse hyperbolic sine (IHS). Panel (a) reports total trades, and Panel (b) reports simple trades, Panel (c) reports other complex trades, panel (d) reports SME loans. The analysis sample is restricted to workers in branches that are also treated by the Task Intervention.




(a) Tech: Reallocation in Space (Branch)



(b) Tech: Reallocation in Jobs (Branch)




(c) Task: Reallocation in Space (Branch)




(d) Task: Reallocation in Jobs (Branch)

Figure A.4: Effects of Tech Roll-out and Task Intervention on worker reallocation at the branch level. All outcomes are in levels. Row 1 reports the Tech Roll-out and Row 2 reports the Task Intervention. *Reallocation in space* is the number of workers transferred across branches in a given week. *Reallocation in tasks* is the number of workers transitioning into customer-facing (front office) activities in a given week. These transitions measure movement into front office roles, not changes in the product mix within a role; product-level reallocation is measured by the trade outcomes. The absence of an effect indicates that neither intervention operates through the reassignment of workers across roles. To exclude temporary training-related movements, the sample is restricted to workers with at least three months of tenure prior to the transition.

Reallocation ratio: complex over simple trades




(a) Tech Intervention (Worker)




(b) Task Intervention (Worker)

Share of SME loans in complex trades




(c) Tech Intervention (Worker)




(d) Task Intervention (Worker)

Specialization: worker-weighted Herfindahl-Hirschman Index



(e) Tech Intervention (Branch)



(f) Task Intervention (Branch)

Figure A.5: Effects of Tech Roll-out and Task Intervention on worker-level reallocation ratios and task specialization. Row 1 reports reallocation between complex and simple trades at the worker level. Row 2 reports the share of SME loans in complex trades at the worker level. Row 3 reports task specialization, measured by the Herfindahl–Hirschman Index at the branch level. The HHI is computed based on the distribution of trades across task categories, with higher values indicating greater specialization, weighted by worker productivity. Results on the left pertain to the Tech Roll-out and include all affected workers, while results on the right pertain to the Task Intervention and include account managers.




Figure A.6: Tech Roll-out: heterogeneity in tenure and baseline productivity in complex trades by task type. Tenure is measured at the treatment date and grouped into 0–3, 4–7, 8–11, and 12+ years. The plotted coefficients correspond to effects from week 31 onward after the intervention, relative to the full pre-treatment period. The analysis includes all workers. Complex productivity is computed as the residual of a regression with complex trades (excluding SME loans) on week and branch fixed effects. Residuals are then averaged over the baseline period (before treatment) for each worker, and split along the median. Task categories are all trades, simple trades, other complex trades (excluding SMEs) and SME loans.




Figure A.7: **Tech Roll-out: tenure-related frictions by task type.** Rows correspond to task categories (simple, complex tasks, and SME loans). Columns report heterogeneity by worker age, worker ability, and peer tenure exposure, measured at baseline. Tenure is measured at the treatment date and grouped into 0–3, 4–7, 8–11, and 12+ years. The plotted coefficients correspond to effects from week 31 onward after the intervention, relative to the full pre-treatment period. The analysis includes all workers. Age is computed at the time of the intervention and split at the median in the cross-section of workers. Ability is computed as the average weekly trades prior to the intervention. High tenure peers are workers whose tenure is above median at the time of the intervention. Shares are computed over the number of workers in a branch (excluding the worker over whom we are computing the split). Low- and high-shares indicate shares of high-tenure peers that are below and above median.






Figure A.8: **Task Intervention: tenure-related frictions by task type.** Rows correspond to task categories (simple, complex, and SME loans). Columns report heterogeneity by worker age, worker ability, and peer tenure exposure, measured at baseline. Tenure is measured at the treatment date and grouped into 0–3, 4–7, 8–11, and 12+ years. The plotted coefficients correspond to effects from week 31 onward after the intervention, relative to the full pre-treatment period. The analysis sample is restricted to account managers. Age is computed at the time of the intervention and split at the median in the cross-section of workers. Ability is computed as the average weekly trades prior to the intervention. High tenure peers are workers whose tenure is above median at the time of the intervention. Shares are computed over the number of workers in a branch (excluding the worker over whom we are computing the split). Low- and high-shares indicate shares of high-tenure peers that are below and above median.




(a) Tech Intervention



(b) Task Intervention

Figure A.9: Tech and Task Interventions: non-parametric age heterogeneity in total trades. Panel (a) reports the Tech Roll-out; Panel (b) reports the Task Intervention. All outcomes are transformed using the inverse hyperbolic sine (IHS). Age is measured at the treatment date and grouped into ≤ 24 , 25–29, 30–34, 35–39, ≥ 40 years. The plotted coefficients correspond to effects from week 31 onward after the intervention, relative to the full pre-treatment period.



(a) Tech Roll-out



(b) Task Intervention

Figure A.10: Non-parametric heterogeneity in Manager Quality on Total Trades after the Tech Roll-out and Task Intervention. The plotted coefficients correspond to long run of each intervention on total trades, by quartiles of manager quality. Panel (a) reports the Tech Roll-out; long run effects are defined from week 13 onward, relative to the full pre-treatment period. Panel (b) reports the Task Intervention; long run effects are defined from week 31 onward, relative to the full pre-treatment period. Managers are sorted into quartiles of manager AKM fixed effects. Both panels use the full sample of branch workers. Standard errors are clustered at the branch level. 83% confidence intervals are reported, such that non-overlapping intervals indicate a statistically significant difference between means at the 5% level (Goldstein and Healy, 1995).




Figure A.11: Heterogeneity in responses to the Tech Roll-out by manager quality and worker tenure. Panels report total trades, simple trades, other complex trades, and SME loans. All outcomes are transformed in inverse hyperbolic sine (IHS), "MQ" denotes manager quality. Manager quality is assigned based on an AKM specification with branch and manager fixed effects, and branch productivity as the outcome. Manager fixed effects are then split at the median of the distribution and are assigned to workers based on the branch in which they are observed at the intervention date. Worker tenure is measured at baseline and split at the median in the cross-section of workers. The sample includes all workers. Standard errors are clustered at the branch level. The figure shows 83% confidence intervals, such that non-overlapping intervals indicate a statistically significant difference between means at the 5% level (Goldstein and Healy, 1995).

A.2. Appendix tables

Variable	Branch Level			Worker Level		
	(1)	(2)	(3)	(4)	(5)	(6)
	Early	Late	Diff.	Early	Late	Diff.
Panel A: Tech Roll-Out						
Average No. Workers	6.513 (5.190)	7.805 (4.969)	1.347*** (0.341)			
Urban Office	0.589 (0.492)	0.952 (0.214)	0.363*** (0.036)			
Trades	4.249 (1.152)	4.475 (1.023)	0.190*** (0.055)	2.554 (1.107)	2.529 (1.104)	-0.018 (0.026)
Complex Trades	3.323 (1.347)	3.640 (1.166)	0.288*** (0.063)	1.724 (1.164)	1.708 (1.157)	-0.011 (0.022)
Complex Trades (Excl. SME)	3.123 (1.470)	3.433 (1.319)	0.263*** (0.063)	1.525 (1.259)	1.490 (1.261)	-0.028 (0.031)
Simple Trades	3.630 (1.139)	3.779 (1.056)	0.104** (0.053)	1.849 (1.224)	1.820 (1.224)	-0.021 (0.034)
SME Loans	1.040 (1.089)	1.200 (1.141)	0.199*** (0.050)	0.221 (0.611)	0.249 (0.628)	0.026 (0.021)
Productivity (Trades/Workers)	2.643 (0.789)	2.604 (0.704)	-0.075*** (0.021)			
Observations	16,846	18,795	35,641	189,904	35,498	225,402
Panel B: Task Intervention						
Average No. Workers	8.488 (4.144)	7.210 (5.704)	-1.227*** (0.441)			
Urban Office	0.899 (0.301)	0.801 (0.399)	-0.086** (0.037)			
Trades	4.708 (0.755)	4.552 (0.744)	-0.191*** (0.049)	2.947 (0.885)	2.911 (0.916)	-0.022 (0.030)
Complex Trades	3.275 (0.972)	3.062 (1.003)	-0.261*** (0.061)	1.495 (0.971)	1.431 (0.989)	-0.067** (0.026)
Complex Trades (Excl. SME)	3.088 (1.084)	2.894 (1.102)	-0.255*** (0.066)	1.483 (0.979)	1.415 (1.000)	-0.071** (0.028)
Simple Trades	4.365 (0.779)	4.228 (0.760)	-0.169*** (0.050)	2.624 (0.944)	2.593 (0.984)	-0.011 (0.033)
SME Loans	1.053 (0.907)	0.854 (0.870)	-0.212*** (0.044)	0.013 (0.149)	0.017 (0.160)	0.003 (0.007)
Productivity (Trades/Workers)	2.698 (0.513)	2.743 (0.533)	0.015 (0.026)			
Observations	20,827	26,330	47,157	57,463	29,043	86,506

Table A.1: **Pre-intervention difference in trades and tasks between early and late treated branches and individuals.** Branches and workers are split on the median treatment cohort in the cross-section of branches. All outcomes are transformed using the inverse hyperbolic sine (IHS), and the specifications include week fixed effects.

Panel A: Tech Roll-out

	(1)	(2)	(3)	(4)
	Branch: Trades	Branch: Trades/Workers	Worker: Trades	Worker: Trades Excl. SME
Tech Roll-out	0.141*** (0.0204)	0.0967*** (0.0152)	0.184*** (0.0284)	0.163*** (0.0288)
Observations	69,465	69,244	382,050	382,050
Dependent Mean	4.349	1.819	2.586	2.401

Panel B: Task Intervention

	(1)	(2)	(3)	(4)
	Branch: Trades	Branch: Trades/Workers	Worker: Trades	Worker: Trades Excl. SME
Task Intervention	0.1540** (0.0683)	0.1504*** (0.0383)	0.1125*** (0.0257)	0.1066*** (0.0258)
Observations	77,830	77,774	158,697	158,697
Dependent Mean	2.726	1.757	2.935	2.922

Table A.2: **Effect of Technology and Task Intervention on trades and productivity.** Panel A reports the average effect of the Tech Roll-out on trade volume and productivity outcomes at the branch and worker levels. Panel B reports analogous estimates for the Task Intervention. In the context of the Task Intervention, worker level results might be inflated by including the training week in the panel. We re-estimate Panel B columns (3) and (4) excluding week 0, and the coefficients are 0.1197*** (0.0270) for trades, and 0.1163*** (0.0271) for trades excluding SME loans, excluding that our results are driven by the training week. Standard errors are clustered at the branch level.

Panel A: Tech Roll-out				
	(1)	(2)	(3)	(4)
	All Trades	Simple Trades	Complex excl. SME	SME Loans
Tech Roll-out	0.268*** (0.0366)	0.265*** (0.0358)	0.133*** (0.0383)	0.0649*** (0.0151)
Tech $\times$ High Tenure	-0.148*** (0.0150)	-0.128*** (0.0151)	-0.0625*** (0.0132)	-0.0484*** (0.00538)
Observations	311,302	311,302	311,302	311,302
Mean (pre)	2.600	1.862	1.585	0.229

Panel B: Task Intervention				
	(1)	(2)	(3)	(4)
	All Trades	Simple Trades	Complex Excl. SME	SME Loans
Task Intervention	0.137*** (0.0256)	0.141*** (0.0256)	0.0812*** (0.0285)	0.0534*** (0.00512)
Task $\times$ High Tenure	-0.0872*** (0.0230)	-0.0752*** (0.0233)	-0.100*** (0.0264)	-0.00268 (0.00587)
Observations	144,656	144,656	144,656	144,656
Mean dep. Variable	3.072	2.712	1.661	0.0520

Table A.3: Effect of ex-ante high tenure on main outcome variables, urban branches only. The table reports heterogeneous treatment effects by worker tenure, estimated only on workers in urban branches. High tenure is defined as above-median average tenure measured in the pre-intervention period. Urban branch location is a binary indicator equal to one if the branch is located in an urban setting. Panel A reports results for the Tech Roll-out; Panel B reports results for the Task Intervention. All outcomes are expressed in inverse hyperbolic sine (IHS). Week and worker fixed effects included, standard errors are clustered at the branch level.

	Tech Roll-out		Task Intervention	
	(1) Total Exits	(2) Total Hiring	(3) Total Exits	(4) Total Hiring
Intervention	-0.000931 (0.00408)	0.00309 (0.00358)	0.00134 (0.00207)	-0.00148 (0.00140)
Intervention $\times$ High Tenure	-0.00319 (0.00279)	-0.00161 (0.00350)	0.000643 (0.00185)	0.00171 (0.00127)
Observations	69,244	69,244	77,391	77,391

Week and branch FEs included. SE clustered by Branch

*p<.05; **p<.01; ***p<.001

Table A.4: **Heterogeneity in Turnover, by Tenure.** The table reports heterogeneous treatment effects by worker tenure, at the branch level, for both the Tech and the Task intervention. A branch is classified as high-tenure if the average tenure of its account managers at the intervention date exceeds the sample median. This classification is held fixed throughout the panel. Total hiring is the total number of workers hired by the company in a given branch and week. Week and branch fixed effects included.

B. Model appendix

This appendix states the full game and characterizes the continuation equilibrium under each observed performance-weight regime. We first define production, pay, costs, and objectives; we then state the sequence of moves and solve the continuation subgame by backward induction. Throughout, $b_k := \bar{\lambda} p_k$ denotes the private return per unit of output in task $k \in \{R, O, S\}$, and we use the compact notation $A_i := A_R(m, \theta_i, \tau_i^R)$, $D_i := D_i(m)$, $\tilde{c}_i := \tilde{c}(\theta_i, \tau_i^A)$, and $c_{Oi} := c_O(\theta_i, \tau_i^A)$ defined below.

B.1. Setup

Workers, tasks, and production. The firm employs a unit mass of account managers indexed by $i \in [0, 1]$. Workers differ in adaptation cost $\theta_i \in [\underline{\theta}, \bar{\theta}]$, distributed according to F with density $f > 0$. A higher θ_i captures more established routines, a larger analytical skill gap, and a greater cost of moving toward complex work. We proxy θ_i with tenure in the empirical analysis. The firm observes F , while branch managers observe individual workers' types through day-to-day interaction.

There are three tasks: routine retail work (R), other complex work (O), and SME lending (S); every worker allocates effort $(e_{Ri}, e_{Oi}, e_{Si}) \geq 0$ across all three. Outputs are

$$y_{Oi} = e_{Oi}, \quad y_{Si} = e_{Si}, \quad y_{Ri} = D_i(m) g(x_i), \quad x_i := \frac{A_R(m, \theta_i, \tau_i^R) e_{Ri}}{D_i(m)} \in [0, 1), \quad (\text{A.1})$$

where x_i is worker i 's capacity utilization and g is twice continuously differentiable, increasing, and strictly concave on $[0, 1)$, with $g(0) = 0$ and $g'(x) \rightarrow 0$ as $x \uparrow 1$, so the marginal return to routine effort vanishes as utilization approaches one. The technology m enters through two channels. First, a *productivity channel*: $A_R(m, \theta_i, \tau_i^R)$ converts routine effort into capacity units, with $A_{R,m} > 0$. Second, a *client demand pool channel*: $D_i(m)$ is worker i 's client book, with $D'_i(m) > 0$, so the app expands the pool of clients seeking routine services.²⁶ Routine coaching τ_i^R raises A_R , $A_{R,\tau^R} > 0$. Let the capacity-utilization elasticity be

²⁶We refer to $D_i(m)$ as the client demand pool scale, not the literal maximum number of clients servable; the specification does not require $\lim_{x \uparrow 1} g(x) = 1$.

$$\varepsilon(x) := -\frac{g''(x)x}{g'(x)} \geq 0, \quad (\text{A.2})$$

which measures the percent drop in marginal conversion per percent increase in utilization.²⁷

Empirical status of demand. The data contain completed trades and realized client-activity outcomes—total and first-time clients, active accounts, and account balances—but not $D_i(m)$, x_i , potential clients, unmet demand, or workers' demand-exhaustion status. The observed outcomes therefore provide only indirect, supply-side proxies for demand. Whether a worker faces unmet demand or not is therefore a theoretical model object. The results below characterize how behavior differs depending on whether simple-product demand is exhausted; the empirical analysis cannot classify workers on this basis or directly test the state-dependent mechanism in Prediction 2.

KPI Score and pay. The worker's KPI score and expected pay are

$$Sc_i = p_R y_{Ri} + p_O y_{Oi} + p_S y_{Si}, \quad W_i = \bar{\lambda} Sc_i, \quad (\text{A.3})$$

where $p = (p_R, p_O, p_S)$ denotes the KPI vector and $\bar{\lambda} > 0$ is the common performance-pay sensitivity.

Costs. Worker costs are

$$C_i = \frac{e_{Ri}^2}{2} + \frac{c_O(\theta_i, \tau_i^A) e_{Oi}^2}{2} + \frac{\tilde{c}(\theta_i, \tau_i^A) e_{Si}^2}{2} + \frac{\chi}{2} (e_{Ri} + e_{Oi} + e_{Si})^2, \quad (\text{A.4})$$

²⁷A hard-cap specification $y_{Ri} = \min\{A_R e_{Ri}, D_i\}$ has constant marginal productivity below the cap and a discontinuous drop at the cap. The smooth specification allows marginal productivity to decline gradually before full utilization; only the smooth model can attribute a joint rise in routine output and analytical effort to the same technology shock, through a utilization elasticity $\varepsilon(x_i) > 1$. The model includes only intensive-margin worker responses: SME activity is captured by the continuous effort choice $e_{Si} \geq 0$, without a separate participation indicator. The model therefore speaks to changes in SME effort intensity rather than to first-time entry into SME lending or to a mass point at zero SME activity; Section 6.2 discusses the corresponding empirical margins.

where $\chi \geq 0$ captures the fact that all tasks draw on the same workday: when $\chi > 0$, taking on more work in one activity raises the cost of supplying additional effort in the others. The cost coefficients satisfy $\tilde{c}_\theta > 0$, $\tilde{c}_{\tau^A} < 0$, $c_{O,\theta} \geq 0$, $c_{O,\tau^A} \leq 0$, and $\tilde{c}, c_O > 0$ on the feasible coaching domain; for brevity we write $\tilde{c}_i := \tilde{c}(\theta_i, \tau_i^A)$ and $c_{Oi} := c_O(\theta_i, \tau_i^A)$. Analytical coaching τ_i^A lowers the cost of other complex work and the marginal cost of SME effort. Both other complex products and SME lending rely on analytical judgment—interpreting client information, matching customers to financial products, and, for SME lending, evaluating a firm’s cash flows and repayment capacity—which is why we allow the same coaching input to reduce the cost of both.

Given the KPI vector and coaching inputs, worker i chooses effort to solve

$$\max_{e_{Ri}, e_{Oi}, e_{Si} \geq 0} \bar{\lambda} [p_R D_i(m) g(x_i) + p_O e_{Oi} + p_S e_{Si}] - C_i, \quad (\text{A.5})$$

subject to $x_i < 1$. We denote the resulting effort choices by $e_i^*(p, \tau_i^A, \tau_i^R; m)$.

Branch objective and game sequence. branch managers act on behalf of the branch and maximize branch surplus by allocating coaching across workers and coaching types:

$$\Pi^*(p, T_M; m) = \max_{\{\tau_i^A, \tau_i^R\}} \int_0^1 [\alpha_R y_{Ri} + \beta_O y_{Oi} + \beta_S(m) y_{Si} - C_i] di \quad \text{s.t.} \quad \int_0^1 (\tau_i^A + \tau_i^R) di \leq T_M, \quad (\text{A.6})$$

with $\beta_S(m) = \beta_{S0} + \gamma m$, $\gamma > 0$, and $\alpha_R, \beta_O > 0$. The increase in $\beta_S(m)$ captures that the larger client base generated by the app raises the branch value of SME lending. The variable T_M is effective coaching capacity—coaching hours weighted by branch manager quality. Because wages are transfers within the branch and worker effort costs enter the objective directly, (A.6) is a total-surplus objective. The model takes the performance-weight vector as given when branch managers allocate coaching.

The model conditions on the three regimes observed in the data: old technology with p^0 , new technology with p^0 , and new technology with p^1 . The timing of the shift from p^0 to p^1 is taken from the data rather than explained within the framework. Within each regime, branch managers allocate coaching across workers and coaching types, anticipating workers’ subse-

quent effort choices; workers then choose effort across tasks. Conditional on each observed (m, p) pair, we solve this continuation game by backward induction: we first derive workers' effort choices from (A.5) and then solve branch managers' coaching allocation using (A.6). Appendix B.6 shows why the gain from the redesign rises with technology. The staggered rollout of HR training and KPI explanations across branches is an implementation process outside the framework.

Assumptions. Let $\gamma_A(\theta_i) := A_{R,m}/A_R \geq 0$ and $\gamma_D := D'_i/D_i \geq 0$ denote the log-derivatives of routine task productivity and the client demand pool with respect to m . The four assumptions in Section 5 are encoded by the production and cost functions above and by the following restrictions.

Assumption 1 (Technology and simple-product demand). (i) The app raises routine task productivity and expands the client demand pool: $A_{R,m} > 0$ and $D'_i(m) > 0$. (ii) $\partial\gamma_A(\theta_i)/\partial\theta_i \leq 0$: the proportional productivity gain is weakly smaller for higher- θ workers. (iii) γ_D is common across workers.

Assumption 2 (Multitasking and limited time). The parameter $\chi \geq 0$ captures time trade-offs across tasks. When $\chi > 0$, effort in one task raises the marginal cost of effort in the others; when $\chi = 0$, task-specific effort costs are separable.

Assumption 3 (Adaptation and coaching). (i) The analytical-capability indices $q_S := 1/\tilde{c}(\theta_i, \tau_i^A)$ and $q_O := 1/c_O(\theta_i, \tau_i^A)$ are weakly concave in τ_i^A for every worker type, and the marginal return to routine coaching MR_i^R is weakly decreasing in τ_i^R . (ii) Among workers for whom both coaching types are feasible, the ratio MR_i^R/MR_i^A is strictly increasing in θ_i . (iii) Analytical coaching is weakly less effective for higher- θ workers: $\tilde{c}_{\theta\tau^A} \geq 0$ and $c_{O,\theta\tau^A} \geq 0$.²⁸

²⁸For fixed performance weights, technology, and coaching inputs, let π_i^* denote worker i 's contribution to branch surplus evaluated at optimal effort. We define $MR_i^A := \partial\pi_i^*/\partial\tau_i^A$ and $MR_i^R := \partial\pi_i^*/\partial\tau_i^R$; Appendix B derives both from primitives.

Assumption 4 (Old SME reward and regularity bounds). For each performance-weight vector considered and every relevant m , $\bar{\lambda}p_R < 2\alpha_R$, $\bar{\lambda}p_O < 2\beta_O$, and $\bar{\lambda}p_S < 2\beta_S(m)$. In the main-text application, the old SME weight is below the surplus-aligned level, $p_S^0 < p_S^{\text{align}}(m) := \beta_S(m)/\bar{\lambda}$. Proposition A.5(i) records the comparative static outside this focal range.²⁹

Assumption 1 separates the productivity channel, which generates a tenure gradient, from the tenure-neutral client demand pool channel. Assumption 2 determines whether changes in simple work affect the other two tasks through limited time. Assumption 3, together with the cost signs in Eq. (A.4), makes analytical coaching relatively more effective for low- θ workers and routine coaching relatively more attractive for high- θ workers.³⁰ Assumption 4 ensures that each task generates nonnegative surplus over the relevant range.

B.2. Worker behavior

We first characterize the worker's effort choices at any $\chi \geq 0$; every later result builds on this system.

Worker's first-order conditions. The worker maximizes $V_i = \bar{\lambda}p_R y_{Ri} + b_O e_{Oi} + b_S e_{Si} - C_i$ over $(e_{Ri}, e_{Oi}, e_{Si}) \geq 0$. Writing $\tilde{b}_{Ri} := \bar{\lambda}p_R A_i g'(x_i^*)$ for the congestion-adjusted routine return and $\mu_i := \chi(e_{Ri} + e_{Oi} + e_{Si})$ for the shadow cost of total effort, the interior first-order conditions are

$$e_{Ri}^* = \tilde{b}_{Ri} - \mu_i, \quad e_{Oi}^* = \frac{b_O - \mu_i}{c_{Oi}}, \quad e_{Si}^* = \frac{b_S - \mu_i}{\tilde{c}_i}, \quad (\text{A.7})$$

with shadow cost

$$\mu_i = \frac{\chi(\tilde{b}_{Ri} + b_O/c_{Oi} + b_S/\tilde{c}_i)}{1 + \chi(1 + 1/c_{Oi} + 1/\tilde{c}_i)}. \quad (\text{A.8})$$

Equations (A.7)–(A.8) are an implicit system: $\tilde{b}_{Ri}$ depends on $x_i^* = A_i e_{Ri}^*/D_i$, which depends on e_{Ri}^* itself. The routine equation $e_{Ri}^* = \tilde{b}_{Ri} - \mu_i$ is therefore a fixed-point condition; the other-complex and SME equations are closed form conditional on μ_i , which in turn depends on $\tilde{b}_{Ri}$

²⁹ p^0 and p^1 are observed institutional primitives.

³⁰ Because each worker's contribution to branch surplus is linear in q_S and q_O under Assumption 4, and the task blocks are additively separable across workers and coaching types at $\chi = 0$, Assumption 3(i) makes branch managers' coaching-allocation problem concave at $\chi = 0$.

and hence on the routine equilibrium. When any unconstrained effort is negative it is set to zero and the shadow cost is recomputed from the active tasks' conditions.

B.3. Technology under old performance priorities

Holding $p = p^0$ fixed, the app raises routine output through higher productivity and a larger client book. Because neither the SME weight nor the analytical cost coefficients change, technology alone leaves the returns to SME and other complex work unchanged; any response of analytical effort can arise only through the shared time channel. The effect on other complex work depends on initial routine-capacity utilization. The key trade-off is whether the client book expansion absorbs the time saved by the app.

Proposition A.1 (Technology under the old vector: full responses). Fix KPI weights at (p_R^0, p_O, p_S^0) and hold coaching fixed.

- (i) Analytical-effort invariance at $\chi = 0$. When $\chi = 0$, $e_{Oi}^* = \bar{\lambda} p_O / c_{Oi}$ and $e_{Si}^* = \bar{\lambda} p_S^0 / \tilde{c}_i$; both analytical effort choices are independent of m , so other-complex and SME output are unchanged by the technology.
- (ii) Effort reallocation at $\chi > 0$. At an interior optimum, define $\Delta_i := c_{Oi} \tilde{c}_i + \chi(c_{Oi} + \tilde{c}_i) > 0$.
Then

$$\begin{aligned} \text{sign} \left(\frac{\partial e_{Ri}}{\partial m} \right) &= \text{sign} [\gamma_A(\theta_i)(1 - \varepsilon(x_i)) + \gamma_D \varepsilon(x_i)], \\ \frac{\partial e_{Oi}}{\partial m} &= -\frac{\chi \tilde{c}_i}{\Delta_i} \frac{\partial e_{Ri}}{\partial m}, \quad \frac{\partial e_{Si}}{\partial m} = -\frac{\chi c_{Oi}}{\Delta_i} \frac{\partial e_{Ri}}{\partial m}. \end{aligned}$$

For $\chi > 0$, both multipliers are strictly negative, so both analytical effort choices move in the opposite direction from routine effort.

Proposition A.1 isolates two effects of the technology, holding the KPI vector and coaching fixed. Part (i) shows that at $\chi = 0$ the two analytical tasks are fully shielded: because effort costs are separable, neither e_{Oi}^* nor e_{Si}^* depends on m , so any change in other-complex or SME output under the old KPI vector must run through $\chi > 0$. Part (ii) signs the routine-effort

response once $\chi > 0$ and identifies two opposing forces. When utilization is low ($\varepsilon(x_i) < 1$), so the worker has not exhausted demand, both terms in the sign condition are positive: productivity growth and client demand pool growth both raise routine effort, which crowds out the other two tasks through the shared workload cost. When utilization is high ($\varepsilon(x_i) > 1$), so the worker has exhausted demand, and $\gamma_A(\theta_i)$ is large relative to γ_D , the first term turns negative: the productivity gain lowers the routine effort needed to serve the existing book, and the released time flows into other-complex and SME work. This configuration is unavailable under a hard-cap specification, in which routine effort and output move together below the cap. The multipliers $\kappa_{Oi} = \chi \tilde{c}_i / \Delta_i$ and $\kappa_{Si} = \chi c_{Oi} / \Delta_i$ are strictly positive whenever $\chi > 0$, so the two analytical tasks always move together and in the opposite direction from routine effort.

The technology raises routine *output* even when it lowers routine *effort*. In the separable benchmark $\chi = 0$, define $\eta(x) := xg'(x)/g(x) \in (0, 1]$. Holding KPI weights and coaching fixed,

$$\frac{d \ln y_{Ri}}{dm} = \gamma_D \left[1 - \frac{\eta(x_i)}{1 + \varepsilon(x_i)} \right] + 2\gamma_A(\theta_i) \frac{\eta(x_i)}{1 + \varepsilon(x_i)} > 0. \quad (\text{A.9})$$

Routine output therefore rises with the technology for every worker. If $H(x) := \eta(x)/(1 + \varepsilon(x))$ is weakly decreasing in utilization and $2\gamma_A(\theta_i) > \gamma_D$, the proportional routine-output response is smaller for workers with higher initial utilization, and the difference in responses between two workers who differ only in γ_A is proportional to $H(x)$ and so narrows as utilization rises (Corollary A.1).

Corollary A.1 (Routine-output response and utilization). Fix KPI weights at (p_R^0, p_O, p_S^0) , hold coaching fixed, and suppose $\chi = 0$. Let $\eta(x) := xg'(x)/g(x)$ and $H(x) := \eta(x)/(1 + \varepsilon(x))$, as in the main text.

- (i) Routine output rises with technology for every worker: $d \ln y_{Ri} / dm = \gamma_D [1 - H(x_i)] + 2\gamma_A(\theta_i) H(x_i) > 0$, which is Eq. (A.9).
- (ii) If H is weakly decreasing in utilization and $2\gamma_A(\theta_i) > \gamma_D$, the response in part (i) is weakly decreasing in initial utilization: the response is smaller for workers with higher initial utilization.

- (iii) For two workers who differ only in γ_A , with $\theta_L < \theta_H$, the tenure gradient in the routine response, $2[\gamma_A(\theta_L) - \gamma_A(\theta_H)]H(x_i)$, shrinks as utilization rises.

Corollary A.1 is the output-side counterpart to Proposition A.1: routine output rises unconditionally, but by how much depends on utilization and on tenure through $\gamma_A(\theta_i)$. Parts (ii)–(iii) say that the technology response is compressed, both across the population and across tenure, exactly where utilization is highest, that is, for workers who have exhausted their demand. The sign of the average other-complex response remains outside the scope of this corollary: it depends on the distribution of utilization and on the empirical ratio $\gamma_A(\theta_i)/\gamma_D$, neither of which is pinned down by theory alone.

Worker type θ_i shapes the response through two channels. First, it raises the cost of analytical work: at $\chi = 0$, $e_{Si}^* = \bar{\lambda}p_S/\tilde{c}(\theta_i, \tau_i^A)$ and $e_{Oi}^* = \bar{\lambda}p_O/c_O(\theta_i, \tau_i^A)$, so higher- θ workers perform less analytical work at any given incentives and coaching. Second, by Assumption 1(i) the lower- θ worker receives the larger productivity gain and so has the larger routine-output response at every utilization level. The implication for analytical effort depends on utilization: at low utilization ($\varepsilon(x_i) < 1$), where the worker has not exhausted demand, the larger productivity gain raises routine effort by more and, when $\chi > 0$, crowds out both other-complex and SME effort; at high utilization ($\varepsilon(x_i) > 1$), where the worker has exhausted demand, it reduces the routine effort needed to serve the expanded book, releasing more time for both analytical tasks. Conditional on utilization, the other-complex and SME responses are therefore more positive for the lower- θ worker when the worker has exhausted demand (Corollary A.2).

Corollary A.2 (Cross-worker technology responses). Fix KPI weights at (p_R^0, p_O, p_S^0) . Consider two workers at the same utilization level who differ only in $\gamma_A(\theta_i)$, with $\gamma_A(\theta_j) > \gamma_A(\theta_k)$ for $\theta_j < \theta_k$.

- (i) Worker j has a strictly larger routine-output response at every utilization level.
- (ii) If $\chi > 0$, worker j has larger responses in both other-complex and SME effort if and only if $\varepsilon(x_i) > 1$. At low utilization, $\varepsilon(x_i) < 1$, both analytical-effort orderings reverse. When $\chi = 0$, both analytical responses to technology are zero.

Corollary A.2 provides the formal mechanism behind Prediction 2. Its sign change is conditional on utilization and on whether simple-product demand is largely served.

B.4. Performance redesign

The Task Intervention changes the KPI vector from the old to the redesigned value:

$$p^0 = (p_R^0, p_O, p_S^0) \quad \text{to} \quad p^1 = (p_R^1, p_O, p_S^1), \quad p_R^1 < p_R^0, \quad p_S^1 > p_S^0. \quad (\text{A.10})$$

Proposition A.2 (The KPI redesign). Hold technology and coaching fixed and consider the redesign in Eq. (A.10), $\Delta p_S > 0$, $\Delta p_R < 0$, $\Delta p_O = 0$. Parts (i)–(iii) suppose $\chi = 0$. Parts (iv)–(vii) suppose $\chi > 0$ and interior effort choices with an unchanged set of active tasks, and consider differential changes $dp_S > 0$, $dp_R < 0$; at a corner solution the corresponding strict inequalities become weak.

(i) SME effort and output. At the old weight, $y_{Si}^0 = e_{Si}^0 = \bar{\lambda} p_S^0 / \tilde{c}_i$. The redesign implies

$$\Delta e_{Si} = \Delta y_{Si} = \frac{\bar{\lambda} \Delta p_S}{\tilde{c}_i} > 0, \quad \frac{\Delta y_{Si}}{y_{Si}^0} = \frac{\Delta p_S}{p_S^0}. \quad (\text{A.11})$$

The level response is strictly decreasing in θ_i , but the proportional response is common across workers.

(ii) Routine and other complex work. $\partial y_{Ri} / \partial p_R = A_R(m, \theta_i, \tau_i^R) g'(x_i) \partial e_{Ri} / \partial p_R > 0$, so $\Delta p_R < 0$ lowers routine output, with a muted output response at high utilization. Because costs are separable and p_O is unchanged, $\Delta e_{Oi} = 0$.

(iii) Wage effect of the routine-weight reduction. Holding p_S fixed,

$$\frac{dW_i}{dp_R} = \bar{\lambda} \left[y_{Ri} + p_R A_R g'(x_i) \frac{\partial e_{Ri}}{\partial p_R} \right] > 0,$$

so $\Delta p_R < 0$ lowers wages through the routine component.

(iv) Effort responses.

$$\begin{aligned} \frac{\partial e_{Si}}{\partial p_S} &> 0, & \frac{\partial e_{Ri}}{\partial p_S} &< 0, & \frac{\partial e_{Oi}}{\partial p_S} &< 0, \\ \frac{\partial e_{Ri}}{\partial p_R} &> 0, & \frac{\partial e_{Si}}{\partial p_R} &< 0, & \frac{\partial e_{Oi}}{\partial p_R} &< 0. \end{aligned}$$

Consequently the joint redesign raises SME effort and lowers routine effort, $de_{Si} > 0$ and $de_{Ri} < 0$; its effect on other-complex effort is ambiguous, since the higher SME weight crowds it out while the lower routine weight frees time for it.

(v) Routine output. Holding technology fixed, $dy_{Ri} = A_i g'(x_i) de_{Ri} < 0$ under the joint change. The reduction in routine output is small when the worker is close to capacity and $g'(x_i)$ is small, even if routine effort falls substantially.

(vi) Wages. A reduction in p_R lowers the worker's wage whenever

$$y_{Ri} + p_R A_i g'(x_i) \frac{\partial e_{Ri}}{\partial p_R} > -p_O \frac{\partial e_{Oi}}{\partial p_R} - p_S \frac{\partial e_{Si}}{\partial p_R}. \quad (\text{A.12})$$

This condition holds at sufficiently high utilization.

(vii) Differences across workers. Holding the routine and other-complex primitives fixed, the increase in SME effort is larger for workers with lower $\tilde{c}_i$. Since $\tilde{c}_\theta > 0$, the response is therefore larger for lower- θ workers. More generally, for any two types whose SME responses are strictly ordered at $\chi = 0$, that ordering continues to hold for all sufficiently small $\chi > 0$.

Part (i) is the result we use most. Because every worker expands SME activity by the same factor $\Delta p_S / p_S^0$, estimates of proportional responses need not display a tenure gradient on the targeted task, even though level responses decline with adaptation cost.³¹ The distinction between routine effort and routine output in part (ii) also matters empirically: near capacity, $g'(x_i)$ is small, so a sizable reduction in routine effort generates only a modest decline in routine output.

³¹At $\chi > 0$ the exact common-proportional result no longer holds; the level ordering in $\tilde{c}_i$ persists for all sufficiently small $\chi > 0$ by continuity.

The joint change $dp_S > 0, dp_R < 0$ raises SME effort and lowers routine effort by part (iv); its effect on other-complex effort is ambiguous, since the higher SME weight crowds it out while the lower routine weight frees time for it. Condition (A.12) requires the routine-pay loss to exceed the combined reallocation gain toward other-complex and SME work; it holds at sufficiently high utilization, where the induced effort responses vanish as $g'(x_i) \rightarrow 0$ while the direct loss from the lower routine weight remains.

The surplus-aligned SME weight and the SME incentive wedge are:

$$p_S^{\text{align}}(m) := \frac{\beta_S(m)}{\bar{\lambda}}, \quad \omega_S(m) := p_S^{\text{align}}(m) - p_S^0. \quad (\text{A.13})$$

Remark A.1 (Alignment for any χ). If every weight equals its surplus-aligned level, $\bar{\lambda}p_R = \alpha_R$, $\bar{\lambda}p_O = \beta_O$, and $\bar{\lambda}p_S = \beta_S(m)$, the worker's objective $\bar{\lambda}[p_R y_{Ri} + p_O e_{Oi} + p_S e_{Si}] - C_i$ coincides term by term with the worker's contribution to branch surplus $\alpha_R y_{Ri} + \beta_O y_{Oi} + \beta_S(m) y_{Si} - C_i$, for any $\chi \geq 0$; private effort choices then maximize branch surplus. Away from this benchmark, the cross-price signs in Proposition A.2(iv) imply that a lower weight on one task raises effort in the other two: an underpriced task loses the competition for the workday, the property used in Section 5.3.

Remark A.2 (The wage decomposition is an identity). For any two KPI vectors p^0, p^1 and the associated outputs y^0, y^1 ,

$$W_i^1 - W_i^0 = \bar{\lambda}[p^1 \cdot y_i^1 - p^0 \cdot y_i^0] = \bar{\lambda}(p^1 - p^0) \cdot y_i^0 + \bar{\lambda}p^1 \cdot (y_i^1 - y_i^0) \quad (\text{A.14})$$

B.5. Coaching and branch managers

The redesign also changes the return to support by branch managers. Analytical coaching lowers $\tilde{c}_i$, the effective cost of SME effort; a higher SME weight gives workers a stronger incentive to use the capability this coaching creates, increasing the SME production induced by a given unit of branch manager support.

branch managers' marginal returns. At $\chi = 0$, the marginal return to analytical coaching is

$$MR_i^A = \frac{\partial \pi_i^*}{\partial \tau_i^A} = (-\tilde{c}_{\tau^A}) \frac{\bar{\lambda} p_S}{\tilde{c}_i^2} \left[\beta_S(m) - \frac{\bar{\lambda} p_S}{2} \right] + (-c_{O,\tau^A}) \frac{b_O}{c_{O,i}^2} \left[\beta_O - \frac{b_O}{2} \right]. \quad (\text{A.15})$$

The first term captures the SME intensive-effort channel; the second is the capability spillover into other-complex work. Under $\tilde{c}_{\tau^A} < 0$, $c_{O,\tau^A} \leq 0$, $\bar{\lambda} p_S < 2\beta_S(m)$, and $\bar{\lambda} p_O < 2\beta_O$ (Assumption 4), the SME term is strictly positive and the other-complex term is weakly positive, so $MR_i^A > 0$. By Assumption 3, MR_i^A is strictly decreasing in θ_i . The expression is strictly increasing in p_S under the stronger condition $\bar{\lambda} p_S < \beta_S(m)$ imposed in Proposition A.3. Define $\Omega_S(p_S; m) := \beta_S(m) \bar{\lambda} p_S - (\bar{\lambda} p_S)^2/2$, so that branch SME surplus per worker equals $\Omega_S(p_S; m)/\tilde{c}_i$ and MR_i^A inherits the same Ω_S -structure from the derivative with respect to τ_i^A .

The marginal return to routine coaching at $\chi = 0$ is derived by differentiating the branch objective in τ_i^R with the worker's induced routine-effort and utilization responses fully taken into account. Let $q_i := 1 - b_R A_i^2 g''(x_i)/D_i$; since $g'' < 0$, $q_i > 0$. Re-differentiating the routine first-order condition $e_{Ri} = b_R A_i g'(x_i)$ with respect to A_i (holding all other primitives fixed) yields

$$\frac{\partial e_{Ri}}{\partial A_i} = \frac{b_R g'(x_i) [1 - \varepsilon(x_i)]}{q_i}.$$

Applying the chain rule together with the envelope theorem for the worker's remaining decisions gives

$$MR_i^R = \frac{\partial \pi_i^*}{\partial \tau_i^R} = A_{R,\tau^R} b_R A_i [g'(x_i)]^2 \left[\alpha_R + (\alpha_R - b_R) \frac{1 - \varepsilon(x_i)}{q_i} \right]. \quad (\text{A.16})$$

At the $\chi = 0$ worker optimum, x_i satisfies $b_R A_i^2 g'(x_i)/D_i = x_i$; hence $b_R A_i^2 g''(x_i)/D_i = x_i g''(x_i)/g'(x_i)$, which equals $-\varepsilon(x_i)$, and therefore $q_i = 1 + \varepsilon(x_i)$. Substituting collapses (A.16) to the equivalent form

$$MR_i^R = A_{R,\tau^R} \frac{b_R A_i [g'(x_i)]^2}{1 + \varepsilon(x_i)} \left[2\alpha_R - b_R (1 - \varepsilon(x_i)) \right], \quad (\text{A.17})$$

which displays both the induced routine-effort response and the utilization-elasticity adjustment through the single factor $1/(1 + \varepsilon)$. Assumption 3 and Corollary A.3 compare this fully optimized return with MR_i^A , and all subsequent results use the fully optimized measure of MR_i^R .

Proposition A.3 (KPI-coaching complementarity). Suppose $\chi = 0$ and Assumption 4 holds. Hold p_R and p_O fixed; where two SME weights appear, suppose $p_S^0 < p_S^1 \leq p_S^{\text{align}}(m)$, with $p_S^{\text{align}}(m)$ the surplus-aligned weight defined in (A.13).

- (i) Worker-level complementarity. Consider an SME weight satisfying $0 < p_S \leq p_S^{\text{align}}(m)$. The SME component of the marginal return to analytical coaching is

$$MR_{i,S}^A(p_S) = (-\tilde{c}_{\tau^A}) \frac{\bar{\lambda} p_S}{\tilde{c}_i^2} \left[\beta_S(m) - \frac{\bar{\lambda} p_S}{2} \right], \quad \frac{\partial MR_{i,S}^A}{\partial p_S} = (-\tilde{c}_{\tau^A}) \frac{\bar{\lambda}}{\tilde{c}_i^2} [\beta_S(m) - \bar{\lambda} p_S] \geq 0,$$

with strict inequality whenever $p_S < p_S^{\text{align}}(m)$.

- (ii) old return. At p_S^0 , the SME component of the marginal return to analytical coaching is $(-\tilde{c}_{\tau^A}) \Omega_S(p_S^0; m) / \tilde{c}_i^2 > 0$ by Assumption 4.
- (iii) Branch-level complementarity. The function

$$\mathcal{D}_S(T_M) := \Pi^*(p_R, p_O, p_S^1, T_M; m) - \Pi^*(p_R, p_O, p_S^0, T_M; m)$$

is weakly increasing in T_M . The increase is strict on any range of coaching capacity over which the shadow value of capacity is strictly larger under p_S^1 than under p_S^0 ; in particular, this occurs when marginal capacity is used for analytical coaching at a worker for whom $MR_i^A(p_S^1) > MR_i^A(p_S^0)$.

- (iv) Conditional shift in coaching composition. If $MR_i^A(p_S^0) \leq MR_i^R < MR_i^A(p_S^1)$ for a positive measure of workers, the redesign shifts the preferred coaching content from routine toward analytical coaching for that group.

The SME weight and analytical coaching are therefore complements below the surplus-aligned weight $p_S^{\text{align}}(m)$: the return to analytical coaching may already be positive under the

old KPI—both through SME lending and through its spillover into other complex work—and the redesign raises it.

Under the full redesign, in which p_S rises and p_R falls, the same aggregate complementarity holds provided that the gain in analytical-coaching value exceeds the decline in routine-coaching value. The result is a complementarity between the SME incentive and *effective coaching capacity*: additional capacity is more valuable after the SME weight rises whenever the marginal coaching unit is used on the analytical margin.

Proposition A.3(ii) shows that the old return can already be strictly positive because the same coaching input also affects other-complex work through $c_{O,\tau^A} \leq 0$ and coexists with routine coaching. Branch managers can therefore respond to changes in p_S even when the old return is positive.

Corollary A.3 (Allocation and content of coaching). Suppose $\chi = 0$, the coaching-capacity constraint binds, and Assumption 3 holds.

- (i) There exists a cutoff $\tilde{\theta}$ such that workers with $\theta_i < \tilde{\theta}$ receive analytical coaching, workers with $\theta_i > \tilde{\theta}$ do not, and the cutoff type may receive a partial allocation.
- (ii) If routine coaching is weakly preferred at the margin for a worker of type θ_i , it is also weakly preferred for every higher type.

The corollary separates the amount of branch manager attention from its content. Lower- θ workers are the first candidates for analytical coaching because coaching reduces their analytical costs by more; as adaptation difficulty rises, routine-oriented coaching becomes relatively more attractive. A branch manager with greater effective capacity extends coaching further down the ordering, so cross-branch differences in T_M show up as differences in analytical coaching—and in the complex-task responses it supports—concentrated among more adaptable workers. Because $c_{O,\tau^A} \leq 0$, the same coaching that supports SME lending also lowers the cost of other complex work, a capability spillover on a margin whose KPI weight does not change. The framework can therefore generate similar total coaching across tenure groups but different coaching composition. We examine this prediction empirically in Section 6.

Proposition A.4 (Local worker-level KPI-coaching complementarity with $\chi > 0$). In addition to the maintained assumptions, suppose $g \in C^3$ on $[0, 1)$.³² Fix $(p_S^0, \tau_i^{A,0})$ with

$$0 < p_S^0 < p_S^{\text{align}}(m) := \frac{\beta_S(m)}{\bar{\lambda}},$$

and suppose that at $(p_S^0, \tau_i^{A,0}, \chi = 0)$ worker i 's effort solution e_i^* is interior and satisfies $x_i^* < 1$. Then there exist $\bar{\chi} > 0$ and an open rectangle $\mathcal{N} \ni (p_S^0, \tau_i^{A,0})$ such that MR_i^A is well defined and C^1 on $\mathcal{N} \times [0, \bar{\chi})$ and

$$\frac{\partial MR_i^A}{\partial p_S}(p_S, \tau_i^A, \chi) > 0 \quad \text{for all } (p_S, \tau_i^A, \chi) \in \mathcal{N} \times [0, \bar{\chi}).$$

In particular, the worker-level branch surplus π_i^* has strictly increasing differences in (p_S, τ_i^A) on $\mathcal{N}$. Hence the SME KPI weight and analytical coaching are strict local complements for sufficiently small values of the common workload term χ .

Unlike the $\chi = 0$ case, the condition

$$p_S \leq \frac{\beta_S(m)}{\bar{\lambda}}$$

is not by itself sufficient to guarantee complementarity globally when $\chi > 0$, because p_S affects all three equilibrium effort choices through the common workload term.

B.6. The gain from the redesign

The framework characterizes why the redesign becomes more valuable after technology adoption while taking the observed redesign date as given. Let the branch-level redesign gain be

$$\mathcal{D}(m, T_M) \equiv \Pi^*(p^1, T_M; m) - \Pi^*(p^0, T_M; m) \tag{A.18}$$

and recall the SME incentive wedge $\omega_S(m)$ from (A.13). The app raises $\beta_S(m)$ and so widens

³²This strengthens the standing assumption $g \in C^2$, $g' > 0$, $g'' < 0$. It is used only here: $\partial MR_i^A / \partial p_S$ involves the cross derivative $\partial^2 e_i^* / \partial p_S \partial \tau_i^A$, which differentiates M_i along e_i^* and hence involves g''' . No first-order comparative static in the main text requires it.

the wedge: the technology increases the surplus at stake in the redesign. Over the range $p_S^0 < p_S^1 \leq p_S^{\text{align}}(m)$, the SME component of $\mathcal{D}(m, T_M)$ is weakly increasing in m , and strictly so in the representative-worker separable benchmark, where the benchmark gain is proportional to $\omega_S(m)^2$ (Lemma A.1).

Lemma A.1 (Representative-agent SME-component benchmark). In a representative-worker, $\chi = 0$ setting, let $p_S^{\text{align}}(m) = \beta_S(m)/\bar{\lambda}$ be the surplus-aligned SME weight of Eq. (A.13). Let $\mathcal{D}_{\text{bench}}^S(m)$ denote the SME component of the redesign surplus gain, i.e. the difference in branch SME surplus between $p_S^1 = p_S^{\text{align}}(m)$ and p_S^0 , holding p_R and p_O fixed. Over the range in which $p_S^{\text{align}}(m) > p_S^0$,

$$\mathcal{D}_{\text{bench}}^S(m) = \frac{\bar{\lambda}^2}{2\bar{c}}(p_S^{\text{align}}(m) - p_S^0)^2 = \frac{\bar{\lambda}^2}{2\bar{c}}\omega_S(m)^2$$

is strictly increasing in m .

The full redesign gain $\mathcal{D}(m, T_M)$ additionally includes a routine-surplus component driven by $p_R \downarrow$ and, at $\chi > 0$, cross-task interaction terms.

Proposition A.5 (Comparative statics of the branch-level redesign gain). Consider the branch-level surplus gain $\mathcal{D}(m, T_M) = \Pi^*(p^1, T_M; m) - \Pi^*(p^0, T_M; m)$ evaluated at the KPI vectors of Eq. (A.10). Decompose $\mathcal{D}(m, T_M) = \mathcal{D}^S(m, T_M) + \mathcal{D}^R(m, T_M) + \mathcal{D}^X(m, T_M)$ into an SME-channel component (driven by $p_S \uparrow$), a routine-channel component (driven by $p_R \downarrow$), and a cross-task interaction (present only at $\chi > 0$). Then:

- (i) If $p_S^{\text{align}}(m) \leq p_S^0$, the observed increase $p_S^1 > p_S^0$ weakly reduces the SME-channel component: $\mathcal{D}^S(m, T_M) \leq 0$. If the firm were instead free to retain the old SME weight, its optimal SME-side choice in this region would be $p_S^1 = p_S^0$, yielding a counterfactual SME redesign gain of zero.
- (ii) Over the range in which $p_S^{\text{align}}(m) > p_S^0$, $\mathcal{D}^S(m, T_M)$ is weakly increasing in m ; in the representative-worker, $\chi = 0$ benchmark of Lemma A.1, the increase is strict.
- (iii) $\mathcal{D}(m, T_M)$ is weakly increasing in T_M when the surplus gain from additional effective coaching capacity on the SME (analytical) margin exceeds any reduction in the value of

routine coaching caused by $p_R \downarrow$; this condition is the joint-change extension of Proposition A.3(iii), which holds p_R fixed.

The proposition is a comparative-statics statement about the branch-level object $\mathcal{D}(m, T_M)$.

B.7. Proofs

Proof of Proposition A.1. Part (i). At $\chi = 0$, Eq. (A.7) gives $e_{Si}^* = \bar{\lambda} p_S^0 / \tilde{c}_i$, which does not depend on m ; SME output is therefore invariant to the technology. The same argument applies to e_{Oi}^* .

Part (ii). The first-order conditions are

$$\begin{aligned}\bar{\lambda} p_R A_i g'(x_i) &= (1 + \chi) e_{Ri} + \chi e_{Oi} + \chi e_{Si}, \\ b_O &= (c_{Oi} + \chi) e_{Oi} + \chi e_{Ri} + \chi e_{Si}, \\ b_S &= (\tilde{c}_i + \chi) e_{Si} + \chi e_{Ri} + \chi e_{Oi}.\end{aligned}$$

Differentiating the second and third conditions with respect to m gives

$$(c_{Oi} + \chi) \frac{\partial e_{Oi}}{\partial m} + \chi \frac{\partial e_{Ri}}{\partial m} + \chi \frac{\partial e_{Si}}{\partial m} = 0, \quad (\tilde{c}_i + \chi) \frac{\partial e_{Si}}{\partial m} + \chi \frac{\partial e_{Ri}}{\partial m} + \chi \frac{\partial e_{Oi}}{\partial m} = 0.$$

Solving these two equations simultaneously yields

$$\frac{\partial e_{Oi}}{\partial m} = -\frac{\chi \tilde{c}_i}{c_{Oi} \tilde{c}_i + (c_{Oi} + \tilde{c}_i) \chi} \frac{\partial e_{Ri}}{\partial m}, \quad \frac{\partial e_{Si}}{\partial m} = -\frac{\chi c_{Oi}}{c_{Oi} \tilde{c}_i + (c_{Oi} + \tilde{c}_i) \chi} \frac{\partial e_{Ri}}{\partial m},$$

and consequently

$$\chi \frac{\partial e_{Oi}}{\partial m} + \chi \frac{\partial e_{Si}}{\partial m} = -\frac{\chi^2 (c_{Oi} + \tilde{c}_i)}{c_{Oi} \tilde{c}_i + (c_{Oi} + \tilde{c}_i) \chi} \frac{\partial e_{Ri}}{\partial m}.$$

For the first condition, the partial derivative of $A_i g'(x_i)$ with respect to m at fixed e_{Ri} , using $\partial x_i / \partial m|_{e_{Ri}} = (\gamma_A - \gamma_D) x_i$ and $\varepsilon(x) := -g''(x)x/g'(x)$, is $A_i g'(x_i) [\gamma_A (1 - \varepsilon(x_i)) + \gamma_D \varepsilon(x_i)]$. The endogenous response through e_{Ri} contributes $-(\tilde{b}_{Ri} \varepsilon(x_i) / e_{Ri}) \partial e_{Ri} / \partial m$, where $\tilde{b}_{Ri} :=$

$\bar{\lambda} p_R A_i g'(x_i)$. Collecting terms after substituting the other-complex and SME responses:

$$\underbrace{\left[\frac{\tilde{b}_{Ri} \varepsilon(x_i)}{e_{Ri}} + \frac{c_{Oi} \tilde{c}_i (1 + \chi) + (c_{Oi} + \tilde{c}_i) \chi}{c_{Oi} \tilde{c}_i + (c_{Oi} + \tilde{c}_i) \chi} \right]}_{>0} \frac{\partial e_{Ri}}{\partial m} = \tilde{b}_{Ri} [\gamma_A(\theta_i)(1 - \varepsilon(x_i)) + \gamma_D \varepsilon(x_i)].$$

The bracket on the left is strictly positive: its second term equals $[c_{Oi} \tilde{c}_i (1 + \chi) + (c_{Oi} + \tilde{c}_i) \chi] / [c_{Oi} \tilde{c}_i + (c_{Oi} + \tilde{c}_i) \chi] > 0$ for any $c_{Oi} > 0, \tilde{c}_i > 0, \chi \geq 0$. Since the coefficient on the left is positive, the sign condition follows immediately, and the other-complex response inherits the multiplier $\chi \tilde{c}_i / [c_{Oi} \tilde{c}_i + (c_{Oi} + \tilde{c}_i) \chi]$ in absolute value. $\square$

Proof of Corollary A.1. At $\chi = 0$, the worker's routine first-order condition is $e_{Ri} = b_R A_i g'(x_i)$. Using $x_i = A_i e_{Ri} / D_i$ gives $x_i = (b_R A_i^2 / D_i) g'(x_i)$. Log-differentiating with respect to m yields $[1 + \varepsilon(x_i)] d \ln x_i / dm = 2\gamma_A(\theta_i) - \gamma_D$. Since $y_{Ri} = D_i g(x_i)$, $d \ln y_{Ri} / dm = \gamma_D + \eta(x_i) d \ln x_i / dm$, which gives Eq. (A.9). Because g is concave and $g(0) = 0, 0 < \eta(x) \leq 1$; hence both weights in Eq. (A.9) are nonnegative, and the response is strictly positive. Finally, rewrite the response as $\gamma_D + [2\gamma_A(\theta_i) - \gamma_D] H(x_i)$. The utilization and cross-worker orderings follow immediately when $H'(x) \leq 0$. $\square$

Proof of Corollary A.2. Part (i). From $y_{Ri} = D_i g(x_i)$ with $x_i = A_i e_{Ri} / D_i$, the total derivative is

$$\frac{dy_{Ri}}{dm} = \gamma_D y_{Ri} + A_i g'(x_i) \left[(\gamma_{Ai} - \gamma_D) e_{Ri} + \frac{\partial e_{Ri}}{\partial m} \right].$$

At the same initial utilization, the $\gamma_D y_{Ri}$ term is identical for both workers. Substituting $\partial e_{Ri} / \partial m$ from the proof of Proposition A.1(ii):

$$\frac{dy_{Rj}}{dm} - \frac{dy_{Rk}}{dm} = A_i g'(x_i) (\gamma_{Aj} - \gamma_{Ak}) \left[e_{Ri} + \frac{\tilde{b}_{Ri} (1 - \varepsilon(x_i))}{\Lambda} \right],$$

where, taking the coefficient from the proof of Proposition A.1(ii),

$$\Lambda := \frac{\tilde{b}_{Ri} \varepsilon(x_i)}{e_{Ri}} + \frac{c_{Oi} \tilde{c}_i (1 + \chi) + (c_{Oi} + \tilde{c}_i) \chi}{c_{Oi} \tilde{c}_i + (c_{Oi} + \tilde{c}_i) \chi} > 0.$$

The bracket $[e_{Ri} + \tilde{b}_{Ri}(1 - \varepsilon)/\Lambda]$ is positive for all $\varepsilon \geq 0$: trivially when $\varepsilon \leq 1$; when $\varepsilon > 1$, $\Lambda > \tilde{b}_{Ri}\varepsilon/e_{Ri}$ implies $\tilde{b}_{Ri}(\varepsilon - 1)/\Lambda < e_{Ri}(\varepsilon - 1)/\varepsilon < e_{Ri}$. Hence $dy_{Rj}/dm > dy_{Rk}/dm$ at every utilization level.

Part (ii). From Proposition A.1(ii), $\partial e_{Oi}/\partial m = -\kappa_{Oi} \partial e_{Ri}/\partial m$ and $\partial e_{Si}/\partial m = -\kappa_{Si} \partial e_{Ri}/\partial m$, where $\kappa_{Oi} = \chi \tilde{c}_i/\Delta_i$ and $\kappa_{Si} = \chi c_{Oi}/\Delta_i$. The two workers in the comparison differ only in γ_A , so these multipliers are common across them. Moreover, $\text{sign}(\partial e_{Rj}/\partial m - \partial e_{Rk}/\partial m) = \text{sign}[(\gamma_{Aj} - \gamma_{Ak})(1 - \varepsilon)]$. If $\chi > 0$ and $\varepsilon(x_i) < 1$, worker j has the larger routine-effort response and therefore the smaller responses in both analytical tasks. If $\chi > 0$ and $\varepsilon(x_i) > 1$, the routine ordering reverses, so worker j has the larger other-complex and SME responses. At $\chi = 0$, $\kappa_{Oi} = \kappa_{Si} = 0$ and both analytical responses to technology vanish. $\square$

Proof of Proposition A.2. Part (i). At $\chi = 0$, $y_{Si}^0 = \bar{\lambda} p_S^0/\tilde{c}_i$ and $\Delta y_{Si} = \bar{\lambda} \Delta p_S/\tilde{c}_i$. The level response is strictly decreasing in θ_i because $\tilde{c}_\theta > 0$, while $\Delta y_{Si}/y_{Si}^0 = \Delta p_S/p_S^0$ does not depend on worker type.

Part (ii). At $\chi = 0$, $e_{Oi}^* = \bar{\lambda} p_O/c_{Oi}$, so $\Delta e_{Oi} = 0$. Moreover, $\partial y_{Ri}/\partial p_R = A_i g'(x_i) \partial e_{Ri}/\partial p_R$; at high utilization $g'(x_i) \approx 0$, so the output response is negligible even when effort responds.

Part (iii). At $\chi = 0$, SME and other-complex effort do not depend on p_R by separability, so $W_i = \bar{\lambda}[p_R y_{Ri} + p_O e_{Oi} + p_S e_{Si}]$ depends on p_R only through $p_R y_{Ri}$:

$$\frac{dW_i}{dp_R} = \bar{\lambda} \left[y_{Ri} + p_R A_i g'(x_i) \frac{\partial e_{Ri}}{\partial p_R} \right] > 0,$$

where both terms are strictly positive. Hence $\Delta p_R < 0$ strictly reduces W_i . As utilization approaches capacity, $g'(x_i) \rightarrow 0$ makes the induced-effort term negligible and the wage change approaches the limit $\Delta W_i \rightarrow \bar{\lambda} y_{Ri} \Delta p_R < 0$, the high-utilization limit underlying the discussion in the text. The extension of the cross-worker ordering to small $\chi > 0$ is part (vii).

Parts (iv)–(vii). At an interior solution with $e_{Si} > 0$, the worker's first-order conditions are

$$\bar{\lambda} p_R A_i g'(x_i) = e_{Ri} + \chi(e_{Ri} + e_{Oi} + e_{Si}), \quad (\text{A.19})$$

$$\bar{\lambda} p_O = c_{Oi} e_{Oi} + \chi(e_{Ri} + e_{Oi} + e_{Si}), \quad (\text{A.20})$$

$$\bar{\lambda} p_S = \tilde{c}_i e_{Si} + \chi(e_{Ri} + e_{Oi} + e_{Si}). \quad (\text{A.21})$$

Define

$$q_i := 1 - \frac{\bar{\lambda} p_R A_i^2}{D_i} g''(x_i) > 0 \quad \text{and} \quad \Omega_i := 1 + \chi \left(\frac{1}{q_i} + \frac{1}{c_{Oi}} + \frac{1}{\tilde{c}_i} \right).$$

Totally differentiating (A.19)–(A.21) gives

$$\left[\begin{pmatrix} q_i & 0 & 0 \\ 0 & c_{Oi} & 0 \\ 0 & 0 & \tilde{c}_i \end{pmatrix} + \chi \mathbf{1}\mathbf{1}' \right] \begin{pmatrix} de_{Ri} \\ de_{Oi} \\ de_{Si} \end{pmatrix} = \bar{\lambda} \begin{pmatrix} A_i g'(x_i) dp_R \\ 0 \\ dp_S \end{pmatrix}.$$

The matrix on the left is positive definite. Using the Sherman–Morrison formula, the relevant comparative statics are

$$\frac{\partial e_{Si}}{\partial p_S} = \frac{\bar{\lambda} [1 + \chi(q_i^{-1} + c_{Oi}^{-1})]}{\tilde{c}_i \Omega_i} > 0, \quad (\text{A.22})$$

$$\frac{\partial e_{Ri}}{\partial p_S} = -\frac{\bar{\lambda} \chi}{q_i \tilde{c}_i \Omega_i} < 0, \quad (\text{A.23})$$

$$\frac{\partial e_{Oi}}{\partial p_S} = -\frac{\bar{\lambda} \chi}{c_{Oi} \tilde{c}_i \Omega_i} < 0, \quad (\text{A.24})$$

$$\frac{\partial e_{Ri}}{\partial p_R} = \frac{\bar{\lambda} A_i g'(x_i) [1 + \chi(c_{Oi}^{-1} + \tilde{c}_i^{-1})]}{q_i \Omega_i} > 0, \quad (\text{A.25})$$

$$\frac{\partial e_{Oi}}{\partial p_R} = -\frac{\bar{\lambda} A_i g'(x_i) \chi}{c_{Oi} q_i \Omega_i} < 0, \quad (\text{A.26})$$

$$\frac{\partial e_{Si}}{\partial p_R} = -\frac{\bar{\lambda} A_i g'(x_i) \chi}{\tilde{c}_i q_i \Omega_i} < 0. \quad (\text{A.27})$$

Equations (A.22)–(A.27) establish part (iv). In particular, $de_{Si} = (\partial e_{Si} / \partial p_S) dp_S + (\partial e_{Si} / \partial p_R) dp_R > 0$ and $de_{Ri} = (\partial e_{Ri} / \partial p_S) dp_S + (\partial e_{Ri} / \partial p_R) dp_R < 0$ when $dp_S > 0$ and $dp_R < 0$. For other-

complex effort,

$$de_{Oi} = -\frac{\bar{\lambda}\chi}{c_{Oi}\Omega_i} \left[\frac{dp_S}{\tilde{c}_i} + \frac{A_i g'(x_i)}{q_i} dp_R \right]. \quad (\text{A.28})$$

Hence $de_{Oi} > 0 \iff (A_i g'(x_i)/q_i)(-dp_R) > dp_S/\tilde{c}_i$: the effect on other-complex effort depends on whether the time released by the lower routine weight is larger than the crowd-out generated by the higher SME weight.

For part (v), the routine production function gives $dy_{Ri} = A_i g'(x_i) de_{Ri}$. Since the joint KPI change lowers routine effort by part (iv), routine output also falls; its response approaches zero as $g'(x_i)$ approaches zero.

For part (vi), the wage is $W_i = \bar{\lambda}[p_R y_{Ri} + p_O e_{Oi} + p_S e_{Si}]$; the p_S -dependent piece is unchanged under $dp_R < 0$, so

$$\frac{dW_i}{dp_R} = \bar{\lambda} \left[y_{Ri} + p_R A_i g'(x_i) \frac{\partial e_{Ri}}{\partial p_R} + p_O \frac{\partial e_{Oi}}{\partial p_R} + p_S \frac{\partial e_{Si}}{\partial p_R} \right]. \quad (\text{A.29})$$

Substituting the closed-form partials from (A.25)–(A.27) and using $g'(x_i) \rightarrow 0$ as $x_i \rightarrow 1$, the induced-effort terms vanish and $dW_i/dp_R \rightarrow \bar{\lambda} y_{Ri} > 0$. Condition (A.12) generalizes this to interior utilization by requiring the direct routine-pay term to dominate any reallocation gain.

Finally, part (vii). Holding q_i , c_{Oi} , and $A_i g'(x_i)$ fixed, let $z_i := 1/\tilde{c}_i$ and $B_i := 1 + \chi(q_i^{-1} + c_{Oi}^{-1})$. Using (A.22) and (A.27), the SME response to the joint KPI change can be written as

$$de_{Si} = \bar{\lambda} \frac{z_i}{B_i + \chi z_i} \left[B_i dp_S + \frac{\chi A_i g'(x_i)}{q_i} (-dp_R) \right].$$

The derivative of this expression with respect to z_i is

$$\bar{\lambda} \frac{B_i}{(B_i + \chi z_i)^2} \left[B_i dp_S + \frac{\chi A_i g'(x_i)}{q_i} (-dp_R) \right] > 0.$$

Thus the response is larger when $\tilde{c}_i$ is lower; since $\tilde{c}_\theta > 0$, it is larger for lower- θ workers, holding the other primitives fixed. More generally, at $\chi = 0$ the SME response is $de_{Si} = (\bar{\lambda}/\tilde{c}_i) dp_S$, which is strictly decreasing in θ_i . Because the optimal efforts and their comparative statics are continuous in χ at an interior solution, this strict ordering persists for sufficiently small $\chi > 0$. $\square$

Proof of Proposition A.3. Part (i). From Eq. (A.15), the SME component of MR_i^A at $\chi = 0$ is

$$(-\tilde{c}_{\tau^A}) \frac{\bar{\lambda} p_S}{\tilde{c}_i^2} \left[\beta_S(m) - \frac{\bar{\lambda} p_S}{2} \right].$$

Differentiating with respect to p_S gives

$$(-\tilde{c}_{\tau^A}) \frac{\bar{\lambda}}{\tilde{c}_i^2} [\beta_S(m) - \bar{\lambda} p_S] \geq 0,$$

where the weak inequality follows from $\tilde{c}_{\tau^A} < 0$ and $p_S \leq p_S^{\text{align}}(m)$. It is strict whenever $p_S < p_S^{\text{align}}(m)$.

Part (ii). The expression follows directly from Eq. (A.15). It measures the SME component of the marginal return to analytical coaching.

Part (iii). Let $\nu(p_S, T_M)$ denote the shadow value of the binding coaching-capacity constraint. Under Assumption 3(i), each coaching-return schedule is weakly decreasing in its own dose, so aggregate coaching demand at a shadow price ν is weakly decreasing in ν . At $\chi = 0$, routine-coaching demand is independent of p_S , while part (i) implies that analytical-coaching demand at any given ν is weakly increasing in p_S over the stated range. The binding capacity constraint therefore implies $\nu(p_S^1, T_M) \geq \nu(p_S^0, T_M)$. By the envelope theorem,

$$\frac{d\mathcal{D}_S(T_M)}{dT_M} = \nu(p_S^1, T_M) - \nu(p_S^0, T_M) \geq 0, \quad \text{equivalently,}$$

$$\mathcal{D}_S(T_M) = \mathcal{D}_S(0) + \int_0^{T_M} [\nu(p_S^1, t) - \nu(p_S^0, t)] dt.$$

The inequality is strict whenever the two shadow values differ, including intervals in which marginal capacity is used for analytical coaching at a worker for whom the worker-level complementarity is strict.

Part (iv). The conclusion follows directly from the stated crossing condition. $\square$

Proof of Corollary A.3. At $\chi = 0$, the analytical and routine coaching blocks are separable. Assumption 3(i) makes the marginal return to each coaching type weakly decreasing in its own

dose, so branch managers' problem is concave and the Kuhn–Tucker conditions are necessary and sufficient. Each coaching input that receives a positive allocation has a marginal return equal to the shadow value of capacity; each unused coaching input has a weakly lower return.

By Eq. (A.15), the cost signs in Eq. (A.4), and Assumption 3, MR_i^A is decreasing in θ_i at every common analytical-coaching dose. Together with diminishing returns in τ_i^A , this implies that shifting analytical coaching from a higher- θ worker toward an otherwise comparable lower- θ worker cannot reduce branch surplus. An optimal analytical-coaching allocation is therefore weakly decreasing in θ_i , and its positive support is a lower interval. This establishes part (i).

Assumption 3(ii) implies that the relative ranking of routine and analytical coaching changes monotonically with θ_i . Consequently, if routine coaching is weakly preferred at the margin for one type, it is also weakly preferred for every higher type, establishing part (ii). Neither condition orders $\tau_i^A + \tau_i^R$, so the model imposes no monotonic prediction on total coaching time. $\square$

Proof of Proposition A.5. Write $z = (p_S, \tau_i^A, \chi)$ and $z^0 = (p_S^0, \tau_i^{A,0}, 0)$, and let $F_i(e_i, z) := \nabla_e u_i(e_i; z)$ denote the first-order-condition map. The utility u_i depends on e_{Ri} through $g(x_i)$ with x_i affine in e_{Ri} ; on p_S only through the linear term $\bar{\lambda} p_S e_{Si}$; on τ_i^A only through the C^2 functions $c_O(\theta_i, \cdot)$ and $\tilde{c}(\theta_i, \cdot)$; and polynomially on (e_i, χ) . Hence $\nabla_e u_i$ involves g' , and each further derivative in e_i raises the order of the derivative of g by one. Since $g \in C^3$ on $[0, 1)$, F_i is C^2 jointly in (e_i, z) on the open set where $x_i < 1$. Because u_i is polynomial in χ , the same formula defines F_i for χ in an open interval around 0, so z^0 is an interior point of the domain.

At z^0 , interiority gives $F_i(e_i^*(z^0), z^0) = 0$, and $\partial F_i / \partial e_i = \nabla_e^2 u_i = -M_i$ is nonsingular by Lemma ?? . By the implicit function theorem, there exist an open neighborhood W of z^0 and a unique C^2 map $\hat{e}_i$ on W with $F_i(\hat{e}_i(z), z) = 0$ and $\hat{e}_i(z^0) = e_i^*(z^0)$. Shrinking W if necessary, the continuity of $\hat{e}_i$ keeps every effort component strictly positive and keeps $x_i < 1$. For $z \in W$ with $\chi \geq 0$, Lemma ?? makes $u_i(\cdot; z)$ strictly concave, so the interior critical point $\hat{e}_i(z)$ is the unique maximizer. Hence $e_i^* = \hat{e}_i$ is C^2 on $W \cap \{\chi \geq 0\}$, with

$$\frac{\partial e_i^*}{\partial \tau_i^A} = M_i^{-1} \frac{\partial F_i}{\partial \tau_i^A}.$$

The worker-level branch surplus

$$\pi_i(e_i, \tau_i^A, \chi) = \alpha_R y_{Ri}(e_i) + \beta_O e_{Oi} + \beta_S(m) e_{Si} - C_i(e_i, \tau_i^A, \chi)$$

is C^2 in (e_i, τ_i^A, χ) on $\{x_i < 1\}$, because $y_{Ri} = D_i g(x_i)$ is C^3 there and C_i is polynomial in (e_i, χ) with C^2 dependence on τ_i^A . By the chain rule, $\pi_i^*(z) := \pi_i(e_i^*(z), \tau_i^A, \chi)$ is C^2 on $W \cap \{\chi \geq 0\}$. Therefore

$$MR_i^A(z) = \frac{d\pi_i^*}{d\tau_i^A} = \pi_{i,\tau^A} + \pi'_{i,e} \frac{\partial e_i^*}{\partial \tau_i^A}$$

is C^1 there, and $\partial MR_i^A / \partial p_S$ is jointly continuous in z .

Since $0 < p_S^0 < \beta_S(m) / \bar{\lambda}$, Proposition A.3 gives $\partial MR_i^A / \partial p_S(z^0) > 0$. By joint continuity, there exist an open rectangle $\mathcal{N} \ni (p_S^0, \tau_i^{A,0})$ and $\bar{\chi} > 0$ with $\mathcal{N} \times [0, \bar{\chi}) \subset W$ on which $\partial MR_i^A / \partial p_S > 0$. Finally, for $p'_S < p''_S$ with $(p'_S, \tau_i^A), (p''_S, \tau_i^A) \in \mathcal{N}$ and $\chi \in [0, \bar{\chi})$,

$$MR_i^A(p''_S, \tau_i^A, \chi) - MR_i^A(p'_S, \tau_i^A, \chi) = \int_{p'_S}^{p''_S} \frac{\partial MR_i^A}{\partial p_S}(s, \tau_i^A, \chi) ds > 0,$$

which is strictly increasing differences of π_i^* in (p_S, τ_i^A) . □

Proof of Lemma A.1. At $\chi = 0$, SME effort is $e_S^* = \bar{\lambda} p_S / \tilde{c}$ and branch SME surplus is

$$\pi_S(p_S; m) = \beta_S(m) \frac{\bar{\lambda} p_S}{\tilde{c}} - \frac{(\bar{\lambda} p_S)^2}{2\tilde{c}} = \frac{\Omega_S(p_S; m)}{\tilde{c}};$$

other-complex and routine surplus do not depend on p_S at $\chi = 0$. Completing the square and evaluating at $p_S^{\text{align}}(m) = \beta_S(m) / \bar{\lambda}$ gives $\mathcal{D}_{\text{bench}}^S(m) = (\bar{\lambda}^2 / 2\tilde{c})(p_S^{\text{align}}(m) - p_S^0)^2$. Over the range in which $p_S^{\text{align}}(m) > p_S^0$, both $p_S^{\text{align}}(m) - p_S^0 > 0$ and $dp_S^{\text{align}}(m) / dm = \gamma / \bar{\lambda} > 0$, so $\mathcal{D}_{\text{bench}}^S$ is strictly increasing in m . The routine-surplus component of the full redesign gain is $\pi_R(p_R^1; m) - \pi_R(p_R^0; m)$; at $\chi = 0$ this depends on m only through $A_R(m, \cdot)$ and $D_i(m)$, and its sign depends on whether p_R^0 exceeds its own surplus-aligned value $\alpha_R / \bar{\lambda}$ (from Assumption 4, $p_R^0 < 2\alpha_R / \bar{\lambda}$, but the position of p_R^0 relative to $\alpha_R / \bar{\lambda}$ is not restricted). $\mathcal{D}_{\text{bench}}^S$ therefore isolates the SME channel underlying the m -monotonicity claim of Proposition A.5(ii); it does not exhaust $\mathcal{D}(m, T_M)$. □

Proof of Proposition A.5. Part (i). When $p_S^{\text{align}}(m) \leq p_S^0$, the SME-surplus coefficient satisfies $\partial\Omega_S(p_S; m)/\partial p_S = \bar{\lambda}[\beta_S(m) - \bar{\lambda}p_S] \leq 0$ for all $p_S \geq p_S^0$. Thus, for every feasible coaching allocation, increasing the SME weight from p_S^0 to $p_S^1 > p_S^0$ weakly reduces the SME component of branch surplus, and maximization over coaching allocations preserves the inequality $\mathcal{D}^S(m, T_M) \leq 0$. If retaining p_S^0 were an available redesign choice, the firm could obtain an SME-side gain of zero by making no change, and that choice would weakly dominate any increase in p_S in this region.

Part (ii). Over the range $p_S^{\text{align}}(m) > p_S^0$, Lemma A.1 establishes strict monotonicity of the SME component $\mathcal{D}^S$ in m in the representative-worker, $\chi = 0$ benchmark. For the heterogeneous population, the SME margin of $\Pi(p^1, T_M; m)$ is weakly increasing in m pointwise because $\beta_S(m)$ is strictly increasing and every optimizing worker's SME contribution to Π^* is weakly increasing in β_S . Whenever the old weight p_S^0 keeps some workers' SME effort low, the SME margin of $\Pi^*(p^0, T_M; m)$ increases at a weakly smaller rate. The SME component $\mathcal{D}^S(m, T_M)$ is therefore weakly increasing in m over this range.

Part (iii). By Proposition A.3(iii), $\mathcal{D}_S(T_M)$ is weakly increasing in T_M under the stated increasing-differences condition, holding p_R fixed. The full redesign raises p_S and lowers p_R simultaneously; the lower p_R can reduce the marginal value of routine coaching. The composite object $\mathcal{D}(m, T_M)$ is therefore weakly increasing in T_M when the increase in analytical-coaching value from $p_S \uparrow$ exceeds the reduction in routine-coaching value from $p_R \downarrow$. $\square$