

Task assignment and restaurant output: a simulation analysis

Ernesto F. L. Amaral

Department of Sociology, Texas A&M University, College Station, Texas, USA, and

Raymond Robertson

University of Minnesota, Twin Cities, Minnesota, USA

Abstract

Purpose – This study isolates the output contribution of productivity-based task assignment in a fixed workforce and examines how that contribution changes with worker heterogeneity, cross-group comparative advantage and the correlation between task productivities.

Design/methodology/approach – The analysis simulates fixed restaurant workforces of 1,000 workers in five 3,000-run Monte Carlo designs. Within each run, the output-maximizing cutoff assignment is paired with a random assignment that preserves exactly the same front- and back-of-house headcounts. Bivariate-normal productivity draws vary task dispersion, group means and within-worker task correlation; additional calculations vary the Cobb–Douglas task weights. A formula-driven Excel realization reproduces Figure 1, and the Stata materials reproduce the remaining simulation evidence.

Findings – Relative to random assignment with the same task headcounts, optimal sorting raises mean output by 7.6% in the low-dispersion design and 12.1% when front-of-house dispersion is doubled. Greater dispersion changes random-assignment output by only –0.2% but raises optimally assigned output by 4.0%. Assignment gains decline from 10.0 to 5.4% as within-worker task correlation rises from –0.5 to 0.5. A separate cross-advantage design produces incomplete group specialization without materially changing output relative to the fixed group-neutral benchmark.

Research limitations/implications – The results are comparative statics from assumed bivariate-normal productivity distributions. They remain conditional on fixed employment, perfect observation, costless reassignment, additive worker contributions within tasks and Cobb–Douglas aggregation across tasks.

Originality/value – This study does not propose a new theory of assignment, labor demand or immigration. Its value is a controlled computational decomposition that separates sorting from dispersion and group labels, makes the task-level mechanism visible and provides reproducible correlation and task-weight sensitivities.

Keywords Worker heterogeneity, Task assignment, Occupational specialization, Firm productivity, Restaurant industry, Simulation modelling

Paper type Research article

1. Introduction

How much output is attributable to assigning a fixed workforce according to relative task productivity? This study answers that question with a deliberately simple restaurant simulation. It pairs the output-maximizing allocation with a random allocation that has the same task headcounts, then varies productivity dispersion, within-worker task correlation and group-specific comparative advantage. The exercise is not intended to develop a new theory of task assignment, immigration or labor demand. It uses controlled simulation to expose a familiar mechanism and the assumptions governing its magnitude.

The organizing idea comes from the Roy model of occupational selection (Roy, 1951). Sattinger (1993) surveys assignment models in which heterogeneous workers and positions



JEL Classification — D24, J23, J24, C63, F22

© Ernesto F. L. Amaral and Raymond Robertson. Published by Emerald Publishing Limited. This article is published under the Creative Commons Attribution (CC BY 4.0) licence. Anyone may reproduce, distribute, translate and create derivative works of this article (for both commercial and non-commercial purposes), subject to full attribution to the original publication and authors. The full terms of this licence may be seen at <http://creativecommons.org/licences/by/4.0/>

Funding statement: The authors received no external funding for this research.

Conflict of interest disclosure: The authors report no competing interests.

are matched, while [Heckman and Honoré \(1990\)](#) clarify the empirical content of Roy-type selection. [Costinot and Vogel \(2010\)](#) develop a continuous assignment framework in which relative productivity supports cutoff-based matching. Related research studies specialization, diversity and task allocation across workers and economies ([Rosen, 1978](#); [Grossman and Maggi, 2000](#); [Peri and Sparber, 2009](#); [Acemoglu and Autor, 2011](#); [Ottaviano *et al.*, 2013](#); [Deming, 2017](#)). The contribution here is narrower than that literature: it supplies a reproducible decomposition of assignment, dispersion and task correlation in a compact service-sector example.

The production structure should be distinguished from models of team-member complementarity. Worker contributions are summed within each task, so workers assigned to the same task are perfect substitutes in production units. Complementarity appears only between the front-of-house aggregate F and the back-of-house aggregate B . Restaurants provide an intuitive setting because customer-facing activities such as service and order management must be combined with food preparation and other back-of-house activities, while workers may differ in their relative productivity across those tasks.

The native/immigrant distinction is utilized only as an illustrative grouping. Descriptive Current Population Survey Annual Social and Economic Supplement (CPS ASEC) data show different front/back occupational ratios by nativity, but those national person-level patterns do not identify latent productivity or an efficient allocation within one restaurant. The simulation therefore does not calibrate task weights or productivity gaps. It instead rounds the observed foreign-born share among economically active front/back workers to set group size, and uses a transparent cross-advantage design in which each group has a higher mean in one task.

The analysis makes three distinctions explicit. First, the difference between optimal and random assignment isolates sorting worker identities while holding task headcounts fixed. Second, the high-versus-low dispersion comparison is a result about group-neutral worker heterogeneity, not immigration. Third, the cross-group design reports assignment patterns and output separately, so occupational specialization is not mistaken for the source of the dispersion result. Employment and wages remain fixed, and the exercise establishes no labor-demand response.

The rest of the article proceeds as follows. [Section 2](#) sets out the assignment problem and defines separate assignment and common-benchmark ratios. [Section 3](#) documents the CPS context, simulation designs, sorting decomposition and sensitivity analyses. [Section 4](#) discusses interpretation and limitations.

2. Task-assignment framework

Consider a restaurant with a fixed number of workers L . Output Y is produced by combining front-of-house task output F and back-of-house task output B . A simple representation is:

$$Y = F^\varphi B^\gamma \quad (1)$$

where φ and γ are positive task-output weights. The main simulations use $\varphi = \gamma = 0.45$. Equal weights avoid imposing relative front- or back-of-house intensity in the baseline; their sum of 0.90 implies decreasing returns when the scale of both task aggregates changes. Because the workforce is fixed, this value is a calibration rather than an estimate of returns to scale. The sensitivity analysis also uses equal weights summing to 0.60 and 1.00 and a front-intensive specification with $\varphi = 0.50$ and $\gamma = 0.30$.

Each worker i supplies one unit of labor and has task-specific productivities α_{if} and α_{ib} . If assigned to the front of the house, the worker contributes $f_i = \alpha_{if}$; if assigned to the back, the worker contributes $b_i = \alpha_{ib}$. Contributions are additive within each task. Productivities are affine transformations of bivariate standard-normal shocks and are observed before assignment. To keep relative-productivity ratios and Cobb–Douglas output defined, the

code replaces a nonpositive productivity draw with 0.01. The baseline within-worker correlation is $\rho = 0$, while $\rho = -0.5$ and $\rho = 0.5$ are examined as comparative statics.

The worker-specific ratio $\theta_i = \alpha_{ib}/\alpha_{if}$ summarizes relative productivity. Lower values indicate comparative advantage in front-of-house work, while higher values indicate comparative advantage in back-of-house work. After sorting workers from low to high θ_i , the firm can assign the first k workers to the front and the remaining $L-k$ workers to the back. Aggregate task outputs are then:

$$F(k) = \sum_{i \in \text{front}(k)} \alpha_{if}, B(k) = \sum_{i \in \text{back}(k)} \alpha_{ib}. \tag{2}$$

The output-maximizing assignment is the cutoff k^* that maximizes $Y(k) = F(k)^\varphi B(k)^\gamma$ over all feasible cutoffs. At the margin, the firm compares a worker’s relative productivity with the common shadow-value ratio implied by Equation (1):

$$MRS_{FB} = \frac{\partial Y / \partial F}{\partial Y / \partial B} = \frac{\varphi B}{\gamma F}. \tag{3}$$

Figure 1 illustrates this logic using the supplementary Excel realization (Supplementary Material 1). The blue curve plots the relative shadow value of front-of-house output, $(\varphi/\gamma) (B/F)$, as the cutoff changes. The orange curve plots θ_i after workers are ordered from lower to




Figure 1. Task-assignment cutoff in the illustrative cross-advantage realization. Note: Author-generated Excel realization with 1,000 workers (770 native and 230 immigrant), $\varphi = \gamma = 0.45$ and $\rho = 0$. Natives have the higher front-of-house mean and immigrants the higher back-of-house mean; all task standard deviations equal 1.5. The horizontal axis reports worker percentile rank after sorting by $\theta_i = \alpha_{ib}/\alpha_{if}$. The blue curve plots $(\varphi/\gamma) (B/F)$, the orange curve plots θ_i and the dashed line marks the output-maximizing cutoff. The display ends at percentile 99.5 and caps the vertical axis at 7 for readability; all 1,000 formula rows and uncapped values remain in Replicate_Figure_1.xlsx (Supplementary Material 1). Source: Authors’ own work

higher relative back-of-house productivity. Their intersection, marked by the dashed line, identifies k^* . Searching all ordered cutoffs yields the global optimum in this two-task additive assignment problem; the horizontal axis is a worker rank, not a temporal ordering of restaurant activities.

2.1 Output benchmarks

Let d index a simulation design. We use two ratios to distinguish the output gain from sorting worker identities within a design from differences in optimally assigned output across designs. First, define the assignment ratio A_d as:

$$A_d = \frac{\text{mean}(Y_d^*)}{\text{mean}(Y_d^R)} \quad (4)$$

For design d , Y_d^* denotes output under the optimal cutoff and Y_d^R denotes output when exactly k^* workers are selected for the front at random, independently of productivity. The remaining workers are assigned to the back. Thus, $A_d - 1$ is the proportional gain from sorting worker identities, conditional on the same design and task headcounts. The reported ratios divide mean output across runs rather than averaging run-specific ratios.

Because A_d compares two assignments within the same design, it does not provide a common basis for comparing output levels across designs. For that purpose, define the common-benchmark ratio H_d as:

$$H_d = \frac{\text{mean}(Y_d^*)}{\text{mean}(Y_0^*)} \quad (5)$$

The fixed common-distribution benchmark Y_0^* is the optimally assigned, group-neutral low-dispersion design: both labels draw both task productivities with mean 10 and standard deviation 1.5, $\rho = 0$ and $\varphi = \gamma = 0.45$. H_d uses one explicit denominator for every design. Under Cobb–Douglas aggregation, any comparison also satisfies $\Delta \ln Y = \varphi \Delta \ln F + \gamma \Delta \ln B$. Percentage output differences consequently depend on the task weights; reporting F and B and varying the weights separates the assignment mechanism from that scaling choice.

2.2 Nativity as an illustrative grouping

The cross-group design contains 770 native and 230 immigrant workers. Natives have the higher mean front-of-house productivity (10.14935 vs 9.5), while immigrants have the higher mean back-of-house productivity (10.5 versus 9.85065). With a 77/23 workforce split, the pooled mean equals 10 in each task. Both groups have a standard deviation of 1.5 in each task and $\rho = 0$. These assumptions create opposite-task group comparative advantages while preserving overlapping distributions; they are not estimates of productivity by nativity.

Figure 2 is descriptive context, not a calibration target. It uses 23 annual CPS ASEC cross-sections from 2000 through 2022. The plotted series are ASEC-weighted percentages of each nativity-year sample reporting front-of-house, back-of-house or managerial restaurant occupations. The civilian labor-force front/back subsample is used to set group size: variable LABFORCE equals to 2, excluding Armed Forces, which are defined in variable EMPSTAT codes 30–36. This procedure results in 96,863 person-year observations yielding a weighted foreign-born share of 22.78%, rounded to 23% in the simulations (variable ASECWT is the person-level survey weight). National occupational shares reflect worker supply, employer demand, institutions, location and selection across establishments; they cannot identify the within-restaurant productivity draws or task weights required to fit the model.




Figure 2. Weighted percentage of U.S.-born and foreign-born CPS ASEC samples reporting restaurant occupations, by task category, 2000–2022. Note: Each series is 100 times the ASECWT-weighted share of all respondents in the indicated nativity group and year reporting a harmonized 2010 occupation in the task category. Front-of-house codes are 4,040, 4,110, 4,120 and 4,150; back-of-house codes are 4,000, 4,030 and 4,140; manager is 4,010. The figure uses 4,462,254 person-year records from 23 annual cross-sections. The simulation share is calculated separately among 96,863 front/back observations in the civilian labor force (LABFORCE = 2, excluding Armed Forces EMPSTAT codes 30–36), for which the ASECWT-weighted foreign-born share is 22.78%. Source: Authors’ calculations using IPUMS CPS: Version 13.0 (Flood *et al.*, 2025), 2000–2022 Current Population Survey Annual Social and Economic Supplement

3. Simulation design and results

3.1 Reproducible simulation design

The supplementary Excel workbook (Supplementary Material 1) contains one formula-driven 1,000-worker realization of the cross-advantage design. It records task-specific productivity, θ_i , the output from every cutoff, the optimal and random assignments, task aggregates, total output and the fixed group-neutral common benchmark. Its native Excel chart reproduces

Figure 1. The random assignment uses fixed ranks independent of productivity and exactly the optimal number of workers in each task. The workbook is an auditable illustration, not the Monte Carlo dataset.

The Stata replication files comprise the do-file ([Supplementary Material 2](#)), execution log ([Supplementary Material 3](#)), run-level dataset ([Supplementary Material 4](#)), design-level summary ([Supplementary Material 5](#)), assignment-dispersion summary ([Supplementary Material 6](#)), task-weight summary ([Supplementary Material 7](#)), correlation summary ([Supplementary Material 8](#)) and paired worker-reassignment summary ([Supplementary Material 9](#)). The Stata analysis contains five design blocks with 3,000 runs each, producing 15,000 run-level observations. Every run contains 1,000 workers, draws front- and back-of-house productivity, searches all ordered cutoffs and records the optimum and a same-headcount random assignment. The designs are: group-neutral low dispersion; group-neutral high front-of-house dispersion; cross-group comparative advantage and low-dispersion correlation designs with $\rho = -0.5$ and $\rho = 0.5$. The $\rho = 0$ low-dispersion design is the correlation reference. Common random numbers pair latent worker shocks across designs. A further 3,000-run paired calculation traces worker movements when front-of-house dispersion changes. All main blocks use $\varphi = \gamma = 0.45$.

3.2 CPS context and cross-group assignment

The pooled CPS front/back employment ratio is 0.951 for U.S.-born workers and 0.381 for foreign-born workers. [Figure 2](#) establishes a descriptive difference in direction and magnitude, but it does not identify efficient firm-level assignment. We do not select φ/γ or the group mean gaps to fit these ratios: doing so would attach structural meaning to national occupation counts that lack establishment identifiers, worker productivity, wages, staffing constraints and the relevant choice set. The simulation uses the CPS only to motivate the task categories and set the rounded foreign-born share.

[Figure 3](#) uses the cross-advantage design specified in [Section 2.2](#). Under random assignment, mean native and immigrant front/back ratios are 1.002 and 1.006, and immigrant shares are approximately 23% in both tasks. Optimal assignment changes the native ratio to 1.247 and the immigrant ratio to 0.468; immigrants then constitute 14.6% of front-of-house workers and 31.4% of back-of-house workers. Both groups remain in both tasks, so specialization is incomplete. These model-generated ratios were not calibrated to the CPS ratios reported above and do not match them; they arise from the assumed productivity distributions and optimal assignment rule.

The cross-advantage design raises mean output by 7.83% relative to its same-headcount random assignment ($A_d = 1.0783$), but its optimal output is only 0.23% above the fixed group-neutral benchmark ($H_d = 1.0023$). The specialization pattern and the output comparison are distinct: the design reallocates groups sharply across tasks without generating a large output difference relative to a benchmark with the same pooled task means. The design-level results are provided in [Supplementary Material 5](#).

3.3 Isolating assignment from dispersion

The two dispersion designs are group-neutral: both labels share task means of 10, and group identity has no effect. Both use a back-of-house standard deviation of 1.5 and $\rho = 0$. The low-dispersion design sets the front-of-house standard deviation to 1.5; the high-dispersion design raises it to 3. Each random counterfactual preserves the optimal front/back headcounts for its run. This four-cell design separates the contribution of sorting from the contribution of dispersion under sorting.

The simulation isolates a simple result: heterogeneity creates a potential gain, but productivity-based assignment realizes it ([Table 1](#)). Under low dispersion, optimal sorting raises mean output by 7.58% relative to random assignment with the same task headcounts. Under high front-of-house dispersion, the sorting gain is 12.11%. High dispersion changes




Figure 3. Optimal task assignment under cross-group comparative advantage. Note: Results are from 3,000 runs with 1,000 workers per run (770 native and 230 immigrant), $\varphi = \gamma = 0.45$ and independent task draws ($\rho = 0$). Native and immigrant front-of-house means are 10.14935 and 9.5; back-of-house means are 9.85065 and 10.5. All standard deviations equal 1.5, and the pooled mean is 10 in each task. Panel A plots optimal front/back employment ratios; Panel B plots immigrant shares within each task. Random assignment uses the same task headcounts but worker identities independent of productivity. *Source:* Authors' own work

mean random-assignment output by only -0.17% , from 2,133.7 to 2,129.9, but raises mean optimally assigned output by 4.03% , from 2,295.4 to 2,387.9. The dispersion comparison is secondary to the within-design optimal/random contrasts.

The task aggregates make the mechanism visible. As front-of-house dispersion rises, the mean optimal cutoff falls from 500.0 to 468.1 workers. Relative to low dispersion, optimally

Table 1. Assignment and front-of-house productivity dispersion

Dispersion	Assignment	Front workers	Front output F	Back output B	Total output Y	Optimal/random	High/low
Low	Random	500.0	5000.8	5000.7	2133.7	–	1.0000
Low	Optimal	500.0	5423.4	5423.7	2295.4	1.0758	1.0000
High	Random	468.1	4683.4	5319.3	2129.9	–	0.9983
High	Optimal	468.1	5743.1	5591.9	2387.9	1.1211	1.0403

Note(s): Entries are means across 3,000 runs per dispersion design. Random assignment preserves the optimal task headcounts in each run but selects worker identities independently of productivity. All groups have mean 10 in both tasks; the front-of-house standard deviation is 1.5 (low) or 3 (high), while the back-of-house standard deviation is 1.5. Task draws are independent and $\varphi = \gamma = 0.45$. Optimal/random compares assignments within a dispersion design; high/low compares dispersion within an assignment rule. The underlying assignment-dispersion summary is provided in [Supplementary Material 6](#)

Source(s): Authors' own work

assigned front output rises 5.89% and back output rises 3.10%. Under random assignment, front output falls 6.35% and back output rises 6.37%, changes that approximately offset in the symmetric Cobb–Douglas aggregator. In the paired worker audit, an average of 67.0 workers move from front to back and 35.2 move from back to front, so 10.2% change task and the net front-of-house decline is 31.9 workers. The paired worker-reassignment statistics are provided in [Supplementary Material 9](#).

The identities of the movers also accord with comparative advantage. Workers sent from front to back have below-average front productivity (mean standardized front draw -0.617); doubling front dispersion raises their mean relative-productivity ratio from 0.950 to 1.073. Workers sent from back to front have above-average front productivity (mean standardized draw 0.857), and their mean relative-productivity ratio falls from 1.033 to 0.930. Greater dispersion stretches relative productivity in both directions, allowing the firm to improve both task aggregates through gross reallocation even though the front-house headcount declines.

[Figure 4](#) displays all four output distributions. The random densities nearly overlap across dispersion levels, while the optimal density moves farther right under high dispersion. This visual comparison and [Table 1](#) jointly show that dispersion creates additional output in this exercise through selective assignment rather than through a mechanical change in mean productivity.

3.4 Correlation and task-weight sensitivity

The correlation comparative static holds both marginal task distributions fixed. When ρ rises from -0.5 to 0 and 0.5 , the sorting gain falls from 10.02% to 7.58% and 5.37%. Mean output under random assignment remains approximately 2,133.5–2,133.7, while optimal output falls as positive correlation compresses the distribution of relative productivity. This result gives the illustration a comparative static that depends on joint, not merely marginal, productivity. [Table 2](#) reports both ratios for all five designs, separating the within-design assignment gain (A_d) from optimal output relative to the fixed common benchmark (H_d). The design-level and correlation summaries are provided in [Supplementary Materials 5](#) and [8](#).

Task weights affect output magnitudes exactly as the Cobb–Douglas identity implies. With equal weights summing to 0.60, 0.90 and 1.00, the low-dispersion sorting gain is 4.99, 7.58 and 8.46%, while the high-dispersion gain is 7.92, 12.11 and 13.54% ([Table 3](#)). Proportional rescaling leaves the cutoff unchanged because φ/γ is unchanged. The front-intensive $\varphi = 0.50$, $\gamma = 0.30$ specification changes that ratio, moves the cutoff and yields gains of 6.39 and 10.63%. No single percentage should be read as technology-free. The task-weight summary data are provided in [Supplementary Material 7](#).

The results separate three channels. Marginal dispersion governs the available tails, within-worker correlation governs the spread of relative productivity and task weights translate




Figure 4. Total output by assignment and front-of-house productivity dispersion. Note: Each panel contains 3,000 runs with 1,000 workers, $\varphi = \gamma = 0.45$ and $\rho = 0$. Both group labels have task means of 10 and identical distributions. The low-dispersion panel sets both task standard deviations to 1.5. The high-dispersion panel raises only the front-of-house standard deviation to 3. In every run, the random assignment uses the same task headcounts as the optimum but selects worker identities independently of productivity. Common random numbers pair latent worker shocks across panels. *Source:* Authors' own work

Table 2. Design-level assignment and common-benchmark ratios

Design	ρ	Mean cutoff	Optimal Y	Random Y	A_d	H_d
Low dispersion	0	500.0	2295.4	2133.7	1.0758	1.0000
High front dispersion	0	468.1	2387.9	2129.9	1.1211	1.0403
Cross-group advantage	0	500.2	2300.6	2133.7	1.0783	1.0023
Low dispersion, negative correlation	-0.5	506.6	2347.3	2133.5	1.1002	1.0226
Low dispersion, positive correlation	0.5	500.0	2248.2	2133.6	1.0537	0.9794

Note(s): Each design contains 3,000 runs. A_d is mean optimal output divided by mean same-headcount random-assignment output. H_d is mean optimal output divided by 2,295.4, the mean optimum in the fixed group-neutral, low-dispersion, $\rho = 0$ benchmark. All rows use $\varphi = \gamma = 0.45$. Marginal distributions are identical across the three correlation rows

Source(s): Authors' own work

changes in F and B into total output. Group labels affect assignment only in the cross-advantage design, where the group-specific distributions differ.

4. Discussion and conclusion

The simulation distinguishes heterogeneity from the assignment process that uses it. Random assignment supplies a direct counterfactual, the paired dispersion analysis traces the workers

Table 3. Task-weight sensitivity of assignment gains

Dispersion	φ	γ	$\varphi + \gamma$	Mean cutoff	A_d	Gain (%)
Low	0.30	0.30	0.60	500.0	1.0499	4.99
Low	0.45	0.45	0.90	500.0	1.0758	7.58
Low	0.50	0.50	1.00	500.0	1.0846	8.46
Low	0.50	0.30	0.80	617.6	1.0639	6.39
High	0.30	0.30	0.60	468.1	1.0792	7.92
High	0.45	0.45	0.90	468.1	1.1211	12.11
High	0.50	0.50	1.00	468.1	1.1354	13.54
High	0.50	0.30	0.80	581.7	1.1063	10.63

Note(s): Entries are means across the same 3,000 runs used in Table 1. A_d compares the optimal and same-headcount random assignments under the indicated task weights. Proportional rescaling of equal weights leaves φ/γ and the cutoff unchanged; the front-intensive 0.50/0.30 specification changes both the weight ratio and the cutoff

Source(s): Authors' own work

and task aggregates behind the output change and the correlation experiment shows that the joint distribution of task skills matters. These analyses remain comparative statics within a familiar assignment model rather than a new theory.

The main interpretive boundary concerns labor demand and managerial implementation. Employment and wages are fixed, so higher output does not establish that the firm would hire more workers. Measuring task-specific productivity, screening workers and changing assignments may also be costly. The simulation identifies the gross output available under full information and costless reassignment, not the net return to a particular management practice.

A second boundary concerns immigration. Figure 3 shows how two overlapping groups with opposite mean comparative advantages can sort incompletely, but the group distributions are assumed rather than estimated. It does not show that nativity causes productivity, that observed CPS allocations are efficient or that immigration raises restaurant output. The group-neutral dispersion and correlation results would be unchanged if the labels were removed.

Further work could estimate joint task-productivity distributions from establishment data, incorporate imperfect observation and reassignment costs and compare Cobb–Douglas with CES or Leontief aggregation. Endogenous hiring and wages would be required to connect output to labor demand. Establishment identifiers and direct measures of task performance would also be needed to calibrate the model to observed assignment ratios.

Within those limits, the note provides a reproducible benchmark for separating four objects that are easily conflated: worker heterogeneity, productivity-based assignment, group composition and the technology that aggregates task output. The central finding is not a universal output percentage. It is that dispersion increases output in this exercise only when assignment can select worker identities, and that the size of this gain varies predictably with task correlation and task weights.

Ethics approval statement

This study does not involve human subjects research requiring institutional review board approval. The descriptive evidence uses publicly available Current Population Survey data and the main analysis is simulation-based.

Data availability statement

The nine supplementary files provided with this article are cited as follows. Replicate_Figure_1.xlsx (Supplementary Material 1) is the formula-driven 1,000-worker Excel realization used for Figure 1. Replicate_Figures_3_4.do, version 1.0 (Supplementary Material 2), contains the Stata replication code, and Replicate_Figures_3_4.log, version 1.0 (Supplementary Material 3), records its execution. Replicate_Figures_3_4.csv (Supplementary Material 4) contains the 15,000 run-level observations, while Simulation_Design_Summary.csv (Supplementary Material 5) contains the five design-level summaries. Table_1_assignment_dispersion.csv (Supplementary Material 6), Table_2_task_weight_sensitivity.csv (Supplementary Material 7), Table_3_correlation_sensitivity.csv (Supplementary Material 8) and Table_4_dispersion_reassignment.csv (Supplementary Material 9) contain the assignment-dispersion, task-weight summaries, correlation and paired worker-reassignment, respectively. Supplementary Materials 1–9 reproduce the simulation evidence reported in the manuscript. Figure 2 uses publicly available CPS ASEC microdata; its exact years, sample restrictions, weight, occupation codes and source citation are reported in the manuscript and workbook, but its source microdata and construction code are not included in the simulation supplement.

Declaration of AI-assisted editing

During revision of this manuscript, the authors used OpenAI ChatGPT (GPT-5.5 with the High reasoning setting) and OpenAI Codex (GPT-5 and GPT-5.5 with the Extra High reasoning setting), accessed June–September 2026, to support copy-editing, restructuring, condensation and the organization and formatting of supplementary Excel and Stata materials. AI assistance was also used to prepare and check code, figures, tables and manuscript wording. The authors developed the economic model, selected the assumptions and parameter designs, interpreted the results, ran the software and verified all formulas, code, outputs, figures and text. The authors take full responsibility for the content of the publication.

Permission to reproduce material from other sources

No copyrighted materials from other sources are reproduced in this manuscript. All figures are original to this study or derived from publicly available data.

Supplementary material

The supplementary material for this article can be found online.

References

- Acemoglu, D. and Autor, D. (2011), “Skills, tasks and technologies: implications for employment and earnings”, in Ashenfelter, O. and Card, D. (Eds), *Handbook of Labor Economics*, Elsevier, Amsterdam, Vol. 4B, pp. 1043–1171. doi: [10.1016/S0169-7218\(11\)02410-5](https://doi.org/10.1016/S0169-7218(11)02410-5).
- Costinot, A. and Vogel, J. (2010), “Matching and inequality in the world economy”, *Journal of Political Economy*, Vol. 118 No. 4, pp. 747–786, doi: [10.1086/655858](https://doi.org/10.1086/655858).
- Deming, D.J. (2017), “The growing importance of social skills in the labor market”, *The Quarterly Journal of Economics*, Vol. 132 No. 4, pp. 1593–1640, doi: [10.1093/qje/qjx022](https://doi.org/10.1093/qje/qjx022).
- Flood, S., King, M., Rodgers, R., Ruggles, S., Warren, J.R., Backman, D., Breton, E., Cooper, G., Rivera Drew, J.A., Richards, S., Van Riper, D. and Williams, K.C.W. (2025), *IPUMS CPS: Version 13.0 [Dataset]*, IPUMS, Minneapolis, MN, doi: [10.18128/D030.V13.0](https://doi.org/10.18128/D030.V13.0).
- Grossman, G.M. and Maggi, G. (2000), “Diversity and trade”, *The American Economic Review*, Vol. 90 No. 5, pp. 1255–1275, doi: [10.1257/aer.90.5.1255](https://doi.org/10.1257/aer.90.5.1255).
- Heckman, J.J. and Honoré, B.E. (1990), “The empirical content of the Roy model”, *Econometrica*, Vol. 58 No. 5, pp. 1121–1149, doi: [10.2307/2938303](https://doi.org/10.2307/2938303).
- Ottaviano, G.I.P., Peri, G. and Wright, G.C. (2013), “Immigration, offshoring, and American jobs”, *American Economic Review*, Vol. 103 No. 5, pp. 1925–1959, doi: [10.1257/aer.103.5.1925](https://doi.org/10.1257/aer.103.5.1925).
- Peri, G. and Sparber, C. (2009), “Task specialization, immigration, and wages”, *American Economic Journal: Applied Economics*, Vol. 1 No. 3, pp. 135–169, doi: [10.1257/app.1.3.135](https://doi.org/10.1257/app.1.3.135).

Rosen, S. (1978), "Substitution and division of labour", *Economica*, Vol. 45 No. 179, pp. 235-250, doi: [10.2307/2553070](https://doi.org/10.2307/2553070).

Journal of
Economic Studies

Roy, A.D. (1951), "Some thoughts on the distribution of earnings", *Oxford Economic Papers*, Vol. 3 No. 2, pp. 135-146, doi: [10.1093/oxfordjournals.oep.a041827](https://doi.org/10.1093/oxfordjournals.oep.a041827).

Sattinger, M. (1993), "Assignment models of the distribution of earnings", *Journal of Economic Literature*, Vol. 31 No. 2, pp. 831-880, available at: <http://www.jstor.org/stable/2728516>

363

Corresponding author

Ernesto F. L. Amaral can be contacted at: amaral@tamu.edu