Technology Bottlenecks

Anders Humlum*

March 8, 2026

Abstract

A technology bottleneck is a task that (*i*) is not itself technologically feasible, but (*ii*) is complementary to tasks that are. We develop an empirical measure of bottleneck pressure based on a simple assumption: that occupations bundle complementary tasks. In a standard task-based model, this assumption implies that the bottleneck pressure on a task is the product of its own infeasibility and the technology exposure of the occupations in which it is performed. Applying this to Generative AI, the tasks under greatest bottleneck pressure are predominantly interpersonal and supervisory—supervising personnel, training staff, and liaising with clients—in otherwise highly exposed white-collar occupations. Bottlenecks are economically important: for the tasks under greatest pressure, the within-occupation bottleneck effect exceeds the economy-wide productivity gain from AI. This pressure is highly concentrated: just 4% of infeasible tasks account for half the total. Our measure identifies where labor demand will shift at the onset of technological change and where R&D would unlock the greatest gains from AI.

*University of Chicago, Booth School of Business, and NBER, anders.humlum@chicagobooth.edu. I thank Claude Code (Anthropic, Opus 4.6) for assistance with writing code for the empirical analysis and drafting and revising manuscript text, and Refine.Ink for editing assistance.

Feasibility measures—technical assessments of which tasks a technology can assist with—have become the standard tool for predicting where workers are exposed to AI (Eloundou et al., 2024; Felten, Raj and Seamans, 2021; Webb, 2020), and have proven successful in predicting where adoption occurs (Bick, Blandin and Deming, 2026; Tomlinson et al., 2025). But what a technology *cannot* do also matters. When a technology can perform some tasks in an occupation but not others, the infeasible tasks become bottlenecks that can prevent task-level improvements from translating into aggregate economic gains (Gans and Goldfarb, 2026; Jones and Tonetti, 2026). At the same time, complementarity between feasible and infeasible tasks can increase the demand for labor in bottleneck tasks by more than technology displaces it elsewhere (Autor and Thompson, 2025). Feasibility measures, on their own, are silent on both of these forces, and while theory establishes that bottlenecks matter in principle, it does not tell us which specific tasks are bottlenecks, where they concentrate, or how large their effects are.

This paper develops a simple and operational measure of technology bottlenecks. We start from the observation that occupations bundle complementary tasks (Becker and Murphy, 1992). In a standard task-based production model, this observation implies that when a technology raises productivity in some tasks within an occupation, the marginal value of the remaining infeasible tasks rises through complementarity. We show that the resulting *bottleneck pressure* on a task can be measured using a simple expression: the interaction between a task's own technological infeasibility and the technology exposure of the occupations in which it is performed. The importance of bottleneck effects relative to economy-wide productivity gains depends on the degree of task complementarity within occupations and the exposure of infeasible tasks relative to the economy-wide average—if feasible and infeasible tasks were perfectly segregated across occupations, there would be no bottlenecks. The measure is general-purpose: it requires only feasibility assessments and task-content data, both readily available for many technologies.

We apply our measure to Generative AI, using the feasibility assessments from Eloundou et al. (2024) for Large Language Models (LLMs). The tasks under the greatest bottleneck pressure are predominantly interpersonal and supervisory—supervising personnel, training staff, and liaising with clients—in white-collar occupations. By contrast, equally infeasible physical tasks in unexposed occupations—sweeping floors, washing dishes, sorting trash—carry zero bottleneck pressure, because there is no AI-driven productivity to constrain. Bottleneck effects are economically large: under a conservative calibration, the within-occupation bottleneck effect for the tasks under greatest pressure is 1.7 times the economy-wide productivity gain from AI. This pressure is highly concentrated: just 4% of infeasible tasks account for half the total. This points to where labor demand will shift—toward interpersonal and supervisory tasks within AI-exposed occupations—and where targeted R&D could unlock disproportionate gains. That the binding constraints on Generative AI are interpersonal and supervisory—rather than physical or cognitive—is a finding about the nature of work that only emerges from measurement. Feasibility measures tell us what a technology can do; bottleneck measures tell us what it cannot do *that matters*.

1 Framework

We use a standard task-based production model in the tradition of Acemoglu and Autor (2011) and Acemoglu and Restrepo (2018) to formalize the concept of a technology bottleneck and derive an empirical measure for it.

Final output combines occupations $o \in \mathcal{O}$ via CES with elasticity $\rho > 0$. Each occupation combines tasks $t \in \mathcal{T}$ via CES with elasticity $\sigma > 0$:

$$Y = \left[\sum_o \alpha_o \, Y_o^{\frac{\rho-1}{\rho}}\right]^{\frac{\rho}{\rho-1}}, \qquad Y_o = \left[\sum_t \alpha_{ot} \, Y_{ot}^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}. \tag{1}$$

Without technology ($F_t = 0$), tasks are produced by labor alone: $Y_{ot} = L_{ot}$. We remain agnostic about the production function in technologically feasible tasks ($F_t = 1$), and instead

parameterize the task-level productivity gain: holding labor fixed at baseline levels, technology increases the output of a feasible task by a factor $\delta > 1$.[1]

Assumption 1. *Tasks are more complementary within occupations than across:* $\sigma < \rho$.

Assumption 1 captures the idea that occupations bundle complementary tasks (Becker and Murphy, 1992)—a standard assumption in task-based models of automation and technological change (Acemoglu and Restrepo, 2022; Humlum, 2021; Jones and Tonetti, 2026). Empirically, Deming (2017) shows that cognitive and social skills are complements within occupations. Consistent with this, Humlum and Vestergaard (2025) find that workers reallocate time savings from AI chatbots toward infeasible tasks rather than scaling up feasible ones—direct evidence of within-occupation complementarities and of infeasible tasks acting as bottlenecks.

Definition 1 (Technology bottleneck)**.** *A technology bottleneck is a task that (i) is not itself technologically feasible, but (ii) is complementary to tasks that are. The* bottleneck pressure *on a task is the resulting increase in its marginal value product of labor (MVPL).*

Bottleneck effect. We compare a *baseline* where no task is feasible ($F_t = 0$ for all t) to a *technology scenario* where some tasks become feasible, holding labor allocations fixed at baseline levels. To a first-order approximation, the change in MVPL for infeasible task t in occupation o is:

$$\Delta \ln \text{MVPL}_{ot} = \underbrace{\ln \delta \times \frac{1}{\rho} \times \bar{E}}_{\text{Scale Effect}} + \underbrace{\ln \delta \times \left(\frac{1}{\sigma} - \frac{1}{\rho} \right) \times E_o}_{\text{Bottleneck Effect}} \qquad \text{for} \qquad F_t = 0, \tag{2}$$

where $E_o \equiv \sum_{t'} \pi_{ot'} F_{t'}$ is the *technology exposure* of occupation o, π_{ot} is the baseline share of task t in occupation o, and $\bar{E} = \sum_o \pi_o E_o$ is aggregate technology exposure. Appendix B

[1]This parameterization mirrors how feasibility is assessed in practice. For example, Eloundou et al. (2024) asks whether access to ChatGPT can at least halve the time required to perform task t, corresponding to $\delta \geq 2$. The productivity gain δ could arise from technology augmenting labor or from technology substituting for it in feasible tasks; Appendix C shows both cases formally. Because bottleneck pressures operate on infeasible tasks only, we can remain agnostic about the sources of productivity effects in feasible tasks.

derives Equation (2) formally.

Interpretation. The first term in Equation (2) is a scale effect: as feasible tasks become more productive, aggregate output expands, raising demand for all labor—including in infeasible tasks. The second term is the bottleneck effect within occupations: under Assumption 1, tasks are stronger complements within than across occupations, so technology exposure in an occupation will raise the marginal value of the infeasible tasks that remain. The ratio of the bottleneck effect to the scale effect is $(\rho/\sigma - 1) \times E_o/\bar{E}$, so the relative importance of bottlenecks is governed by the interaction of task complementarity (ρ/σ) and the exposure of infeasible tasks relative to the economy-wide average ($E_o/\bar{E}$). If feasible and infeasible tasks were perfectly segregated across occupations, E_o would be zero wherever infeasible tasks occur and there would be no bottlenecks.

Empirical measure. Since the structural parameters in Equation (2) are positive and common across all task-occupation pairs, we define the *empirical bottleneck pressure*:

$$B_{ot} \equiv (1 - F_t) \times E_o, \tag{3}$$

All inputs—F_t from feasibility assessments, E_o from task-content surveys and employment data—are directly observable. By construction, $B_{ot} = 0$ whenever task t is feasible or occupation o has zero exposure—only infeasible tasks in exposed occupations are technology bottlenecks. Thus B_{ot} suffices to *identify* which tasks are bottlenecks and to quantify their relative importance, but their absolute economic magnitude depends on the size of the productivity gain $\ln \delta$ and the complementarity gap $(1/\sigma - 1/\rho)$.

2 Generative AI

We apply our measure to Generative AI, using the feasibility assessments from Eloundou et al. (2024) for LLMs. Under their most permissive threshold ($F_t = \mathbf{1}[\gamma_t > 0]$, combining LLMs

with complementary tools), 56% of tasks are feasible; the remainder are credible bottleneck candidates. Task contents come from O*NET (Appendix E).

Which tasks are (not) bottlenecks? Table 1 classifies tasks along the two dimensions that determine B_{ot}: feasibility (F_t) and occupation exposure (E_o). Only the lower-right cell—infeasible tasks in occupations where other tasks *are* feasible—constitutes technology bottlenecks ($B_{ot} > 0$). Feasible tasks (top row) are never bottlenecks regardless of exposure, and infeasible tasks in zero-exposure occupations (lower left) have $B_{ot} = 0$: there is no AI-driven productivity for them to constrain.

Table 1: Technology bottleneck classification

	$E_o = 0$ (No exposure)	$E_o > 0$ (Exposure)
$F_t = 1$ (Feasible)	*Not a bottleneck* *Empty by construction: if $E_o = 0$, no tasks are feasible*	*Not a bottleneck* Develop pricing strategies, balancing firm objectives and customer satisfaction — Marketing Managers ($E_o = 0.95$) Analyze financial information to determine strategies for meeting clients' objectives — Personal Financial Advisors ($E_o = 0.95$)
$F_t = 0$ (Infeasible)	*Not a bottleneck* Sweep, mop, or vacuum floors — Janitors ($E_o = 0$) Wash dishes or pans by hand — Dishwashers ($E_o = 0$)	*Technology bottleneck* Direct the hiring, training, or performance evaluations of marketing staff — Marketing Managers ($B_{ot} = 0.95$) Meet with clients' other advisors to fully understand clients' financial goals — Personal Financial Advisors ($B_{ot} = 0.95$)

Notes: Each cell shows example task-occupation pairs, using LLM feasibility assessments from Eloundou et al. (2024). Parenthetical values report occupation exposure E_o or bottleneck pressure $B_{ot} = (1 - F_t) \times E_o$.

The distinction between bottlenecks and non-bottlenecks is sharp in the data. Sweeping floors and washing dishes are beyond LLM capabilities, but they sit in occupations—Janitors ($E_o = 0$), Dishwashers ($E_o = 0$)—where no tasks are feasible. Directing the hiring and evaluation of staff and meeting with a client's attorneys and accountants are equally infeasible, yet they sit in Marketing Managers ($E_o = 0.95$) and Personal Financial Advisors ($E_o = 0.95$), where nearly all other tasks are feasible. Only the latter are technology bottlenecks. Table A.1

lists the infeasible tasks under greatest and least bottleneck pressure.[2]

Composition and concentration of bottlenecks. Figure 1 reveals two patterns. First, the *composition* of bottleneck tasks shifts markedly with exposure (Panel (a)): in low-exposure occupations, bottlenecks are overwhelmingly physical (75% in the bottom quintile), while in the most exposed occupations, interpersonal and supervisory tasks dominate (80%). While bottleneck pressures on individual tasks are by definition highest in the most exposed occupations, *total* bottleneck pressure concentrates at middling exposure (Panel (b)), where occupations retain many infeasible tasks alongside meaningful exposure: the second and third quintiles together account for two-thirds of the total. Across all infeasible tasks, just 4% account for half of total bottleneck pressure (Figure A.1).

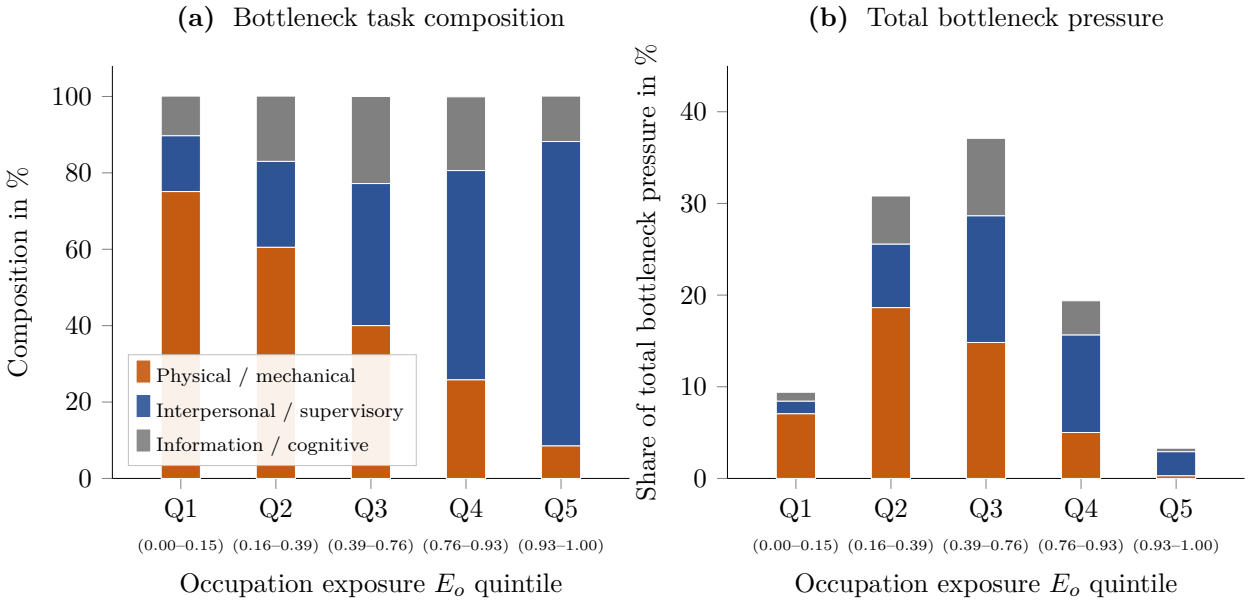


Figure 1: Bottleneck pressure by occupation exposure

Notes: Occupations are binned into employment-weighted quintiles by $E_o = \sum_t \pi_{ot} F_t$, so that each quintile contains approximately 20% of total employment. Panel (a) shows the composition of $\sum_{o,t} \pi_o \pi_{ot} B_{ot}$ by broad task type within each quintile. Panel (b) shows the share of total $\sum_{o,t} \pi_o \pi_{ot} B_{ot}$ within each quintile. Tasks are classified into three broad types—physical/mechanical, interpersonal/supervisory, and information/cognitive—using O*NET's Generalized Work Activity taxonomy (Appendix F).

[2]Bottleneck pressures for all task-occupation pairs are available at `www.andershumlum.com/s/bottlenecks.xlsx`.

The economic importance of bottlenecks. While B_{ot} measures relative bottleneck pressures, quantifying their economic effects requires structural parameters: the productivity gain $\ln \delta$ and the complementarity gap $(1/\sigma - 1/\rho)$. Under conservative parameter values,[3] the scale effect—common across all infeasible tasks—raises MVPL by 0.39 log points (47%). For the tasks under greatest bottleneck pressure, with E_o near 0.97, the within-occupation bottleneck pressure adds a further 0.67 log points (96%)—1.7 times the scale effect. Bottleneck pressures also matter in aggregate: among infeasible tasks (44% of employment), the total employment-weighted bottleneck pressure is about half the corresponding scale effect, because the average E_o among those tasks (0.28) is well below the top-ranked values. Figure D.1 shows the sensitivity of these magnitudes to ρ/σ.

3 Implications and Limitations

Predicting labor demand shifts. Applied to AI, bottleneck pressure predicts rising demand for interpersonal and supervisory capabilities in exposed occupations—the infeasible tasks that become more valuable as technology raises productivity in complementary tasks within the same occupation. Consistent with this, Deming (2017) documents that the return to social skills has grown in occupations where cognitive tasks have become more technologically augmented.

Guiding education and training. Workforce development in AI-exposed occupations should target interpersonal and supervisory capabilities rather than additional analytical skills, enabling workers to move into roles augmented by AI rather than potentially displaced by it. In the most exposed occupations, interpersonal tasks account for 80% of bottleneck pressure, and the highest-pressure tasks are specific interpersonal capabilities such as supervising loan personnel and meeting with clients' advisors.

[3]We set $\delta = 2$ (the lower bound in Eloundou et al. (2024)), $\sigma = 0.5$ (as in Acemoglu and Restrepo (2022); Humlum (2021); Jones and Tonetti (2026)), and $\rho = 1$. Appendix D details the calculations.

Directing R&D. Solving a high-pressure bottleneck not only creates direct value but also releases the productivity gains from the feasible tasks it currently constrains. Since just 4% of infeasible tasks account for half of total bottleneck pressure, targeted capability investments in a small number of areas could unlock a disproportionate share of these gains—and our measure identifies exactly which areas those are.

Evolving capabilities. Rapidly evolving AI capabilities are a feature, not a limitation, of the bottleneck measure. Updated feasibility assessments can be combined with the same task-content data to trace how bottlenecks shift—identifying which constraints have eased and which new ones have emerged.

Limitations. Our measure captures bottlenecks arising from task complementarity *within* occupations; bottlenecks across occupations and along supply chains (Acemoglu, Autor and Patterson, 2024) are outside its scope. It inherits the limitations of its inputs: feasibility assessments are subjective, and O*NET may not reflect how work is actually organized. Finally, B_{ot} captures initial pressure on labor demand holding labor allocations fixed, not equilibrium outcomes.

References

Acemoglu, Daron, and David Autor. 2011. "Skills, Tasks and Technologies: Implications for Employment and Earnings." In *Handbook of labor economics*. Vol. 4, 1043–1171. Elsevier.

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.

Acemoglu, Daron, and Pascual Restrepo. 2022. "Tasks, Automation, and the Rise in U.S. Wage Inequality." *Econometrica*, 90(5): 1973–2016.

Acemoglu, Daron, David Autor, and Christina Patterson. 2024. "Bottlenecks: Sectoral Imbalances and the US Productivity Slowdown." In *NBER Macroeconomics Annual 2023*. Vol. 38. University of Chicago Press.

Autor, David, and Neil Thompson. 2025. "Expertise." *Journal of the European Economic Association*, 23(4): 1203–1271.

Becker, Gary S, and Kevin M Murphy. 1992. "The division of labor, coordination costs, and knowledge." *The Quarterly journal of economics*, 107(4): 1137–1160.

Bick, Alexander, Adam Blandin, and David J Deming. 2026. "The Rapid Adoption of Generative AI." *Management Science*.

Deming, David J. 2017. "The Growing Importance of Social Skills in the Labor Market." *The Quarterly Journal of Economics*, 132(4): 1593–1640.

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock. 2024. "GPTs are GPTs: Labor Market Impact Potential of LLMs." *Science*, 384(6702): 1306–1308.

Felten, Edward, Manav Raj, and Robert Seamans. 2021. "Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses." *Strategic Management Journal*, 42(12): 2195–2217.

Gans, Joshua S, and Avi Goldfarb. 2026. "O-Ring Automation." National Bureau of Economic Research w34639.

Humlum, Anders. 2021. "Robot Adoption and Labor Market Dynamics."

Humlum, Anders, and Emilie Vestergaard. 2025. *Large Language Models, Small Labor Market Effects.* National Bureau of Economic Research.

Jones, Charles I, and Christopher Tonetti. 2026. "Past Automation and Future AI: How Weak Links Tame the Growth Explosion." Working paper, Stanford GSB and NBER.

Tomlinson, Kiran, Sonia Jaffe, Will Wang, Scott Counts, and Siddharth Suri. 2025. "Working with AI: measuring the applicability of generative AI to occupations." *arXiv preprint arXiv:2507.07935*.

Webb, Michael. 2020. "The Impact of Artificial Intelligence on the Labor Market." *Working paper, Stanford University*.

Supplementary Materials
Technology Bottlenecks

Anders Humlum
University of Chicago

A Additional Figures and Tables **2**

B Derivation of the MVPL Expression **3**

C Microfoundations for the Productivity Gain **4**

D Calibrating Economic Magnitudes **5**

E Data **7**

F Generalized Work Activity Classification **8**

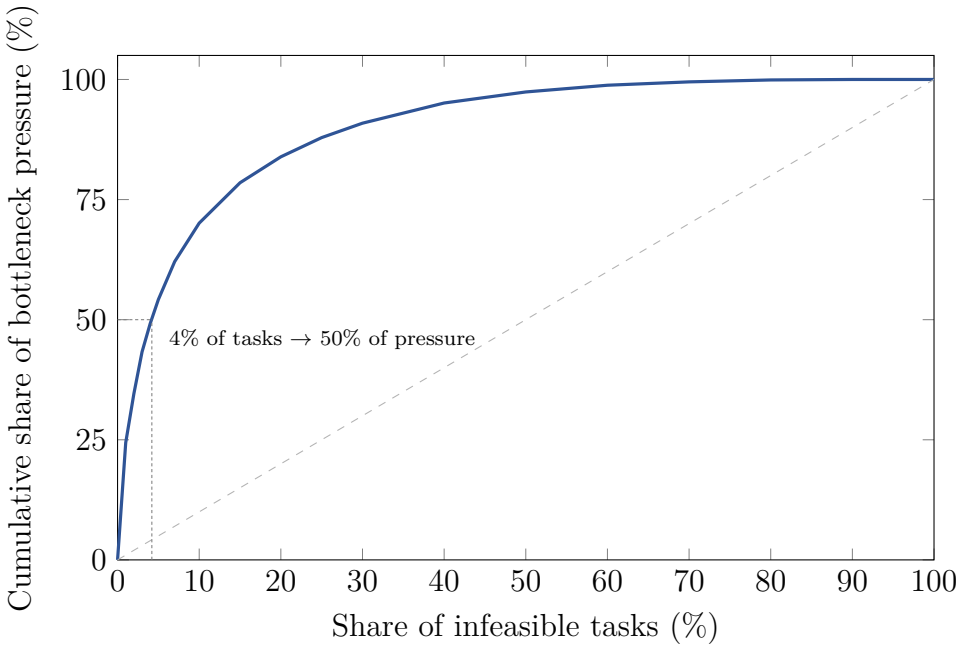
A Additional Figures and Tables

Table A.1: Highest and lowest bottleneck pressures among infeasible tasks

	Occupation	Task description	Type	Emp. (000s)	E_o	B_{ot}
Panel A: Highest bottleneck pressure ($B_{ot} > 0$)						
1	Regulatory Affairs Specialists	Direct the collection and preparation of laboratory samples as requested by regulatory agencies.	I	48	0.97	0.97
2	Loan Officers	Supervise loan personnel.	I	340	0.97	0.97
3	Court & Municipal Clerks	Train other workers or coordinate their work, as necessary.	I	150	0.97	0.97
4	Compliance Managers	Serve as a confidential point of contact for employees to communicate with management...	I	83	0.97	0.97
5	Supply Chain Managers	Appraise vendor manufacturing capabilities through on-site observations...	C	72	0.97	0.97
6	Radiologists	Administer radioisotopes to clinical patients or research subjects.	I	30	0.97	0.97
7	Biomedical Engineers	Conduct training or in-services to educate clinicians and other personnel on proper use of equipment.	I	17	0.97	0.97
8	Claims Adjusters & Examiners	Attend mediations or trials.	C	278	0.97	0.97
9	Architectural & Civil Drafters	Supervise and train other technologists, technicians, and drafters.	I	101	0.96	0.96
10	Environmental Engineers	Assess, sort, characterize, or pack known or unknown materials.	P	43	0.96	0.96
Panel B: Lowest bottleneck pressure ($B_{ot} = 0$)						
1	Janitors & Cleaners	Service, clean, or supply restrooms.	P	2,037	0	0
2	Janitors & Cleaners	Gather and empty trash.	P	2,037	0	0
3	Cooks, Fast Food	Maintain sanitation, health, and safety standards in work areas.	P	768	0	0
4	Cooks, Fast Food	Clean food preparation areas, cooking surfaces, and utensils.	P	768	0	0
5	Dishwashers	Wash dishes, glassware, flatware, pots, or pans, using dishwashers or by hand.	P	377	0	0
6	Dishwashers	Place clean dishes, utensils, or cooking equipment in storage areas.	P	377	0	0
7	Maids & Housekeeping	Carry linens, towels, toilet items, and cleaning supplies...	P	723	0	0
8	Maids & Housekeeping	Clean rooms, hallways, lobbies, lounges, restrooms, corridors, elevators...	P	723	0	0
9	Packaging Machine Ops.	Observe machine operations to ensure quality and conformity...	P	359	0	0
10	Packaging Machine Ops.	Adjust machine components and machine tension and pressure...	P	359	0	0

Notes: Panel A ranks task-occupation pairs by B_{ot} (descending), with occupation employment π_o as tiebreaker, limited to two tasks per occupation. $B_{ot} = (1 - F_t) \times E_o$ is the empirical bottleneck pressure. Panel B lists infeasible tasks with $B_{ot} = 0$ because their occupations have zero exposure ($E_o = 0$), also limited to two per occupation. Type indicates the O*NET Generalized Work Activity classification: P = Physical/mechanical, I = Interpersonal/supervisory, C = Information/cognitive. Emp. is occupation employment (in thousands).

Figure A.1: Concentration of bottleneck pressure



Notes: Infeasible tasks are ranked by employment-weighted bottleneck pressure $\pi_o \pi_{ot} B_{ot}$ (descending). The curve plots the cumulative share of total bottleneck pressure against the share of infeasible tasks. The dashed line shows a uniform distribution for reference.

B Derivation of the MVPL Expression

We derive the change in the marginal value product of labor for infeasible tasks stated in the main text. The MVPL of task t in occupation o is:

$$\text{MVPL}_{ot} = \frac{\partial Y}{\partial L_{ot}} = \frac{\partial Y}{\partial Y_o} \cdot \frac{\partial Y_o}{\partial Y_{ot}} \cdot \frac{\partial Y_{ot}}{\partial L_{ot}}. \tag{4}$$

From the CES aggregators in Equation (1):

$$\frac{\partial Y}{\partial Y_o} = \alpha_o \left(\frac{Y}{Y_o}\right)^{1/\rho}, \qquad \frac{\partial Y_o}{\partial Y_{ot}} = \alpha_{ot} \left(\frac{Y_o}{Y_{ot}}\right)^{1/\sigma}. \tag{5}$$

For an infeasible task ($F_t = 0$), the production function is $Y_{ot} = L_{ot}$ regardless of whether technology is available, so $\partial Y_{ot}/\partial L_{ot} = 1$. Taking logs:

$$\ln \text{MVPL}_{ot} = \ln \alpha_o + \ln \alpha_{ot} + \frac{1}{\rho} \ln Y + \left(\frac{1}{\sigma} - \frac{1}{\rho}\right) \ln Y_o - \frac{1}{\sigma} \ln Y_{ot}. \tag{6}$$

We compare a baseline where no task is feasible ($F_t = 0$ for all t, so $Y_{ot} = L_{ot}$) to a technology scenario where some tasks become feasible. Holding labor allocations fixed, infeasible tasks have $\Delta \ln Y_{ot} = 0$, while feasible tasks have $\Delta \ln Y_{ot} = \ln \delta$. Therefore:

$$\Delta \ln \text{MVPL}_{ot} = \frac{1}{\rho} \Delta \ln Y + \left(\frac{1}{\sigma} - \frac{1}{\rho} \right) \Delta \ln Y_o, \qquad F_t = 0. \tag{7}$$

It remains to compute $\Delta \ln Y_o$ and $\Delta \ln Y$. In the CES, the output elasticity with respect to input t equals the baseline expenditure share. Denoting these baseline shares by π_{ot} for tasks within occupations and π_o for occupations in aggregate output, a first-order approximation gives:

$$\Delta \ln Y_o \approx \sum_{t'} \pi_{ot'} \Delta \ln Y_{ot'} = \ln \delta \sum_{t'} \pi_{ot'} F_{t'} = \ln \delta \cdot E_o, \tag{8}$$

$$\Delta \ln Y \approx \sum_{o} \pi_o \Delta \ln Y_o = \ln \delta \sum_{o} \pi_o E_o = \ln \delta \cdot \bar{E}, \tag{9}$$

where $E_o \equiv \sum_{t'} \pi_{ot'} F_{t'}$ is occupation-level technology exposure and $\bar{E} \equiv \sum_o \pi_o E_o$ is aggregate technology exposure. Substituting (8)–(9) into (7):

$$\Delta \ln \text{MVPL}_{ot} = \ln \delta \times \frac{1}{\rho} \times \bar{E} + \ln \delta \times \left(\frac{1}{\sigma} - \frac{1}{\rho} \right) \times E_o, \qquad F_t = 0. \qquad \square$$

C Microfoundations for the Productivity Gain

The main text parameterizes that technology raises feasible-task output by a factor $\delta > 1$, holding labor fixed. Here we show that this is consistent with both polar cases of technology capital augmenting and substituting for labor.

For feasible tasks ($F_t = 1$), let $Y_{ot} = G(L_{ot}, K_{ot})$ where K_{ot} denotes technology capital (e.g., AI systems).

Perfect complements (augmentation). Technology augments labor but cannot function without it:

$$Y_{ot} = \min(\delta\, L_{ot},\ K_{ot}). \tag{10}$$

Optimal adoption is $K^*_{ot} = \delta\, L_{ot}$, yielding $Y_{ot} = \delta\, L_{ot}$. Here δ is a technological parameter, naturally uniform across tasks and occupations.

Perfect substitutes (substitution). Technology directly replaces human effort:

$$Y_{ot} = L_{ot} + K_{ot}. \tag{11}$$

The firm optimally adopts K^*_{ot} units, and we define $\delta_{ot} \equiv 1 + K^*_{ot}/L_{ot}$, so that $Y_{ot} = \delta_{ot}\, L_{ot}$. Optimal capital allocation implies δ_{ot} is uniform within each occupation: baseline labor mobility gives $L_{ot} \propto \alpha_{ot}^{\sigma}$, and the first-order condition for K_{ot} gives $L_{ot} + K_{ot} \propto \alpha_{ot}^{\sigma}$, so their ratio is constant across tasks. Across occupations, δ_o varies with exposure E_o but is pinned down by E_o and (σ, ρ).

In both cases, the task-level output gain takes the form $\delta_o\, L_{ot}$. The empirical bottleneck measure $B_{ot} = (1 - F_t) \times E_o$ is δ-free: it is invariant to which case applies and to any cross-occupation variation in δ_o. This agnosticism mirrors how feasibility is assessed in practice—Eloundou et al. (2024) evaluate whether technology can reduce the time to perform a task, without distinguishing augmentation from substitution.

D Calibrating Economic Magnitudes

We calibrate Equation (2) using $\delta = 2$ (the lower bound in Eloundou et al. (2024)), $\sigma = 0.5$ (following Acemoglu and Restrepo (2022); Humlum (2021); Jones and Tonetti (2026)), and $\rho = 1$, with aggregate exposure $\bar{E} = 0.56$.

Scale effect. Common to all infeasible tasks:

$$\ln \delta \times \tfrac{1}{\rho} \times \bar{E} = \ln 2 \times 1 \times 0.56 = 0.39 \text{ log points } (47\%).$$

Bottleneck effect. For the highest-ranked bottleneck tasks ($E_o \approx 0.97$):

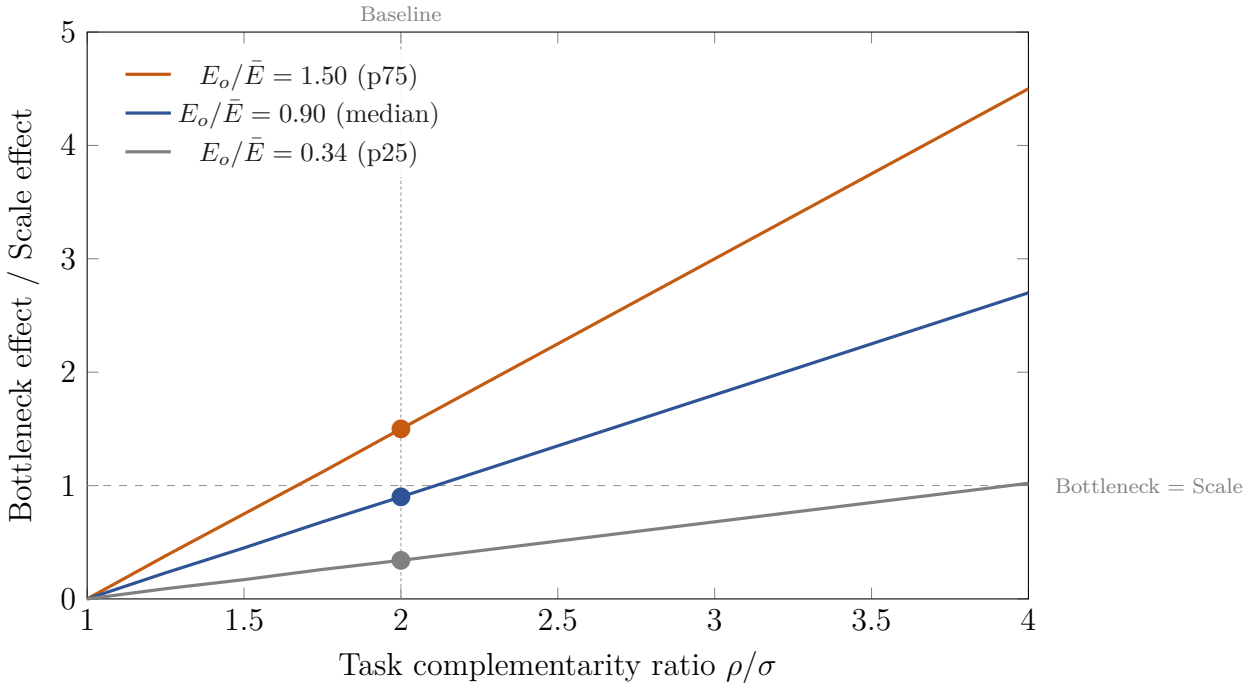
$$\ln \delta \times \left(\tfrac{1}{\sigma} - \tfrac{1}{\rho}\right) \times E_o = \ln 2 \times 1 \times 0.97 = 0.67 \text{ log points } (96\%),$$

or 1.7 times the scale effect.

Aggregate. Among infeasible tasks (44% of employment), the employment-weighted average bottleneck-to-scale ratio is 0.51, reflecting the employment-weighted average $E_o/\bar{E}$ of 0.51 among those tasks.

Sensitivity. The ratio of the bottleneck to the scale effect is $(\rho/\sigma - 1) \times E_o/\bar{E}$, which does not depend on δ. Figure D.1 plots this ratio against ρ/σ at different percentiles of relative exposure $E_o/\bar{E}$ across occupations containing infeasible tasks. The relationship is linear: the bottleneck effect is zero when tasks are equally substitutable within and across occupations ($\rho/\sigma = 1$) and grows proportionally as within-occupation complementarity increases. At the baseline calibration ($\rho/\sigma = 2$), bottleneck pressure exceeds the scale effect for occupations above the 75th percentile of exposure, and remains substantial at the median.

Figure D.1: Sensitivity of bottleneck-to-scale ratio to task complementarity



Notes: Each line plots the ratio of the bottleneck effect to the scale effect, $(\rho/\sigma - 1) \times E_o/\bar{E}$, as a function of the task complementarity ratio ρ/σ. Lines correspond to employment-weighted percentiles of $E_o/\bar{E}$ across occupations containing infeasible tasks. The vertical dotted line marks the baseline calibration ($\rho/\sigma = 2$). The horizontal dashed line marks parity between the bottleneck and scale effects.

E Data

Task contents of occupations (π_{ot}) come from O*NET 30.2, which lists occupation-specific tasks—concrete activities such as "supervise loan personnel" or "sweep, mop, or vacuum floors"—for each of 892 occupations (8-digit O*NET-SOC codes), yielding 18,648 task-occupation pairs. Following Eloundou et al. (2024), we assign equal weight to all tasks within an occupation. Feasibility assessments (F_t) are at the level of O*NET Detailed Work Activities, which map one-to-one to these tasks. Employment shares (π_o) come from the BLS Occupational Employment and Wage Statistics for May 2021, allocated equally across O*NET occupations within each 6-digit SOC code. Analysis code was written with assistance from Claude Code (Anthropic, Opus 4.6).

F Generalized Work Activity Classification

We classify tasks into three broad types using O*NET's Generalized Work Activity (GWA) taxonomy. The O*NET database maps each task to one or more Detailed Work Activities (DWAs), which aggregate to Intermediate Work Activities (IWAs), which in turn aggregate to 41 Generalized Work Activities. We assign each of the 41 GWAs to one of three categories. Table F.1 provides the full mapping.

Table F.1: Generalized Work Activity classification

Physical / mechanical	**Interpersonal / supervisory**	**Information / cognitive**
Performing General Physical Activities	Training and Teaching Others	Getting Information
Handling and Moving Objects	Guiding, Directing, and Motivating Subordinates	Analyzing Data or Information
Controlling Machines and Processes	Coaching and Developing Others	Making Decisions and Solving Problems
Operating Vehicles, Mechanized Devices, or Equipment	Coordinating the Work and Activities of Others	Processing Information
Inspecting Equipment, Structures, or Materials	Developing and Building Teams	Evaluating Information to Determine Compliance
Monitoring Processes, Materials, or Surroundings	Staffing Organizational Units	Thinking Creatively
Repairing and Maintaining Mechanical Equipment	Assisting and Caring for Others	Updating and Using Relevant Knowledge
Repairing and Maintaining Electronic Equipment	Performing for or Working Directly with the Public	Scheduling Work and Activities
Drafting, Laying Out, and Specifying Technical Devices	Establishing and Maintaining Interpersonal Relationships	Organizing, Planning, and Prioritizing Work
	Communicating with Supervisors, Peers, or Subordinates	Interpreting the Meaning of Information
	Communicating with People Outside the Organization	Estimating the Quantifiable Characteristics
	Resolving Conflicts and Negotiating with Others	Judging the Qualities of Objects, Services, or People
	Selling or Influencing Others	Documenting/Recording Information
	Monitoring and Controlling Resources	Providing Consultation and Advice
		Developing Objectives and Strategies
		Identifying Objects, Actions, and Events
		Performing Administrative Activities
		Working with Computers

Notes: Each of O*NET's 41 Generalized Work Activities is assigned to one of three broad categories. When a task maps to multiple GWAs across categories, its bottleneck pressure is split equally across assignments.