Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Shared-Capacity Constraints in Occupational Transition Networks: A Network-Flow Audit of U.S. Workforce Transition Capacity

Ghassan Malkawi

gmalkawi@hct.ac.ae

Higher Colleges of Technology

Ahmed Elsayed

Canadian University Dubai

Mohammed Alhagyan

United Arab Emirates University

Bakeel Hussein

Sana'a Community College

Research Article

Keywords:

Posted Date: November 13th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11308901/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Shared-Capacity Constraints in Occupational Transition Networks: A Network-Flow Audit of U.S. Workforce Transition Capacity

Ghassan Malkawi*[1], Ahmed Elsayed[2], Mohammed Alhagyan[3], and Bakeel Hussein[4]

[1]Faculty of Computer Information Science, Higher Colleges of Technology, Al Ain, Abu Dhabi, United Arab Emirates

[2]Department of Computer Engineering and Computational Sciences, Canadian University Dubai, Dubai, United Arab Emirates

[3]Department of Mathematical Sciences, College of Science, United Arab Emirates University, Al Ain, United Arab Emirates

[4]Department of Applied Engineering, Sana'a Community College, Sana'a, Yemen

Abstract

Occupational-transition studies often identify plausible origin–destination links without testing whether many routes can scale at the same time when they draw on

*Corresponding author: gmalkawi@hct.ac.ae; ORCID: 0000-0002-5320-7561

the same receiving occupations. We formulate this aggregation problem as a directed capacitated occupational network and compare two quantities: an optimistic separate-origin ceiling, which lets each origin reuse destination capacity independently, and a joint maximum-flow allocation, which shares each destination capacity across all origins. In the primary U.S. route network (System A), the 169-origin denominator is embedded in a 661-occupation analytical frame. At the declared 50% opening-access stress, separate-origin routability is 83.65%, whereas joint routability is 49.40%, giving a 34.25 percentage-point shared-capacity gap. A fixed-design 2025–35 update yields a 32.97 percentage-point gap. Alternative route constructions show that route semantics and accessible capacity govern the gap and can eliminate it under sufficiently permissive assumptions. A continuous parametric audit over $0 \leq \rho \leq 1$ shows that the gap is nonmonotone in accessible openings: in the 2024–34 released/no-backfill case it peaks at 45.93 percentage points at $\rho = 0.234068$, while the headline $\rho = 0.50$ value remains 34.25 points. Optimal-face analysis shows that the aggregate shortfall is not uniquely assigned to a fixed origin subset, whereas all-minimum-cut analysis identifies persistent destination-side capacity arcs. A separately constructed Occupational Information Network (O*NET) route system provides a sensitivity check based on a different edge source. Linked December 2023–December 2024 Survey of Income and Program Participation (SIPP) records provide a downstream behavioral boundary: strict cross-major annual mobility is 9.72%, System-A route concordance is bounded between 3.97% and 8.10%, and only two of five persistent receiving destinations receive strict annual support. The analysis distinguishes route adjacency from jointly scalable transition capacity.

Keywords: occupational transition networks; network flow; shared capacity; occupational mobility; workforce planning; minimum cut; reskilling

1 Introduction

Occupational mobility can be represented as a directed network in which occupations are nodes and plausible moves are links, enabling analysis of redeployment opportunities, skill proximity, and mobility constraints (Schmutte, 2014; del Rio-Chanona et al., 2021; Moro et al., 2021; Joyez and Laffineur, 2022; Waters and Shutters, 2022; Frank et al., 2024). Technological change can generate transition recommendations for many occupations simultaneously, so route feasibility must also be assessed at the system level. A route that is feasible for one origin in isolation may become capacity-constrained when several origins depend on the same destinations.

This paper focuses on that network-scaling problem. We transform occupation-to-occupation route evidence into a directed capacitated bipartite flow network. Origin occupations carry modeled supply T_i, destination occupations carry accessible capacity $C_j(\rho)$, and admitted edges encode candidate transition routes. We compare two network objects built on the same origins, edges, and capacities: an intentionally optimistic separate-origin ceiling that allows each origin to reuse the full destination-capacity vector, and a joint maximum-flow allocation in which each destination capacity can be consumed only once across all origins. Their difference is the shared-capacity competition gap.

The contribution is a network-science measurement of how routability changes when many occupational pathways are evaluated as one capacity-constrained system rather than as isolated origin-specific routes. Existing occupational-network studies have used worker flows, skills, and relatedness to map occupation spaces, centrality, and retraining proximity (del Rio-Chanona et al., 2021; Joyez and Laffineur, 2022; Waters and Shutters, 2022; Frank et al., 2024). Those analyses establish that mobility is structured by network position and edge semantics. The present study adds a distinct question: how much of the apparent opportunity in a transition network survives simultaneous use of a common destination-capacity budget, and which origin- and destination-side constraints are invariant across the full set of maximum-flow optima?

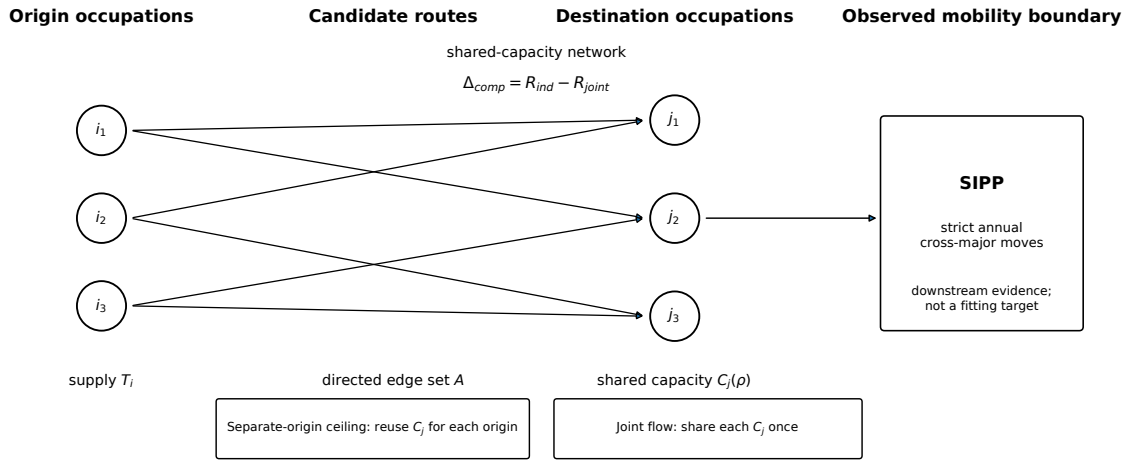


Figure 1: Network representation and shared-capacity estimand. Candidate occupation-to-occupation edges form a directed bipartite flow network with origin supply T_i and destination capacity $C_j(\rho)$. The separate-origin ceiling reuses destination capacity independently across origins, whereas joint routing shares each destination capacity once. Survey of Income and Program Participation (SIPP) data remain downstream observed-mobility evidence rather than a fitting target.

To answer that question, we use two occupational route networks built from distinct edge sources. System A, the primary test, combines U.S. Department of Labor (DOL)/Emsi ranked transition evidence with U.S. Bureau of Labor Statistics (BLS) occupational transfers and openings. System B uses Occupational Information Network (O*NET) Related Occupations under explicit relatedness, education, wage, and directionality screens. The two systems are not averaged because they encode different edge semantics. Across both, projected transfers represent modeled origin-supply pressure and openings serve as a planning-capacity proxy. Figure 1 summarizes the network representation and the distinction between separate and joint capacity accounting.

The primary System-A denominator contains 169 origins drawn from a common 661-occupation analytical frame. The reference and later-vintage analyses both yield positive competition gaps at $\rho = 0.50$, with magnitude varying across route definitions and opening-access assumptions. Longitudinal SIPP is analyzed separately as observed-mobility evidence.

The study is organized around two scientific questions. **RQ1:** When plausible routes from

many origins must use one shared destination-capacity vector, how much apparent routability is lost and which constraints remain structural across alternative maximum-flow allocations? Route-menu, opening-access, and BLS-vintage analyses test the dependence of this network result on edge semantics and capacity assumptions. **RQ2:** Do declared System-A routes and persistent receiving destinations appear in one-year worker mobility histories once job identity, survey uncertainty, and set-valued occupation mapping are treated conservatively?

Three findings organize the contribution. First, shared destination capacity materially reduces routability in the primary network, and the gap remains positive across all three route constructions at the primary opening-access stress in both BLS vintages. Second, the aggregate shortfall is allocation-sensitive across the optimal face rather than structurally assigned to a fixed origin subset, while persistent destination-capacity arcs limit additional total flow. Third, observed annual mobility and route concordance are substantially narrower than structural route feasibility, with strict annual support concentrated in two of the five persistent receiving destinations. System B and the stress analyses assess sensitivity to alternative network representations and capacity assumptions.

2 Occupational Transition Networks and Research Gap

2.1 Exposure is upstream of transition

Capability-based, task-based, and use-informed AI exposure measures differ conceptually (Felten et al., 2021; Eloundou et al., 2024; Tomlinson et al., 2025; Massenkoff and McCrory, 2026; OECD, 2026). Multi-model replication shows that exposure scores can vary materially even under matched rubrics and that differences can propagate into downstream regression results (Yin et al., 2026). BLS combines five sources of theoretical and current evidence to classify detailed occupations into four relative exposure categories, while explicitly documenting sensitivity to source choice, crosswalks, imputation, normalization, and categorization (U.S. Bureau of Labor Statistics, 2026b). For the present paper, exposure is therefore an

origin-screening input, not a displacement probability.

This separation matters because adoption and outcomes add constraints absent from exposure. Lindenlaub et al. show a large gap between exposure and observed AI adoption when comparative advantage is ignored (Lindenlaub et al., 2026); Bick et al. similarly find widespread but shallow generative AI use and substantial within-occupation heterogeneity (Bick et al., 2026). Hyman et al. study actual training spells and post-training earnings (Hyman et al., 2025); Manning and Aguirre study adaptive capacity conditional on displacement (Manning and Aguirre, 2026), while cross-country evidence links exposure and occupational mobility without collapsing the two constructs (Cazzaniga et al., 2024). Our estimand is upstream and narrower: whether declared occupational routes and destination openings can accommodate a simultaneous origin set.

2.2 Mobility is a network object

Occupation-to-occupation movement is structured by skill relatedness, labor-market adjacency, search frictions, and network position (Schmutte, 2014; Bechichi et al., 2019; del Rio-Chanona et al., 2021; Moro et al., 2021). Applied network studies map occupation spaces from observed mobility and skill similarity, identify retraining pathways, and document network constraints on worker mobility (Joyez and Laffineur, 2022; Waters and Shutters, 2022; Frank et al., 2024; Dawson et al., 2021). Workforce-technology studies further use O*NET/BLS evidence to characterize adjacent occupations and reskilling opportunities (Christenko, 2022; Chang and Seo, 2026; Tekic et al., 2026). These studies establish the importance of edge construction and network position, but they generally do not impose a single destination-capacity budget shared by multiple origins. The present analysis addresses that aggregation constraint.

Related work models capacity scarcity in retraining pipelines and technology adoption (Levy Yeyati, 2026; Emadi, 2026), while firm- and worker-level studies document diffusion, hiring, task reallocation, heterogeneous labor-market responses, and mobility constraints (Bonney et al., 2026; Hampole et al., 2025; Humlum and Vestergaard, 2025; Tucker, 2026; Song

et al., 2026). Here, projected destination openings form the shared capacity resource. The resulting problem is a capacitated network-allocation problem in which route admissibility and receiving capacity are evaluated jointly, and optimal-face analysis distinguishes persistent destination use from substitutable allocations.

2.3 Observed mobility remains a different estimand

Maximum flow characterizes feasible allocation under a specified graph and capacity vector. Observed worker transitions additionally reflect geography, preferences, hiring, credentials, firm demand, timing, training access, family constraints, search, and macroeconomic conditions. SIPP is therefore used as a downstream concordance layer: it assesses whether observed cross-major changes from compatible origins align with declared routes and whether persistent receiving occupations appear in contemporary worker flows.

3 Data and Research Design

3.1 Common occupational frame and upstream screen

The common analytical frame contains 661 U.S. occupations defined using a deterministic taxonomy and common-support inclusion criteria. Upstream inputs include AI Occupational Exposure (AIOE), two Microsoft applicability channels, O*NET occupational requirements, and BLS employment/projection fields (Felten et al., 2021; Tomlinson et al., 2025; U.S. Department of Labor, Employment and Training Administration, 2026; U.S. Bureau of Labor Statistics, 2025). The present analysis takes the upstream exposure-requirement construction as given rather than re-estimating it. The baseline design uses AIOE-cardinal exposure with equal weight across the 12 declared O*NET requirement descriptors. Here, baseline-positive means that an occupation has a positive exposure-minus-requirement contrast under that fixed baseline design. That design yields 372 positive occupations, of which 17 remain positive across the 30-design sensitivity family described in Additional file 1. These counts describe

sensitivity of the admitted positive-origin set, not a separate exposure result.

The analysis retains BLS 2024–34 as the primary projection vintage. BLS released 2025–35 projections on August 27, 2026 (U.S. Bureau of Labor Statistics, 2026a). The later-vintage stability analysis holds the measurement design, Standard Occupational Classification (SOC) mapping, route ordering, cross-major screen, Top-5/no-backfill rule, and opening-access stress ρ fixed while replacing the complete model-relevant BLS occupational-transfer and opening vector with 2025–35 values.

3.2 System A: ranked transition evidence

System A uses the public DOL/Emsi career-pathways transition data (U.S. Department of Labor, Office of the Assistant Secretary for Policy, 2022; Schwartz et al., 2021). After deterministic SOC normalization and the primary cross-major screen, 408 of 661 occupations have ranked route support, representing 68.19% of analytical-frame employment. Unsupported occupations remain part of the frame but cannot contribute to route-specific claims.

Candidate-route definition is itself an estimand choice. The released/no-backfill rule takes the Top-k routes released by the source, where k is the retained route count, and then applies analytical-frame and cross-major filters. The conditioned/no-backfill rule first filters for cross-major eligibility, retains the Top-k ranked destinations, and then applies analytical-frame joinability without replacement. The fully joinable/backfill rule searches deeper in the ranked route list until k joinable destinations are available. These alternatives are reported separately rather than averaged.

System A defines its primary supply denominator over baseline-positive occupations with any joinable cross-major DOL/Emsi support. There are 169 such origins in this denominator, carrying 3,096.0 thousand projected transfers per year in the 2024–34 vintage. The released Top-5/no-backfill graph retains 335 directed edges from 146 of those origins to 71 destinations; 145 of the 146 edge-bearing origins have positive 2024–34 transfer supply, while one has zero. The remaining 23 denominator origins have no retained released/no-backfill Top-5 edge and

therefore contribute zero routed flow rather than disappearing from the denominator.

For denominator origin i, T_i denotes BLS projected annual occupational transfers. For destination j, O_j denotes projected annual openings and $C_j(\rho) = \rho O_j$, with $\rho \in [0, 1]$, is the assumed accessible-capacity proxy. Projected transfers and openings do not constitute a bilateral BLS origin–destination accounting identity: the former are treated as modeled origin-supply pressure and the latter as a destination capacity ceiling for the declared structural allocation. The primary result uses gross cross-major routes, $k = 5$, and $\rho = 0.50$. This denominator convention is held unchanged in the 2025–35 fixed-design later-vintage stability check.

3.3 System B: O*NET related-occupation network

System B is constructed from O*NET Related Occupations rather than the DOL/Emsi route ranking. The candidate ledger contains 10,415 directed relatedness records before declared feasibility filters. The same 372 baseline-positive origins supply 4,763.6 thousand projected occupational transfers per year. Three nested specifications form an interpretable capacity spectrum: strict (513 admitted directed edges; 272 destinations), representative (1,658 edges; 458 destinations), and loose (3,716 edges; 562 destinations). The edge screens combine relatedness tier, preparation/education compatibility, wage retention, and directionality; exact definitions are tabulated in Additional file 1. System-B percentages are reported separately from System-A percentages because the edge source and origin eligibility differ.

3.4 Longitudinal SIPP mobility layer

The primary observed-mobility layer links the 2024 and 2025 SIPP public-use releases to compare December 2023 with December 2024 (U.S. Census Bureau, 2025, 2026). Persons are matched by the public-use person key and panel consistency is verified across releases. At each endpoint, the primary job is the valid active job among jobs 1–6 with the highest reported monthly earnings; ties across different occupational codes are excluded. The primary age

rule requires age 18–64 at both December endpoints. Because occupation codes can change while the same survey job spell remains active, an annual occupation-changing job transition additionally requires a nonmissing job identifier at both endpoints, a different job identifier, and a different public-use occupation code. The strict cross-major subset is reported only where the official mapping identifies different SOC major groups without ambiguity.

The two-year longitudinal weight FINYR2 provides population weights; REPWGT1–REPWGT240 provide the 240 replicate weights used to estimate sampling uncertainty. Variance estimates use the Census Fay balanced repeated replication (Fay-BRR) construction with coefficient 0.5. Survey-weighted rates are descriptive population estimates for the declared longitudinal event and are reported separately from BLS structural flows measured in thousands of workers per projected year.

Public-use occupation codes can correspond to multiple detailed SOCs. Official Census bridges are therefore represented as sets, with no title-similarity disambiguation, equal-share allocation, or forced one-to-one choice. Public-use code 7640 remains unresolved because its listed official SOC constituent differs across official source materials; it is not manually repaired. The 2024-reference single-release successive-change SIPP analysis is reported only as a short-horizon sensitivity in Additional file 1.

Table 1 summarizes the evidence sources, units, and scientific roles of the three analytical layers.

4 Methods

4.1 Joint capacity and the independent-origin ceiling

To isolate the cost of sharing the same receiving opportunities, we first define the maximum modeled supply that can be assigned when all origins must use one common destination-capacity vector. For either transition system, let I and J be the admitted origin and destination sets, let $A \subseteq I \times J$ be the directed route set, let T_i be origin supply, and let C_j be

Table 1: Analytical layers, evidence sources, scales, and roles.

Layer	Edge/evidence source	Scale/capacity	Scientific role
System A	DOL/Emsi ranked transition evidence	BLS projected transfers and destination openings	Primary shared-capacity comparison, route-menu sensitivity, persistent receiving core.
System B	O*NET Related Occupations plus preparation/wage/directional screens	BLS projected transfers and openings	Alternative route semantics, exact all-optima destination certificate, and structural controls.
SIPP layer	Linked Dec-2023 to Dec-2024 public-use worker histories	FINYR2-weighted rates with 240 replicate weights	Observed-mobility triangulation and route support.

destination capacity. In Systems A and B, T_i, C_j, and the flow f_{ij} are measured in thousands of workers per year. The standard capacitated bipartite maximum-flow problem is (Ford and Fulkerson, 1956; Ahuja et al., 1993; Bertsimas and Tsitsiklis, 1997)

$$F^\star = \max_f \sum_{i\in I}\sum_{j\in J} f_{ij} \tag{1}$$

$$\text{s.t.} \quad \sum_{j\in J} f_{ij} \le T_i, \quad \forall i \in I, \tag{2}$$

$$\sum_{i\in I} f_{ij} \le C_j, \quad \forall j \in J, \tag{3}$$

$$f_{ij} \ge 0, \qquad f_{ij} = 0 \quad \text{if } (i,j) \notin A. \tag{4}$$

The optimum $F^\star$ is the maximum jointly assignable modeled supply, and $R_{\text{joint}} = F^\star / \sum_{i\in I} T_i$ is its dimensionless routability ratio. Neither quantity is a behavioral transition probability.

The comparison ceiling asks a deliberately more permissive question: how much could each origin route if it were allowed to reuse the full destination-capacity vector independently of every other origin? For origin i, that ceiling is

$$F_i^{\text{ind}} = \min\left(T_i, \sum_{j:(i,j)\in A} C_j\right), \qquad F_{\text{ind}} = \sum_{i\in I} F_i^{\text{ind}}, \qquad R_{\text{ind}} = \frac{F_{\text{ind}}}{\sum_{i\in I} T_i}. \tag{5}$$

Because the same C_j can be reused separately by different origins in this calculation, F_{ind} is an intentionally non-joint ceiling. The competition gap is

$$\Delta_{\text{comp}} = R_{\text{ind}} - R_{\text{joint}}. \tag{6}$$

A zero gap means that sharing capacity does not reduce the best routable fraction under the declared graph and capacities; a positive gap is the additional structural loss attributable to joint competition under the declared graph and capacities. Total deficiency then decomposes exactly as

$$1 - R_{\text{joint}} = (1 - R_{\text{ind}}) + \Delta_{\text{comp}}. \tag{7}$$

The competition share of deficiency is $\Delta_{\text{comp}}/(1 - R_{\text{joint}})$ when $R_{\text{joint}} < 1$.

4.2 Destination certificates over the full optimal face

A maximum-flow value need not imply a unique allocation. Here $F^\star$ denotes only the scalar optimal total-flow value from the maximum-flow problem above, while $\mathcal{F}^\star$ denotes the set of feasible allocations attaining that value:

$$\mathcal{F}^\star = \left\{ f : \ f \text{ satisfies the flow constraints and } \sum_i \sum_j f_{ij} = F^\star \right\}. \tag{8}$$

For each destination,

$$m_j = \min_{f \in \mathcal{F}^\star} \sum_{i \in I} f_{ij}, \qquad M_j = \max_{f \in \mathcal{F}^\star} \sum_{i \in I} f_{ij}. \tag{9}$$

Numerical classifications use a tolerance of 10^{-6} in units of thousands of workers per year, applied uniformly across the numerical calculations. A destination is necessarily used if $m_j > 10^{-6}$, possibly used if $M_j > 10^{-6}$, saturated in every optimum if $|m_j - C_j| \leq 10^{-6}$, saturable if $|M_j - C_j| \leq 10^{-6}$, and never used when $M_j \leq 10^{-6}$. Each endpoint is found by a separate linear program with total flow fixed at $F^\star$. This avoids interpreting one arbitrary optimum as the system's unique route allocation.

4.3 Origin-side and minimum-cut certificates

The same optimal face can be interrogated from the origin side. For origin i, define

$$s_i^- = \min_{f \in \mathcal{F}^\star} \sum_{j \in J} f_{ij}, \qquad s_i^+ = \max_{f \in \mathcal{F}^\star} \sum_{j \in J} f_{ij}. \tag{10}$$

Using the same separate-origin ceiling F_i^{ind} defined above,

$$L_i^{\mathrm{comp}} = \max\Big(0, F_i^{\mathrm{ind}} - s_i^+\Big) \tag{11}$$

is the minimum competition-induced loss that origin i must bear across all joint optima. If $L_i^{\mathrm{comp}} = 0$ but $s_i^- < F_i^{\mathrm{ind}}$, competition can affect that origin in some optima but is not unavoidable for it.

To localize aggregate capacity scarcity, we also use the max-flow/min-cut equivalence. Let V denote the network node set, with s and t the source and sink. After computing a maximum flow, let $\mathcal{R}$ be the nodes reachable from s in the positive-residual graph and let $\mathcal{Q}$ be the nodes that can reach t in that residual graph. Then $\mathcal{R}$ is the smallest source side and $V \setminus \mathcal{Q}$ the largest source side among all minimum cuts. A destination-capacity arc $d \to t$ is therefore present in every minimum cut if $d \in \mathcal{R}$, present in some but not all minimum cuts if $d \in (V \setminus \mathcal{Q}) \setminus \mathcal{R}$, and absent from every minimum cut if $d \in \mathcal{Q}$. Both canonical cut capacities are verified to equal $F^\star$ within the same 10^{-6} model-flow tolerance. This classification distinguishes a destination that is necessarily used in every maximum-flow allocation from a capacity arc that is itself structurally responsible for limiting additional flow.

4.4 Route-definition and representation stresses

For System A, the 30-design sensitivity family provides the joint-routability and receiving-core sensitivity results. The separate-versus-joint competition-gap analysis across route menus

and opening-access stresses is restricted to the fixed baseline exposure-origin design because separate-origin ceilings cannot be reconstructed for every upstream design from the available row-level records. For System B, strict/representative/loose networks are supplemented by point-nonworsening, full-preparation, entry-wage, and half-opening-capacity stresses. Network representation is varied as directed, reciprocal-only, and symmetrized union. Randomized controls are reported in Additional file 1 as structural diagnostics.

4.5 Continuous accessibility frontier and planning-stress certificate

The declared opening-access grid treats ρ as a deterministic stress parameter. To test whether the four displayed grid values hide relevant structure, we also evaluate the System-A problem continuously over $0 \leq \rho \leq 1$ for each route menu and both BLS projection vintages. Destination capacities remain $C_j(\rho) = \rho O_j$; no worker-level probability, vacancy-reservation probability, or behavioral interpretation is introduced.

For a fixed graph, the joint maximum-flow value can be written as the lower envelope of affine minimum-cut functions in ρ. The independent-origin ceiling is the sum of originwise piecewise-linear ceilings. The competition gap $\Delta_{\text{comp}}(\rho) = R_{\text{ind}}(\rho) - R_{\text{joint}}(\rho)$ is therefore evaluated on the combined set of exact breakpoints rather than by interpolating between the displayed grid values. When a nested accessibility support is declared, the same frontier yields interval and possibilistic planning-stress certificates by taking minima and maxima over the corresponding ρ cuts. Those certificates remain structural planning quantities rather than probabilities.

4.6 Set-valued longitudinal route concordance and survey uncertainty

For a strict annual SIPP event e, let S_e and D_e be the detailed-SOC sets compatible with its public-use source and destination codes after intersection with the 661-frame endpoints;

Census classification-change guidance is incorporated into the mapping ledger (U.S. Census Bureau, 2020). Let O_A be the fixed 169-origin System-A denominator and A_A the prespecified released Top-5/no-backfill edge set. Define the origin-admitted pair set

$$P_e^A = (S_e \cap O_A) \times D_e. \tag{12}$$

Origin admission is *necessary* when S_e is nonempty and $S_e \subseteq O_A$, *possible* when $S_e \cap O_A \neq \varnothing$, and *exact* when S_e is a singleton inside O_A. Route compatibility is possible when $P_e^A \cap A_A \neq \varnothing$ and necessary, conditional on origin admission, when P_e^A is nonempty and $P_e^A \subseteq A_A$. The lower endpoint resolves mapping ambiguity against concordance: it credits a route only when all admitted pairs are routes and includes denominator-only admissions whenever a nonroute admitted mapping is available. The upper endpoint resolves ambiguity in favor of concordance: it credits any event with at least one route-compatible pair and excludes optional denominator-only admissions when origin membership can be resolved outside O_A. The resulting interval is a sharp crosswalk-identification bound for the declared mapping sets, not a sampling confidence interval.

Survey sampling uncertainty is evaluated separately. FINYR2 supplies the reported point estimates; REPWGT0 is its full-sample counterpart in the replicate-weight file, and REPWGT1–REPWGT240 are the 240 Fay-BRR replicates with Fay coefficient 0.5. The mapping bound and the replicate-weight confidence interval therefore address different uncertainty sources and are reported separately. The single-release successive-change System-B mapping bounds use the 372-origin O*NET design as a supplementary sensitivity; no longitudinal System-B bound is estimated because the corresponding row-level 513/1,658/3,716 edge ledgers are not available for a design-identical longitudinal calculation.

5 Results

5.1 Separate-origin plans overstate joint route opportunity across BLS vintages

Figure 2 and Table 2 show the 2024–34 reference result and its fixed-design later-vintage stability check. Under the 2024–34 baseline gross cross-major released Top-5 specification with $\rho = 0.50$, $R_{\text{ind}} = 83.65\%$ but $R_{\text{joint}} = 49.40\%$. The support-conditioned competition gap is therefore 34.25 percentage points. Total joint deficiency is 50.60 percentage points: 16.35 points reflect separate-origin route/capacity scarcity and 34.25 points reflect joint competition. Thus shared competition accounts for 67.69% of deficiency.

Replacing only the BLS supply/opening vector with 2025–35 values while holding the baseline positive-origin screen and released Top-5/no-backfill DOL/Emsi graph fixed reduces modeled supply from 3,096.0 to 2,870.2 thousand projected transfers per year. The resulting values are $R_{\text{ind}} = 82.72\%$, $R_{\text{joint}} = 49.74\%$, and a 32.97-percentage-point competition gap; shared competition accounts for 65.61% of joint deficiency. Relative to 2024–34, separate-origin routability decreases by 0.93 points and joint routability increases by 0.35 points despite a 7.3% reduction in modeled supply.

Table 2: System-A later-vintage stability check at $\rho = 0.50$. Supply is in thousands of projected occupational transfers per year; gap is computed from full-precision ratios before display rounding.

Vintage	Supply	Separate	Joint	Gap (pp)
2024–34	3096.0	83.65%	49.40%	34.25
2025–35	2870.2	82.72%	49.74%	32.97

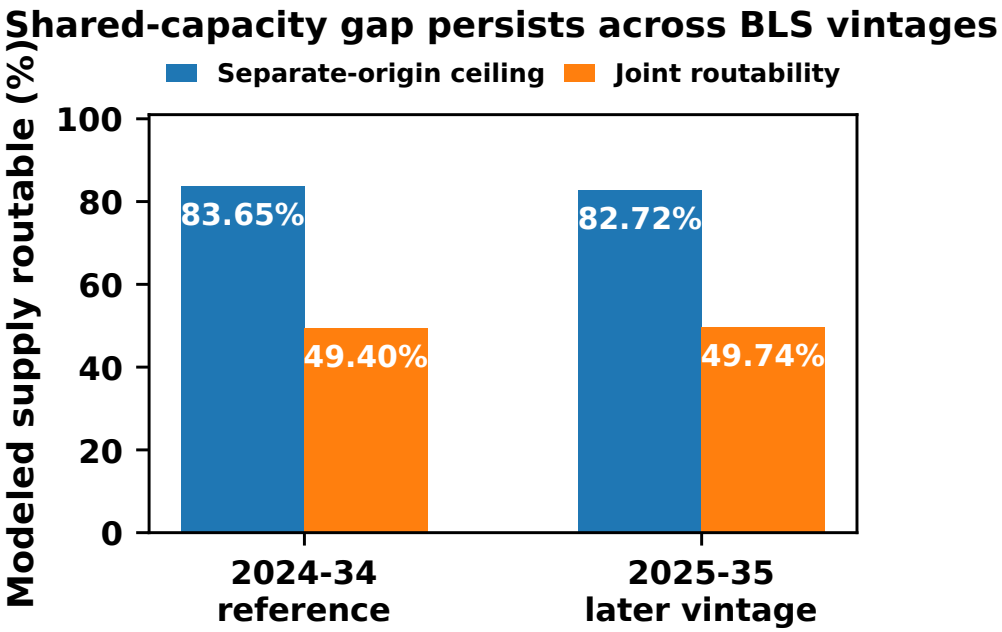


Figure 2: 2024–34 System-A reference comparison and later-vintage 2025–35 fixed-design stability check. The measurement design and DOL/Emsi released Top-5/no-backfill graph are fixed; the later result replaces the complete model-relevant BLS transfer/opening vector. Percentages are support-conditioned structural routing ratios.

5.2 The primary System-A estimand uses a declared support-conditioned denominator

The 661 occupations define the analytical frame, not the denominator of the primary flow ratio. DOL/Emsi contains at least one joinable cross-major route for 408 frame occupations. Intersecting that evidence with the 372 baseline-positive occupations yields the 169-origin System-A denominator. The released Top-5/no-backfill graph then retains at least one edge for 146 denominator origins; the other 23 remain in the denominator with zero routed flow. Figure 3 separates these nested scopes so the reported routability ratios are interpreted relative to the 169-origin denominator rather than the full analytical frame.

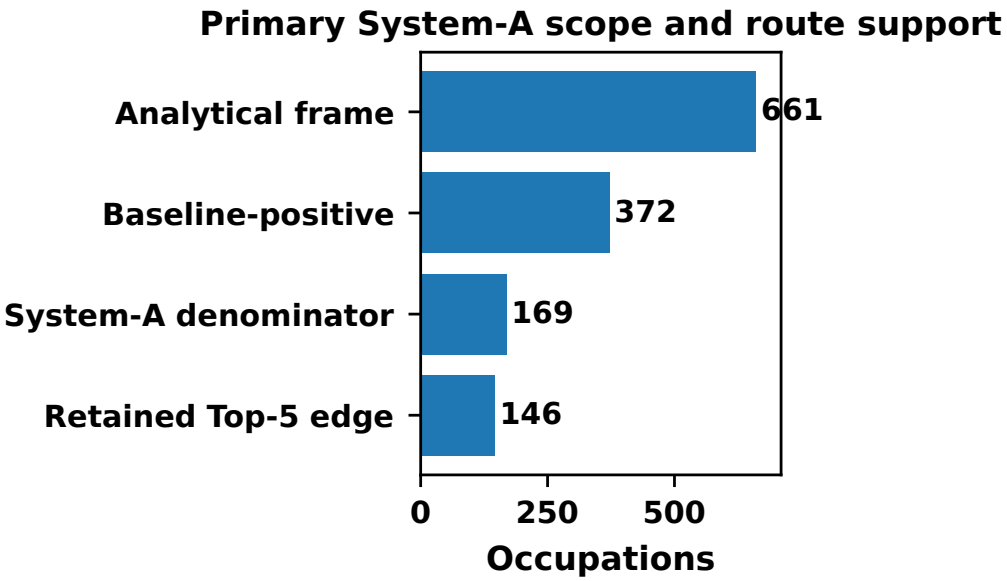


Figure 3: Nested scopes for the primary System-A estimand. The 169-origin denominator is conditioned on any joinable cross-major DOL/Emsi support; the released Top-5/no-backfill graph retains an edge for 146 of those origins, while 23 remain in the denominator with zero routed flow.

5.3 Route construction and opening access determine gap magnitude

The 30-design sensitivity analysis shows that upstream measurement choices change absolute joint routability across the declared design family. The full ranked DOL/Emsi route list was used to construct the three route definitions. The released/no-backfill route set contains exactly the prespecified 335 edges.

For the fixed baseline exposure-origin design at the declared primary $\rho = 0.50$, the competition gap remains positive and exceeds 23 percentage points under all three route constructions (Table 3). In 2024–34, the competition gap is 34.25, 33.93, and 23.52 percentage points under released/no-backfill, conditioned/no-backfill, and fully joinable/backfill, respectively. The 2025–35 vector yields 32.97, 32.87, and 23.34 percentage points in the same order. Thus the magnitude depends on route semantics, while the competition gap remains above 23 points across all three route constructions at the primary opening-access stress in both vintages.

Table 3: System-A competition-gap robustness across route constructions at $\rho = 0.50$. Percentages use the same fixed 169-origin denominator within each vintage. Gaps are computed from full-precision ratios before display rounding.

Vintage	Route construction	Separate	Joint	Gap (pp)
2024–34	Released/no-backfill	83.65%	49.40%	34.25
	Conditioned/no-backfill	96.83%	62.90%	33.93
	Fully joinable/backfill	100.00%	76.48%	23.52
2025–35	Released/no-backfill	82.72%	49.74%	32.97
	Conditioned/no-backfill	96.21%	63.34%	32.87
	Fully joinable/backfill	100.00%	76.66%	23.34

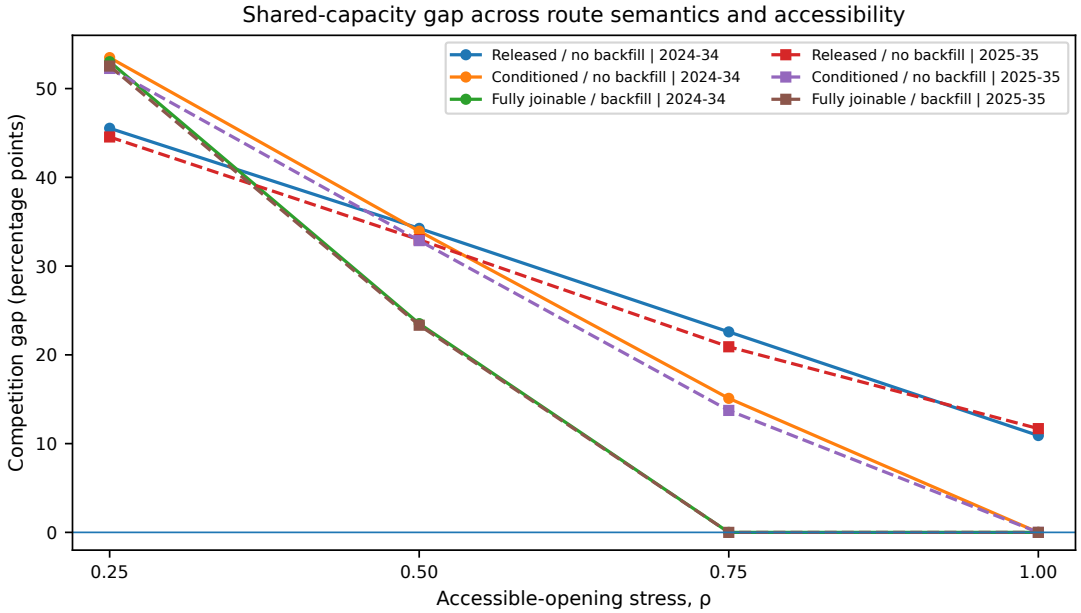


Figure 4: Competition gap across route definitions and the declared opening-access stress. For the fixed baseline exposure-origin design, the primary $\rho = 0.50$ gap remains positive and exceeds 23 percentage points under all three route menus, but the declared ρ grid also shows that more permissive route/capacity combinations can eliminate the gap. Solid and dashed traces denote the 2024–34 and 2025–35 BLS vintages, respectively.

Figure 4 summarizes the opening-access sensitivity profile across the declared route menus. With more accessible capacity, the gap can decrease to zero: the fully joinable/backfill model reaches 100% joint routability by $\rho = 0.75$ in both vintages, and conditioned/no-backfill reaches 100% by $\rho = 1.00$. The competition gap is therefore conditional on route semantics and accessible capacity, while remaining positive under all declared route menus at the primary $\rho = 0.50$.

5.4 Continuous accessibility frontier and support-wide bottlenecks

The declared four-point grid is a sensitivity display, not a complete description of the function $\Delta_{\text{comp}}(\rho)$. Exact parametric evaluation over $0 \leq \rho \leq 1$ shows that the gap is piecewise linear and nonmonotone. In every System-A route menu and both BLS vintages, the gap first increases, reaches an interior maximum, and then declines as destination capacity becomes sufficiently permissive.

For the primary 2024–34 released/no-backfill case, the maximum gap is 45.93 percentage points at $\rho = 0.234068$ (Figure 5). The previously reported $\rho = 0.50$ result remains 34.25 points. The accessibility stress required to reach 50% joint routability is $\rho = 0.508449$, while 75% joint routability requires $\rho = 0.876192$; 90% joint routability is not reached in released/no-backfill for $\rho \leq 1$. The corresponding 2025–35 peak is 45.06 points at $\rho = 0.219983$.

The continuous audit also clarifies the route-menu comparison. The conditioned/no-backfill and fully joinable/backfill menus have larger interior peak gaps, 56.89 and 63.78 points in 2024–34, but they also reach high joint routability at lower accessibility because their route sets are more permissive. Thus the magnitude of the shared-capacity penalty is not ordered by a single grid point; it depends on where along the accessibility frontier the system is evaluated.

A declared nested accessibility support can be placed on the same frontier. Under the planning-stress coordinate $\tilde{\rho} = (0.25, 0.50, 1.00)$ in the primary released/no-backfill case, the event $\Delta_{\text{comp}} \geq 20$ percentage points has possibility 1.00 and necessity 0.614; the event $\Delta_{\text{comp}} \geq 30$ points has possibility 1.00 and necessity 0.187; and the event $\Delta_{\text{comp}} \geq 40$ points has possibility 0.512 and necessity 0.000. These are declared planning-stress certificates, not vacancy-access probabilities.

Across the full declared support $\rho \in [0.25, 1.00]$ for the released/no-backfill graph, seven destination-capacity arcs are present in every minimum cut in both BLS vintages: Chief Executives, Sales Managers, Health and Safety Engineers except Mining Safety Engineers and Inspectors, Travel Agents, First-Line Supervisors of Office and Administrative Support

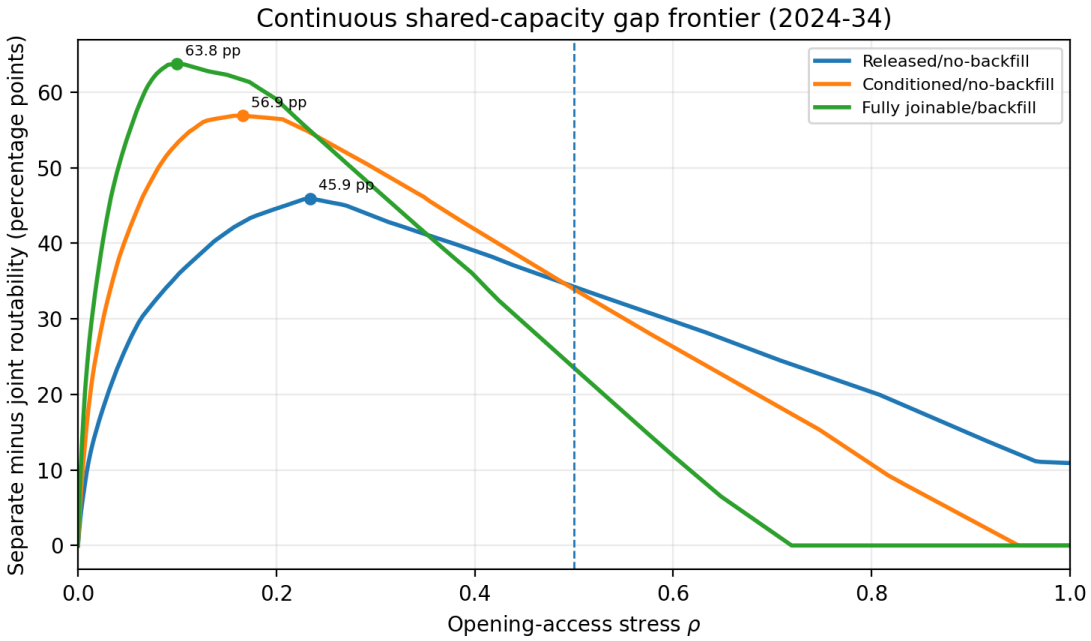


Figure 5: Continuous System-A accessibility frontier for the 2024–34 BLS vintage. Curves trace the exact separate-origin ceiling, joint routability, and competition gap over $0 \leq \rho \leq 1$ for the three route menus. The primary manuscript point $\rho = 0.50$ remains a declared stress; the continuous frontier shows that the maximum shared-capacity penalty can occur at lower accessibility.

Workers, Legal Secretaries and Administrative Assistants, and Secretaries and Administrative Assistants except Legal, Medical, and Executive. Three of these – Sales Managers, Office/Administrative Support Supervisors, and Secretaries/Administrative Assistants – are also part of the five-occupation receiving-core result. Their bottleneck status is therefore not an artifact of selecting $\rho = 0.50$ alone.

5.5 Origin-side allocation sensitivity

The origin-side optimal face characterizes origin-level service variation under shared capacity. In the 2024–34 primary graph at $\rho = 0.50$, 23 denominator origins have no retained route capacity and one has zero projected transfer supply. Among the remaining origins, 68 are competition-neutral and 77 are allocation-sensitive: their service can fall below the separate-origin ceiling in some joint optima. However, no origin has positive unavoidable competition

loss, because every origin's maximum service over the joint optimal face reaches its own separate-origin ceiling. The same result holds in 2025–35, with 70 competition-neutral and 75 allocation-sensitive origins.

The aggregate competition gap coexists with zero unavoidable origin-specific competition loss. The shortfall can be redistributed across alternative maximum-flow allocations and is therefore not structurally pinned to a unique set of origin occupations.

5.6 Where does the limiting capacity reside?

The minimum-cut analysis distinguishes destinations that are necessary receivers from capacity arcs that actually limit additional system flow. At $\rho = 0.50$, 18 destination-capacity arcs are present in every minimum cut in 2024–34. The later-vintage graph has 16 such mandatory capacity arcs, all 16 of which are retained from the reference set (Jaccard similarity $16/18 = 0.889$). Aggregate persistence therefore coexists with a small change in bottleneck composition.

Three of the five cross-design receiving-core destinations are also mandatory capacity bottlenecks in both vintages: Sales Managers (11-2022), First-Line Supervisors of Office and Administrative Support Workers (43-1011), and Secretaries and Administrative Assistants, Except Legal, Medical, and Executive (43-6014). Bookkeeping, Accounting, and Auditing Clerks and Hotel, Motel, and Resort Desk Clerks remain necessary receivers but are not mandatory minimum-cut capacity arcs. Necessary destination use and mandatory minimum-cut capacity status are therefore distinct properties.

5.7 A small receiving core is necessary across the primary System-A family

Five destinations are used in every specification of the 30-design sensitivity family and every maximum-flow allocation under the primary gross cross-major Top-5 case at $\rho = 0.50$. Two

Table 4: Persistent System-A receiving core in the 30-design sensitivity analysis.

SOC	Receiving occupation	Necessary designs	Saturated designs
43-6014	Secretaries and Administrative Assistants, Except Legal, Medical, and Executive	30/30	30/30
43-1011	First-Line Supervisors of Office and Administrative Support Workers	30/30	30/30
11-2022	Sales Managers	30/30	25/30
43-3031	Bookkeeping, Accounting, and Auditing Clerks	30/30	0/30
43-4081	Hotel, Motel, and Resort Desk Clerks	30/30	0/30

are saturated in all 30 designs (Table 4). These destinations are priorities for additional evidence on geography, credentials, training access, and employer demand.

The number of necessary destinations is 46 in each BLS vintage, with 45 destination identities shared across vintages. All five destinations in Table 4 remain necessary under the 2025–35 baseline optimal face; Secretaries and Administrative Assistants, Office/Administrative Support Supervisors, and Sales Managers are saturated in every 2025–35 baseline optimum. This is a fixed-baseline temporal certificate; the 30-design family was not re-evaluated under the new BLS vintage.

5.8 System B evaluates sensitivity under alternative route evidence

System B asks the same general question with different edge evidence. The System-B capacity spectrum uses the full-opening baseline ($\rho = 1.00$) and is reported separately from the primary System-A result at $\rho = 0.50$ because the systems use different route evidence and origin eligibility. With modeled supply of 4,763.6 thousand projected transfers per year, the strict network routes 2,704.8 thousand (56.78%), the representative network 4,199.6 thousand (88.16%), and the loose network 4,626.5 thousand (97.12%) (Figure 6). This spread shows that relatedness semantics and feasibility screens materially change the structural ceiling.

The representative model admits many maximum-flow allocations: 51 destinations are saturated in every optimum, while many others are substitutable. Additional feasibility and capacity stresses change the representative 88.16% ceiling substantially; the all-optima counts,

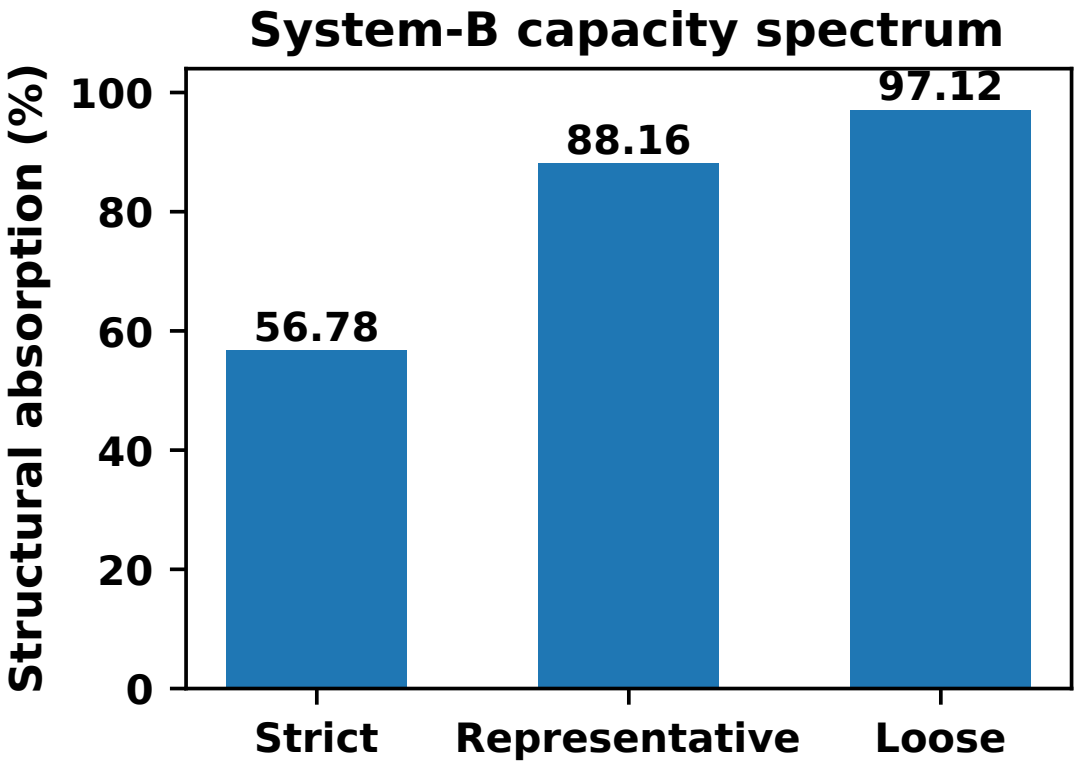


Figure 6: System-B capacity spectrum under strict, representative, and loose O*NET-related-occupation specifications. Values use projected annual transfer/opening units.

stress values, and randomized network diagnostics are reported in Additional file 1. System B quantifies sensitivity to an alternative route representation.

5.9 Longitudinal worker mobility is narrower than structural route feasibility

The age-18–64 longitudinal sample contains 6,556 persons with valid primary occupations at both endpoints and 5,867 with nonmissing job identifiers at both endpoints. The raw occupation-code change rate is 58.33% unweighted and 59.56% FINYR2-weighted, but this diagnostic is not interpreted as worker mobility because many code changes occur while the job identifier is unchanged. Requiring both a different job identifier and a different occupation code reduces the weighted annual rate to 12.19% (95% Fay-BRR interval 11.19–13.19%). The weighted rate for strict, definitely cross-major transitions is 9.72% (8.83–10.61%). Table 5

Table 5: Primary longitudinal SIPP mobility estimates, Dec-2023 to Dec-2024, ages 18–64 at both endpoints. Replicate-weight intervals use Fay-BRR and do not incorporate mapping ambiguity.

Event definition	Raw n	FINYR2 weighted	95% interval
Different job identifier	905/5867	16.66%	15.39–17.93%
Different job identifier + different occupation code	664/5867	12.19%	11.19–13.19%
Strict definitely cross-major	527/5867	9.72%	8.83–10.61%

reports the principal estimates.

Among the 527 strict annual cross-major events, 349 have nonempty candidate SOC sets at both endpoints inside the 661 frame. Relative to the prespecified System-A graph, 146 events are necessarily associated with the 169-origin denominator, 190 are possibly associated, and 111 have exact singleton origin mappings. Fourteen events can match at least one prespecified released System-A route under a compatible mapping. The resulting FINYR2-weighted origin-conditioned route-concordance bound is 3.97–8.10%; the exact-origin bound is 5.99–9.07%. These are crosswalk-identification bounds, not worker-transition probabilities. Conditional Fay-BRR intervals for the fixed lower and upper endpoint rules are reported in Additional file 1.

The five persistent System-A receiving-core destinations receive 11 strict annual cross-major events in total, accounting for 2.13% of strict cross-major FINYR2 mass. Seven events enter Office/Administrative Support Supervisors and four enter Bookkeeping, Accounting, and Auditing Clerks; Sales Managers, Hotel/Motel/Resort Desk Clerks, and Secretaries/Administrative Assistants receive no strict annual cross-major events in this sample. Strict annual support appears for two of the five destinations, whereas the single-release successive-change sensitivity observed all five. The two designs are reported separately.

Because the structural and longitudinal analyses use different estimands, denominators, and units, their percentages are not combined into a model-error statistic. No longitudinal System-B concordance is estimated because the required row-level graphs are not available for a design-identical calculation.

6 Discussion and Interpretation

6.1 Interpretation of the shared-capacity gap

The independent-origin ceiling permits each origin to reuse the full destination-capacity vector, whereas the joint model applies one common capacity vector across all origins. Their difference therefore measures the reduction in routability attributable to shared destination capacity under an otherwise identical network specification.

For the fixed baseline exposure-origin design, the gap remains positive across the three route constructions at $\rho = 0.50$ but can reach zero under more permissive route and opening-access combinations. The computed gap is therefore conditional on the specified network and capacity assumptions.

The origin-side certificate shows that the aggregate competition loss is not uniquely attributable to a fixed set of origins because service can shift across maximum-flow optima. Joint routability is a structural network quantity, and ρ is a deterministic access stress rather than an estimated behavioral parameter.

6.2 Sensitivity across two route systems

System A and System B are analyzed separately because they use different edge evidence and origin eligibility. System A evaluates ranked transition routes under shared projected capacity, whereas System B applies explicit feasibility screens to an O*NET related-occupation graph. Differences between their results quantify sensitivity to route representation rather than validation of one system by the other.

The formulation applies to any capacitated bipartite origin–destination network for which origin supply, destination capacity, and admissible directed routes are defined.

Table 6: Condensed interpretation boundaries for the main reported objects.

Output	Permitted interpretation	Not supported
Joint routability	Maximum assignable modeled supply under declared shared capacities	Probability that a worker changes occupation
Competition gap	Additional structural loss created by sharing destination capacity	Causal loss due to technology exposure
Minimum-cut bottleneck	Capacity arc limiting additional model flow across all minimum cuts	Automatic recommendation to create jobs in that occupation
Longitudinal SIPP concordance	Observed-route compatibility under set-valued mapping and survey uncertainty	Model validation, BLS-capacity validation, or technology-exposure causality

6.3 Implications for workforce planning and technology governance

For workforce agencies, training systems, employers, and regional planners, occupation-specific pathway lists should be evaluated jointly against a common destination-capacity budget before they are combined into a portfolio. The assessment should report the origin definition, route semantics, opening-access stress, competition gap, and structural bottlenecks.

The bottleneck certificate identifies destinations where additional evidence on geography, credentials, training availability, employer demand, wages, and hiring standards would refine the capacity assumptions. The origin-side certificate shows that system scarcity does not determine a unique set of affected origins. Table 6 summarizes the interpretation boundaries.

7 Limitations and Claim Boundaries

The analysis is national rather than local, whereas labor markets are geographically segmented. BLS openings and origin transfers are occupation-level annual projections, and the DOL/Emsi and O*NET route sources provide occupation-level adjacency evidence. Geography, employer-specific demand, credentials, training access, and worker heterogeneity are therefore not represented directly in the network capacities or edge definitions.

The SIPP analysis compares two December endpoints and depends on public-use occupa-

tion aggregation. Mappings from public-use occupation codes to detailed SOCs are set-valued, small receiving-destination cells limit occupation-specific inference, and the one-year design does not observe intermediate jobs or training. Fay-BRR replicate weights and set-valued mapping bounds quantify sampling and classification uncertainty separately.

The route-menu and 2025–35 analyses are conditional robustness checks. The competition gap can disappear under sufficiently permissive route and opening-access assumptions, and the later-vintage calculation does not re-evaluate the upstream 30-design family. The reported network quantities therefore remain conditional on the specified route and capacity construction.

The continuous accessibility extension does not calibrate an empirical distribution for vacancy access. Its role is to expose the exact structural frontier of the declared network and, where a nested accessibility support is explicitly declared, to report planning-stress certificates. The triangular support $\widetilde{\rho} = (0.25, 0.50, 1.00)$ reuses the paper's lower stress, primary stress, and full-access bound; it is not a probability model for job openings, hiring, worker eligibility, training seats, or technology-induced movement. A future policy application could replace this declared support with externally elicited or region-specific accessibility information.

8 Reproducibility

The reproducibility materials identify the DOL/Emsi route source (256,831 rows, 28 variables) by its recorded SHA-256 hash and acquisition rule and provide the fixed 169-origin denominator, the BLS 2024–34 and 2025–35 model-relevant vectors, and the prespecified 335-edge released/no-backfill baseline. The route-construction script reproduces that baseline, constructs the conditioned/no-backfill and fully joinable/backfill menus, and recomputes the separate-origin and joint-flow results over the declared ρ grid. A full route reconstruction requires reacquiring the public DOL/Emsi file and matching the recorded hash.

The reproducibility materials contain the optimal-face and minimum-cut calculations for

both BLS vintages, machine-readable origin and destination certificates, aggregate longitudinal SIPP outputs, analysis code, and integrity manifests. Raw SIPP microdata, person-level analytical extracts, event-level person records, the upstream 30-design measurement matrix, and row-level System-B graphs are not included.

9 Conclusion

Evaluating occupational-transition routes jointly under shared destination capacity produces a different network quantity from origin-by-origin route feasibility. Across both BLS vintages, the competition gap remains positive under all three route constructions at the primary opening-access stress but can disappear under more permissive route and capacity assumptions. The continuous accessibility frontier further shows that the gap is nonmonotone and can peak below the primary $\rho = 0.50$ stress. Optimal-face analysis shows that the aggregate shortfall is not uniquely assigned to a fixed origin subset, while minimum-cut analysis identifies persistent destination-capacity constraints.

The longitudinal SIPP analysis provides a separate observed-mobility comparison and shows sparse route concordance relative to structural network feasibility. Taken together, the results support shared-capacity accounting when multiple occupational origins draw on common receiving destinations and identify a system-level network constraint rather than an individual transition forecast.

Additional files

Additional file 1. PDF. *Supplementary Information.* Detailed unit register; System-A route-menu and opening-access grid; origin-side optimal-face and minimum-cut certificates; System-B specifications and structural controls; longitudinal SIPP uncertainty and route-concordance details; later-vintage stability checks; and interpretation and reproducibility boundaries.

Additional file 2. ZIP. *Reproducibility materials.* Code, analytical ledgers, machine-readable outputs, integrity manifests, and aggregate longitudinal SIPP outputs.

Abbreviations

AI, artificial intelligence; AIOE, AI Occupational Exposure; BLS, U.S. Bureau of Labor Statistics; DOL, U.S. Department of Labor; Fay-BRR, Fay balanced repeated replication; O*NET, Occupational Information Network; SIPP, Survey of Income and Program Participation; SOC, Standard Occupational Classification.

Declarations

Availability of data and materials

The datasets and repository materials supporting the reported calculations are available through the public *AI Workforce Shared Capacity* GitHub repository (Malkawi, 2026): `https://github.com/shaikhamalkawi-ux/AI-Workforce-Shared-Capacity`. Raw SIPP microdata and person-level analytical records are not redistributed, and a full System-A route rebuild requires reacquisition of the public DOL/Emsi source and verification against the recorded SHA-256 hash.

Software information: project name, *AI Workforce Shared Capacity*; project home page, `https://github.com/shaikhamalkawi-ux/AI-Workforce-Shared-Capacity`; archived version, `https://github.com/shaikhamalkawi-ux/AI-Workforce-Shared-Capacity/commit/bb258195eda72e676d251d95c905ffe69dcae811`; version, 1.0.0; operating system, platform-independent; programming language, Python; principal requirements, pandas, NumPy, NetworkX, and SciPy; license, MIT for original repository code, with third-party data remaining subject to their source terms; restrictions on non-academic use, none for original MIT-licensed repository code.

Competing interests

The authors declare that they have no competing interests.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Use of AI tools

OpenAI ChatGPT was used for editorial and consistency assistance during manuscript preparation. All scientific content, analyses, results, interpretations, and conclusions were reviewed and approved by the authors, who take full responsibility for the submitted work.

Authors' contributions

GM: Conceptualization, Methodology, Formal analysis, Investigation, Data curation, Software, Validation, Visualization, Writing – original draft, Writing – review & editing, Project administration. AE: Methodology, Formal analysis. MA: Validation, Formal analysis. BH: Investigation, Validation.

Acknowledgements

Not applicable.

References

Ravindra K. Ahuja, Thomas L. Magnanti, and James B. Orlin. *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, 1993.

Nagui Bechichi, Stéphanie Jamet, Gustave Kenedi, Robert Grundke, and Mariagrazia Squicciarini. Occupational mobility, skills and training needs. Technical Report 70, OECD, 2019. URL https://doi.org/10.1787/30a12738-en.

Dimitris Bertsimas and John N. Tsitsiklis. *Introduction to Linear Optimization*. Athena Scientific, 1997.

Alexander Bick, Adam Blandin, David J. Deming, and Tyler R. Schumacher. What work does generative AI do? Working Paper 35677, National Bureau of Economic Research, 2026. URL https://doi.org/10.3386/w35677.

Kathryn Bonney, Cory Breaux, Emin Dinlersoz, Lucia Foster, John Haltiwanger, and Aditya Pande. The microstructure of AI diffusion: Evidence from firms, business functions, and worker tasks. Working Paper 26-25, U.S. Census Bureau, Center for Economic Studies, 2026. URL https://www.census.gov/library/working-papers/2026/adrm/CES-WP-26-25.html.

Mauro Cazzaniga, Carlo Pizzinelli, Emma J. Rockall, and Marina Mendes Tavares. Exposure to artificial intelligence and occupational mobility: A cross-country analysis. Working Paper 2024/116, International Monetary Fund, 2024. URL https://doi.org/10.5089/9798400278631.001.

Yu-Ling Chang and Beomseok Seo. Using explainable artificial intelligence to envision the future of work and the workforce—from prediction to explanation. *Technology in Society*, 84:103139, 2026. doi: 10.1016/j.techsoc.2025.103139. URL https://doi.org/10.1016/j.techsoc.2025.103139.

Aleksandr Christenko. Automation and occupational mobility: A task and knowledge-based approach. *Technology in Society*, 70:101976, 2022. doi: 10.1016/j.techsoc.2022.101976. URL https://doi.org/10.1016/j.techsoc.2022.101976.

Nikolas Dawson, Mary-Anne Williams, and Marian-Andrei Rizoiu. Skill-driven recommendations for job transition pathways. *PLOS ONE*, 16(8):e0254722, 2021. doi: 10.1371/journal.pone.0254722. URL https://doi.org/10.1371/journal.pone.0254722.

R. Maria del Rio-Chanona, Penny Mealy, Mariano Beguerisse-Diaz, François Lafond, and J. Doyne Farmer. Occupational mobility and automation: A data-driven network model. *Journal of the Royal Society Interface*, 18:20200898, 2021. doi: 10.1098/rsif.2020.0898. URL https://doi.org/10.1098/rsif.2020.0898.

Tyna Eloundou, Sam Manning, Pamela Mishkin, and Daniel Rock. GPTs are GPTs: Labor market impact potential of LLMs. *Science*, 384(6702):1306–1308, 2024. doi: 10.1126/science.adj0998. URL https://doi.org/10.1126/science.adj0998.

Seyed Morteza Emadi. The capacity-adoption paradox: Managing workforce bottlenecks during rapid technology deployment. Working Paper 6654059, SSRN, 2026. URL https://doi.org/10.2139/ssrn.6654059.

Edward W. Felten, Manav Raj, and Robert Seamans. Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, 42(12):2195–2217, 2021. doi: 10.1002/smj.3286. URL https://doi.org/10.1002/smj.3286.

L. R. Ford, Jr. and D. R. Fulkerson. Maximal flow through a network. *Canadian Journal of Mathematics*, 8:399–404, 1956. doi: 10.4153/CJM-1956-045-5. URL https://doi.org/10.4153/CJM-1956-045-5.

Morgan R. Frank, Esteban Moro, Tobin South, Alex Rutherford, Alex Pentland, Bledi Taska, and Iyad Rahwan. Network constraints on worker mobility. *Nature Cities*, 1:94–104, 2024. doi: 10.1038/s44284-023-00009-1. URL https://doi.org/10.1038/s44284-023-00009-1.

Menaka Hampole, Dimitris Papanikolaou, Lawrence D. W. Schmidt, and Bryan Seegmiller. Artificial intelligence and the labor market. Working Paper 33509, National Bureau of Economic Research, 2025. URL https://doi.org/10.3386/w33509. Revised September 2025.

Anders Humlum and Emilie Vestergaard. Still waters, rapid currents: Early labor market transformation under generative AI. Working Paper 33777, National Bureau of Economic Research, 2025. URL https://doi.org/10.3386/w33777. Revised March 2026.

Benjamin G. Hyman, Benjamin Lahey, Karen Ni, and Laura Pilossoph. How retrainable are AI-exposed workers? Working Paper 34174, National Bureau of Economic Research, 2025. URL https://doi.org/10.3386/w34174. Revised August 2026.

Charlie Joyez and Catherine Laffineur. The occupation space: Network structure, centrality and the potential of labor mobility in the french labor market. *Applied Network Science*, 7:16, 2022. doi: 10.1007/s41109-022-00453-3. URL https://doi.org/10.1007/s41109-022-00453-3.

Eduardo Levy Yeyati. Too fast to adjust: Adoption speed and the permanent cost of AI transitions. Working Paper IDB-WP-01822, Inter-American Development Bank, 2026. URL https://doi.org/10.18235/0014025.

Ilse Lindenlaub, Ryungha Oh, Maria Alejandra Rodriguez, and Laura Veldkamp. Beyond exposure: Predicting AI adoption based on comparative advantage. Working Paper 35271, National Bureau of Economic Research, 2026. URL https://doi.org/10.3386/w35271.

Ghassan Malkawi. AI Workforce Shared Capacity: Reproducibility materials. GitHub repository, version 1.0.0, 2026. URL https://github.com/shaikhamalkawi-ux/AI-Workforce-Shared-Capacity. Accessed 7 October 2026.

Sam J. Manning and Tomás Aguirre. How adaptable are American workers to AI-induced

job displacement? Working Paper 34705, National Bureau of Economic Research, 2026. URL https://doi.org/10.3386/w34705.

Maxim Massenkoff and Peter McCrory. Labor market impacts of AI: A new measure and early evidence. Anthropic Economic Research, 2026. URL https://maximmassenkoff.com/. 5 March 2026.

Esteban Moro, Morgan R. Frank, Alex Pentland, Alex Rutherford, Manuel Cebrian, and Iyad Rahwan. Universal resilience patterns in labor markets. *Nature Communications*, 12:1972, 2021. doi: 10.1038/s41467-021-22086-3. URL https://doi.org/10.1038/s41467-021-22086-3.

OECD. The OECD AI exposure measure: Mapping the OECD AI capability indicators to occupations. Technical Report 59, OECD Publishing, 2026. URL https://doi.org/10.1787/f3da0f0a-en.

Ian M. Schmutte. Free to move? A network analytic approach for learning the limits to job mobility. *Labour Economics*, 29:49–61, 2014. doi: 10.1016/j.labeco.2014.05.003. URL https://doi.org/10.1016/j.labeco.2014.05.003.

Deena Schwartz, Andrew Clarkwest, Marissa Hashizume, Tresa Kappil, and Julie Strawn. Building better pathways: An analysis of career trajectories and occupational transitions. Technical report, Abt Associates for the U.S. Department of Labor, 2021. URL https://www.dol.gov/agencies/oasp/evaluation/completedstudies/career-pathways-descriptive-and-analytical-study.

Xi Song, Jennie E. Brand, Sukie Xiuqi Yang, and Michael Lachanski. AI-accelerated occupational decline and the mobility trap. *AEA Papers and Proceedings*, 116:246–250, 2026. doi: 10.1257/pandp.20261004. URL https://doi.org/10.1257/pandp.20261004.

Z. Tekic, Y. Omelchenko, and A. Dunaeva. Generative AI and the uneven disruption of freelance labor: Exposure, diffusion, and adaptation under constraint. *Technology in*

Society, 88:103464, 2026. doi: 10.1016/j.techsoc.2026.103464. URL https://doi.org/10.1016/j.techsoc.2026.103464.

Kiran Tomlinson, Sonia Jaffe, Will Wang, Scott Counts, and Siddharth Suri. Working with AI: Measuring the applicability of generative AI to occupations. arXiv:2507.07935, 2025. URL https://doi.org/10.48550/arXiv.2507.07935.

Lee C. Tucker. You're (not) hired: Artificial intelligence and early career hiring in the Quarterly Workforce Indicators. Working Paper 26-27, U.S. Census Bureau, Center for Economic Studies, 2026. URL https://www.census.gov/library/working-papers/2026/adrm/CES-WP-26-27.html.

U.S. Bureau of Labor Statistics. Occupational separations and openings, projected 2024–34. Employment Projections Program, 2025. URL https://www.bls.gov/emp/tables/occupational-separations-and-openings.htm.

U.S. Bureau of Labor Statistics. Occupational separations and openings, projected 2025–35, table 1.10. Employment Projections Program, 2026a. URL https://www.bls.gov/emp/tables/occupational-separations-and-openings.htm. Released 27 August 2026.

U.S. Bureau of Labor Statistics. Artificial intelligence (AI) exposure categories. Employment Projections Program, 2026b. URL https://www.bls.gov/emp/publications/ai-exposure-categories.htm. Released 27 August 2026.

U.S. Census Bureau. Recent changes in the census industry and occupation classification systems. Technical Paper 78, U.S. Census Bureau, 2020. URL https://www.census.gov/programs-surveys/demo/guidance/industry-occupation/technical-paper-78.html.

U.S. Census Bureau. 2024 survey of income and program participation data. Public-use microdata, 2025. URL https://www.census.gov/programs-surveys/sipp/data/

datasets/2024-data.html. Reference period January–December 2023; accessed 4 October 2026.

U.S. Census Bureau. 2025 survey of income and program participation data. Public-use microdata, 2026. URL https://www.census.gov/programs-surveys/sipp/data/datasets/2025-data.html. Reference period January–December 2024; accessed 4 October 2026.

U.S. Department of Labor, Employment and Training Administration. O*NET 30.3 database. O*NET Resource Center, 2026. URL https://www.onetcenter.org/database.html. May 2026 release.

U.S. Department of Labor, Office of the Assistant Secretary for Policy. Career pathways descriptive and analytical study data. Public-use data release, 2022. URL https://www.dol.gov/agencies/oasp/evaluation/completedstudies/career-pathways-descriptive-and-analytical-study. 8 September 2022.

Keith Waters and Shade T. Shutters. Skills-approximate occupations: Using networks to guide jobs retraining. *Applied Network Science*, 7:43, 2022. doi: 10.1007/s41109-022-00487-7. URL https://doi.org/10.1007/s41109-022-00487-7.

Michelle Yin, Hoa Vu, and Claudia Persico. How (un)stable are LLM occupational exposure scores? Evidence from multi-model replication. Working Paper 35110, National Bureau of Economic Research, 2026. URL https://doi.org/10.3386/w35110.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- AIWorkforceR5BAppliedNetworkScienceAdditionalfile1.pdf