The Compressed Apprenticeship

AI, Task Recomposition, and the Economics of Human Judgment

JUNE 2026 · AI, JOBS & THE FUTURE OF WORK

AUTHORS

Christophe Kolb
Chief Executive Officer, Taller
Christophe.Kolb@tallertechnologies.com

Jim Caron
Chief Investment Officer, Morgan Stanley Investment Management
Jim.Caron@morganstanley.com

ESSAY

The Compressed Apprenticeship

Artificial intelligence enters the economy of knowledge work through tasks, then remakes occupations from the inside. A job title represents a bundle of these tasks assembled within the context of information costs, coordination technologies, supervision practices, liability rules, and training conventions. When machine cognition changes the relative price of search, drafting, summarization, code generation, classification, and first-pass analysis, the name on the organization chart may survive while the function beneath it changes. The analyst, associate, paralegal, accountant, consultant, and junior engineer remain recognizable, yet the human hours inside their roles migrate away from routinized cognitive production toward verification, exception handling, judgment, persuasion, and accountability.

This is the economics of task recomposition. Firms purchase portfolios of labor services: produce, interpret, certify, and bear responsibility. Generative AI lowers the cost of a subset of these services because their inputs are digital, their outputs follow recognizable forms, and their quality can often be reviewed after production. The most exposed activities begin with filings, tickets, records, codebases, transcripts, or prior work product. They end as memos, chart shells, code stubs, issue lists, contract markups, or reconciliations. They reward retrieval, transformation, pattern completion, and linguistic fluency. They also carry bounded error costs when a competent human reviewer remains in the loop.

The surviving human tasks possess a different anatomy. They require tacit organizational context, fiduciary or legal responsibility, negotiated trust, interpretation of silence, resolution of conflicting objectives, and authority to decide under ambiguity. They draw on verification capital: the accumulated capacity to see why a plausible answer may be false, dangerous, incomplete, or strategically misdirected. Verification capital is built through repeated exposure to mistakes, institutional idiosyncrasies, edge cases, and feedback from seniors. AI raises the return to this capital precisely because it increases the supply of plausible drafts. When first versions become abundant, the scarce factor becomes knowing which version deserves confidence.

An economy of cheap cognition puts a premium on inquisitiveness as well as judgment. Most discussions treat AI as an answer engine, a device for producing drafts, summaries, code, analysis, and retrieval at high speed. As the marginal cost of producing answers falls, value shifts toward asking the right questions: identifying which problems matter, which hypotheses deserve attention, which data are meaningful, and which line of inquiry can create insight. Expertise moves from information retrieval toward problem framing, intellectual curiosity, hypothesis generation, and accountable judgment. In a world of abundant answers, the capacity to ask intelligent questions becomes a scarce form of human capital.

That scarcity changes the employment situation. The evidence consistent with current adoption points more toward transformation than broad job destruction. AI changes task composition, workflow design, skill requirements, and organizational architecture, while preserving a central role for human judgment. If the value of AI depends on better questions, the technology can expand the opportunity set for labor. New questions create new products, services, occupations, and fields of expertise. Economic history repeatedly shows that technological progress expands the feasible frontier of problems society can address. AI may lower the cost of routine cognitive production while increasing demand for workers who identify opportunities, define problems, interpret results, and translate insight into action.

Viewed through this lens, AI is a multiplier of human curiosity and accountable judgment. The central labor-market challenge is institutional: Can firms, schools, and professional systems cultivate the ability to ask better questions, verify better answers, and carry responsibility for consequential decisions? The productivity gain from cheap cognition becomes durable when it is joined to human capital formation. The danger arises when firms treat machine output as a substitute for the developmental work through which novices learn to exercise judgment.

The problem of compressed apprenticeship

This danger takes the form of bottom-up compression. The initial pressure falls on lower rungs of professional hierarchies because junior workers historically performed structured cognitive production. They collected comparables, summarized cases, reconciled accounts, and formatted slides. These tasks were output and apprenticeship at once. The grind produced deliverables, then quietly built taste, speed, intuition, professional discipline, and error recognition. AI directly targets the output component of the grind. The learning component remains socially necessary, while its private value to the firm becomes easier to underweight when software can produce the immediate deliverable.

The result is a wedge between short-run productivity and long-run capability. A bank, law firm, software company, accounting practice, or consulting shop can produce more output with fewer junior hours per transaction, case, sprint, close, or project. Over time, the same organization may discover that fewer workers have accumulated the tacit experience required for responsibility. A profession can enjoy a productivity dividend today while thinning the bench of experts tomorrow. The biggest institutional risk of AI in knowledge work is compressed apprenticeship: Production accelerates faster than the learning architecture is rebuilt.

The production function has three stages: question selection, answer generation, and accountable use. Traditional professional labor bundled all three because the cost of answers was high and the same worker often searched, drafted, checked, and presented. AI unbundles the stages. It makes answer generation cheaper, faster, and more abundant, thereby raising the shadow price of the upstream and downstream stages. Upstream value lies in choosing the right problem, scoping the relevant comparison set, defining the decision criterion, and designing the test that can change a mind. Downstream value lies in verification, interpretation, persuasion, and responsibility for action. The wage premium should therefore migrate toward workers who can connect curiosity to consequences.

This framework also explains why automation and task creation are jointly determined. A fall in the price of answers reduces labor required for a fixed volume of tasks. It can also raise the volume of viable projects. A law firm can investigate more claims, a bank can evaluate more transactions, a software company can test more product variants, and a consulting team can explore more strategic options. The employment outcome depends on whether cheaper cognition is absorbed as cost saving or converted into broader problem coverage. Scale expansion turns answer abundance into demand for framing, verification, and execution. Task creation turns curiosity into jobs.

New foundations for junior work

The organizational frontier will therefore be shaped by governance as much as by model capability. Firms need protocols that assign question rights, review duties, escalation triggers, and responsibility for final decisions. They need repositories of verified work product, feedback loops from errors, and career ladders that expose juniors to consequential ambiguity in protected ways. Apprenticeship must become more intentional as incidental learning from routine production declines. Simulations, supervised AI audits, rotating exception teams, after-action reviews, and structured red-team exercises can make the hidden curriculum explicit. The best organizations will convert AI from a shortcut around learning into a machine for producing more informative mistakes under safer conditions.

The same logic gives a principled account of entrepreneurship. Cheap answers lower the cost of exploring product space, testing customer hypotheses, preparing prototypes, and translating technical ideas into commercial narratives. This can broaden entry for small firms and independent workers whose main scarcity was the expense of preliminary analysis and production. Yet durable advantage will accrue to those who ask questions that competitors miss, validate answers against reality, and organize trust around reliable execution. AI expands the search space; human judgment selects a path through it.

A task-based model clarifies why exposure statistics alone are insufficient. Technical exposure measures what AI can reach inside work content. Employment outcomes depend on demand elasticity, liability, regulation, organizational redesign costs, market power, product innovation, and new-task creation, as well as what society believes about work. Software engineering contains highly exposed tasks, yet demand for software, cybersecurity, automation, analytics, and AI-enabled products may expand enough to absorb productivity gains. Other domains with standardized output and weaker demand expansion may experience tighter labor demand. The same exposure can generate expansion, compression, wage polarization, or new occupational ladders depending on market structure and the elasticity of final demand.

The new model begins with a professional-service output produced by routine cognition, specialized judgment, coordination, and accountability. AI lowers the unit cost of routine cognition. Where routine cognition substitutes readily for junior labor, first-pass human production shrinks. Where routine cognition complements expert judgment, cheaper drafts raise demand for review, tailoring, escalation, and decision authority. In most knowledge-work settings, both forces operate together. The low-discretion production component contracts, and the review-and-accountability component becomes denser:

- A first-year banking analyst spends fewer hours pulling comparables, drafting decks, and scaffolding models, while remaining work shifts toward reconciliation, client-specific framing, escalation, and judgment under transaction pressure.
- A junior software engineer receives assistance with boilerplate, routine tests, documentation, and search, while requirements ambiguity, integration, debugging in living systems, maintenance, and product tradeoffs grow more central.
- A paralegal spends fewer hours on first-pass research, drafting, e-discovery, and issue spotting, while citation checking, source grounding, attorney support, procedural accountability, client coordination, and strategic tailoring become dominant.
- An accounting associate sees automation of classification, records work, routine reconciliation, rule retrieval, and workpaper drafting, while exceptions, variance logic, client follow-up, auditability, and sign-off risk absorb attention.
- A junior consultant relies on AI for desk research, slide outlines, boilerplate visuals, first-pass synthesis, and phrasing variations, while framing, interview quality, organizational diagnosis, stakeholder management, and persuasion become the human center of gravity.

This pattern resembles skill-biased technical change, although its mechanism differs from earlier computerization. Previous digital waves rewarded workers who could operate formal systems. Generative AI reaches semi-formal cognition: language, synthesis, drafting, and code. Its force therefore lands on tasks that once served as entry pathways into elite labor markets. The distributional consequence may arrive through access, timing, and career mobility. Incumbents with domain authority gain leverage. Entrants face fewer routine openings, higher expected productivity on day one, and stronger pressure to demonstrate judgment before they have lived through enough mistakes to earn it.

The ultimate effect on wages may be irregular. Reduced demand for routine junior hours weakens the bottom of the professional wage ladder. Higher productivity among capable workers strengthens complementarities between expertise and software. Better templates may compress quality differences inside some tasks, narrowing dispersion for standardized production. Scarcity of client trust, reputation, and final accountability can raise returns at the top. The same firm may therefore experience wage compression in rote production, wage premia for judgment, and superstar returns where responsibility remains scarce.

Managing the new architecture of expertise

The most powerful complement to AI is institutional judgment. A firm gains little from turning every employee into an isolated prompt operator. The larger gains come from workflow design: deciding which tasks belong to AI, which outputs require verification, which error types deserve escalation, which data sources are safe, and which decisions require named accountability. AI is elevating workflow from an operational concern to a primary workforce design consideration, forcing organizations to rethink how work is allocated across humans, machines, and human-machine teams. Productivity depends on the match among model, task, organization, and liability regime. Competence is jagged across tasks and contexts. A system that accelerates bounded drafting can degrade performance in environments with ambiguous goals, hidden constraints, or weak review. Disciplined adoption will outperform indiscriminate experimentation.

AI also changes management. Professional hierarchies were built around delegation: Seniors assigned routine pieces to juniors, then reviewed the work. AI changes the span of control because one experienced worker can supervise more draft output. This may reduce the need for large cohorts of pure junior producers. It may also create bottlenecks in attention, review fatigue, and accountability. A system that generates endless plausible artifacts increases the value of deciding which artifacts matter. Firms will need more editors, auditors, integrators, explainers, and question framers. The managerial scarce resource becomes the allocation of attention under abundance.

Policy should start from capability preservation. The public problem is the maintenance of human expertise when private incentives to train novices weaken. Training subsidies, apprenticeship requirements, professional standards, public procurement rules, and disclosure norms can help align firm incentives with social returns. Education should teach students to supervise machine outputs, reason from first principles, and ask economically meaningful questions. The disciplines historically dismissed as impractical—philosophy, rhetoric, history—turn out to be training in the one skill cheap cognition cannot supply. The goal is higher cognitive leverage joined to preserved expertise formation. Cheap drafts are valuable; cheap drafts without accountable judgment create liability.

The risk of social inequality depends on the distribution of complements. Workers in elite firms and institutions may receive AI-augmented apprenticeship with proprietary data, mentoring, feedback, governance, and rigorous review. Workers outside those systems may receive tools absent the social infrastructure that turns tools into expertise. AI can democratize high-quality scaffolding, while organizational inequality can harden when data, reputation, and senior oversight remain concentrated. Compute and software are one side of the production function; trusted data, feedback, governance, and mentorship are the other.

A mature economics of AI in knowledge work centers on three questions. Which tasks become cheap? Which human complements become scarce? Which institutions preserve the formation of those complements? The answer points toward disciplined urgency. Routinized cognitive work will compress, especially at the junior end. Judgment, verification, persuasion, coordination, problem framing, and accountable curiosity will gain relative value. The decisive labor-market variable will be the evolution of apprenticeship under compressed production.

The future of knowledge work is a contest over the architecture of expertise. The old model used routine production as both output and training. The new model can separate these functions. With deliberate institutional redesign, AI can raise productivity and accelerate learning. With simple removal of the grind, the economy risks a thinner bench of professionals capable of bearing responsibility. The task bundle is the unit of disruption. The apprenticeship system is the unit of renewal. The firms and countries that understand both can turn AI from a labor-saving shock into a capability-building institution.

DATA APPENDIX · AI, JOBS & THE FUTURE OF WORK

Task-Decomposition Atlas

Each chart converts residual shares into scaled human-hours by applying the midpoint of the gross compression range. The left bar shows the pre-AI task bundle. The right bar shows the residual human role after compression. The translucent guides trace task migration from production toward residual bottlenecks.

Role	Gross compression range	Midpoint residual hours	Largest residual concentration
First-year banking analyst	35–45%	60.0h per 100h	Client context, escalation, and judgment (30%)
Junior software engineer	20–35%	72.5h per 100h	Debugging, integration, and maintenance (25%)
Paralegal or legal assistant	35–50%	57.5h per 100h	Citation checking, attorney support, client coordination, and strategic tailoring (45%)
Audit or accounting associate	30–45%	62.5h per 100h	Sign-off, judgment, and quality review (30%)
Management analyst or junior consultant	20–35%	72.5h per 100h	Stakeholder management, interviews, and change persuasion (30%)

First-year banking analyst

Gross human-hour compression: 35–45% | midpoint residual: 60.0 human-hours per 100 pre-AI hours

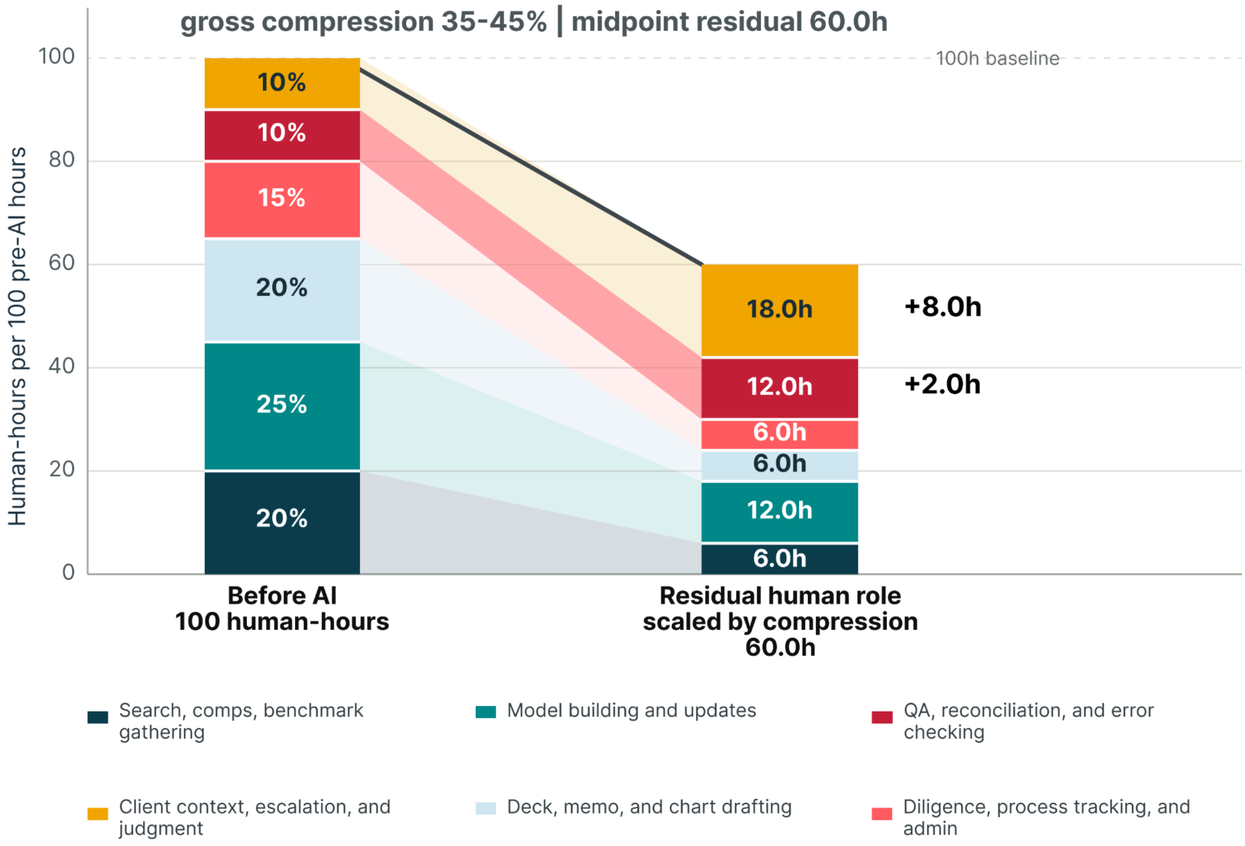

AI attacks comps pulls, deck drafting, chart scaffolds, routine diligence synthesis, and first-pass model support. Residual work concentrates in reconciliation, client-specific framing, escalation, and judgment under transaction pressure.

Exhibit 2. First-year banking analyst — task migration from production to residual role

Segment	Before AI share	Residual role share	Scaled residual hours	Hour delta
Search, comps, benchmark gathering	20%	10%	6.0h	−14.0h
Model building and updates	25%	20%	12.0h	−13.0h
Deck, memo, and chart drafting	20%	10%	6.0h	−14.0h
Diligence, process tracking, and admin	15%	10%	6.0h	−9.0h
QA, reconciliation, and error checking	10%	20%	12.0h	**+2.0h**
Client context, escalation, and judgment	10%	30%	18.0h	**+8.0h**

Junior software engineer

Gross human-hour compression: 20-35% | midpoint residual: 72.5 human-hours per 100 pre-AI hours

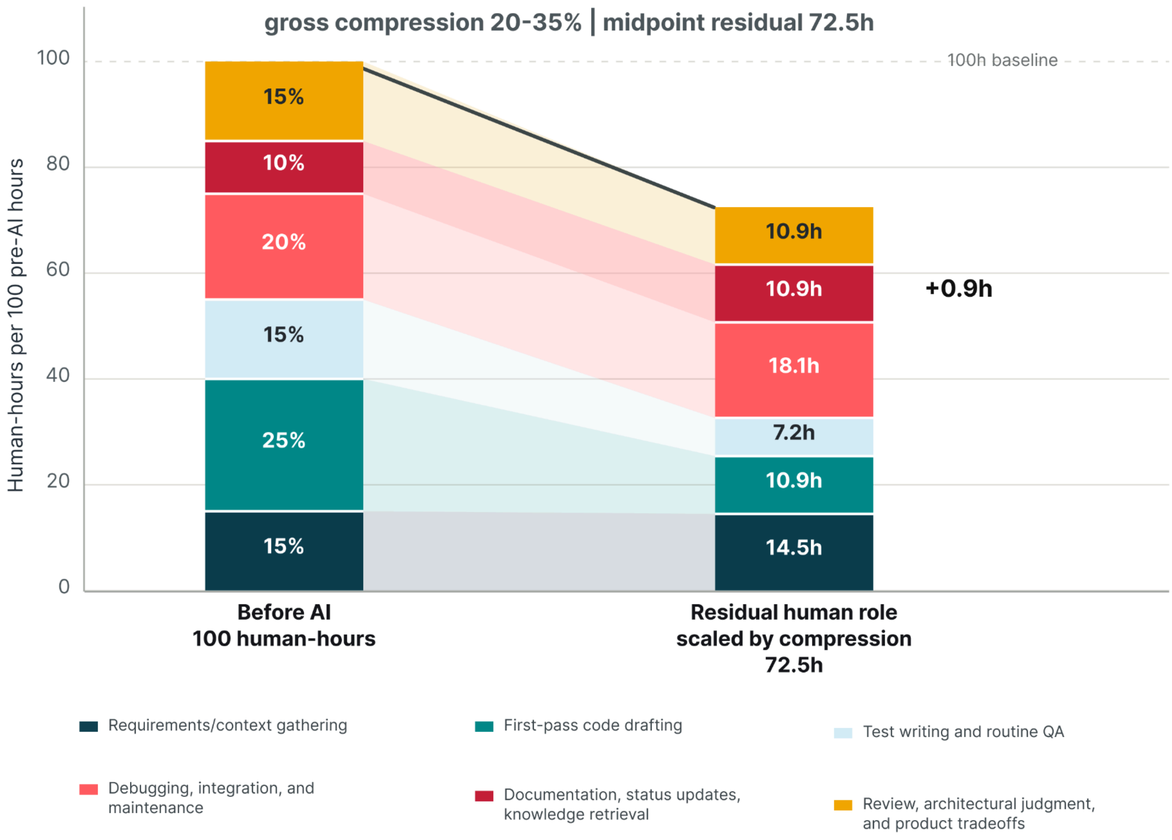

AI compresses code scaffolding, boilerplate, routine QA, documentation, and search. Residual work tilts toward requirements, integration, debugging in living systems, maintenance, and judgment over product tradeoffs.

Exhibit 3. Junior software engineer — task migration from production to residual role

Segment	Before AI share	Residual role share	Scaled residual hours	Hour delta
Requirements/context gathering	15%	20%	14.5h	−0.5h
First-pass code drafting	25%	15%	10.9h	−14.1h
Test writing and routine QA	15%	10%	7.2h	−7.8h
Debugging, integration, and maintenance	20%	25%	18.1h	−1.9h
Documentation, status updates, knowledge retrieval	10%	15%	10.9h	**+0.9h**
Review, architectural judgment, and product tradeoffs	15%	15%	10.9h	−4.1h

Paralegal or legal assistant

Gross human-hour compression: 35-50% | midpoint residual: 57.5 human-hours per 100 pre-AI hours

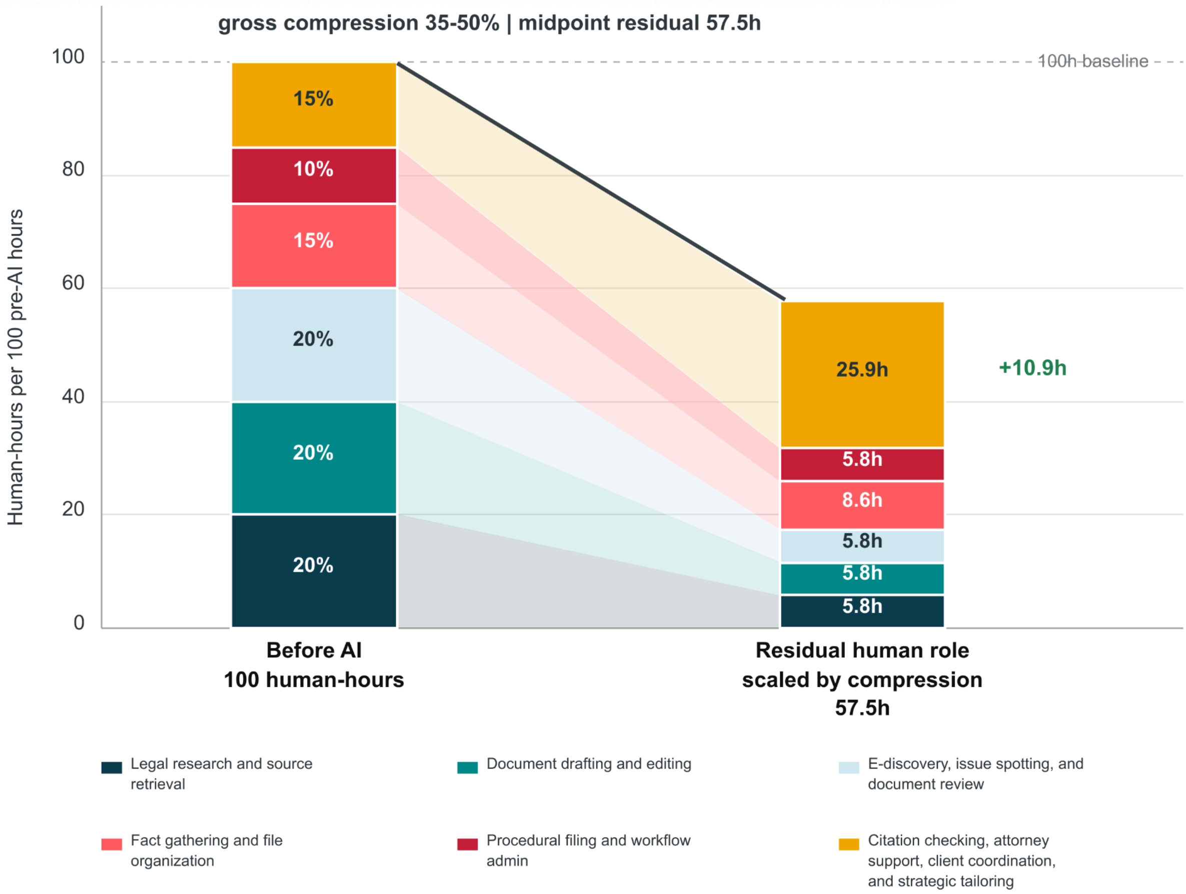

AI compresses research, drafting, e-discovery, and first-pass issue spotting. Residual work shifts sharply toward citation checking, source grounding, attorney support, procedural accountability, client coordination, and strategic tailoring.

Exhibit 4. Paralegal or legal assistant — task migration from production to residual role

Segment	Before AI share	Residual role share	Scaled residual hours	Hour delta
Legal research and source retrieval	20%	10%	5.8h	−14.2h
Document drafting and editing	20%	10%	5.8h	−14.2h
E-discovery, issue spotting, and document review	20%	10%	5.8h	-14.2h
Fact gathering and file organization	15%	15%	8.6h	−6.4h
Procedural filing and workflow admin	10%	10%	5.8h	−4.2h
Citation checking, attorney support, client coordination, and strategic tailoring	15%	45%	25.9h	**+10.9h**

Audit or accounting associate

Gross human-hour compression: 30-45% | midpoint residual: 62.5 human-hours per 100 pre-AI hours

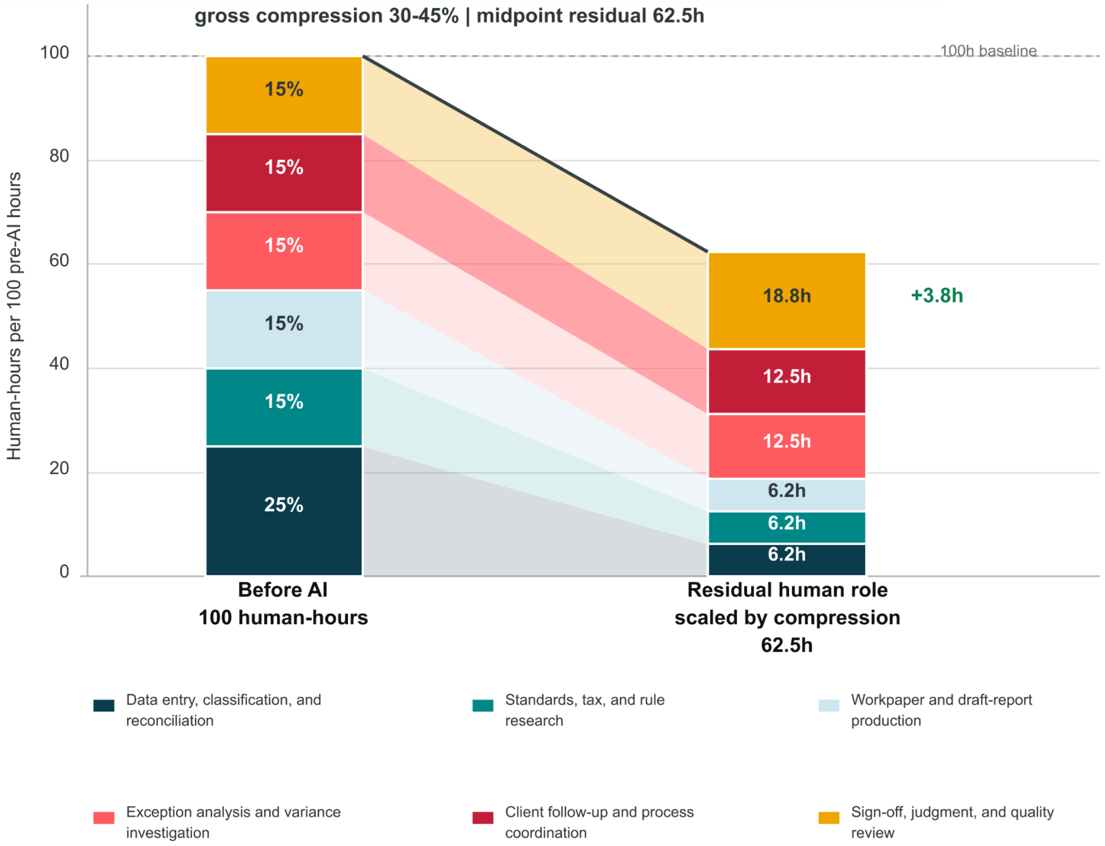

AI compresses structured records work, classification, routine reconciliation, rules retrieval, and workpaper drafting. Residual work concentrates in exceptions, variance logic, client follow-up, auditability, and sign-off risk.

Exhibit 5. Audit or accounting associate — task migration from production to residual role

Segment	Before AI share	Residual role share	Scaled residual hours	Hour delta
Data entry, classification, and reconciliation	25%	10%	6.2h	−18.8h
Standards, tax, and rule research	15%	10%	6.2h	−8.8h
Workpaper and draft-report production	15%	10%	6.2h	−8.8h
Exception analysis and variance investigation	15%	20%	12.5h	−2.5h
Client follow-up and process coordination	15%	20%	12.5h	−2.5h
Sign-off, judgment, and quality review	15%	30%	18.8h	**+3.8h**

Management analyst or junior consultant

Gross human-hour compression: 20–35% | midpoint residual: 72.5 human-hours per 100 pre-AI hours

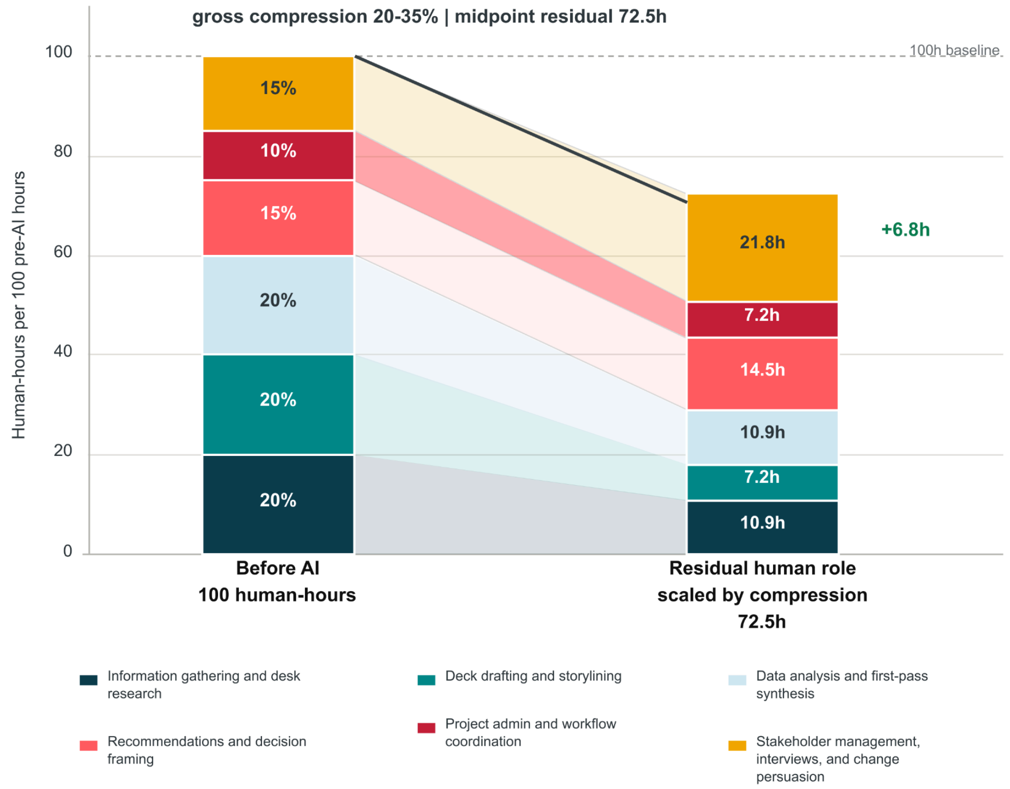

AI compresses desk research, slide outlines, boilerplate visuals, first-pass synthesis, and alternative phrasings. Residual work moves toward framing, interview quality, organizational diagnosis, stakeholder management, and persuasion.

Exhibit 6. Management analyst or junior consultant — task migration from production to residual role

Segment	Before AI share	Residual role share	Scaled residual hours	Hour delta
Information gathering and desk research	20%	15%	10.9h	−9.1h
Deck drafting and storylining	20%	10%	7.2h	−12.8h
Data analysis and first-pass synthesis	20%	15%	10.9h	−9.1h
Recommendations and decision framing	15%	20%	14.5h	−0.5h
Project admin and workflow coordination	10%	10%	7.2h	−2.8h
Stakeholder management, interviews, and change persuasion	15%	30%	21.8h	**+6.8h**

Methodological Note:

Producing the Atlas Bars

Estimand and interpretation

The atlas estimates gross human-hour compression for a fixed unit of professional-service output at an unchanged quality, liability, and accountability standard. Its unit is 100 pre-AI human-hours for a role-specific junior knowledge-work bundle. The object of interest is a central-scenario accounting decomposition: which portions of the original human-hour bundle become machine-assisted, and which portions remain scarce because they require verification, exception handling, coordination, persuasion, or accountable judgment. Employment, wages, and occupational entry are downstream equilibrium outcomes, so the decomposition is read as a workflow-production estimate.

Each bar is a calibrated model object, rather than a regression coefficient. The sources identify tasks and triangulate plausible compression magnitudes; the residual bars enforce exact arithmetic. Uncertainty enters through model specification: source weights, review add-back floors, and the discount between theoretical exposure and observed workflow coverage.

1. Baseline task allocation

The baseline begins with occupation proxies, because junior titles vary across firms while official occupation systems preserve stable task inventories. O*NET supplies the richest task list; BLS supplies official duty language; employer program pages supply early-career workflow specificity. Each source is coded separately, normalized to 100, and combined as an explicit prior over information value:

$$s0\_rc = 100 \times [0.60\ O\_rc + 0.25\ B\_rc + 0.15\ P\_rc]$$

Here r indexes the role, c indexes a task segment, O_rc is the normalized O*NET share, B_rc is the normalized BLS share, P_rc is the normalized employer-program or posting share, and s0_rc is the pre-AI human-hour share per 100 role hours. O*NET receives the largest weight because it offers the broadest task inventory; BLS receives intermediate weight because its duties are official and coarser; employer pages receive the smallest weight because they are role-specific and volatile.

Within O*NET, task-importance scores receive the primary weight. Where public task displays omit task-level importance, the ordered-task proxy assigns 100 to the first task, 95 to the second, 90 to the third, continuing in five-point steps with a floor of 50. Core tasks receive twice the weight of supplemental tasks. BLS duty bullets and employer clauses receive equal weights within source after clause splitting.

Category	Verb-object rule	AI-compressed work	Residual human work
Research and retrieval	research; investigate; gather; monitor; search; collect source material	source search, summarization, precedent or market retrieval	source choice, grounding, provenance, relevance
Analysis and modeling	analyze; evaluate; model; value; forecast; test hypotheses; inspect patterns	model scaffolds, spreadsheet interpretation, first-pass analysis	materiality, causal logic, scenario discipline
Synthesis and drafting	prepare reports; draft; create presentations; summarize; communicate recommendations	memos, slide text, chart captions, first drafts	tone, audience fit, strategic framing, accountability
Production, execution, and coding	build code; prepare workpapers; file pleadings; update models; compile exhibits; maintain records	boilerplate code, workpapers, filings, exhibits, routine updates	repository context, process control, regulated execution
QA, review, and compliance	test; validate; reconcile; audit; review; proofread; verify; compliance-check	assisted testing, checklist coverage, comparison passes	evidence review, source validation, sign-off reliability
Coordination and communication	meet; interview; confer; collaborate; liaise; present; obtain evidence	meeting notes, follow-up drafts, status summaries	trust, elicitation, conflict resolution, client alignment
Judgment and exception handling	recommend; decide; interpret ambiguity; handle exceptions; prioritize under liability	narrow direct substitution	taste, escalation, fiduciary or legal responsibility

Table M1. Clause-level coding rules used to map official and employer task text into mutually exclusive segments.

2. AI retiming priors and human add-backs

Post-AI residual hours combine direct machine compression with human review, coordination, quality-assurance, and judgment add-backs:

$$h\_rc = s0\_rc(1 - m\_rc) + a\_rc$$

The prior m_rc is the fractional direct compression of segment c in role r. The term a_rc is measured in original human-hours and captures verification, source checking, model reconciliation, code review, client tailoring, attorney or partner supervision, and supervisory judgment created by AI-assisted first passes. This add-back is central to high-liability workflows: more plausible drafts increase the need for triage and accountable review.

Evidence layer	Sources used	Role in calibration
Task economics	Acemoglu and Restrepo; Autor and Thompson	Frame occupations as task bundles and separate automation, task creation, and retained expertise.
Theoretical exposure bounds	Eloundou et al.; ILO Working Paper 140	Bound tasks that can plausibly receive large language model assistance at constant quality.
Observed use and applicability	Anthropic Economic Index; Microsoft Copilot occupational-applicability study; Bick, Blandin, and Deming	Discount theoretical exposure by real usage, adoption intensity, and workflow coverage.
Cross-metric uncertainty	Budget Lab at Yale	Treat exposure magnitude as model-uncertain while preserving broad exposure rankings.
Productivity and field evidence	Noy and Zhang; Brynjolfsson, Li, and Raymond; Peng et al.; Cui et al.; METR; Choi et al.; Choi and Xie; Dell'Acqua et al.	Anchor writing, support, coding, legal, accounting, and consulting compression priors.
Professional governance and duties	ABA Formal Opinion 512; Magesh et al.; BLS and O*NET occupation sources	Set verification floors where legal, audit, fiduciary, and product responsibility constrain automation.

Table M2. Evidence hierarchy used to calibrate direct compression, usage discounts, and review add-backs.

Task family	Central prior range for m_rc	Calibration logic
Research/retrieval; synthesis/drafting in text-heavy junior roles	0.45–0.65	High language-model fit for search, summarization, first drafts, and structured professional writing.
Analysis and modeling	0.20–0.45	Meaningful assistance for model scaffolds, spreadsheet reading, first-pass analysis, and hypothesis tests; context verification caps savings.
Production, execution, and coding	0.15–0.50	High for boilerplate and templateable artifacts; lower for live processes, mature repositories, filings, workpapers, and regulated outputs.
QA, review, and compliance	0.00–0.15	Assistance with tests and checks; final evidence, source validation, and auditability preserve human floors.
Coordination and communication	0.00–0.10	Drafting and summarization savings; interpersonal coordination, elicitation, and alignment remain bottlenecks.

Table M3. Direct machine-compression priors before human add-backs.

3. Formulaic construction of the stacked bars

For each role, the compression interval [gL_r, gU_r] defines the central gross compression gbar_r = (gL_r + gU_r) / 2. The residual human-hour total is H_r = 100 − gbar_r. The residual-share vector q_rc sums to 100 within each role, and raw residual hours equal h_rc = H_r × q_rc / 100. Category-level compression is d_rc = s0_rc − h_rc, and the atlas hour delta is Delta_rc = h_rc − s0_rc.

The plotted residual bar is therefore drawn on the original 100-hour scale. A segment can gain residual share while losing raw hours, because residual shares are compositional. Raw hours require multiplication by H_r. A segment can expand in absolute human time when review, coordination, or judgment add-backs exceed direct machine compression.

Role	Compressed center of gravity	Residual bottleneck and add-back logic	Central residual total
First-year banking analyst	Comps pulls, benchmark gathering, deck and chart drafting, first-pass model support, routine diligence synthesis.	Reconciliation, client-specific framing, escalation, fiduciary judgment, and transaction-pressure review.	60.0h
Junior software engineer	First-pass code, boilerplate, routine tests, documentation, and knowledge search.	Requirements ambiguity, integration, debugging in living systems, maintenance, review, architecture, and product tradeoffs.	72.5h
Paralegal or legal assistant	Legal research, first drafts, e-discovery triage, issue spotting, file organization, and workflow admin.	Citation checking, source grounding, attorney supervision, procedural accountability, client coordination, and strategic tailoring.	57.5h
Audit or accounting associate	Structured records work, classification, routine reconciliation, standards retrieval, workpaper drafting, and draft reports.	Exceptions, variance logic, client evidence requests, auditability, professional skepticism, and sign-off risk.	62.5h
Management analyst or junior consultant	Desk research, slide outlines, boilerplate visuals, first-pass synthesis, and alternative phrasings.	Problem framing, interview quality, organizational diagnosis, recommendations, stakeholder management, and persuasion.	72.5h

Table M4. Role-level calibration logic behind the five atlas decompositions.

Methodology:

Bibliography

Source	Use in methodology
Acemoglu, D., and Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. Journal of Economic Perspectives, 33(2), 3–30. https://doi.org/10.1257/jep.33.2.3	Task-based labor-demand frame; displacement and reinstatement logic.
Autor, D. H., and Thompson, N. (2025). Expertise. NBER Working Paper No. 33941. https://doi.org/10.3386/w33941	Scarcity value of retained expertise after task automation.
Eloundou, T., Manning, S., Mishkin, P., and Rock, D. (2024). GPTs are GPTs: Labor market impact potential of LLMs. Science, 384(6702), 1306–1308. https://doi.org/10.1126/science.adj0998	Theoretical occupational and task exposure bounds.
Gmyrek, P., Berg, J., Kaminski, K., Konopczynski, F., Ladna, A., Nafradi, B., Roslaniec, K., and Troszynski, M. (2025). Generative AI and jobs: A refined global index of occupational exposure. ILO Working Paper 140. https://doi.org/10.54394/HETP0387	Refined exposure index and task-level exposure discipline.
Gimbel, M., Kendall, J., and Kulsakdinun, R. (2026). Labor market AI exposure: What do we know? The Budget Lab at Yale. https://budgetlab.yale.edu/research/labor-market-ai-exposure-what-do-we-know	Cross-metric exposure uncertainty and exposure-versus-outcome boundary.
Handa, K., Tamkin, A., McCain, M., Huang, S., Durmus, E., Heck, S., Mueller, J., Hong, J., Ritchie, S., Belonax, T., Troy, K. K., Amodei, D., Kaplan, J., Clark, J., and Ganguli, D. (2025). Which economic tasks are performed with AI? Evidence from millions of Claude conversations. arXiv. https://doi.org/10.48550/arXiv.2503.04761	Observed usage patterns mapped to O*NET task structure.
Tomlinson, K., Jaffe, S., Wang, W., Counts, S., and Suri, S. (2025). Working with AI: Measuring the applicability of generative AI to occupations. arXiv. https://doi.org/10.48550/arXiv.2507.07935	Observed Copilot work activities and occupation applicability scores.
Bick, A., Blandin, A., and Deming, D. J. (2024, rev. 2025). The rapid adoption of generative AI. NBER Working Paper No. 32966. https://doi.org/10.3386/w32966	Adoption-intensity boundary for translating capability into workflow coverage.

Table M5. Task economics, exposure, observed-use, and adoption sources. The bibliography above includes sources used to construct the task baseline, calibrate compression priors, set verification floors, and define audit boundaries.

Source	Use in methodology
Noy, S., and Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. Science, 381(6654), 187–192. https://doi.org/10.1126/science.adh2586	Writing, summarization, and professional drafting priors.
Brynjolfsson, E., Li, D., and Raymond, L. R. (2025). Generative AI at work. Quarterly Journal of Economics, 140(2), 889–942. https://doi.org/10.1093/qje/qjae044	Suggested-reply, customer-support, and novice-skewed productivity calibration.

Source	Use in methodology
Peng, S., Kalliamvakou, E., Cihon, P., and Demirer, M. (2023). The impact of AI on developer productivity: Evidence from GitHub Copilot. arXiv. https://doi.org/10.48550/arXiv.2302.06590	Coding-boilerplate and greenfield-task compression prior.
Cui, K. Z., Demirer, M., Jaffe, S., Musolff, L., Peng, S., and Salz, T. (2026). The effects of generative AI on high-skilled work: Evidence from three field experiments with software developers. Management Science. https://doi.org/10.1287/mnsc.2025.00535	Field evidence for software-development productivity and experience heterogeneity.
Becker, J., Rush, N., Barnes, E., and Rein, D. (2025). Measuring the impact of early-2025 AI on experienced open-source developer productivity. arXiv. https://doi.org/10.48550/arXiv.2507.09089.	Cap on software compression in mature, repository-specific tasks with review burden.
Dell'Acqua, F., McFowland, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K. C., Rajendran, S., Krayer, L., Candelon, F., and Lakhani, K. R. (2026). Navigating the jagged technological frontier. Organization Science. https://doi.org/10.1287/orsc.2025.21838	Consulting and knowledge-worker evidence for jagged task-specific gains and losses.
Choi, J. H., Monahan, A. B., and Schwarcz, D. (2024). Lawyering in the age of artificial intelligence. Minnesota Law Review, 109, 147–218. https://doi.org/10.24926/265535.4225	Legal research and drafting speed gains with quality safeguards.

Table M6. Productivity, domain-specific field, and verification sources.

Role	Official occupation sources	Employer/program source used for junior-role specificity
First-year banking analyst	O*NET OnLine. Financial and Investment Analysts, 13-2051.00. U.S. Bureau of Labor Statistics. Financial Analysts, Occupational Outlook Handbook.	JPMorgan Chase & Co. Investment Banking Full-Time Analyst Program.
Junior software engineer	O*NET OnLine. Software Developers, 15-1252.00. U.S. Bureau of Labor Statistics. Software Developers, Quality Assurance Analysts, and Testers, Occupational Outlook Handbook.	Microsoft Careers. Recent Graduate Opportunities; Explore Microsoft program pages.
Paralegal or legal assistant	O*NET OnLine. Paralegals and Legal Assistants, 23-2011.00. U.S. Bureau of Labor Statistics. Paralegals and Legal Assistants, Occupational Outlook Handbook.	Skadden, Arps, Slate, Meagher & Flom LLP. Paralegals career page.
Audit or accounting associate	O*NET OnLine. Accountants and Auditors, 13-2011.00. U.S. Bureau of Labor Statistics. Accountants and Auditors, Occupational Outlook Handbook.	EY Careers. Audit Associate / Assurance Audit Associate program and posting pages.
Management analyst or junior consultant	O*NET OnLine. Management Analysts, 13-1111.00. U.S. Bureau of Labor Statistics. Management Analysts, Occupational Outlook Handbook.	McKinsey & Company. Business Analyst career page.

Table M7. Occupation and employer pages used for baseline task-source coding. The sources above supply the role proxies, duty statements, and employer-specific clauses that feed the 0.60/0.25/0.15 baseline task allocation.